# Best Guesses and Surprises

William Poole

I have a simple message today—that anyone interested in monetary policy should spend less time on economic forecasts and more time on implications of forecast surprises. If you are in the forecasting business, it makes good sense to write at length about the forecast and the analysis behind it. For the rest of us, the forecast provides the baseline for examining the most important policy issues. The true art of good monetary policy is in managing forecast surprises and not in doing the obvious things implied by the baseline forecast.

I'll proceed by outlining the consensus outlook and then will discuss how I view the job of dealing with surprises. I'll emphasize that the key issue is that monetary policy responses to surprises ought not to be random, but as predictable as possible. There are some principles of good responses that make it easier for students of monetary policy to predict what the Federal Reserve will do.

Before digging into the substance of my subject, I want to emphasize that the views I express do not necessarily reflect official positions of the Federal Reserve System. I thank my colleagues at the Federal Reserve Bank of St. Louis for their comments—especially Bob Rasche, senior vice president and Director of Research, who provided special assistance. However, I retain full responsibility for errors.

## CONSENSUS OUTLOOK TODAY

What is the consensus outlook for the U.S. economy today? Numerous forecasts are in the public domain, from government and private sources.[1] Direct comparison of these forecasts is not straightforward, as there are differences among the variables for which the forecasts are presented, differences in the forecasting time horizons, and differences in averaging, with some forecasts presented on a fourth-quarter to fourth-quarter basis and others presented on an annual-average over annual-average basis.

Nevertheless, at present a remarkably uniform picture from a perusal of these various sources emerges for the major economic indicators. Real gross domestic product (GDP) is forecast to grow in the 4 to 4 1/2 percent range from the fourth quarter of 2003 to the fourth quarter of 2004. Inflation as measured by the consumer price index (CPI) is forecast in the 1 1/2 to 2 percent range and as measured by the GDP chain price index is forecast in the 1 to 1 1/2 percent range over that horizon. The unemployment rate is forecast to be around 5 1/2 percent by the fourth quarter of 2004.

My colleagues around the Federal Open Market Committee (FOMC) table are on average slightly more bullish than the above picture: The midpoint of the range of forecasts of real GDP growth included in the *Monetary Policy Report to the Congress* submitted two weeks ago is 4 1/2 percent for the fourth quarter of 2004 over the fourth quarter of 2003. The midpoint of the range of inflation forecasts (measured by the chained price index for personal consumption expenditures) in that report is 1.13 percent, and the midpoint of the forecast for the unemployment rate in the fourth quarter of 2004 is 5.38 percent. I'll refer to this forecast as the "FOMC members' forecast." The forecast reflects a survey of FOMC members,

[1] A partial list includes *Blue Chip Economic Indicators* (published each month); *Survey of Professional Forecasters* (compiled quarterly by the Federal Reserve Bank of Philadelphia); *Wall Street Journal Forecasting Survey* (published early January and July of each year); Congressional Budget Office, *Budget and Economic Outlook* (published each February and August); Council of Economic Advisers economic outlook (published each February in the *Economic Report of the President* and updated each July); and Federal Reserve System, *Monetary Policy Report to the Congress* (published each February and July), containing the economic projections of the Federal Reserve governors and Reserve Bank presidents.

William Poole is the president of the Federal Reserve Bank of St. Louis. This article was adapted from a speech of the same title presented at the Charlotte Economics Club, Charlotte, North Carolina, February 25, 2004. The author thanks colleagues at the Federal Reserve Bank of St. Louis. Robert H. Rasche, senior vice president and Director of Research, provided special assistance. The views expressed do not necessarily reflect official positions of the Federal Reserve System.

Federal Reserve Bank of St. Louis *Review*, January/February 2004, *86*(3), pp. 1-7.

but is not an FOMC forecast per se because the Committee does not debate and vote on the forecast to make it a Committee forecast as such. Nor is it the Board of Governors staff forecast prepared for each FOMC meeting and reproduced in the Greenbook; the Greenbook is released with the FOMC meeting transcript only after a five-year lag.

Those forecasters who risk interest rate forecasts (the *Wall Street Journal Forecasting Survey*, the *Blue Chip Economic Indicators*, and the Congressional Budget Office *Economic Outlook*) expect Treasury bill rates around 1.2 percent, and ten-year Treasury bond rates around 4.6 percent either on an annual average basis or at the middle of 2004. I should note that FOMC members do not make public forecasts of interest rates.

## HOW RELIABLE IS THE CONSENSUS OUTLOOK?

The small dispersion among forecasts today is not unusual and should not be interpreted as a measure of likely forecast accuracy. Over the years, numerous studies have investigated the forecast accuracy of private forecasters. More recently several studies have compared the accuracy of both the FOMC members' forecasts and the Greenbook forecasts with those of private forecasters. One recent analysis was produced by William Gavin and Rachel Mandal in the Research Department at the Federal Reserve Bank of St. Louis. Their paper was published last year in the *International Journal of Forecasting*.[2] Because of the lag in releasing the Greenbook, this study analyzes the Greenbook forecasting record up through 1996. The other comparisons include forecasts through 2001.

The authors compared the Blue Chip forecasts, the Greenbook forecasts, and the FOMC members' forecasts against a naive, same-change forecast beginning in 1980 for both real output growth and inflation. Three different forecasting horizons were examined: six, twelve, and eighteen months.[3] Not surprisingly, the accuracy of the forecasts deteriorates as the forecasting horizon is lengthened. For a one-year-ahead forecast, the root-mean-squared forecast error (a measure of the dispersion of the forecasts around the realized value) for real output growth is on the order of 1.4 percentage points for all the three sets of forecasts considered. The root-mean-squared forecast error for the naive forecast is considerably larger, on the order of 2.2 percentage points.

Clearly, the forecast accuracy of the forecasters is substantially better than that of the naive forecast, but still leaves a lot of room for surprises. To make this point clear in today's context, if for convenience we say that the GDP growth forecast is 4 percent over the four quarters ending 2004:Q4, then one standard error leaves us with a forecast band of 3 to 6 percent growth over this period. If we were to have a 3 percent outcome, everyone would fear that the recovery is faltering; if we were to have a 6 percent outcome, the most likely characterization would be that we have a boom on our hands.

Moreover, keep in mind that one standard deviation on either side of the expected value does not by any means exhaust the range of possible outcomes. As a rough approximation, one time out of three, the one-year-ahead forecast of real output growth will fall outside a range of plus or minus 1.4 percentage points of the stated forecast number. Assuming a symmetrical distribution of forecast errors, which seems reasonable, there is a probability of about 0.16 that real output growth over the next four quarters will exceed 6 percent and a probability of about 0.16 that output growth will fall short of 3 percent.

Clearly the range of error associated with the current state of the forecasting art fails to distinguish between a really strong expansion—a boom—and a faltering recovery. And the accuracy of inflation forecasts is not much better. On a one-year-ahead forecasting horizon, the root-mean-squared error of inflation forecasts is in the range of 0.75 to 0.9 percentage points. This forecasting record is not much better than would have been achieved with a naive forecasting model. As with real GDP, there is a significant probability that the outcome could fall outside the one-standard-deviation band. And it is also true that an inflation outcome outside the band would create considerable concern.

Published forecasts are repeatedly updated on ever shorter and shorter forecasting horizons. The results reported by Gavin and Mandal indicate that as the horizon becomes shorter the uncertainty surrounding the forecast realization is reduced—though, perhaps surprisingly, not by a particularly large amount. Their analysis suggests that for real

---

[2] William T. Gavin and Rachel J. Mandal, "Evaluating FOMC Forecasts," *International Journal of Forecasting*, 2003, *19*(4), pp. 655-67.

[3] Details can be found in Tables 4 and 5 of Gavin and Mandal.

GDP growth the root-mean-squared forecast error on an eighteen-month horizon is between 1.5 and 1.9 percentage points while at a six-month horizon it is reduced to only 1.3 percentage points. For inflation, the root-mean-squared forecast error at the eighteen-month horizon is between 1.1 and 1.3 percentage points but is substantially reduced to around 0.5 percentage points at a six-month horizon.

As we go through the year, the forecast for 2004 will be updated as results for each quarter come in. An example of this process is provided by the monthly *Blue Chip* consensus forecast for 2003. The initial release of the *Blue Chip* forecast for last year was in January 2002; thus, we have a record of 24 successive *Blue Chip* forecasts for 2003. The initial forecast was for a year-over-year growth rate of 3.4 percent. Through the first half of 2002 the consensus forecast was revised up slightly, reaching a peak of 3.6 percent in June. Thereafter the consensus was fairly steadily revised downward over the next year, reaching a trough of 2.3 percent in August 2003. The final forecast, in December 2003, was 3.1 percent, which is the currently published number for real growth in 2003 over 2002. The initial release of the consensus CPI inflation forecast for 2003 over 2002 in January 2002 was 2.4 percent. This forecast changed very little over the following 24 months, increasing to 2.5 percent in mid-2002 and then settling down at 2.3 percent, the CPI inflation rate that was realized for 2003 over 2002.

The relatively low variability of the consensus forecast for 2003 masks the heterogeneity among the individual survey respondents that reflects the inherent uncertainty of economic forecasts. In early 2002, the range of forecasts for real growth in 2003 across the *Blue Chip* respondents was 2.0 to 6.0 percent. By the beginning of 2003 this range had narrowed to about 2.5 to 4.5 percent. Only after the middle of 2003, when data from six of the eight quarters involved in the computation of a year-over-year growth rate were available, did the range of individual forecasts drop below 1 percentage point.

A similar dispersion is observed among the individual forecasts of CPI inflation for 2003. At the beginning of 2002 the range of forecasts was from less than 1.0 percent to almost 4.5 percent. By early 2003 this range had narrowed to less than 1.5 to slightly more than 2.5 percent. It was only after September 2003 that the range of forecasts shrunk to less than 1 percentage point.

## SOME EXAMPLES OF FORECAST SURPRISES

Forecast surprises, or forecast errors, are a standard part of the policy landscape. It is very easy to criticize forecasts and extremely difficult to come up with better forecasts. The fact is that good forecasters produce state-of-the-art forecasts. Policymakers must deal with forecast surprises. What are the sources of those surprises?

The difficulty of forecasting turning points in economic activity is most significant. Whatever creates a recession also creates a forecast surprise. For example, the October 2000 *Blue Chip* consensus for real growth over the five quarters ending 2001:Q4 was for a very steady quarter-to-quarter expansion in real GDP in the range of 3.3 to 3.6 percent at an annual rate. The business cycle dating committee of the National Bureau of Economic Research later dated a cycle peak in March 2001 and a trough in November 2001. Actual quarter-to-quarter real growth during this period ranged from –1.3 to 2.1 percent. Thus, five months before the onset of the 2001 recession, the *Blue Chip* consensus forecast missed the recession completely!

My point here is not to pick on the *Blue Chip* respondents. My colleagues on the FOMC had no greater foresight. In the minutes of the FOMC meeting in October 2000 we can read that, "[l]ooking ahead, they [FOMC members] generally anticipated that the softening in equity prices and the rise in interest rates that had occurred earlier in the year would contribute to keeping growth in demand at a more subdued but still relatively robust pace."[4]

A second noteworthy example is October 2001. In the immediate aftermath of the 9/11 attacks, forecasters turned extremely bearish on the near-term prospects. The *Blue Chip* consensus for real growth in 2001:Q4 in the October 10, 2001, survey was –1.3 percent, with a range of forecasts from –3.2 to 0.8 percent. The *Blue Chip* respondents were particularly pessimistic about prospects for the manufacturing sector; the consensus was for growth of –3.1 percent, with a range from –7.4 to 0.6 percent. We now know that in 2001:Q4 the economy rebounded to a 2.1 annual rate of growth in real GDP, led by an all-time record rate of light vehicle sales. Keep in mind that this GDP growth rate was above the forecast of every single one of the 50 plus *Blue Chip* respondents at the beginning of the quarter.

[4] Minutes of the FOMC Meeting of October 3, 2000, *Federal Reserve Bulletin*, January 2001, p. 23.

A third example is the history of real growth and inflation forecasts in the second half of the 1990s after the now-apparent increase in trend productivity growth. Consider the midpoint of the range of forecasts of real growth and inflation by FOMC members prepared each February for the four quarters ending in the fourth quarter of each year from 1996 through 1999. On average these forecasts underestimated real growth by 2.1 percentage points for these four years. The range of forecast errors was from 2.4 to 1.9 percentage points. The errors were all in the same direction and all of significant size. During the same four years the CPI inflation forecast error averaged 0.0 percent—right on the button. However the forecast errors for the individual years ranged from –1.2 to 0.6 percentage points.

The reasons for forecast errors are many. Some reflect incomplete understanding of how the economy works, such as the errors in projecting productivity growth, or consumer behavior right after the 9/11 terrorist attacks. Some reflect unpredictable shocks, such as a sharp change in energy prices or the 9/11 terrorist attacks themselves. Some reflect financial disturbances, such as the 1987 stock market crash. Whatever may be the reasons for forecast errors, they are a fact of life.

## THE POLICY SIGNIFICANCE OF FORECAST UNCERTAINTY

What are the implications of the documented uncertainty surrounding forecasts of future economic activity? Some dismiss forecasts altogether and view them as irrelevant for policy because their errors are so large. To me, that response is completely wrong. Instead, policy needs to be informed by the best guesses incorporated in forecasts and by knowledge of forecast errors. Forecast errors create risk, and that risk needs to be managed as efficiently as possible. And the surprises that create forecast errors also create the need for policy changes that cannot be anticipated in advance because the surprises cannot be anticipated.

Given the size of forecast errors, we will frequently observe the economy evolving along a substantially different path from that portrayed by consensus forecasts only a short time earlier. With newly available information, forecasters will adjust their prognostications, and policymakers, such as the FOMC, will adjust their view of the appropriate policy stance. If the revised view of the appropriate policy stance is sufficiently changed, policymakers can and should implement the changes in policy settings, such as the intended federal funds rate, that they believe are consistent with the new information.

Such policy actions should be implemented whether or not they will come as surprises to market participants and the general public. Here it is important to be clear about the distinction between a *policy* and a *policy action*. A policy should be viewed as a rule or response regularity that links policy actions, such as adjustments in the intended federal funds rate, to the state of the world. A policy is like a decision to drive 65 miles per hour; given that policy, a policy action is the adjustment of the accelerator to maintain the target speed. If the effects of wind and hills on speed cannot be anticipated, then neither can the policy actions of accelerator adjustments. A good policy requires clarity about policy objectives and as much precision as possible as to how policy actions will respond to new information to best achieve the policy objectives.

In the monetary policy context, anticipated policy actions are naturally tied closely to the forecast. To maintain the policy of achieving low and stable inflation, unanticipated policy actions must often accompany forecast surprises. Should an inflationary shock hit the economy, for example, that shock and an increase in the FOMC's target federal funds rate would both be surprises.

On numerous occasions I have stated my view that the FOMC should communicate its intention about monetary policy as precisely as possible to get markets in "synch" with policy. My view should not be interpreted to imply that the FOMC can only act after it has "prepared" market participants for a change in the intended federal funds rate. There will be times when significant unforeseen economic news will be revealed very suddenly—events that can be appropriately described as "shocks." If, in the judgment of the FOMC, such news calls for policy actions even though market participants could not have been forewarned of such actions, the FOMC would be derelict in its responsibilities if it failed to act. *Given* the shock, the FOMC's action ought not to be a surprise. The real surprise would arise if the FOMC were to do nothing in the face of a shock calling for action.

A couple of examples are illustrative. Consider the Asian debt crisis, the Russian default, and the collapse of Long Term Capital Management (LTCM) in August-September 1998. No one foresaw this combination of events, nor was the impact of these events on the liquidity of major financial markets predictable. At the time of the FOMC meeting on August 18, 1998, federal funds futures for contracts as far out as December 1998 were trading within a couple of basis points of the prevailing 4.50 percent intended funds rate. Nevertheless, by the conclusion of the FOMC meeting on November 17, 1998, the intended funds rate had been reduced in three steps by a total of 75 basis points, including a cut of 25 basis points at an unscheduled conference call meeting on October 15.

These rate cuts in the fall of 1998 were a surprise from the vantage point of early August. But the real surprise would have been if the FOMC had totally ignored the Russian default and collapse of LTCM. Holding the intended funds rate constant given the market turmoil would not have been consistent with the Fed's responsibility to contribute to financial stability.

I could walk through numerous other examples to drive home the point that a key feature of monetary policy is measured and appropriate responses to the constant stream of surprises. In discussions of monetary policy, I would like to see much more emphasis on appropriate policy responses to surprises and potential surprises and much less emphasis on forecasts. An overemphasis on consensus forecasts may lead market participants to a false precision in their views about the federal funds rate going forward. It is much more productive to think through the sorts of things that might happen and the appropriate response to such events. A careful analysis of risks helps to prepare the mind for dealing with surprises when they occur. Market participants and the FOMC should not focus on the predictability of a particular path for the funds rate, but instead on the predictability of the response of the FOMC to new information about the economy.

## PRINCIPLES OF FOMC RESPONSES TO "SHOCKS"

For at least 40 years, economists have been trying to develop a quantitative characterization of FOMC policy actions—a policy reaction function.[5] A review published in 1990 analyzed 42 published examples of attempts at characterizing FOMC behavior.[6] Since 1993, the prevalent framework to quantify FOMC action is the "Taylor rule."[7] None of these efforts have achieved their objective.[8] In my judgment, it is not possible at the current state of knowledge to define a precise reaction function of the FOMC, and perhaps it never will be possible.

It is possible, however, to describe some general principles that guide FOMC behavior and that can be applied by market participants to form expectations about how the FOMC will respond to new and unexpected information. I believe that these principles are fairly widely accepted, but different FOMC members will apply them in different ways at different times. And the principles always involve a degree of judgment.

The first principle is that the FOMC will not respond to "shocks" that are seen as very transitory. Policy should only react to "shocks" that are longer lasting—highly persistent. The reason for this principle is quite straightforward—nothing that the FOMC can do will offset the impact on the economy of a "here today, gone tomorrow" event. While economists continue to debate exactly how "long and variable" the response of the economy is to a policy action, there is a consensus of professional opinion that it takes at least several months before the economy responds. Of course, judgment is always necessary to determine whether any particular shock is likely to be transitory.

Some transitory shocks occur because "news" does not provide accurate information. Many data releases are subject to several revisions, and often the revised data reveal a different picture than that portrayed by the initial release. Quarterly GDP data are revised twice until the "final estimates," and these "final estimates" are subject to annual benchmark revisions and comprehensive revisions at roughly five-year intervals. For a recent example of significant data revisions, initial measurement of the 2001 recession suggested negative real GDP growth in the second and third quarters of 2001, whereas the currently available data measure negative real growth as early as the third quarter of 2000 and for the first three quarters of 2001.

5 An early analysis is William G. Dewald and Harry G. Johnson, "An Objective Analysis of the Objectives of American Monetary Policy, 1952-61," in Deane Carson (ed.), *Banking and Monetary Studies*, Homewood, IL: Richard D. Irwin, 1963, pp. 171-89.

6 Salwa S. Khoury, "The Federal Reserve Reaction Function: A Specification Search," in Thomas Meyer (ed.), *The Political Economy of American Monetary Policy*, New York: Cambridge University Press, 1990, pp. 27-50.

7 John B. Taylor, "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy*, 1993, *39*(0), pp. 195-214.

8 An illustration of the deviations of the predicted from actual funds rate from a Taylor rule with a constant target inflation rate over the past decade can be found in Federal Reserve Bank of St. Louis, *Monetary Trends*, p. 10, available at http://research.stlouisfed.org/publications/mt/.

The monthly payroll employment data are revised in the two months following the initial release and are revised again at the beginning of the following calendar year to incorporate benchmarks to the unemployment insurance system records. With the initial release of the payroll employment numbers for October 2003 at the beginning of last November, the measured monthly increase of 126,000 workers generated the hope that the transition from the two-year "jobless recovery" to a period of rapid employment growth was at hand. The October data as currently revised indicate an increase in payroll employment of only 88,000 workers. We now believe that only in January 2004 did month-to-month payroll employment growth exceed 100,000 jobs—and only an anemic 112,000 at that.

The presence of measurement error in individual economic data series, particularly in the initial releases of such data, requires that analysts and policymakers examine multiple data sources for a consistent picture of the underlying trends in economic activity. There is no way to generalize about this issue. Different series have different sources of error and different frequencies of large revisions. Some series are more subject to special disturbances, such as bad weather, than others. The Federal Reserve has tremendous staff expertise and access to statistical agencies that permit it to form the best judgments possible on these tricky issues.

As an aside, let me offer another observation. Currently, the Federal Reserve enjoys very high credibility. Among other benefits, that credibility enables the Fed to react to its best judgment about what incoming data mean. The Fed does not have to act to maintain appearances. For example, my impression is that there were times in the 1970s when the Fed failed to react to accumulating evidence of economic weakness for fear that to react before inflation declined could be interpreted as a lack of inflation-fighting resolve. Policy actions designed primarily to attempt to affect expectations, even though contrary to the fundamentals, ultimately increase uncertainty as it becomes clear that the action did not fit the fundamentals. Success in bringing down inflation and keeping it down means that the Fed can ignore a surge in price indices or any other troubling information if its best judgment is that the data reflect a statistical aberration or a transitory event.

I've argued that the Federal Reserve must analyze the data for potential statistical problems and that it must do its best to sort out transitory disturbances from longer lasting ones. Another dimension of the problem is that a central part of making such judgments is to collate information from a variety of sources. Employment data provide an excellent example. The establishment and household surveys are quite distinct statistically, as the surveys' coverage and methods do not overlap. Data on initial claims for unemployment insurance supplement the message from the two main surveys. In addition, the Federal Reserve accumulates a wealth of anecdotal information on the labor market from contacts across the country; much of this information appears in the Beige Book. When data from diverse sources point in the same direction, confidence in the direction indicated is increased. Conversely, when the signals are conflicting, there is often good reason to reserve judgment and delay policy action. Analysis of the strength of household demand, business investment demand, inflationary pressures, and all other key elements of the picture can and does proceed the same way.

Once FOMC members have reached a conclusion on where the economy is and where it is heading, there are situations where the decision on the appropriate policy response is straightforward and other cases where the appropriate response is problematic. Consider some examples of easy cases first. Suppose the economy has shown robust growth with low inflation for a period of time, and information accumulates that leads to a reasonable interpretation that both real growth and inflation pressures are increasing, or both decreasing. Faced with such information, central bankers with a dual mandate (such as the FOMC) are likely to respond by raising or lowering, as appropriate, the nominal interest rate target. When credibility is high, moreover, the decision need not be a quick one. But when the issue is clear, the central bank must act vigorously enough to ensure maintenance of a non-inflationary equilibrium.

Indeed, such responses are, qualitatively, exactly those predicted by the Taylor rule. Under these conditions, market participants and private agents can likely accurately anticipate the direction, if not the timing and magnitude, of FOMC actions. The FOMC practice since February 1994 of generally restricting changes in the intended federal funds rate to regularly scheduled meetings and making changes in multiples of 25 basis points has demonstratively improved the predictability of the timing and magnitude of changes in the intended funds rate in such cases.[9]

The appropriate policy response in other cases is less clear. Suppose the economy has shown robust growth with low inflation for a period of time and information arrives that leads to a reasonable interpretation that real growth is decreasing and inflation is increasing. A historical episode of this sort is the "oil shock" in late 1973 and 1974. Here, one component of a "dual mandate" signals a policy action in one direction and the other component in the opposite direction. This is a dilemma case in which the behavior of the economy is pulling the policymakers to be both easier and tighter. A weighting of objectives and careful attention to long-run concerns is necessary. Even if a central bank were to follow a Taylor rule approach to implementing policy changes, in the absence of disclosure of the exact reaction function, outside observers would be unable to predict the policy action. It is unrealistic to believe that a central bank can provide the transparency required for outsiders to accurately predict policy actions in all such circumstances.

Predictability in the dilemma cases can be improved and the appropriate policy response facilitated when a central bank has a credible commitment to maintaining low inflation in the long run. In these circumstances, the central bank can likely pursue short-run stabilization objectives without significant influence on expectations of long-run inflation. In current parlance, inflationary expectations are "well anchored." In such environments, policy actions aimed at short-run stabilization are likely to be more effective. Under conditions of high credibility, policy actions are likely more predictable.

The Federal Reserve has policy responsibilities beyond a narrow interpretation of the "dual mandate." In particular, a fundamental responsibility envisioned by the architects of the Federal Reserve Act was that the new central banking structure avoid recurrence of episodes of financial instability and banking panics such as those that occurred regularly in the late 19th and early 20th century. The Great Depression was a Federal Reserve policy failure of the first order. Recent episodes, such as the 1987 stock market crash, the financial market upset in the fall of 1998, Y2K, and 9/11 provide evidence that the Federal Reserve has learned lessons from the 1930s well and can deal effectively with systemic threats to financial stability.

It is important to understand, however, that concerns about financial stability require that the Federal Reserve sort out shocks that raise such concerns from those that do not. Not every large event creates risks for the financial system as a whole. The large stock market decline that started in early 2001 did tend to depress economic activity but, unlike the crash of 1987, never raised issues of systemic stability.

## CONCLUDING COMMENT

Forecast uncertainty is a fact of life. Forecasts are like newspapers. Just as last week's newspaper is of little value in understanding today's news, last month's forecast is of little value in determining today's policy stance. Old newspapers and old forecasts are primarily of historical interest.

The obvious fact that we insist on using the most up-to-date forecast available indicates that forecasts change, sometimes substantially, with new information. Forecasts are valuable in formulating monetary policy, but it is of critical importance that we not allow today's policy settings to become entrenched in our minds.

---

[9] William Poole, Robert H. Rasche, and Daniel L. Thornton, "Market Anticipations of Monetary Policy Actions," Federal Reserve Bank of St. Louis *Review*, July/August 2002, *84*(4), pp. 65-93; and William Poole and Robert H. Rasche, "The Impact of Changes in FOMC Disclosure Practices on the Transparency of Monetary Policy: Are Markets and the FOMC Better 'Synched'?" Federal Reserve Bank of St. Louis *Review*, January/February 2003, *85*(1), pp. 1-9.

# Monetary Policy Actions, Macroeconomic Data Releases, and Inflation Expectations

**Kevin L. Kliesen and Frank A. Schmid**

Do surprises in macroeconomic data releases and monetary policy actions of the Federal Reserve lead economic agents to update their beliefs about the rate of inflation? If so, which macroeconomic data releases matter for inflation expectations? Does Federal Reserve communication (for example, speeches and testimonies by Fed officials) affect the uncertainty surrounding the rate of inflation that economic agents expect for the following 10 years? These are some of the questions we are trying to answer in this paper.

We gauge inflation expectations by two different concepts of inflation compensation, both of which are derived from the market valuation of the expected cash flows of nominal Treasury securities versus inflation-indexed Treasury securities (TIIS). We look at 35 macroeconomic data series (for example, the monthly change in nonfarm payroll employment) and determine whether daily changes in inflation compensation are associated with the surprise component in these series. We define the surprise component as the difference between the expected and the actually released value of the series, normalized by the degree of uncertainty surrounding these expectations. The time period of our analysis is January 31, 1997, to June 30, 2003, for one concept of inflation compensation and January 4, 1999, to June 30, 2003, for the other. For the 35 macroeconomic data series, we find the following: for 17, the surprise component in the announcement has no bearing on inflation expectations; for 5, the surprise component has an effect on one measure of inflation expectations, but not on the other; for 13, the surprise component in the release consistently affects inflation expectations, independent of the employed concept of inflation compensation. Further, we show that monetary policy actions that are tighter or easier than expected by the federal funds futures market have a statistically significant effect on the expected rate of inflation, independent of the employed concept of inflation compensation. Finally, we provide evidence that Federal Reserve communication and surprises in monetary policy actions bear on the uncertainty surrounding the expected rate of inflation. For one concept of inflation expectations, we find that Federal Reserve communication reduces uncertainty about the future rate of inflation, while surprises in monetary policy actions increase this uncertainty. For the other concept of inflation compensation, we find no such effects.

## RELATED LITERATURE

The effect of macroeconomic announcements on inflation compensation embedded in the market valuation of expected cash flows of nominal Treasury securities versus TIIS has been investigated before.

The studies most closely related to our work are Sack (2000) and Gürkaynak, Sack, and Swanson (2003). Using daily data, Sack (2000) analyzes how surprises in the releases of six monthly macroeconomic data series affect the inflation compensation embedded in Treasury securities for the period 1997-99. These macroeconomic data series are the consumer price index (CPI), the CPI excluding food and energy (core CPI), the producer price index for finished goods (PPI), nonfarm payroll employment, retail sales, and the NAPM index.[1] Sack matches the on-the-run 10-year TIIS with a portfolio of nominal Treasury STRIPS (Separate Trading of Interest and Principal Securities) that replicates the pattern of expected payments of the TIIS.[2] The author finds a

[1] The National Association for Purchasing Managers' (NAPM) index is now simply the Purchasing Managers Index; it is released by the Institute for Supply Management.

[2] The Treasury STRIPS program, which was introduced in January 1985, "lets investors hold and trade the individual interest and principal components of eligible Treasury notes and bonds as separate securities" (www.publicdebt.treas.gov/of/ofstrips.htm).

Kevin L. Kliesen is an economist and Frank A. Schmid is a senior economist at the Federal Reserve Bank of St. Louis. Jason Higbee and Thomas Pollmann provided research assistance.

Federal Reserve Bank of St. Louis *Review*, May/June 2004, *86*(3), pp. 9-21.

statistically significant response of inflation expectations to surprises in the CPI, the core CPI, retail sales, and the NAPM index.

Gürkaynak, Sack, and Swanson (2003) analyze the response of inflation compensation embedded in one-year nominal versus inflation-adjusted *forward* interest rates to surprises in macroeconomic data releases and Federal Reserve monetary policy actions—that is, changes to the Federal Open Market Committee's (FOMC) federal funds target rate. The forward rates are derived from the yields of 10-year Treasury securities—the TIIS and nominal securities. The pair of one-year forward rates that the authors study is the one for the 12-month time window between the maturity dates of the on-the-run 10-year TIIS and the (off-the-run) TIIS issued 12 months earlier. Prior to July 2002, and starting in 1997, 10-year TIIS were issued only once per year, in January. This implies that the authors analyze changes to the inflation rate that is expected to prevail during a 12-month time window that starts, on average, 8.5 years from the time of the data release. The sample period runs from January 1997 through July 2002 and covers 39 macroeconomic data series. The authors show that the rate of inflation expected to prevail in about nine years' time correlates positively with surprises in consumer confidence, consumer credit, the employment cost index (ECI), gross domestic product (GDP, advance), new home sales, and retail sales and negatively with surprises in the federal funds target rate. The authors conclude (on p. 2) that their "empirical findings suggest that private agents adjust their expectations of the long-run inflation rate in response to macroeconomic and monetary policy surprises."

Like Gürkaynak, Sack, and Swanson (2003), we analyze Treasury securities with a 10-year maturity. Unlike those authors, we study the average inflation rate expected for the next 10 years rather than the value expected for a 12-month window late in this time period. In other words, we make no statement about whether (and, if so, the degree to which) surprises in monetary policy actions and macroeconomic data releases affect the rate of inflation that economic agents expect to prevail in 8 to 9 years.

Another study related to ours is by Kohn and Sack (2003), who study the effect of Federal Reserve communication on financial variables but make no attempt to gauge the influence that Chairman Greenspan's speeches and testimonies have on the *level* of Treasury yields. Rather, the authors measure the effect of Fed communication on Treasury yield *volatility*. These authors investigate the effect that Federal Reserve communication has on various financial variables by using daily observations for the period January 3, 1989, through April 7, 2003. Federal Reserve communication comprises statements released by the FOMC and, since June 1996, congressional testimonies and speeches delivered by Chairman Greenspan. Among the financial variables Kohn and Sack analyze are the yields (to maturity) of the 2-year and 10-year Treasury notes. These authors find that statements of the FOMC and testimonies of Chairman Greenspan have a statistically significant impact on the variance of 2-year and 10-year Treasury note yields; no such effect was found for Chairman Greenspan's speeches. We follow Kohn and Sack (2003) and study the effect of Federal Reserve communication on the (conditional) volatility of inflation compensation—that is, on the uncertainty that surrounds the future rate of inflation.

## MEASURES OF INFLATION COMPENSATION

In 1997, the U.S. Treasury introduced TIIS. These securities are issued alongside traditional (nominal) Treasury securities. Both types of securities are endowed with a fixed coupon yield—that is, the coupon payment per annum as a percent of the principal. Unlike the principal of the nominal Treasury, which is fixed for the lifetime of the security, the principal of the TIIS is adjusted daily to past changes in the rate of inflation, as measured by the (not seasonally adjusted) CPI for all urban consumers. The coupon payments of TIIS are made off of the inflation-adjusted principal. If the rate of inflation turns negative, the principal is adjusted downward, possibly dropping below the par amount at issue. At maturity, TIIS are redeemed at their inflation-adjusted principal or par value at issue, whichever is greater.

It has become common practice to gauge inflation expectations—that is, expectations about the future, *average* rate of change in the CPI—from the inflation rate at which the market prices of comparable TIIS and nominal Treasury securities break even. To illustrate the theoretical motivation of this concept, consider, as an example, two default-free securities with annual coupon payments, identical coupon yield, an original principal of $1, and a time to maturity of one year. One is a nominal Treasury security and the other is a TIIS. For the two securities to deliver the same return to an investor who is indifferent to inflation risk (but not to inflation, of course), the following equation must hold:

$$\frac{1+c}{p^n} = \frac{(1+\pi)\cdot(1+c)}{p^i}, \tag{1}$$

where $c$ is the coupon yield, $\pi$ is the expected rate of inflation, and $p^n$ and $p^i$ are the prices of the nominal security and the TIIS, respectively. Solving equation (1) for the expected rate of inflation delivers

$$\hat{\pi} = \frac{\frac{1+c}{p^n}}{\frac{1+c}{p^i}} - 1 \equiv \frac{y^n - y^i}{1+y^i} \approx y^n - y^i, \tag{2}$$

where $y^n$ and $y^i$ are the yields to maturity of the nominal security and the TIIS, respectively. Equation (2) states that the break-even inflation rate, $\hat{\pi}$—that is, the expected rate of inflation at which the two securities trade at the same price—can be approximated by the difference in the yields to maturity between a nominal and an inflation-indexed security. For securities of more than one year to maturity, matters are more complex but the same principle applies.

Although the break-even rate of inflation makes up the bulk of the inflation compensation embedded in the market valuation of the expected cash flows of nominal securities versus TIIS, there is also compensation for inflation risk. Hence, the embedded inflation compensation exceeds the expected rate of inflation. Further, for positive rates of inflation, the payment stream on TIIS is back-loaded compared with the cash flow of nominal Treasury securities. This is because the TIIS principal and, hence, the coupon payments grow with the price level. Because the payment stream of TIIS is back-loaded, their duration with respect to the real (that is, inflation-adjusted) term structure of interest rates is longer than the duration of nominal Treasury securities. In other words, the two types of securities do not have the same price sensitivity to real interest rates. Hence, the real-interest-rate risk (and thus the amount of compensation for this type of risk) might differ between the two securities; this may distort the measured inflation compensation. For technical details on the differences between the two securities, see Emmons (2000) and Sack (2000).

A simple and readily available concept of inflation compensation is the difference in yields to maturity between the on-the-run nominal Treasury security and the on-the-run TIIS of the same *original* time to maturity. This concept of gauging inflation expectations has four major disadvantages. First, and most importantly, the market for nominal Treasury securities is more liquid than the market for TIIS, which may cause the yield spread to understate the expected rate of inflation by a liquidity premium (Sack, 2000). Second, nominal Treasury securities and TIIS might not have the same duration with respect to real interest rates, which might cause the compensation for real-interest-rate risk in the yields of the two securities to differ. Third, the two types of securities might not be issued at the same dates and with the same frequency. Hence, the two on-the-run securities might not have the same remaining time to maturity. For instance, whereas nominal 10-year Treasuries are issued several times per year, the corresponding inflation-indexed securities are issued only once (1997-2001) or twice (2002) per year. Fourth, the remaining time to maturity of on-the-run securities varies because new securities are not issued every trading day.

To avoid some of the drawbacks that come with the simple difference in yields between the on-the-run nominal and inflation-indexed securities, we derive inflation compensation from a smoothed zero-coupon yield curve estimated from off-the-run nominal Treasury coupon securities. We use two different measures for the inflation-indexed yield. The first measure uses the on-the-run TIIS—the OTR measure. The other measure uses a smoothed zero-coupon yield curve estimated from TIIS, which allows us to compare the nominal and inflation-indexed yields at constant maturity—the CM measure.

## THE DATA

Our analysis covers the period from January 31, 1997, to June 30, 2003, for the OTR measure, and from January 5, 1999, to June 30, 2003, for the CM measure. Each period begins on the starting date of the respective daily data series, provided by the Board of Governors of the Federal Reserve System. Both series contain missing values (OTR: 53; CM: 16)—that is, have trading days for which no OTR or CM observation is on record. The macroeconomic data releases are from Money Market Services. The dataset comprises median polled forecast values for 38 macroeconomic data series, along with the sample standard deviations of these forecast values. The Money Market Services survey is conducted every Friday morning among senior economists and bond traders with major commercial banks, brokerage houses, and some consulting firms mostly in the greater New York, Chicago, and San Francisco areas. Among these 38 series in the survey, there are three items—CPI, PPI, and retail sales—for which there also exists a "core" concept. Whereas the comprehensive items of the CPI and the PPI include food and energy items, the respective core measures do not. For retail sales, the narrowly defined concept excludes automotive sales. In the regression analysis, we exclude the three "core" concepts, which leaves us with 35 macroeconomic variables.[3] Data that were released on days when the markets were closed were moved to the next trading day—the day on which this information could be priced in the marketplace.

We try to relate daily changes in inflation compensation to the surprise component in macroeconomic data releases. We define the surprise component as the difference between the actual and the median forecast values, normalized by the sample standard deviation of the individual forecasts.[4] We also control for Federal Reserve communication and actions. Our concept of Federal Reserve communication comprises (i) Chairman Greenspan's semi-annual testimony to Congress (formerly known as Humphrey-Hawkins Testimony) and (ii) speeches and other testimonies of Chairman Greenspan. Consistent with the macroeconomic data releases, we moved Federal Reserve communication to the next trading day if this communication occurred after-hours (that is, after the data input for the inflation compensation measures were recorded) or on days on which there was no trading. Finally, we control for the surprise component in changes (or the absence thereof) of the federal funds target rate, which we measure as suggested by Kuttner (2001) and discussed by Watson (2002). For each scheduled and unscheduled FOMC meeting, we scaled up by $30/(k+1)$ the change of the price of the federal funds futures contract for the current month on the day of the FOMC meeting, $t$, where $t+k$ denotes the last calendar day of the month.[5] Note that this variable is not on the same scale as the surprise component in the macroeconomic data releases. In a sensitivity analysis, we use an alternative measure of the surprise component in monetary policy actions. This alternative measure, devised by Poole and Rasche (2000), rests on price changes of federal funds futures contracts also.[6]

Table 1 shows the frequency with which releases of the 38 macroeconomic data series match recorded OTR and CM inflation compensation during the two respective sample periods. The numbers of data releases during the sample period are in parentheses, and the differences in the two numbers are due to missing values. We also report matches for scheduled and unscheduled FOMC meetings—the federal funds target variable, the surprise component of which was calculated as outlined above—and the two Federal Reserve communication variables defined above—(i) semi-annual testimony to Congress and (ii) Greenspan speeches, testimonies other than semi-annual testimony to Congress. The only weekly series in the dataset, initial jobless claims, has the highest frequency. The next-to-highest frequency is observed for testimonies other than semi-annual testimony to Congress, followed by monthly data releases, FOMC actions (federal funds target), quarterly data releases, and the Chairman's semi-annual testimonies to Congress. An exception is nonfarm productivity, which entered the Money Market Services dataset during the sample period; the first surveyed number refers to the first quarter of 1999.

Table 2, center column, offers a frequency distribution for the coincidence of surprises in macroeconomic data releases (Money Market Services survey) and monetary policy actions. For instance, for the OTR measure of inflation expectations, 453 of the 1,555 trading days analyzed had no surprises in data releases or monetary actions, possibly because no data were released or no action taken; 615 trading days (40 percent) had more than one surprise; and 270 (17 percent) had more than two surprises. Table 2, right column, offers a frequency distribution with Federal Reserve communication included.

## EMPIRICAL FINDINGS

Our empirical approach rests on the following regression equation:

---

[3] As will be discussed, our exclusion of the "core" variables did not materially change our empirical results.

[4] Fleming and Remolona (1997, 1999) calculate the surprise component by normalizing the difference between the actual and the forecast values by the mean absolute difference observed for the respective variable during the sample period. Balduzzi, Elton, and Green (2001) normalize the difference between the actual and the forecast values by the standard deviation of this difference during the sample period. Gürkaynak, Sack, and Swanson (2002) do not normalize their variables.

[5] Following Gürkaynak, Sack, and Swanson (2003), we use the (unscaled) change in the price of the federal funds futures contract due to expire in the following month if the FOMC meeting took place within the last seven calendar days of the month.

[6] For a discussion of measures of market expectations concerning monetary policy actions, see Gürkaynak, Sack, and Swanson (2002).

## Table 1

### Number of Data Releases That Match Inflation Compensation Observations

| Data series (FOMC communication and actions) | OTR measure | CM measure |
|---|---|---|
| Auto sales | 75 (78) | 51 (54) |
| Business inventories | 76 (78) | 54 (54) |
| Capacity utilization | 77 (78) | 54 (54) |
| Civilian unemployment rate | 75 (78) | 54 (54) |
| Construction spending | 75 (78) | 51 (54) |
| Consumer confidence | 76 (78) | 54 (54) |
| Consumer credit | 77 (78) | 54 (54) |
| Consumer price index (CPI-U) | 75 (78) | 54 (54) |
| CPI excluding food and energy | 75 (78) | 54 (54) |
| Durable goods orders | 75 (78) | 54 (54) |
| Employment cost index (Q) | 25 (26) | 18 (18) |
| Existing home sales | 76 (78) | 54 (54) |
| Factory orders | 75 (78) | 53 (54) |
| Federal funds target: unscheduled FOMC meetings* | 4 (4) | 2 (3) |
| Federal funds target: scheduled FOMC meetings* | 51 (52) | 36 (36) |
| GDP price index (advance) (Q) | 25 (26) | 18 (18) |
| GDP price index (final) (Q) | 25 (26) | 18 (18) |
| GDP price index (preliminary) (Q) | 25 (26) | 18 (18) |
| Goods and services trade balance (surplus) | 76 (78) | 54 (54) |
| Greenspan speeches, testimonies other than semi-annual testimony to Congress* | 130 (135) | 89 (90) |
| Hourly earnings | 75 (78) | 54 (54) |
| Housing starts | 76 (78) | 53 (54) |
| Industrial production | 77 (78) | 54 (54) |
| Initial jobless claims (W) | 323 (334) | 224 (234) |
| Leading indicators | 75 (78) | 53 (54) |
| Purchasing Managers' Index | 75 (78) | 51 (54) |
| New home sales | 75 (78) | 53 (54) |
| Nonfarm payrolls | 75 (78) | 54 (54) |
| Nonfarm productivity (preliminary) | 16 (17) | 17 (17) |
| Nonfarm Productivity (revised) | 17 (17) | 17 (17) |
| Personal consumption expenditures | 75 (78) | 53 (54) |
| Personal income | 75 (78) | 53 (54) |
| Producer price index (PPI) | 76 (78) | 54 (54) |
| PPI excluding food and energy | 76 (78) | 54 (54) |
| Real GDP (advance) (Q) | 25 (26) | 18 (18) |
| Real GDP (final) (Q) | 25 (26) | 18 (18) |
| Real GDP (preliminary) (Q) | 25 (26) | 18 (18) |
| Retail sales | 76 (78) | 54 (54) |
| Retail sales excluding autos | 76 (78) | 54 (54) |
| Semi-annual testimony to Congress* | 12 (13) | 8 (9) |
| Treasury budget (surplus) | 74 (78) | 53 (54) |
| Truck sales | 75 (78) | 51 (54) |

NOTE: Monthly series if not indicated otherwise (Q, quarterly; W, weekly). Numbers in parentheses indicate total number of observation, not all of which are used because of missing observations for the measures of inflation compensation.

*Variables not included in the dataset of macroeconomic data releases.

**Table 2**

**Frequency Distribution of Coincidence in Surprises**

| Number of surprises per trading day | Number of observations: MMS survey and federal funds target — OTR | MMS survey and federal funds target — CM | MMS survey, federal funds target, and Federal Reserve communication — OTR | MMS survey, federal funds target, and Federal Reserve communication — CM |
|---|---|---|---|---|
| 0 | 453 | 323 | 415 | 298 |
| 1 | 487 | 335 | 488 | 329 |
| 2 | 345 | 249 | 355 | 265 |
| 3 | 147 | 112 | 159 | 113 |
| 4 | 86 | 68 | 96 | 79 |
| 5 | 21 | 10 | 25 | 13 |
| 6 | 10 | 9 | 10 | 8 |
| 7 | 3 | 2 | 3 | 3 |
| 8 | 1 | 1 | 2 | 1 |
| 9 | 2 | 1 | 2 | 1 |
| Total | 1,555 | 1,110 | 1,555 | 1,110 |

$$\hat{\pi}_t - \hat{\pi}_{t-1} = \alpha + \beta \cdot D + \sum_{k=1}^{35} \delta_k \cdot x_t^k + \gamma \cdot ff_t + \varepsilon_t \ , \quad (3)$$

where $\hat{\pi}_t - \hat{\pi}_{t-1}$ is change in the inflation compensation from trading day $t-1$ to trading day $t$, $D$ is an indicator variable that is equal to 1 if all explanatory variables are equal to 0 (and is equal to 0 otherwise), $x_t^k$ is the surprise component in the macroeconomic data release, $ff_t$ is the surprise component in the Federal Reserve action (the federal funds target variable), and $\varepsilon_t$ is an identically and independently distributed error term with mean 0 and finite variance $\sigma^2$.[7] The dependent variable is, alternately, the OTR and the CM measures of inflation compensation.[8]

We expect signs on the surprises in the macroeconomic data releases to be consistent with the conventional macroeconomic theory as taught in the classroom. Bernanke (2003) presents a brief summary of this theory using the expectations-augmented Phillips curve. In this model, inflation depends on inflation expectations, the output gap, and supply shocks. The most important determinant of inflation is the expectation component. Anything that causes people to expect more inflation is likely to lead to behavior that causes higher inflation. Forecasting models and studies using surveys of inflation expectations show that the most important variable in predicting future inflation is past inflation.[9] Therefore, we expect surprise increases in price indices (such as the CPI or PPI) to lead to higher inflation expectations.

We assume that real variables affect inflation indirectly through the output gap, that potential output is relatively stable (as measured by the Congressional Budget Office, CBO), and that most news about real activity is news about aggregate demand. Therefore a surprise increase in any subset of real economic activity may lead to expectations of higher inflation. The Phillips curve model also performs better if we account for supply shocks. A surprise in real activity may be associated with a supply shock rather than an aggregate demand shock. The real variable included in this study that

[7] The intercept indicator variable, $D$, eliminates the influence on the observed mean of the dependent variable of those observations for which none of the explanatory variables contains information pertinent to the measured inflation compensation.

[8] Using kernel estimation, we verified that the dependent variables are close to normally distributed. Yet, there is mild (and statistically significant) skewness (OTR: 0.402; CM: 0.409) and excess kurtosis (OTR: 1.673; CM: 1.114).

[9] See Stock and Watson (1999) for evidence about the role of past inflation and real variables in predicting inflation.

is explicitly thought to measure changes in aggregate supply is labor productivity.

The expected sign of the federal funds target variable is negative. Remember that this variable measures the surprise in the choice of the federal funds target rate at scheduled and unscheduled FOMC meetings. We predict economic agents to revise down (up) their expected rate of inflation in response to FOMC actions that are indicative of monetary policy tighter (easier) than expected.

Tables 3 and 4 show the regression results for the OTR and CM measures, respectively.[10] The regression coefficients of the federal funds target variable have the expected sign and are statistically significant. A surprise in the expected federal funds target rate of 10 basis points reduces the expected rate of inflation by 2.2 basis points (OTR measure) and 1 basis point (CM measure), respectively. The macroeconomic data series that prove statistically significant are ranked in Table 5 by the magnitude of their influence on each measure's (OTR and CM) inflation expectations. Remember that all macroeconomic data items are on the same scale (scaled by their standard deviation). Data releases whose surprise component proved statistically significant in the Gürkaynak, Sack, and Swanson (2003, Table 3) study, are noted with an asterisk. All statistically significant regression coefficients have the expected sign, except consumer credit and housing starts.[11] Greater than expected numbers for consumer credit and housing starts numbers might be thought of as "bullish"; yet, the regression coefficient is negative. The six *common* variables whose surprise components bear most heavily on both measures of inflation expectations are the employment cost index, the Purchasing Managers' Index, CPI, retail sales, factory orders, and personal income. The capacity utilization rate was the most significant mover of inflation expectations using the CM measure, while the ECI mattered the most in the OTR measure.

We repeat the regression analysis (3) for the core measures of CPI, PPI, and retail sales. That is, we replace the respective comprehensive (total) numbers in model (3) by CPI excluding food and energy, PPI excluding food and energy, and retail sales excluding automotive. These regression results, which are not shown, have less explanatory power than the regression equations with the comprehensive numbers, as judged by the $R^2$. Further, while CPI retains its statistical significance, retail sales do not; PPI remains statistically significant only for the OTR measure. The results for the federal funds target variable are nearly unchanged.

Poole, Rasche, and Thornton (2002) argue that monetary policy surprises as gauged by the change in federal funds futures prices are measured with error. This is because federal funds futures prices not only change in response to monetary policy actions, but also respond to other information pertinent to the future path of the federal funds rate. Because of the measurement error introduced by such ambient price changes of federal funds futures contracts, the regression coefficient of the federal funds target variable is biased toward 0. We deal with the error-in-variables problem by employing an instrumental-variables approach. As an instrument for the federal funds target, we use an indicator that is equal to 1 if the federal funds target exceeds its median positive value, equal to –1 if it falls short of its median negative value, and 0 otherwise.[12]

Table 6 shows the regression results of the instrumental-variables approach applied to equation (3). We use two different definitions of the surprise component of monetary policy actions (the federal funds target variable). First, we provide results for the concept that we have used throughout the paper—the measure suggested by Gürkaynak, Sack, and Swanson (2003), which is denoted federal funds target (GSS) in the table. Second, we present results for the surprise measure devised by Poole and Rasche (2000), which is denoted federal funds target (PR) in the table. Unlike the GSS measure, which rests on the scaled price change of the current month's federal funds futures contract (unless the monetary policy surprise happens within the last seven days of the month), the PR measure always uses the price change of the next month's federal funds futures contract. One of the regression coefficients for the federal funds target (GSS) is indeed larger (in absolute value) than without the error-in-variables correction (Table 4, CM measure), while the other is smaller (Table 3, OTR measure); also, both federal funds

[10] The $t$-statistics and the $F$-statistics are based on Newey-West (1987) corrected standard errors. We applied a standard lag length of floor($4(T/100)^{2/9}$), where floor(.) indicates rounding down to the nearest integer and $T$ is the number of observations. This lag length is also used for the Ljung-Box statistic shown in the tables.

[11] Note that deficits in the Treasury budget and the trade balance are recorded as negative numbers. In other words, a positive surprise component in these variables indicates a deficit that is smaller than expected in absolute value or a surplus that is greater than expected.

[12] For details on this error-in-variables approach, see Greene (2002).

## Table 3

### Regression Results for the 10-year OTR Measure

| Explanatory variable | Coefficient | *t*-statistic |
|---|---|---|
| Auto sales | $-3.250 \cdot 10^{-5}$ | –0.011 |
| Business inventories | $2.928 \cdot 10^{-3}$ | 0.808 |
| Capacity utilization | $3.395 \cdot 10^{-3}$ | 0.514 |
| Civilian unemployment rate | $-2.759 \cdot 10^{-3}$ | –1.612 |
| Construction spending | $-1.604 \cdot 10^{-3}$ | –0.892 |
| Consumer confidence | $3.643 \cdot 10^{-3}$ | 2.237** |
| Consumer credit | $-2.157 \cdot 10^{-3}$ | –1.658* |
| Consumer price index (CPI-U) | $1.016 \cdot 10^{-2}$ | 2.109** |
| Durable goods orders | $1.760 \cdot 10^{-3}$ | 0.800 |
| Employment cost index | $1.223 \cdot 10^{-2}$ | 2.673*** |
| Existing home sales | $-5.957 \cdot 10^{-7}$ | –0.497 |
| Factory orders | $8.538 \cdot 10^{-3}$ | 3.210*** |
| Federal funds target | $-2.194 \cdot 10^{-1}$ | –2.424** |
| GDP price index (advance) | $8.201 \cdot 10^{-3}$ | 1.602 |
| GDP price index (final) | $3.793 \cdot 10^{-3}$ | 1.900* |
| GDP price index (preliminary) | $1.617 \cdot 10^{-3}$ | 0.816 |
| Goods and services trade balance (surplus) | $-4.242 \cdot 10^{-3}$ | –2.764*** |
| Hourly earnings | $4.269 \cdot 10^{-3}$ | 1.773* |
| Housing starts | $-3.087 \cdot 10^{-3}$ | –1.806* |
| Industrial production | $7.385 \cdot 10^{-3}$ | 1.297 |
| Initial jobless claims | $-1.840 \cdot 10^{-3}$ | –1.996** |
| Leading indicators | $2.114 \cdot 10^{-3}$ | 0.526 |
| Purchasing Managers' Index | $1.045 \cdot 10^{-2}$ | 4.757*** |
| New home sales | $4.375 \cdot 10^{-3}$ | 2.310** |
| Nonfarm payrolls | $4.099 \cdot 10^{-3}$ | 2.673*** |
| Nonfarm productivity (preliminary) | $-2.588 \cdot 10^{-3}$ | –0.545 |
| Nonfarm productivity (revised) | $-2.398 \cdot 10^{-3}$ | –0.334 |
| Personal consumption expenditures | $-6.524 \cdot 10^{-4}$ | –0.211 |
| Personal income | $5.176 \cdot 10^{-3}$ | 1.730* |
| Producer price index (PPI) | $-1.634 \cdot 10^{-3}$ | –0.954 |
| Real GDP (advance) | $1.010 \cdot 10^{-3}$ | 0.185 |
| Real GDP (preliminary) | $-2.305 \cdot 10^{-3}$ | –0.641 |
| Real GDP (final) | $-7.614 \cdot 10^{-3}$ | –1.042 |
| Retail sales | $1.015 \cdot 10^{-2}$ | 2.748*** |
| Treasury budget (surplus) | $-1.121 \cdot 10^{-2}$ | –2.117** |
| Truck sales | $6.420 \cdot 10^{-4}$ | 0.247 |
| Intercept indicator variable (*D*) | $-1.434 \cdot 10^{-3}$ | –0.587 |
| Intercept | $-6.330 \cdot 10^{-4}$ | –0.436 |
| *F*-statistic (1) | 3.621*** | |
| *F*-statistic (2) | 3.790*** | |
| $R^2$ | 0.081 | |
| $R^2$ adj. | 0.059 | |
| Ljung-Box statistic | 18.02** | |
| Rao's score test | 13.16*** | |
| Number of nonzero observations | 1,100 | |
| Number of observations | 1,555 | |

NOTE: ***/**/* Indicate significance at the 1/5/10 percent levels (*t*-tests are two-tailed). *F*-statistics and *t*-statistics are Newey and West (1987) corrected. Federal funds target is not included in the MMS survey. *F*-statistic (1): all MMS survey variables and federal funds target; *F*-statistic (2): all MMS survey variables. The number of nonzero observations indicates the number of trading days where Federal Reserve communication or the surprise in a monetary policy action was priced.

## Table 4

### Regression Results for the 10-year CM Measure

| Explanatory variable | Coefficient | *t*-statistic |
|---|---|---|
| Auto sales | $-4.363 \cdot 10^{-4}$ | –0.152 |
| Business inventories | $2.325 \cdot 10^{-3}$ | 0.528 |
| Capacity utilization | $1.241 \cdot 10^{-2}$ | 2.082** |
| Civilian unemployment rate | $-1.669 \cdot 10^{-3}$ | –0.734 |
| Construction spending | $-2.740 \cdot 10^{-3}$ | –1.441 |
| Consumer confidence | $3.046 \cdot 10^{-3}$ | 2.012** |
| Consumer credit | $-2.680 \cdot 10^{-3}$ | –1.896* |
| Consumer price index (CPI-U) | $7.698 \cdot 10^{-3}$ | 1.723* |
| Durable goods orders | $2.075 \cdot 10^{-3}$ | 0.743 |
| Employment cost index | $9.162 \cdot 10^{-3}$ | 2.038** |
| Existing home sales | $8.273 \cdot 10^{-4}$ | 0.591 |
| Factory orders | $6.343 \cdot 10^{-3}$ | 2.008** |
| Federal funds target | $-9.907 \cdot 10^{-2}$ | –3.071*** |
| GDP price index (advance) | $1.466 \cdot 10^{-2}$ | 1.193 |
| GDP price index (final) | $4.147 \cdot 10^{-3}$ | 2.350** |
| GDP price index (preliminary) | $3.469 \cdot 10^{-3}$ | 2.049** |
| Goods and services trade balance (surplus) | $-2.121 \cdot 10^{-3}$ | –1.534 |
| Hourly earnings | $5.752 \cdot 10^{-3}$ | 1.998* |
| Housing starts | $-4.455 \cdot 10^{-3}$ | –2.149** |
| Industrial production | $2.740 \cdot 10^{-3}$ | 0.681 |
| Initial jobless claims | $-2.160 \cdot 10^{-3}$ | –2.048** |
| Leading indicators | $7.567 \cdot 10^{-4}$ | 0.210 |
| Purchasing Managers' Index | $1.086 \cdot 10^{-2}$ | 5.221*** |
| New home sales | $2.574 \cdot 10^{-3}$ | 1.505 |
| Nonfarm payrolls | $3.921 \cdot 10^{-3}$ | 1.818* |
| Nonfarm productivity (preliminary) | $-6.456 \cdot 10^{-3}$ | –1.073 |
| Nonfarm productivity (revised) | $-1.675 \cdot 10^{-3}$ | –0.246 |
| Personal consumption expenditures | $-1.184 \cdot 10^{-5}$ | –0.025 |
| Personal income | $6.184 \cdot 10^{-3}$ | 1.696* |
| Producer price index (PPI) | $-1.845 \cdot 10^{-3}$ | –0.878 |
| Real GDP (advance) | $9.222 \cdot 10^{-3}$ | 1.500 |
| Real GDP (preliminary) | $-2.707 \cdot 10^{-3}$ | –0.668 |
| Real GDP (final) | $-4.609 \cdot 10^{-3}$ | –0.700 |
| Retail sales | $8.745 \cdot 10^{-3}$ | 2.127** |
| Treasury budget (surplus) | $-6.082 \cdot 10^{-3}$ | –1.208 |
| Truck sales | $-4.671 \cdot 10^{-4}$ | –0.153 |
| Intercept indicator variable (*D*) | $-2.896 \cdot 10^{-4}$ | –1.070 |
| Intercept | $1.676 \cdot 10^{-3}$ | 1.025 |
| *F*-statistic (1) | 2.749*** | |
| *F*-statistic (2) | 2.867*** | |
| $R^2$ | 0.087 | |
| $R^2$ adj. | 0.055 | |
| Ljung-Box statistic | 14.47** | |
| Rao's score test | 1.352 | |
| Number of nonzero observations | 785 | |
| Number of observations | 1,110 | |

**NOTE: ***/**/* Indicate significance at the 1/5/10 percent levels (*t*-tests are two-tailed). *F*-statistics and *t*-statistics are Newey and West (1987) corrected. Federal funds target is not included in the MMS survey. *F*-statistic (1): all MMS survey variables and federal funds target; *F*-statistic (2): all MMS survey variables. The number of nonzero observations indicates the number of trading days where Federal Reserve communication or the surprise in a monetary policy action was priced.**

## Table 5

**Ranking of Macroeconomic Data Releases by Impact on Inflation Expectations**

| Data release | OTR sign (+/–) | CM sign (+/–) | OTR rank | CM rank |
|---|---|---|---|---|
| Employment cost index* | + | + | 1 | 4 |
| Capacity utilization | 0 | + | N/A | 1 |
| Treasury budget (surplus) | – | 0 | 2 | N/A |
| Purchasing Managers' Index | + | + | 3 | 2 |
| Consumer price index (CPI-U) | + | + | 4 | 5 |
| Retail sales* | + | + | 5 | 3 |
| Factory orders | + | + | 6 | 6 |
| Personal income | + | + | 7 | 7 |
| New home sales* | – | 0 | 8 | N/A |
| Hourly earnings | + | 0 | 9 | N/A |
| Goods and services trade balance (surplus) | + | + | 10 | 8 |
| Nonfarm payrolls | + | + | 11 | 11 |
| GDP price index (final) | + | + | 12 | 10 |
| Consumer confidence* | + | + | 13 | 13 |
| Housing starts | – | – | 14 | 9 |
| Consumer credit* | – | – | 15 | 14 |
| Initial jobless claims | – | – | 16 | 15 |
| GDP price index (preliminary) | 0 | + | N/A | 10 |

NOTE: Variables in Tables 3 and 4 that are significant at the 10 percent level (based on two-tailed *t*-tests) are ranked above; 0 is not significant.

*Five of the six variables that turned out statistically significant in Gürkaynak, Sack, and Swanson (2003, Table 3).

target (PR) coefficients are larger in magnitude when corrected (original estimates not shown).

Finally, we turn to the influence of Federal Reserve communication. As discussed above, the surprise component in Federal Reserve communication is next to impossible to ascertain. Yet, following Kohn and Sack (2003), we can analyze the effect of Federal Reserve communication on the (conditional) volatility of the dependent variable. Specifically, we are interested in whether Federal Reserve communication and surprises in monetary policy actions bear on inflation rate uncertainty. Intuitively, one might expect Federal Reserve communication to decrease the uncertainty surrounding the future rate of inflation as the chairman of the Federal Reserve offers guidance about the future path of monetary policy. Also, one might expect that monetary policy actions that take the market by surprise will increase uncertainty about future inflation. Note that, if Federal Reserve communication and surprises in monetary policy actions bear on inflation uncertainty, then the error term of the regression equation (3) is heteroskedastic. Rao's score test on heteroskedasticity shows that the null hypothesis of no heteroskedasticity is indeed rejected for the OTR measure of inflation compensation (Table 3) but not for the CM inflation compensation measure (Table 4).[13]

To address the issue of inflation uncertainty, we use the squared residuals from regression equation (3)—the regression results of which are shown in Tables 3 and 4—in an estimation approach suggested by Amemiya (1977, 1978). We regress these squared residuals on the federal funds target variable, an indicator variable that is equal to 1 on days where Federal Reserve communication was priced in the market (and 0 otherwise), and the previously introduced intercept indicator variable (*D*). The regression

[13] For Rao's score test, see Amemiya (1985).

## Table 6

### Instrumental-Variables Approach

**Panel A: OTR measure**

| Explanatory variable | Coefficient | *t*-statistic |
|---|---|---|
| Federal funds target (GSS) | $-2.016 \cdot 10^{-1}$ | –2.710** |
| Federal funds target (PR) | $-2.659 \cdot 10^{-1}$ | –3.191*** |

**Panel B: CM measure**

| Explanatory Variable | Coefficient | *t*-statistic |
|---|---|---|
| Federal funds target (GSS) | $-1.671 \cdot 10^{-1}$ | –2.975*** |
| Federal funds target (PR) | $-1.239 \cdot 10^{-1}$ | –2.575** |

NOTE: ***/** Indicate significance at the 1/5 percent levels (*t*-tests are two-tailed; *t*-statistics are Newey and West (1987) corrected). GSS and PR indicate the federal funds market measure for monetary policy surprises as suggested by Gürkaynak, Sack, and Swanson (2002) and Poole and Rasche (2000), respectively.

## Table 7

### Inflation Uncertainty

**Panel A: OTR measure**

| Explanatory variable | Coefficient | *t*-statistic |
|---|---|---|
| Federal Reserve communication | $-8.632 \cdot 10^{-3}$ | –3.989*** |
| Federal funds target | $1.547 \cdot 10^{-3}$ | 4.009*** |
| Intercept indicator variable (*D*) | $9.189 \cdot 10^{-4}$ | 4.017*** |
| Intercept | $1.049 \cdot 10^{-3}$ | 4.157*** |
| Number of nonzero observations | 166 | |
| Number of observations | 1,555 | |

**Panel B: CM measure**

| Explanatory variable | Coefficient | *t*-statistic |
|---|---|---|
| Federal Reserve communication | $3.021 \cdot 10^{-3}$ | 1.295 |
| Federal funds target | $5.160 \cdot 10^{-4}$ | 1.107 |
| Intercept indicator variable (*D*) | $4.173 \cdot 10^{-4}$ | 1.104 |
| Intercept | $1.381 \cdot 10^{-3}$ | 3.740*** |
| Number of nonzero observations | 117 | |
| Number of observations | 1,110 | |

NOTE: *** Indicates significance at the 1 percent level (*t*-tests are two-tailed). The variable Federal Reserve communication equals 1 on trading days on which Chairman Greenspan's semi-annual testimony to Congress (formerly known as Humphrey-Hawkins testimony) or speeches and other testimonies of Chairman Greenspan were priced in the market. The number of nonzero observations indicates the number of trading days where Federal Reserve communication or the surprise in a monetary policy action was priced.

results, which are presented in Table 7, indicate that neither Federal Reserve communication nor monetary policy surprises bear on the conditional variance of the CM measure of inflation compensation. This finding is not surprising, given that Rao's score test does not suggest heteroskedasticity. Things are different with the OTR measure of inflation compensation. Here, Rao's score test indicates heteroskedasticity, and, indeed, the coefficients for the variables Federal Reserve communication and federal fund target are statistically significant and have the expected sign. Hence, we conclude that, at least judged by one of our two measures of inflation compensation, Federal Reserve communication diminishes the uncertainty surrounding the future rate of inflation, while surprises in monetary policy actions increase it.

## CONCLUSION

Do monetary policy actions that are tighter or easier than expected by the federal funds futures market bear on the average rate of inflation that economic agents expect to prevail over the next 10 years? Moreover, do surprises in macroeconomic data releases lead economic agents to update their beliefs about the average rate of inflation they expect for the next 10 years; if so, which data series matter the most? We gauged inflation expectations by two different measures of inflation compensation, both of which are derived from the market valuation of the expected cash flows of nominal and inflation-indexed Treasury securities.

We find that surprises in monetary policy actions bear on both measures of inflation expectations. Monetary policy actions that are viewed as tighter (easier) than expected by the market lead economic agents to revise down (up) their expected rate of inflation. Further, one measure indicates that Federal Reserve communication reduces uncertainty about the future rate of inflation, while surprises in monetary policy actions increase uncertainty about the path the rate of inflation is going to take. We also show that surprises in macroeconomic data releases matter. In particular, we show that the surprise components in data releases for the employment cost index, the Purchasing Managers' Index, CPI, retail sales, factory orders, and personal income bear most heavily on both measures of inflation expectations.

## REFERENCES

Amemiya, Takeshi. "A Note on a Heteroscedastic Model." *Journal of Econometrics*, November 1977, *6*(3), pp. 365-70.

Amemiya, Takeshi. "Corrigenda." *Journal of Econometrics*, October 1978, *8*(2), p. 275.

Amemiya, Takeshi. *Advanced Econometrics*. Cambridge, MA: Harvard University Press, 1985.

Balduzzi, Pierluigi; Elton, Edwin J. and Green, Clifton T. "Economic News and Bond Prices: Evidence from the U.S. Treasury Market." *Journal of Financial and Quantitative Analysis*, December 2001, *36*(4), pp. 523-43.

Bernanke, Ben S. "An Unwelcome Fall in Inflation?" Remarks before the Economics Roundtable, University of California–San Diego, La Jolla, California, July 23, 2003. www.federalreserve.gov/boarddocs/speeches/2003/20030723/default.htm.

Emmons, William R. "The Information Content of Treasury Inflation-Indexed Securities." Federal Reserve Bank of St. Louis *Review*, November/December 2000, *82*(6), pp. 25-37.

Fleming, Michael J. and Remolona, Eli M. "What Moves Bond Prices?" *Journal of Portfolio Management*, Summer 1999, *25*(4), pp. 28-38.

Fleming, Michael J. and Remolona, Eli M. "What Moves the Bond Market?" Federal Reserve Bank of New York *Economic Policy Review*, December 1997, *3*(4), pp. 31-50.

Greene, William H. *Econometric Analysis*. 5th edition. Upper Saddle River, NJ: Prentice-Hall, August 2002.

Gürkaynak, Refet S.; Sack, Brian P. and Swanson, Eric T. "The Excess Sensitivity of Long-Term Interest Rates: Evidence and Implications for Macroeconomic Models." Working paper, Division of Monetary Affairs, Board of Governors of the Federal Reserve System, April 4, 2003. www.clevelandfed.org/CentralBankInstitute/conf2003/august/sensitivity_apr4.pdf.

Gürkaynak, Refet S.; Sack, Brian P. and Swanson, Eric T. "Market-Based Measures of Monetary Policy Expectations." Working paper, Division of Monetary Affairs, Board of Governors of the Federal Reserve System, August 1, 2002. www.federalreserve.gov/pubs/feds/2002/200240/200240pap.pdf.

Kohn, Donald L. and Sack, Brian P. "Central Bank Talk: Does It Matter and Why?" Paper presented at the Macroeconomics, Monetary Policy, and Financial Stability Conference in honor of Charles Freedman, Bank of Canada, Ottawa, Canada, June 20, 2003.

www.federalreserve.gov/boarddocs/speeches/2003/20030620/paper.pdf.

Kuttner, Kenneth N. "Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market." *Journal of Monetary Economics*, June 2001, *47*(3), pp. 523-44.

Newey, Whitney K. and West, Kenneth D. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica*, May 1987, pp. 703-08.

Poole, William and Rasche, Robert H. "Perfecting the Market's Knowledge of Monetary Policy." *Journal of Financial Services Research*, December 2000, *18*(2-3), pp. 255-98.

Poole, William; Rasche, Robert H. and Thornton, Daniel L. "Market Anticipations of Monetary Policy Actions." Federal Reserve Bank of St. Louis *Review*, July/August 2002, *84*(4), pp. 65-93.

Sack, Brian P. "Deriving Inflation Expectations from Nominal and Inflation-Indexed Treasury Yields." *Journal of Fixed Income*, September 2000, *10*(2), pp. 6-17.

Stock, James H. and Watson, Mark W. "Forecasting Inflation." *Journal of Monetary Economics*, October 1999, *44*(2), pp. 293-335.

Watson, Mark W. Commentary. Federal Reserve Bank of St. Louis *Review*, July/August 2002, *84*(4), pp. 95-97.

# A Rational Pricing Explanation for the Failure of the CAPM

Hui Guo

Fama and French (2003), among many others, show that the capital asset pricing model (CAPM) does not explain stock returns. These results should not be a surprise because the model has some strong assumptions, and the failure of any one of them may cause the model to fail. In particular, the CAPM is a static model in which expected stock returns are assumed to be constant. However, if expected returns are time-varying, Merton (1973) and Campbell (1993), among others, show that the return on an asset is determined not only by its covariance with stock market returns, as in the CAPM, but also by its covariance with variables that forecast stock market returns. In this article, I estimate a variant of Campbell's intertemporal CAPM (ICAPM), using forecasting variables advocated in recent research. I find that the CAPM fails to explain the predictability of stock market returns because covariances with the forecasting variables are also important determinants of stock market returns. Therefore, consistent with some recent authors, for example, Brennan, Wang, and Xia (forthcoming) and Campbell and Vuolteenaho (2002), the failure of the CAPM is related to time-varying expected returns.

The remainder of the article is organized as follows. I first briefly summarize the recent developments of the asset pricing literature and then present evidence that stock market returns and volatility are predictable. For illustration, I discuss and estimate a variant of Campbell's ICAPM and show that changing investment opportunities have important effects on stock prices.[1]

## A BRIEF REVIEW OF THE LITERATURE

In the past two decades, financial economists have documented many anomalies in financial markets. For example, contrary to the market efficiency hypothesis by Fama (1970), Fama and French (1989) argue that stock market returns are predictable. There is also evidence of the predictability in the cross section of stock returns, which casts doubt on the widely accepted CAPM by Sharpe (1964) and Lintner (1965). In particular, Fama and French (1992, 1993) report that value stocks, stocks of high book-to-market value ratio, have much higher risk-adjusted returns than growth stocks, stocks of low book-to-market value ratio. Also, Jegadeesh and Titman (1993) show that the momentum strategy of buying the past winners and selling the past losers is quite profitable.

The advocates of the market efficiency hypothesis argue quite convincingly that many of these anomalies can be attributed to data snooping. For example, if we experiment with a large number of macro variables, it should not be a surprise that a few of them might be correlated with future stock market returns by chance. However, investors cannot profit from such an ex post relation if it does not persist in the future. Indeed, Bossaerts and Hillion (1999), among others, find that, although the variables uncovered by the early authors forecast stock returns in sample, their out-of-sample predictive power is negligible. Similarly, Schwert (2002), among others, finds that many trading strategies, which have been found to generate abnormal returns, were unprofitable in the past decade. Overall, Malkiel (2003) asserts that there is no reliable evidence of persistent stock return predictability and the U.S. equity market is remarkably efficient in the sense that abnormal returns disappear quickly after they are discovered.

Some anomalies, however, cannot be easily discarded as data snooping. Jegadeesh and Titman (2001) and Schwert (2002) find that the momentum strategy remained highly profitable in the 1990s, one decade after it was published in academic jour-

[1] Stock market returns and volatility are measures of investment opportunities.

Hui Guo is an economist at the Federal Reserve Bank of St. Louis. William Bock and Jason Higbee provided research assistance.

Federal Reserve Bank of St. Louis *Review*, May/June 2004, *86*(3), pp. 23-33.

**Table 1**

**Forecasting Quarterly Stock Market Returns and Volatility**

| Intercept | $r_m$ | *cay* | $\sigma_m^2$ | *rrel* | Adjusted $R^2$ |
|---|---|---|---|---|---|
| | | | Return | | |
| –1.50 | 0.02 | 2.42 | 7.19 | –5.63 | 0.20 |
| (–5.12) | (0.34) | (5.13) | (4.73) | (–2.95) | |
| | | | Volatility | | |
| 0.05 | 0.01 | –0.09 | 0.39 | 0.68 | 0.24 |
| (3.49) | (1.32) | (–3.37) | (4.17) | (0.93) | |

NOTE: White-corrected *t*-statistics are in parentheses.

nals. Also, recent authors, for example, Lettau and Ludvigson (2001), show that the consumption-wealth ratio, especially when combined with realized stock market volatility, has statistically and economically significant out-of-sample forecasting power for stock market returns. It is reasonable to believe, as argued by Campbell (2000), that stock returns have some predictable variations. Moreover, the excess stock volatility puzzle (Shiller, 1981), the equity premium puzzle (Mehra and Prescott, 1985), and large fluctuations of stock market volatility (Schwert, 1989) remain unexplained by conventional theories (e.g., Lucas, 1978).

The failure of rational expectations theories leads some researchers to be skeptical about the assumption that individual investors are fully rational. They try to incorporate various well-documented cognitive biases into asset pricing models and find that such combinations have some success in explaining the anomalies mentioned above. Behavioral finance has developed rapidly since the 1990s, and Shiller (2003), among others, has stressed its important role in rebuilding modern finance. However, in my view, we should be at least cautious about it. The main criticism is that a long list of cognitive biases gives researchers so many degrees of freedom that anything can be explained. But financial economists are more interested in the out-of-sample forecast than in explaining what has happened. Also, it is difficult to believe that the investors who frequently misinterpret fundamentals can survive in an arbitrage-driven financial market. Barberis and Thaler (2003) provide a comprehensive survey of the behavioral finance literature and come to this conclusion: "First, we will find that most of our current theories, both rational and behavioral, are wrong. Second, substantially better theories will emerge."

In this article, I want to emphasize the promising role of another alternative hypothesis—*stock return predictability does not necessarily contradict rational expectations theories.*[2] In particular, Campbell and Cochrane (1999) recently proposed a novel explanation using a habit-formation model that investors are more risk tolerant and thus require a smaller equity premium during economic expansions than during economic recessions. Their model not only replicates stock return predictability, but also resolves many other outstanding issues, including the equity premium puzzle and the excess volatility puzzle.

As mentioned, stock return predictability has important implications for asset pricing. Fama (1991) also conjectures that we should relate the cross-section properties of expected returns to the variation of expected returns through time. Consistent with these theories, some recent authors (e.g., Brennan, Wang, and Xia, forthcoming; and Campbell and Vuolteenaho, 2002) find that the predictability of stock market returns and volatility indeed helps explain the cross section of stock returns.

## FORECASTING STOCK MARKET RETURNS AND VOLATILITY

The early authors, for example, Campbell (1987) and Fama and French (1989), find that the short-

[2] This point has been well understood in theory; for example, Lucas (1978) shows that predictable variations of stock returns should be explained by predictable variations of consumption growth. Similarly, in the CAPM, the equity premium is predictable because of predictable variations of stock market volatility, a measure of stock market risk. However, these models cannot generate sizable predictable variations of stock returns as observed in the data.

## Figure 1

### Expected Excess Stock Market Returns

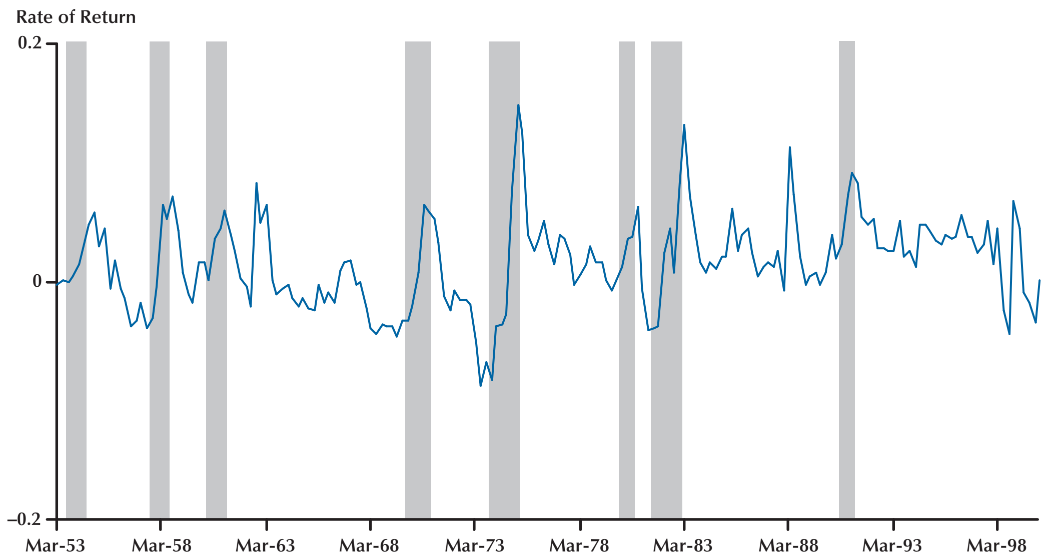


NOTE: Shaded bars indicate recessions as determined by the National Bureau of Economic Research.

term interest rate, the dividend yield, the default premium, and the term premium forecast stock market returns. Recently, Lettau and Ludvigson (2001) report that the consumption-wealth ratio, *cay*—the error term from the cointegration relation among consumption, net worth, and labor income—forecasts stock market returns in sample and out of sample. Interestingly, I (Guo, 2003a) find that the predictive power of *cay* improves substantially if past stock market variance, $\sigma_m^2$, is also included in the forecasting equation and the stochastically detrended risk-free rate, *rrel*, provides additional information about future returns.[3]

I replicate this result in the upper panel of Table 1, with the White-corrected (White, 1980) *t*-statistics reported in parentheses. It shows that all three variables are statistically significant in the forecasting equation of real stock market return, $r_m$, and the adjusted $R^2$ is about 20 percent; however, the lagged dependent variable is insignificant.[4] Moreover, these variables drive out the other commonly used forecasting variables, including the dividend yield, the default premium, and the term premium.[5] Figure 1 plots the fitted value from the forecasting regression of returns and shows that expected returns tend to rise during recessions.

Schwert (1989), among many others, also finds clustering in stock market volatility: When volatility increases, it stays at its high level for an extended period before it reverts to its average level. Research shows that some macro variables predict stock volatility as well. Consistent with Lettau and Ludvigson (2003), the lower panel of Table 1 shows that, while past volatility is positively related to future volatility, *cay* is negatively related to it. Figure 2 plots the fitted value from the forecasting regression of stock market volatility, which also tends to increase during recessions.

I (Guo, 2003b) provide some theoretical insight on these empirical results in a limited stock market participation model. In particular, I argue that, in addition to a market risk premium (as in the standard consumption-based model), investors also require a liquidity premium on stocks because investors cannot use stocks to hedge income risk—due to limited stock market participation. Therefore, stock volatility and the consumption-wealth ratio forecast

[3] The stochastically detrended risk-free rate is the difference between the nominal risk-free rate and its average in the previous 12 months.

[4] See the appendix for data descriptions.

[5] To conserve space, I do not report the results here; but they are available upon request.

**Figure 2**

**Expected Stock Market Volatility**

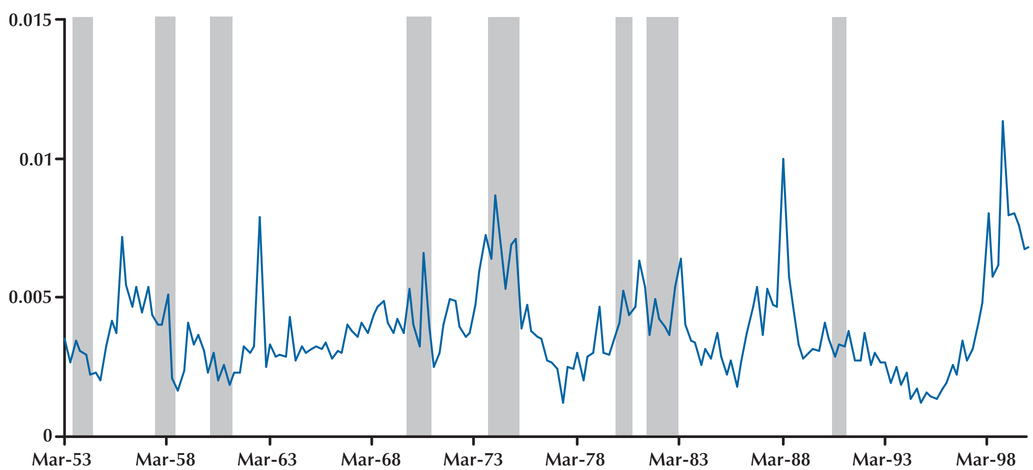


NOTE: Shaded bars indicate recessions as determined by the National Bureau of Economic Research.

stock market returns because they are proxies for the risk and liquidity premiums, respectively.[6]

It should not be a surprise that stock market volatility forecasts returns, which is an important implication of the CAPM.[7] Intuitively, risk-averse investors tend to reduce their holding of equities relative to safe assets such as Treasury bills when volatility is expected to rise. To induce investors to hold a broadly measured stock market index, the expected stock market return has to rise as well. Given that the level of volatility tends to persist over time, we expect that past volatility should provide some indication of future volatility and hence stock market returns. On the other hand, the consumption-wealth ratio measures investors' liquidity conditions. When investors are borrowing constrained because of, for example, a bad income shock, they require a high liquidity premium on stocks and stock prices thus fall. Conversely, the liquidity premium is low and stock prices rise when investors have plenty of liquidity.

It is important to note that the risk and liquidity premiums or their proxies, stock market volatility and the consumption-wealth ratio, could be negatively related to one another in the limited stock market participation model. Intuitively, when investors have excess liquidity, they might be willing to hold stocks when the expected return is low, even though expected volatility is high. This implication is particularly relevant for the stock market boom in the late 1990s, during which investors accepted a low expected return even though volatility rose to a historically high level. Indeed, as shown in Table 1, while volatility and the consumption-wealth ratio are both positively related to future stock market returns, they are negatively related to one another in the post-World War II sample. This pattern explains that, because of an omitted variable problem, the predictive power of the consumption-wealth ratio improves dramatically when past variance is also included in the forecast equation. As shown in Guo and Whitelaw (2003), it also explains why the early authors fail to find conclusive evidence of a positive risk-return relation, as stipulated by the CAPM.

There is an important conceptual issue of using the consumption-wealth ratio as a forecasting variable, because consumption and labor income data are subject to revision. In particular, Guo (2003c) finds that the predictive power of the consumption-wealth ratio deteriorates substantially if we use information available at the time of the forecast. It is

[6] Patelis (1997) suggests that variables such as the stochastically detrended risk-free rate forecast stock returns because these variables reflect the stance of monetary policies, which have state-dependent effects on real economic activities through a credit channel (e.g., Bernanke and Gertler, 1989).

[7] However, it is puzzling that many authors (e.g., Campbell, 1987; Glosten, Jagannathan, and Runkle, 1993; and Lettau and Ludvigson, 2003) report a negative risk-return tradeoff in the stock market. As I will discuss here, these results reflect the fact that the early authors fail to control for the liquidity premium, which may be negatively related to the risk premium.

important to note that the use of real-time data does not call the predictive abilities of the consumption-wealth ratio into question. In fact, we expect the consumption-wealth ratio to have better predictive power in the current vintage data than in the real-time data because the latter is a noisier and potentially biased measure of its “true” value. Moreover, investors may obtain similar information from alternative sources; for example, Guo and Savickas (2003a) show that a measure of the (value-weighted) idiosyncratic volatility, which is available in real time, has forecasting abilities that are very similar to those of the consumption-wealth ratio. Therefore, as stressed by Lettau and Ludvigson (2003), it is appropriate to use the consumption-wealth ratio estimated from the current vintage data in this paper because I address the question of whether expected excess returns are time-varying.[8]

## CAMPBELL’S ICAPM

In this section, I briefly discuss how stock market returns are determined in a rational expectations model (i.e., Campbell’s, 1993, ICAPM) if stock market returns and volatility are predictable. In particular, Campbell argues that the expected return on any asset is determined by its covariance with stock market returns and variables that forecast stock market returns. This simple exercise helps illustrate why the CAPM fails to explain the cross section of stock returns, as mentioned in the introduction.

Campbell’s ICAPM is quite intuitive. For example, because the consumption-wealth ratio is positively related to future stock market returns, a negative innovation in the consumption-wealth ratio indicates a low future expected return or worsened “future investment opportunities.” A stock is thus risky if its return is low when future investment opportunities deteriorate—that is, there is a negative shock to the consumption-wealth ratio. As a result, in addition to compensation for the market risk, investors require additional compensation on this stock because it provides a poor hedge for changing investment opportunities. Below, I briefly discuss the testable implications of Campbell’s ICAPM. Interested readers may look to Campbell (1993, 1996) for details.

Campbell’s ICAPM is a model of an infinite horizon economy, in which a representative agent maximizes an Epstein and Zin (1989) objective function,

(1)
$$U_t = \{(1-\beta\}C_t^{1-(1/\sigma)} + \beta(E_t U_{t+1}^{1-\gamma})^{[1-(1/\sigma)]/(1-\gamma)}\}^{1/[1-(1/\sigma)]}$$
$$= [(1-\beta)C_t^{(1-\gamma)/\theta} + \beta(E_t U_{t+1}^{1-\gamma})^{1/\theta}]^{\theta/(1-\gamma)},$$

subject to the intertemporal budget constraint

(2)
$$W_{t+1} = R_{m,t+1}(W_t - C_t).$$

In the above equations, $C_t$ is consumption, $W_t$ is aggregate wealth, $R_{m,t+1}$ is the return on aggregate wealth, $\beta$ is the time discount factor, $\gamma$ is the relative risk aversion coefficient, $\sigma$ is the elasticity of intertemporal substitution, and $\theta$ is defined as $\theta = (1-\gamma)/[1-(1/\sigma)]$. If we set $\theta$ equal to 1, we obtain the familiar power utility function, in which the relative risk aversion coefficient is equal to the reciprocal of the elasticity of intertemporal substitution.

Suppose that there are $(K-1)$ state variables, $x_{t+1} = [x_{1,t+1},\ldots,x_{K-1,t+1}]'$, lags of which forecast stock market return, $r_{m,t+1}$, and its volatility.[9] Also, $r_{m,t+1}$ and $x_{t+1}$ follow a first-order vector autoregressive (VAR) process:

(3)
$$\begin{bmatrix} r_{m,t+1} \\ x_{t+1} \end{bmatrix} - A_0 - A\begin{bmatrix} r_{m,t} \\ x_t \end{bmatrix} = \begin{bmatrix} \varepsilon_{1,t+1} \\ \vdots \\ \varepsilon_{K,t+1} \end{bmatrix},$$

where $A_0$ is a $K$-by-1 vector of intercepts, $A$ is a $K$-by-$K$ matrix of slope coefficients, and $[\varepsilon_{1,t+1}, \varepsilon_{2,t+1},\ldots, \varepsilon_{K,t+1}]'$ is a $K$-by-1 vector of error terms, which are orthogonal to the lagged state variables. Campbell (1993) shows that, if stock market returns and volatility are predictable, as shown in equation (3), the expected return on any asset, e.g., $r_{i,t+1}$, is determined by its covariance with stock market returns and variables that forecast stock market returns:

(4)
$$E_t r_{i,t+1} - r_{f,t+1} + \frac{V_{ii,t}}{2} = \gamma V_{im,t} + \sum_{k=1}^{K}[(\gamma - 1 - \frac{\theta\psi}{\sigma})]\lambda_{hk} V_{ik,t},$$

where $r_{f,t+1}$ is the risk-free rate, $V$ is conditional variance or covariance, $\Psi$ is the coefficient relating stock market return to volatility, and $\lambda_{hk}$ is a function

[8] I focus on the consumption-wealth ratio in this article because, as mentioned, it is a theoretically motivated variable. In contrast, the idiosyncratic volatility forecasts stock returns because of its co-movements with the consumption-wealth ratio, and such a link has not been well understood. Also, the consumption-wealth ratio is available in a longer sample than the idiosyncratic volatility.

[9] $r_{m,t+1}$ is the log of return on aggregate wealth, $R_{m,t+1}$.

**Table 2**

**Model Specifications**

| Model | Restrictions on equation (5) |
|---|---|
| I | $\lambda_{hk} = 0,\ k = 1, \cdots, K$ |
| II | $\frac{\theta\psi}{\sigma} = 0$ |
| III | No restrictions |
| IV | Coefficients of $\varepsilon_{m,t+1}\varepsilon_{k,t+1},\ k = 1, \cdots, K$ are free parameters |

of $A$.[10] In particular, excess stock market return is given by

$$(5)\quad \begin{aligned} & r_{m,t+1} - r_{f,t+1} + \frac{\varepsilon^2_{m,t+1}}{2} - \gamma\varepsilon^2_{m,t+1} \\ & - \sum_{k=1}^{K} [(\gamma - 1 - \frac{\theta\psi}{\sigma})]\lambda_{hk}\varepsilon_{m,t+1}\varepsilon_{k,t+1} = u_{m,t+1}, \end{aligned}$$

where $\varepsilon_{m,t+1}$ is the shock to the stock market return, which I also denote as $\varepsilon_{1,t+1}$ in equation (3). Equation (5) indicates that stock market return is predictable because its covariances with state variables—for example, stock market volatility, $\varepsilon^2_{m,t+1}$—are predictable.

Campbell's ICAPM of equations (3) and (5) can be estimated using the generalized method of moments (GMM) by Hansen (1982). I use a quarterly sample spanning from 1952:Q4 to 2000:Q4, with a total of 193 observations. To mitigate the potential small sample problem, I follow the advice of Ferson and Foerster (1994) and use the iterative GMM unless otherwise noted. I assume that the error terms in equations (3) and (5) are orthogonal to the lagged state variables and have zero means. Equation (3) is exactly identified. In equation (5), there are two parameters, $\gamma$ and $\frac{\theta\psi}{\sigma}$, and $K+1$ orthogonality conditions. The equation system, therefore, is overidentified with $K-1$ degrees of freedom. Hansen's (1982) $J$-test can be used to test the null hypothesis that the pricing error of equation (5), $u_{m,t+1}$, is orthogonal to the lagged state variables and has a zero mean. We can also back-out the price of risk for each factor using the formula

$$(6)\quad \begin{aligned} & p_1 = \gamma + [(\gamma - 1 - \frac{\theta\psi}{\sigma})]\lambda_{h1} \\ & p_i = [(\gamma - 1 - \frac{\theta\psi}{\sigma})]\lambda_{hi},\ i = 2, \ldots, K, \end{aligned}$$

where $p_1$ is the risk price for stock market returns and $p_i$ is the risk price for forecasting variable $i$.

If we impose the restriction that $\Psi$—the parameter for time-varying stock market volatility—is equal to zero, we obtain the special case analyzed by Campbell (1996), in which volatility changes have no effects on asset prices. It should be noted that, as discussed in footnote 10, equation (4) or (5) is a special case of Campbell's ICAPM with time-varying volatility. Under general conditions (e.g., conditional stock market volatility is a linear function of lagged state variables), I (Guo, 2002) show that conditional stock market return is still a linear function of its covariances with state variables, but the risk prices are complicated functions of the underlying structural parameters. For robustness, I also estimate equation (5) by assuming that the risk prices are free parameters. It should also be noted that equation (5) reduces to the familiar CAPM if we drop the covariances between stock market returns and the forecasting variables. These four specifications are nested, and the $D$-test proposed by Newey and West (1987) can be used to test the restrictions across these specifications.

Table 2 summarizes the specifications of the four models investigated in this paper. Model I is the CAPM, in which I assume that the covariances with variables that forecast stock market returns have no effects on the expected stock market return. Model II is Campbell's ICAPM with constant stock market volatility, in which the parameter $\frac{\theta\psi}{\sigma}$ is restricted to be zero. In Model III, I allow volatility changes to affect the expected stock market return and estimate $\frac{\theta\psi}{\sigma}$ as a free parameter. Model IV is the general case of Campbell's ICAPM, in which I estimate the risk prices as free parameters. In models II and III, I estimate the structural parameters $\gamma$ and/or $\frac{\theta\psi}{\sigma}$ and use equation (6) to back-out the risk prices. The risk prices are estimated directly in models I and IV.

## EMPIRICAL RESULTS

Before presenting the empirical results, I want to emphasize that Campbell's ICAPM is not a general

[10] For example, $V_{i1,t} = E_t(\varepsilon_{1,t+1}\varepsilon_{i,t+1})$ is the conditional covariance between the shock to stock market return, $\varepsilon_{1,t+1}$, and the shock to the return on asset $i$, $\varepsilon_{i,t+1}$. To derive equation (4), I also use a simplifying assumption, $E_t r_{m,t+1} = \Psi V_{mm,t}$, as suggested by Campbell (1993).

## Table 3

### Campbell's ICAPM with Constant $\gamma$

| Model | $\gamma$ | $\frac{\theta\psi}{\sigma}$ | Risk prices for $r_m$ | $cay$ | $\sigma_m^2$ | $rrel$ | *J*-test (*p*-value) |
|---|---|---|---|---|---|---|---|
| I | 6.53 | | 6.53 | | | | $\chi^2(4) = 14.01$ |
| | (5.12) | | (5.12) | | | | (0.01) |
| II* | 32.82 | | 9.50 | 441.67 | 405.79 | –807.87 | $\chi^2(4) = 7.12$ |
| | (3.35) | | (3.13) | (2.79) | (3.37) | (–3.17) | (0.13) |
| III | 31.29 | –14.62 | 8.91 | 449.16 | 447.76 | –719.24 | $\chi^2(3) = 0.73$ |
| | (2.31) | (–1.97) | (2.02) | (2.10) | (2.40) | (–2.58) | (0.87) |
| IV | | | 9.32 | 71.44 | 387.39 | –739.37 | $\chi^2(1) = 0.04$ |
| | | | (1.03) | (0.18) | (1.27) | (–0.94) | (0.85) |

| *D*-test (*p*-value) |
|---|
| I vs. IV: $\chi^2(3) = 19.64$ (0.00) |
| II vs. IV: $\chi^2(3) = 7.19$ (0.07) |
| II vs. III: $\chi^2(1) = 5.60$ (0.02) |
| III vs. IV: $\chi^2(2) = 1.01$ (0.60) |

NOTE: * I use the identity matrix as the initial weighting matrix and use five iterations; *t*-statistics are in parentheses unless otherwise indicated.

equilibrium model, because it takes stock return predictability as given. Therefore, the test of Campbell's ICAPM is a joint test of an equilibrium model, which explains the choice of forecasting variables. This explains, in contrast with my results, that the early authors such as Campbell (1996), Li (1997), and Chen (2002) find little support for Campbell's ICAPM because they use different sets of forecasting variables.

Table 3 presents the empirical results. Consistent with asset pricing theories, the relative risk aversion coefficient, $\gamma$, is statistically positive with a point estimate of about 6.5 in model I.[11] However, the *J*-test rejects the model at the 1 percent significance level, indicating that the stock return predictability cannot be explained solely by predictable variations in volatility. This result, which is consistent with Harvey (1989), should not be a surprise. Table 1 shows that other variables such as the consumption-wealth ratio and the stochastically detrended risk-free rate also forecast stock market returns. Their covariances with stock market returns, therefore, are also components of the expected stock market return, as shown in equation (5).

In model II, the relative risk aversion coefficient is also significantly positive, with a point estimate of about 32.8. Moreover, the risk prices of all factors are statistically significant with expected signs. In particular, the covariance with the consumption-wealth ratio, *cay*, and realized stock market variance, $\sigma_m^2$, is positively priced because these two variables are positively related to future stock market returns, as shown in Table 1. Similarly, the covariance with the stochastically detrended risk-free rate, *rrel*, is negatively priced because it is negatively related to future stock market returns. The price of the market risk, $r_m$, is positive; its point estimate of 9.5, however, is much smaller than that of the relative risk aversion coefficient. This result, as argued by Campbell (1996), reflects the mean-reversion in stock market returns. However, there is only weak support for model II: It is not rejected at the 10 percent significance level by the *J*-test.

In model III, the coefficient $\frac{\theta\psi}{\sigma}$ is negative and statistically significant. This should not be a surprise given mounting evidence of large fluctuations in

[11] The price of the market risk is equal to the relative risk aversion coefficient in the CAPM.

## Figure 3

### A Decomposition of Stock Market Return

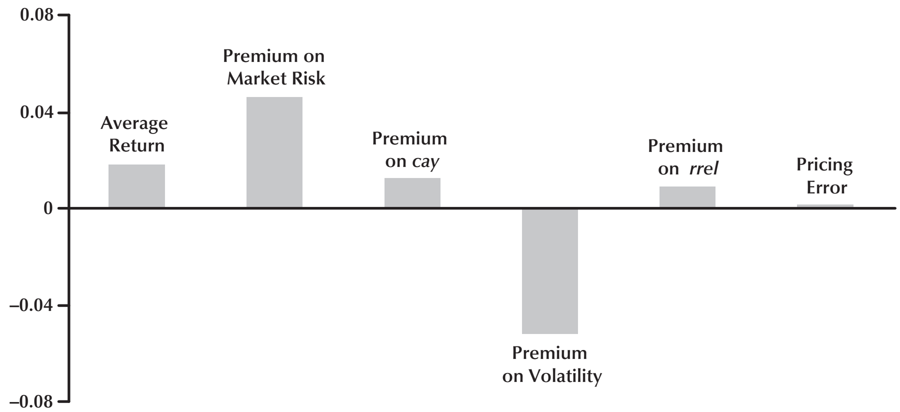


stock market volatility, e.g., Schwert (1989). Moreover, model III fits the data pretty well: It is not rejected at the 80 percent significance level, according to the *J*-test. Nevertheless, the point estimates of the relative risk aversion coefficient and the risk prices are similar to those reported in model II. Finally, we find that model IV, the most general specification, also explains the dynamic of stock market returns well.

The *D*-test reveals a similar pattern. Model I is overwhelmingly rejected relative to model IV, indicating that the CAPM cannot explain the dynamic of stock market returns. Model II is also rejected relative to models III and IV, indicating that changes in volatility have important effects on asset prices.[12] However, we cannot reject model III relative to model IV at the 60 percent significance level. Therefore, equation (4) or (5) provides a good approximation for the effect of return heteroskedasticity. One advantage of model III is that it allows us to estimate the structural parameter of the relative risk aversion coefficient, $\gamma$.

In Figure 3, I use the estimation results of model III in Table 2 to decompose the average stock market return into its covariances with the four risk factors and the pricing error. It shows that, although the market risk is an important determinant of the average return, the risk premiums on *cay*, $\sigma_m^2$, and *rrel* are also substantial. The pricing error, however, is very small relative to the average return, which confirms the *J*-test in Table 3.

Figure 3 sheds light on the failure of the CAPM, as argued by Fama and French (2003), among others: The market risk is not the only determinant of stock returns when conditional stock market return and volatility change over time. For example, as mentioned in the introduction, value stocks earn higher average returns than growth stocks, even though their covariances with stock market returns are similar. This is because value stocks have higher covariances with *cay* than growth stocks; similarly, the momentum profit is explained by the fact that the past winners have higher covariances with $\sigma_m^2$ than the past losers do (Guo, 2002). Brennan, Wang, and Xia (forthcoming) and Campbell and Vuolteenaho (2002) also find that a hedge for changing investment opportunities explains the value premium, although they use different instrumental variables.

Lastly, I want to stress that, although many financial economists agree that the CAPM does not explain the cross section of stock returns, they disagree on the source of the deviation from the CAPM. This is because the early authors test the CAPM using portfolios formed according to various characteristics, such as book-to-market value ratios, and the failure of the CAPM is consistent with a host of alternative hypotheses. For example, while Fama and French (2003) interpret the value premium as being consistent with ICAPM, Lakonishok, Shleifer, and Vishny (1994) and MacKinlay (1995) attribute it to irrational pricing and data mining.[13] Based on

[12] Ang et al. (2003) also find that stock market volatility is a significantly priced risk factor.

[13] The value premium is the return on a portfolio that is long in stocks with high book-to-market value ratios and short in stocks with low book-to-market value ratios.

Campbell's ICAPM or equation (4), Guo and Savickas (2003b) provide some new insight on this issue by forming portfolios on conditionally expected returns. In particular, they use the same variables used in this paper to make out-of-sample forecasts for individual stocks and then sort the stocks into decile portfolios based on the forecast. They show that the decile portfolios, which are motivated directly from the ICAPM and thus not vulnerable to the criticism of data mining or irrational pricing, pose a serious challenge to the CAPM. Their results provide direct support for ICAPM.

## CONCLUSION

In this article, I provide a brief survey of rational pricing explanations for stock return predictability. For illustration purposes, I also estimate and test a variant of Campbell's ICAPM, which allows for time-varying conditional return and volatility. Consistent with the recent authors, ICAPM appears to explain the dynamic of stock market returns better than the CAPM does. The results suggest that stock return predictability is important for understanding asset pricing.

## REFERENCES

Ang, Andrew; Hodrick, Robert J.; Xing, Yuhang and Zhang, Xiaoyan. "The Cross-Section of Volatility and Expected Returns." Working paper, Columbia Business School, 2003.

Barberis, Nicholas C. and Thaler, Robert H. "A Survey of Behavioral Finance," in George Constantinides, Milton Harris, Rene Stulz, eds., *Handbook of the Economics of Finance*. Amsterdam: North-Holland, 2003.

Bernanke, Ben S. and Gertler, Mark. "Agency Costs, Net Worth, and Business Fluctuations." *American Economic Review*, March 1989, *79*(1), pp. 14-31.

Bossaerts, Peter and Hillion, Pierre. "Implementing Statistical Criteria to Select Return Forecasting Models: What Do We Learn?" *Review of Financial Studies*, Summer 1999, *12*(2), pp. 405-28.

Brennan, Michael J.; Wang, Ashley W. and Xia, Yihong. "Estimation and Test of a Simple Model of Intertemporal Capital Asset Pricing" (forthcoming in *Journal of Finance*).

Campbell, John Y. "Stock Returns and the Term Structure." *Journal of Financial Economics*, June 1987, *18*(2), pp. 373-99.

Campbell, John Y. "Intertemporal Asset Pricing without Consumption Data." *American Economic Review*, June 1993, *83*(3), pp. 487-512.

Campbell, John Y. "Understanding Risk and Return." *Journal of Political Economy*, April 1996, *104*(2), pp. 298-345.

Campbell, John Y. "Asset Pricing at the Millennium." *Journal of Finance*, August 2000, *55*(4), pp. 1515-67.

Campbell, John Y. and Cochrane, John H. "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior." *Journal of Political Economy*, April 1999, *107*(2), pp. 205-51.

Campbell, John Y.; Lettau, Martin; Malkiel, Burton G. and Xu, Yexiao. "Have Individual Stocks Become More Volatile? An Empirical Exploration of Idiosyncratic Risk." *Journal of Finance*, February 2001, *56*(1), pp. 1-43.

Campbell, John Y. and Vuolteenaho, Tuomo. "Bad Beta, Good Beta." Working paper, Harvard University, 2002.

Chen, Joseph. "Intertemporal CAPM and the Cross-Section of Stock Returns." Working paper, University of Southern California, 2002.

Epstein, Larry G. and Zin, Stanley E. "Substitution, Risk Aversion, and the Temporal Behavior of Asset Returns." *Journal of Political Economy*, April 1989, *99*(2), pp. 263-86.

Fama, Eugene. "Efficient Capital Markets: A Review of Theory and Empirical Work." *Journal of Finance*, May 1970, *25*(2), pp. 383-417.

Fama, Eugene. "Efficient Capital Markets: II." *Journal of Finance*, December 1991, *46*(5), pp. 1575-618.

Fama, Eugene and French, Kenneth. "Business Conditions and Expected Returns on Stocks and Bonds." *Journal of Financial Economics*, November 1989, *25*(1), pp. 23-49.

Fama, Eugene and French, Kenneth. "The Cross-Section of Expected Stock Returns." *Journal of Finance*, June 1992, *47*(2), pp. 427-65.

Fama, Eugene and French, Kenneth. "Common Risk Factors in the Returns on Stocks and Bonds." *Journal of Financial Economics*, February 1993, *33*(1), pp. 3-56.

Fama, Eugene and French, Kenneth. "The CAPM: Theory and Evidence." Working Paper No. 550, Center for Research in Security Prices, University of Chicago, 2003.

Ferson, Wayne E. and Foerster, Stephen R. "Finite Sample Properties of the Generalized Method of Moments in Tests of Conditional Asset Pricing Models." *Journal of Financial Economics*, August 1994, *36*(1), pp. 29-55.

Guo, Hui. "Time-Varying Risk Premia and the Cross Section of Stock Returns." Working Paper 2002-013B, Federal Reserve Bank of St. Louis, 2002.

Guo, Hui. "On the Out-of-Sample Predictability of Stock Market Returns" (forthcoming in *Journal of Business*, 2003a).

Guo, Hui. "Limited Stock Market Participation and Asset Prices in a Dynamic Economy" (forthcoming in *Journal of Financial and Quantitative Analysis*, 2003b).

Guo, Hui. "On the Real-Time Forecasting Ability of the Consumption-Wealth Ratio." Working Paper 2003-007b, Federal Reserve Bank of St. Louis, 2003c.

Guo, Hui and Savickas, Robert. "Idiosyncratic Volatility, Stock Market Volatility, and Expected Stock Returns." Working Paper 2003-028A, Federal Reverse Bank of St. Louis, 2003a.

Guo, Hui and Savickas, Robert. "On the Cross Section of Conditionally Expected Stock Returns." Working Paper 2003-043A, Federal Reserve Bank of St. Louis, 2003b.

Guo, Hui and Whitelaw, Robert. "Uncovering the Risk-Return Relation in the Stock Market." NBER Working Paper 9927, National Bureau of Economic Research, 2003.

Glosten, Lawrence R.; Jagannathan, Ravi and Runkle, David E. "On the Relation Between the Expected Value and the Variance of the Nominal Excess Return on Stocks." *Journal of Finance*, December 1993, *48*(5), pp. 1779-801.

Hansen, Lars P. "Large Sample Properties of Generalized Method of Moments Estimators." *Econometrica*, July 1982, *50*(4), pp. 1029-54.

Harvey, Campbell R. "Time-Varying Conditional Covariances in Tests of Asset Pricing Models." *Journal of Financial Economics*, October 1989, *24*(2), pp. 289-317.

Jegadeesh, Narasimhan and Titman, Sheridan. "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency." *Journal of Finance*, March 1993, *48*(1), pp. 65-91.

Jegadeesh, Narasimhan and Titman, Sheridan. "Profitability of Momentum Strategies: An Evaluation of Alternative Explanations." *Journal of Finance*, April 2001, *56*(2), pp. 699-720.

Lakonishok, Josef; Shleifer, Andrei and Vishny, Robert W. "Contrarian Investment, Extrapolation, and Risk." *Journal of Finance*, December 1994, *49*(5), pp. 1541-78.

Lettau, Martin and Ludvigson, Sydney. "Consumption, Aggregate Wealth, and Expected Stock Returns." *Journal of Finance*, June 2001, *56*(3), pp. 815-49.

Lettau, Martin and Ludvigson, Sydney. "Measuring and Modeling Variation in the Risk-Return Tradeoff." Working paper, New York University, 2003.

Li, Yuming. "Intertemporal Asset Pricing without Consumption Data: Empirical Tests." *Journal of Financial Research*, Spring 1997, *20*(1), pp. 53-69.

Lintner, John. "Security Prices, Risk and Maximal Gains from Diversification." *Journal of Finance*, December 1965, *20*(4), pp. 587-615.

Lucas, Robert E. Jr. "Asset Prices in an Exchange Economy." *Econometrica*, November 1978, *46*(6), pp. 1429-46.

MacKinlay, Craig A. "Multifactor Models Do Not Explain Deviations from the CAPM." *Journal of Financial Economics*, May 1995, *38*(1), pp. 3-28.

Malkiel, Burton G. "The Efficient Market Hypothesis and Its Critics." *Journal of Economic Perspectives*, Winter 2003, *17*(1), pp. 59-82.

Mehra, Rjnish and Prescott, Edward C. "The Equity Premium: A Puzzle." *Journal of Monetary Economics*, March 1985, *15(*2), pp. 145-61.

Merton, Robert C. "An Intertemporal Capital Asset Pricing Model." *Econometrica*, September 1973, *41*(5), pp. 867-87.

Merton, Robert C. "On Estimating the Expected Return on the Market: An Exploratory Investigation." *Journal of Financial Economics*, December 1980, *8*(4), pp. 323-61.

Newey, Whitney K. and West, Kenneth D. "Hypothesis Testing with Efficient Method of Moments Estimation." *International Economic Review*, October 1987, *28*(3), pp. 777-87.

Patelis, Alex D. "Stock Return Predictability and The Role of Monetary Policy." *Journal of Finance*, December 1997, *52*(5), pp. 1951-72.

Schwert, William G. "Why Does Stock Market Volatility Change Over Time?" *Journal of Finance*, December 1989, *44*(3), pp. 1115-53.

Schwert, William G. "Anomalies and Market Efficiency." *Handbook of the Economics of Finance*, Chap. 15. Amsterdam: North-Holland, 2002.

Sharpe, William F. "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk." *Journal of Finance*, September 1964, *19*(3), pp. 425-42.

Shiller, Robert J. "Do Stock Prices Move Too Much to Be Justified by Subsequent Changes in Dividends?" *American Economic Review*, June 1981, *71*(3), pp. 421-36.

Shiller, Robert J. "From Efficient Markets Theory to Behavioral Finance." *Journal of Economic Perspectives*, Winter 2003, *17*(1), pp. 83-104.

White, Halbert L. "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity." *Econometrica*, May 1980, *48*(4), pp. 817-38.

## Appendix

### DATA DESCRIPTIONS

Because the consumption-wealth ratio, *cay*, is available on a quarterly basis, I analyze a quarterly sample spanning from 1952:Q4 to 2000:Q4, with a total of 193 observations. Following Merton (1980), among many others, realized stock market variance, $\sigma_m^2$, is the sum of the squared deviation of the daily excess stock return from its quarterly average in a given quarter. It should be noted that, as in Campbell et al. (2001), I make a downward adjustment for the realized stock market variance of 1987:Q4, on which the 1987 stock market crash has a compounding effect. The stochastically detrended risk-free rate, *rrel*, is the difference between the risk-free rate and its average over the previous four quarters, and the quarterly risk-free rate is approximated by the sum of the monthly risk-free rate in a quarter. The consumption-wealth ratio data were obtained from Martin Lettau at New York University. I obtain the daily value-weighted stock market return from the Center for Research in Security Prices (CRSP) at the University of Chicago. The daily risk-free rate is not directly available, but I assume that it is constant within a given month. The monthly risk-free rate is also obtained from the CRSP. The real stock market return, $r_m$, is the difference between the CRSP value-weighted market return and the inflation rate of the consumer price index obtained from the Bureau of Economic Analysis.

# How Costly Is Sustained Low Inflation for the U.S. Economy?

James B. Bullard and Steven Russell

## 1. INTRODUCTION

In this paper we describe and analyze a quantitative-theoretic general equilibrium model in which permanent changes in monetary policy have important welfare consequences for households. Our main findings are estimates of the welfare cost of inflation that are an order of magnitude larger than most estimates found in the cost-of-inflation literature. In particular, we find that a permanent, 10-percentage-point increase in the inflation rate—a standard experiment in this literature—imposes an annual welfare loss equivalent to 11.2 percent of output. Most estimates of the cost of inflation place this loss at less than 1 percent of output. Thus, our analysis helps account for the widely held view that the benefits of reducing the inflation rate from the double-digit levels experienced during the 1970s were very large.

The model we employ belongs to a class of models—the overlapping generations (OLG) or general equilibrium life-cycle models—that have rarely been used to study the cost of inflation and have never been used to obtain practical estimates of the magnitude of these costs. The distinctive features of the model allow us to study a source of welfare losses from inflation that has not been described previously. Although our analysis is novel in these respects, in most other ways it is entirely conventional. We make standard assumptions about preferences, production, and capital accumulation. Households and firms have rational expectations, and equilibria occur at prices and interest rates that clear markets. Money demand is introduced through a reserve requirement. Changes in monetary policy take the form of permanent changes in the growth rate of the base money stock that produce permanent changes in the rate of inflation. We follow the bulk of the inflation-cost literature by basing our cost estimates on comparisons of alternative steady states.

We follow the recent trend in applied macroeconomic theory by calibrating our model to increase the empirical credibility of its predictions. The principal goal of our calibration procedure is to produce a steady-state equilibrium that matches certain long-run-average features of U.S. postwar data. We have given the model a variety of characteristics that increase both its overall plausibility and its ability to mimic these data. The characteristics include households that live for a large but finite number of periods, exogenous technological progress, exogenous population growth, costly financial intermediation, and endogenous labor-leisure decisions. The model also includes a fairly elaborate government sector, including real expenditures (government purchases), taxes on labor and capital income, seigniorage revenue, and government debt. The importance of the role played by the government sector is a distinctive feature of our analysis.

A characteristic of the observed public finance system that plays a key role in driving our results is that capital income taxes are levied on net *nominal* income, so that increases in the inflation rate increase effective capital income tax rates. In this respect, our analysis is similar to recent work by Feldstein (1997) and Abel (1997). However, their estimates of the cost of inflation are based largely on the tendency of higher effective capital income tax rates to increase the wedge between the before-tax and

---

James B. Bullard is a vice president and economist at the Federal Reserve Bank of St. Louis. Steven Russell is a professor of economics at Indiana University–Purdue University at Indianapolis (IUPUI).

The authors thank Bruce Smith at the conference "Macroeconomic Theory and Monetary Policy," sponsored by the University of Pennsylvania and the Federal Reserve Bank of Philadelphia; Roger Farmer at the conference "Dynamic Models of Economic Policy," sponsored by the University of Rochester, University of California, Los Angeles, and the Federal Reserve Bank of Minneapolis; Wilbur John Coleman II at a Summer North American Meeting of the Econometric Society; and Frank Gong at a meeting of the Federal Reserve System Committee on Macroeconomics for useful discussant comments. They also thank seminar participants at the Bank of England, Cornell University, Indiana University, Midwest Macroeconomics, Rutgers University, Simon Fraser University, University of Texas, and the University of Western Ontario for helpful feedback. Deborah Roisman provided research assistance.

Federal Reserve Bank of St. Louis *Review*, May/June 2004, *86*(3), pp. 35-67.

after-tax real rates of return on capital. We identify an entirely new channel through which increased effective capital income tax rates contribute to the cost of inflation. Our inflation-cost estimates are attributable mostly to this channel, which we discuss at length below.

We can use our model to identify the portions of our total welfare-cost estimates that are attributable to effects analogous to those studied by other researchers. In particular, the "purely monetary" component of our cost estimate—the portion that is due to the fact that an increase in the inflation rate is a decrease in the real rate of return on money—accounts for somewhere between 1 and 5 percent of our total cost estimate, which makes it roughly as large as most inflation-cost estimates in the literature. The component of the cost that is due to the effects emphasized by Feldstein accounts for roughly 15 percent of our total cost estimate. Thus, the new inflation-cost-generating mechanism we describe is responsible for about 80 percent of our estimate of the total welfare cost of inflation.

## 1.1 Previous Research

Some of the previous research on the cost of inflation has been conducted using partial equilibrium models. One recent contribution to this literature is Feldstein (1997), which is closely related to our analysis. Much of the rest of the recent work on the cost of inflation is based on general equilibrium models—almost invariably, the infinite-horizon representative agent (IHRA) model, which has become the standard model in applied macroeconomic theory. Research of this type includes Cooley and Hansen (1989), İmrohoroğlu and Prescott (1991), Gomme (1993), Lucas (2000), Haslag (1994), Jones and Manuelli (1995), Dotsey and Ireland (1996), and Lacker and Schreft (1996). Another example is Abel (1997), who presents a general equilibrium adaptation of Feldstein's analysis. For our purposes, Feldstein (1997) and Abel (1997) can serve as the representatives of cost-of-inflation research using partial equilibrium models and IHRA models, respectively.

Feldstein follows most other partial equilibrium investigations of the cost of inflation by assuming that the before-tax real interest rate (or real rate of return on capital) is invariant to policy-induced changes in the inflation rate. According to Feldstein, most of the cost of inflation grows out of the fact that it increases the effective tax rate on capital income and consequently reduces the after-tax real rate of return facing savers. Feldstein also emphasizes a closely related problem, which is that inflation affects the after-tax rates of return on some assets (notably, housing capital) more strongly than others, causing capital to be misallocated.

Feldstein also discusses the effect of inflation on the rate of return on money and the opportunity cost of holding money. This effect has been the main one emphasized by most other contributors to the cost-of-inflation literature. The "monetary cost" of inflation grows out of the fact that an increase in the opportunity cost of holding money causes households to overeconomize on transactions balances, while a reduction in the rate of return on money distorts saving and/or labor-leisure decisions by increasing the opportunity cost of future consumption. In Feldstein's analysis the *net* monetary cost of inflation is actually negative (i.e., a welfare benefit), since an increase in the inflation rate produces an increase in currency seigniorage revenues that allows a reduction in other distorting taxes—a reduction whose welfare benefits exceed the costs just described.[1] Although most other analyses of the monetary cost of inflation produce positive cost estimates, these estimates are uniformly small relative to Feldstein's estimates of the total cost of inflation estimates or to the total cost estimates we present in this paper.

Cost-of-inflation analyses using IHRA models do not assume that the before-tax real interest rate is constant. However, in the standard IHRA model the steady-state value of the *after-tax* real interest rate is essentially invariant to monetary or fiscal policy: It is a function of the exogenous output growth rate plus preference parameters such as the rate of time preference and the intertemporal elasticity of substitution in consumption.[2] For this reason, in Abel's general equilibrium adaptation of Feldstein's analysis, an increase in the inflation rate produces an increase in the before-tax real rate of return on capital that is roughly equal to the decrease

[1] Feldstein's estimates of the monetary cost of inflation are based on empirical evidence suggesting that money demand is not very sensitive to changes in the rate of return on money. As a result, the distortions caused by increasing the inflation tax on money balances are relatively modest, and the resulting increases in the volume of currency seigniorage revenue are relatively large.

[2] In stochastic models, uncertainty about asset returns also plays a role in the determination of real interest rates on both safe and risky assets. In a formal model, changes in the average rate of inflation do not in themselves affect the amount of uncertainty of this type. There is, however, a fairly extensive literature on the empirical relationship between the average level of inflation and the variability of the inflation rate.

in the after-tax real rate of return predicted by Feldstein. The increase in the pretax return rate on capital causes a substantial decline in the capital stock—a decline that reduces the marginal product of capital and produces a lower wage rate. It is this decline in household income, rather than a decrease in the rate of return facing savers, that is responsible for most of Abel's estimate of the welfare cost of inflation. Despite this rather profound difference between the inflation-cost-generating mechanisms postulated by Feldstein and Abel, their estimates of the total cost of inflation are very close to one another.

### 1.2 Our Approach

Although our analysis of the cost of inflation shares a number of important features with the analyses conducted by Feldstein and Abel, it differs from these analyses in one centrally important way: In the Feldstein and Abel models, the change in the real rate of return to capital produced by an increase in the inflation rate must be approximately equal to the implied increase in the effective tax rate on the real return to capital. The after-tax real rate falls (Feldstein) or the before-tax real rate rises (Abel) by this amount, leaving the other rate unchanged. In our model, by contrast, an increase in the inflation rate causes a decrease in both the after-tax real rate of return to capital and the before-tax real rate of return to capital, and it also widens the spread between these two rates. The increase in the spread is equal to the increase in the effective tax rate on capital returns, but the total decline in the after-tax real interest rate is considerably larger. For example, in our baseline case, a 10-percentage-point increase in the inflation rate causes the spread between the before-tax real rate of return on capital and the after-tax real rate of return to increase by approximately 1.8 percentage points. However, the total decline in the after-tax real rate of return is 3.6 percentage points.

The large changes in real interest rates that produce our relatively high inflation-cost estimates are driven by a combination of two features of our model. The first feature involves our assumptions about the role of the government budget constraint—more specifically, about the disposition of the substantial increase in capital income tax revenue that an increase in the inflation rate produces (all else held constant) when the government taxes on a nominal basis. In the Feldstein and Abel analyses, this revenue is used to finance proportional decreases in all direct tax rates. In our analysis, on the other hand, the government uses the increase in capital income tax revenue to reduce the amount it borrows from the public. The resulting decrease in aggregate demand for credit produces a substantial decline in the before-tax real interest rate. This decline is possible because of the second distinctive feature of our analysis: In the general equilibrium model we use, it is possible for the government to change the amount it borrows without adjusting future taxes in a way that produces offsetting shifts in the aggregate supply of credit. We discuss both of these features of our model in detail below.

One interesting result we obtain concerns the implications of attempts to reduce the inflation rate from its postwar-average level of approximately 4 percent to a level of 0 percent or lower. As we have indicated, under normal conditions the government can use debt policy to offset the loss of revenue caused by declines in the inflation rate. However, our analysis implies that once the inflation rate reaches a threshold level, further increases in government borrowing no longer succeed in increasing net government revenue. Additional progress in reducing the inflation rate then requires active cooperation from the fiscal authorities, who must be willing to reduce government expenditures and/or increase tax rates.[3]

Our results indicate that the threshold level of the inflation rate is about 2.5 percent. They also indicate that once this threshold is reached, the welfare benefits from further reductions in the inflation rate are much smaller because of the corresponding need to increase direct tax rates. Thus, our analysis suggests that further reductions in the inflation rate for the U.S. economy are likely to be both much more difficult to achieve, and much less beneficial if achieved, than the reductions that have taken place since the early 1980s.

In the next section we present a more complete description of our approach to analyzing the real effects of changes in monetary policy and the welfare costs of inflation. This section also includes a graphical depiction of our mechanism for generating high inflation costs and a discussion of the empirical plausibility of our approach. In section 3 we lay out

[3] We view changes in government borrowing as constituting a passive cooperation by the fiscal authorities. In the United States, the Treasury Department can (and does) respond to most changes in borrowing requirements without seeking authorization from Congress. Changes in expenditures or tax rates, however, require Congressional action.

the model that we use to obtain our welfare cost estimates. Section 4 describes the procedure we use to select values for the parameters of the model, while section 5 describes the baseline steady state associated with our parameter choices and discusses some of its features. In section 6 we calculate the welfare costs of changes in the inflation rate using our baseline parameterization and some alternative parameterizations. Section 7 discusses some qualifications of our results.

## 2. A NEW APPROACH TO THE WELFARE COST OF INFLATION

### *2.1 Theoretical Principles*

Our objective in this section is to explain in general terms the principles behind the model, as well as the mechanisms at work, and to defend the empirical plausibility of our results before proceeding to the formal model in the next section.

For almost a generation, two-period versions of the OLG model have been widely used for theoretical analyses of the real effects of monetary policy. An important reason for this is that in OLG models, unlike the IHRA models that have been favored by empirically oriented macroeconomists, permanent changes in monetary policy can have large permanent effects on real interest rates and other real variables. In our multi-period model, as in many of its two-period predecessors, permanent changes in monetary policy affect real interest rates by influencing the government's demand for credit. This is possible because the finite lives of OLG households make it possible for government credit demand to rise or fall without producing offsetting adjustments in households' supply of credit. In the case of particular interest to us, this can happen because the government's debt is "unbacked": It does not have to be serviced by a stream of future surpluses. Government debt can be unbacked because OLG models, unlike IHRA models, can have steady-state competitive equilibria in which the real interest rate is lower than the output growth rate. These equilibria were first studied by Samuelson (1958) and Diamond (1965) and are sometimes described as "Samuelson case" equilibria following Gale (1973). (Equilibria with higher real interest rates are often called "classical case" equilibria.[4])

Bullard and Russell (1999) use a calibrated multi-period OLG model—a simpler version of the model developed in this paper—to argue that it is empirically plausible to model the postwar U.S. economy as a Samuelson-case steady state and to view the postwar U.S. government debt as unbacked. Their analysis is grounded on the fact that during the postwar period the average real interest rate on U.S. government debt has been substantially lower than the average U.S. output growth rate and the government's average primary surplus has been approximately zero.[5]

If the government issues debt at a real interest rate lower than the output growth rate, then it can earn "bond seigniorage" revenue in a steady state by extending the debt at the rate necessary to keep its share of output constant.[6] Across steady states of this type, a permanent decrease in the debt stock will shift the aggregate credit demand curve to the left along an unchanged aggregate credit supply curve, causing the steady-state real interest rate to fall. The decline in the real interest rate will have permanent effects on investment, output, and other real variables, including the welfare of households.

**2.1.1 The Real Effects of Monetary Policy.** What does the situation just described have to do with monetary policy or the cost of inflation? The answer to this question begins with a seminal insight of Sargent and Wallace (1981): The government's budget constraint enforces a connection between fiscal policy and monetary policy. For the purposes of their analysis, Sargent and Wallace define monetary policy as consisting of the central bank's choice of a combination of base money and government debt policies that is consistent with the fiscal authority's policies regarding taxes and government spending.[7] They assume that these fiscal policy

---

4 Stochastic versions of IHRA models can have stationary equilibria in which the average real interest rate is lower than the average output growth rate. However, these equilibria do not have most of the other distinctive characteristics of low-real-rate steady states in nonstochastic OLG models.

5 This point also has been made by Darby (1984), among others. During 1948-97, the average ex post real (CPI-deflated) yield on three-month Treasury bills was 1.1 percent. This yield is widely used as an empirical proxy for the risk-free real interest rate: The short term presumably ensures that the ex post real rate is quite close to the ex ante rate and that the premium for interest risk is minimal. Of course, the average real interest rate that the government actually paid was higher, because the average term of the bonds it issued exceeded three months. However, even the average ex post real yields on bonds with terms of ten years or more fall well short of the average output growth rate.

6 The term "bond seigniorage" seems to have been coined by Miller and Sargent (1984).

7 An alternative interpretation of the Sargent-Wallace policy assumptions is that the fiscal authority controls debt policy but conducts it in a way that passively accommodates the fiscal authority's active decisions concerning taxes and spending as well as the monetary authority's

decisions leave the government with a fixed real primary deficit. They show that under this assumption a permanent tightening of monetary policy (a permanent decrease in the base money growth rate) is not feasible and a temporary tightening must eventually cause the inflation rate to rise. The logic behind this result is simple. A monetary tightening reduces the amount of inflation tax revenue and leads to a gradual accumulation of government debt. At some point the debt accumulation must cease and the increased debt service costs must be financed by increased revenue from base money seigniorage.

In the Sargent-Wallace (1981) model the real interest rate was assumed to be fixed at a level above the output growth rate. However, Miller and Sargent (1984) and Wallace (1984) conduct similar analyses in models in which the real interest rate is endogenous, and the former analysis encompasses situations in which it is lower than the real growth rate. Miller and Sargent sketch out a model that implies that, under certain conditions, a permanent tightening of monetary policy can lead to a permanent decrease in the inflation rate and a permanent increase in the real interest rate. The explanation for this result grows out of the fact that when the real interest rate is relatively low the government can earn revenue from bond seigniorage. A decrease in the inflation rate reduces the amount of base money seigniorage revenue and forces the government to issue more debt to increase its revenue from bond seigniorage. The increase in the debt stock causes the real interest rate to rise.[8] It is this situation—one where there is an inverse relationship between inflation rates and real interest rates—that we study in this paper.

We now turn to the welfare consequences of permanent changes in inflation in these equilibria.

**2.1.2 The Cost of Inflation.** Across Samuelson-case steady states of models in which the households are intragenerationally homogeneous, such as Espinosa-Vega and Russell (1998a), policy-induced decreases in the real interest rate are welfare-reducing. When the real interest rate is lower than the output growth rate, the tendency of a lower real interest rate to increase household welfare by increasing the real wage (because a lower real interest rate produces a larger capital stock) is more than offset by its tendency to increase the relative price of future consumption.[9] Thus, models of monetary policy that support Samuelson-case steady states seem to provide a potential contribution to the literature on the cost of inflation.[10]

In the research program that led to the current paper, we began by reconstructing the Espinosa-Vega and Russell analysis in a multi-period calibrated model that allows us to talk credibly about magnitudes. Our previous research—Bullard and Russell (1999)—had established that plausible calibrations of a relatively simple, nonmonetary version of this model produce realistic-looking low-real-interest-rate steady states.[11] In the present paper we find that plausible calibrations of a richer version of the model—a version that includes both money and government debt—support realistic-looking steady states across which permanent increases in the base money growth rate produce higher inflation rates, lower real interest rates, and lower levels of welfare for households.

---

active decisions about the growth rate of the base money supply (which determines government revenue from currency seigniorage). This assumption seems consistent with modern U.S. monetary and fiscal arrangements. Congress and the Administration make active decisions about taxes and spending and the Federal Reserve makes active decisions about the base money growth rate. The Treasury then passively issues the bonds necessary to cover any resulting deficit.

8 The Miller-Sargent (1984) analysis has been refined and extended by Espinosa-Vega and Russell (1998a,b), who study the real interest rate effects of permanent changes in monetary policy across Samuelson-case steady states of a pure exchange model (1998a) and in a model with production and capital (1998b). In both models, a permanent, policy-induced increase in the inflation rate can result in a permanent decrease in the real interest rate on government debt. In the latter model, the increase in the inflation rate is also associated with a permanent increase in the level of output.

9 Part of the reason for this is that in Samuelson-case equilibria, capital is overaccumulated: The marginal return to capital is smaller than the marginal cost of maintaining the capital stock via investment. A reduction in the real interest rate increases the capital stock and increases the severity of this “dynamic inefficiency.”

10 The dynamic inefficiency of the steady states in our model helps explain why our estimates of the size of the Feldstein effect are relatively low. In the Feldstein and Abel models, the fact that increases in the inflation rate increase the effective tax rate on capital income is almost unambiguously welfare-reducing. In our model, in contrast, higher capital income tax rates tend to reduce welfare because they increase the distortion of the saving and labor supply decisions (as in Feldstein and Abel), but they tend to increase welfare because a higher tax rate on capital income tends to reduce the degree of capital overaccumulation.

11 In the postwar U.S. economy the average real rate of return on capital seems to be substantially higher than the average output growth rate. In Bullard and Russell (1999) we show that augmenting the basic model by adding capital income taxes and/or intermediation costs at plausibly calibrated levels can produce realistic-looking steady states in which the real rate of return on capital exceeds the output growth rate but the real interest rate on government bonds is lower. The model in the current paper includes both these features and produces a baseline steady state of this type.

**Figure 1**

**Dependence of Government Revenue on the Real Bond Rate**

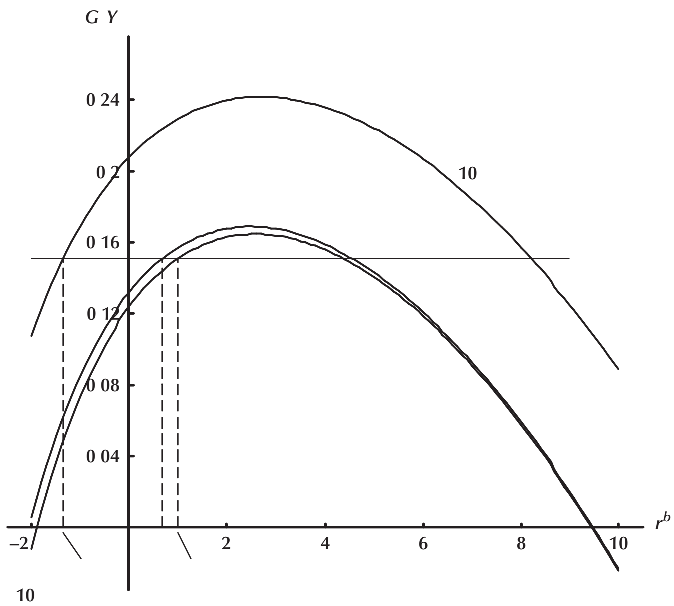


In a basic version of the model presented below—one where we abstract from capital income taxation—the mechanisms work as we have described them but the quantitative real interest rate and welfare effects of moderate increases in the inflation rate are relatively small.[12] The small size of these effects is easily understood. We have seen in the models that, when the inflation rate increases, the government earns additional revenue from currency seigniorage and consequently needs less revenue from bond seigniorage. As a result, the government needs to borrow less, the real stock of bonds declines, and the real interest rate falls to allow private debt to replace some of these bonds in the portfolios of savers. However, in the U.S. economy the stock of base money is quite small relative to output, so the ratio of base money seigniorage to output is also quite small. As a result, in a plausibly calibrated model a moderate increase in the inflation rate will produce a relatively small increase in the amount of revenue from currency seigniorage and thus a relatively small decrease in the amount of bond seigniorage. This small decrease in bond seigniorage can be produced by a small decrease in the real stock of government bonds and thus a small decline in the real interest rate on these bonds.

However, when nominal capital income taxation is included in the model, as it is in this paper, these effects are greatly magnified—they are an order of magnitude larger, in fact. As we have seen, the formal analyses of this argument conducted by Feldstein (1997) and Abel (1997) are based primarily on the fact that increases in effective capital income tax rates increase the spread between the before-tax and after-tax real return rates on capital. The revenue implications of these tax rate increases are a distinctly secondary consideration in those papers. In our analysis, in contrast, the increase in capital income tax revenue that occurs when the inflation rate rises is the principal driving force behind our results: Real interest rates have to fall substantially to restore equilibrium. Under plausible calibrations, this revenue increase is much larger than the inflation-induced increase in revenue from currency seigniorage. It is large enough to cause a substantial decline in the real interest rate and a large decrease in household welfare.

[12] Although the changes in the real interest rate that are produced are small relative to our main results in this paper, the welfare cost these changes produce are about as large as most of the inflation-cost estimates that appear in the literature.

## 2.2 A Graphical Description

A simple graphical description of the mechanism we describe is presented in Figure 1. The lowest curve in the figure describes the ratio of total government revenue to output as a function of the before-tax real interest rate on government debt (hereafter, the real bond rate) in the baseline parameterization of our model. The inflation rate is held fixed at its baseline value of 4 percent. The tax rates on labor and capital income are also held fixed at their baseline values.

One might expect the government revenue curve to be uniformly downward-sloping, since lower real interest rates are associated with higher real wages and the U.S. government gets most of its revenue from labor income taxes. Since the U.S. tax system taxes net rather than gross capital income, however, decreases in the real interest rate cause capital income tax revenues to decline sharply. In addition, at relatively low real bond rates, further declines in the rate produce substantial declines in the revenue from bond seigniorage. As a result, there exists a real bond rate below which further decreases in the rate cause the revenue-output ratio to fall. In our baseline calibration, this rate is roughly 2.5 percent.

Equilibrium in the model occurs at a real bond rate at which the ratio of government revenue to output is equal to the government's target for the expenditures share of output. The expenditures target is indicated by the horizontal line in the figure. The equilibrium real bond rate on which we focus is the one on the left side of the government revenue curve: Its value is 1 percent. Although there is an alternative equilibrium rate on the right side of the curve, its level is counterfactually high and it produces implausible values for other endogenous variables.

Now suppose that the monetary authority increases the base money growth rate, and thus the inflation rate, by 10 percentage points. If capital income tax revenues were indexed to inflation, this would affect total government revenue only by increasing the revenue from currency seigniorage. The increase in revenue would be small and the government revenue curve would shift upward by a small amount, producing the curve in Figure 1 that lies just above the baseline curve. The new equilibrium value of the real bond rate would be slightly lower (roughly 0.3 percentage points) than the original value.

Under the actual U.S. tax system, however, the increase in the inflation rate produces a relatively large increase in the amount of capital income tax revenue. As a result, the total revenue curve shifts upward by a relatively large amount, producing the curve in Figure 1 that lies well above the baseline curve. The left side of the new curve intersects the expenditures target line at a point well to the left of the original intersection point, indicating a decrease of 2.4 percentage points in the before-tax real interest rate on government debt. The after-tax real bond rate falls by a larger amount, 3.6 percentage points, because the increase in the effective tax rate on capital income increases the size of the wedge between the two rates. The spread between the before-tax and after-tax values of the real rate of return on physical capital widens by 1.8 percentage points. As we have indicated, however, only a small part of our estimate of the welfare cost of inflation is attributable to the increase in these rate spreads.

## 2.3 Empirical Plausibility

Is it plausible to believe that permanent increases in the inflation rate can produce substantial permanent declines in before-tax real interest rates? During the past few years, a number of authors have used recent developments in time series analysis to study the long-run relationship between the money growth rate and/or the inflation rate and the levels of real variables such as the real interest rate, output, and investment. The papers in question include King and Watson (1992, 1997), Weber (1994), Bullard and Keating (1995), Ahmed and Rogers (1996), Serletis and Koustas (1998), Koustas (1998), and Rapach (2003). As in the case of most literature that conducts empirical tests of propositions from macroeconomic theory, the results reported are mixed and the evidence cannot be regarded as conclusive. Nevertheless, the literature provides plenty of evidence that indicates that money may not be long-run superneutral and that the direction and magnitude of the long-run effects of inflation on real variables may be consistent with the implications of our model. A key aspect of these papers is that they distinguish permanent from temporary movements in nominal variables. Since the experiment we study (and that most of the inflation cost literature studies) is a permanent change in the inflation rate that alters the steady state of the model, these studies provide the appropriate empirical counterpart.

King and Watson (1992, 1997) use postwar U.S. data to study the long-run relationship between the money growth rate and the level of output, and between the inflation rate and nominal interest rate, under a range of alternative identifying assumptions. They find that there is a broad range of plausible identifying assumptions under which (i) the hypothesis that money is superneutral can be rejected and (ii) the relationship between the rate of money growth and the level of output is significantly positive. When they assume that money is contemporaneously exogenous—the most common identifying assumption in this literature—they find that a permanent 1-percentage-point increase in the money growth rate tends to increase the level of output by 3.8 percent. As we shall see in section 5, this is almost identical to the percentage increase in output that results, in our model, from a permanent 1-percentage-point increase in the money growth and inflation rates starting from the baseline inflation rate. Under the analogous identifying assumption regarding inflation and nominal interest rates, the authors find that the simple Fisher relationship can be rejected easily and that a 1-percentage-point increase in the inflation rate tends to reduce the real interest rate by more than 80 basis points. As we have seen, this estimate is roughly twice the size of the effect our model produces in the vicinity of the baseline steady state. Alternative identifying assumptions produce smaller estimates, however. Weber (1994) uses similar data and methodology to study the relationships between money growth and output/interest rates in the Group of Seven countries. He finds that superneutrality can be rejected, under a wide range of alternative identifying assumptions, for all of these countries except France. Weber also finds that the data for the United States and the United Kingdom are strongly inconsistent with the simple Fisher relationship, with increases in the inflation rate producing substantial decreases in the real interest rate. The evidence for the other countries is less conclusive.[13]

Bullard and Keating (1995) use bivariate models to study the long-run relationship between the rate of inflation and the level of output using postwar data for 58 countries. They find evidence of statistically significant departures from superneutrality for a number of developed countries with relatively low average inflation rates, though not for the United States. In each of these countries an increase in the inflation rate is associated with an increase in the level of output.[14] Rapach (2003) uses trivariate models to study the long-run relationship between the inflation rate, the real interest rate, and the level of output in 14 OECD countries. He finds that for each of these countries, increases in the inflation rate produce statistically significant decreases in the real interest rate. For seven countries, higher inflation rates produce statistically significant increases in the level of output. For the other seven countries (including the United States) the output responses vary in sign but are not statistically significant.

Ahmed and Rogers (1996) use vector autoregressive methods to study the long-run relationship between the inflation rate and the levels of output, consumption, and investment using U.S. data for the past 100 years. They find that permanent increases in the inflation rate are associated with large permanent increases (decreases) in the ratio of investment (consumption) to output. They view these results as inconsistent with superneutrality or the simple Fisher relationship. Their point estimate is that a permanent, 1-percentage-point increase in the inflation rate causes the consumption share of output to fall by 2.5 percentage points and the investment share of output to rise by 1.0 percentage points.[15] The effects generated by our model are qualitatively similar, although the magnitudes are somewhat different. In our model, a permanent increase in the inflation rate from 4 percent to 5 percent causes the consumption share of output to fall by 1.6 percent and the investment share to rise by 1.3 percent.

## 3. A GENERAL EQUILIBRIUM LIFE-CYCLE MODEL

### *3.1 Overview*

Our model can be succinctly described as the result of a hypothetical meeting between Sargent and Wallace and Auerbach and Kotlikoff. Auerbach

[13] There is a large literature that uses more traditional econometric methods to investigate the relationship between ex ante real return rates on assets and the expected rate of inflation. Most of these studies find that these variables are strongly negatively correlated in the short run and/or the long run. A good example is Huizinga and Mishkin (1984). See Marshall (1992) and Boyd, Levine, and Smith (2001) for more complete descriptions of this literature.

[14] For Germany, Austria, and the United Kingdom, the estimated effects are clearly significant. For Japan and Spain, the estimated effects are large but only marginally significant.

[15] When Ahmed and Rogers (1996) confine their analysis to data from the postwar period, they find that the departures from superneutrality are qualitatively similar, and remain statistically significant, but are much smaller in size.

and Kotlikoff (1987) were pioneers in the use of multi-period OLG models to study issues in public finance, but neither they nor their successors have used these models to study issues in monetary economics.[16] The principal differences between our model and the Auerbach-Kotlikoff model are that our model includes monetary elements and that it allows for productivity growth and capital depreciation. As we have noted in the previous section, Sargent and Wallace (1981, 1982, 1985) were pioneers in the use of two-period OLG models to study questions in monetary theory and policy. We have adopted many aspects of their approach, perhaps the most important of which is their emphasis on the role of the government budget constraint in helping determine the real effects of changes in monetary policy. In addition, since the source of money demand in our model is a reserve requirement, our analysis might be considered an application of the legal restrictions theory of money that was developed by Wallace (1983, 1988) and has been applied repeatedly by Sargent and Wallace. However, we do not expect readers to take our reserve requirement assumptions seriously as a deep theory of the demand for money: Instead we view them as providing a proxy for money demand from all sources. In this sense, our money demand specification is similar to the cash-in-advance specifications that are common in the cost-of-inflation literature.[17]

## 3.2 Primitives

**3.2.1 Demographics.** A generation of identical households is born at each discrete date $t = \ldots,-2,-1,0,1,2,\ldots$ and lives for $n$ periods. Successive generations of households are identified by their birthdates and differ from each other only in their populations, which grow at gross rate $\psi \geq 1$ per period. Each household is endowed with a single, perfectly divisible unit of time per period and must allocate this time unit between labor and leisure. A household in its $i$th period of life has an effective labor productivity coefficient $e_i$, $i = 1,\ldots,n$. If this household supplies $l$ units of labor during this period, then its effective labor supply is $e_i l$.

**3.2.2 Goods and Technologies.** There is a single good. Units of the good available at a given date can be consumed or stored. If stored they are called "capital goods" and can be used in production during the following period. Capital goods depreciate at a *net* rate of $\delta \in [0,1]$ per period whether or not they are used in production.

At each date, an arbitrary number of competitive firms have access to a technology that uses capital and effective labor to produce the consumption good. The aggregate stock of capital goods available for use in production at the beginning of date $t$ is denoted $K(t)$. Since the technology exhibits constant returns, it suffices to describe the aggregate production function

$$Y(t) = \lambda^{(1-\alpha)(t-1)} K(t)^{\alpha} L(t)^{1-\alpha}, \quad (1)$$

where $L(t)$ is the aggregate supply of effective labor and $k(t) \equiv K(t)/L(t)$ is the ratio of capital to effective labor. The parameter $\lambda \geq 1$ is the gross rate of growth of labor productivity, and the parameter $\alpha \in (0,1)$ governs the capital share of output.

**3.2.3 Preferences.** A household's preferences are defined over intertemporal bundles that include the quantities of the single good that it consumes during each period of its life and the quantity of leisure enjoyed in each period. The consumption and leisure choices of a member of generation $t$ at date $t+j$ are denoted $c_t(t+j)$ and $\ell_t(t+j)$, $j = 0,\ldots,n-1$, respectively.

The preferences of the households are described by the standard utility function

$$U\left(\{c_t(t+j), \ell_t(t+j)\}_{j=0}^{n-1}\right) = \sum_{j=0}^{n-1} \frac{\beta^j}{1-\gamma}\left[c_t(t+j)^{\eta}\,\ell_t(t+j)^{1-\eta}\right]^{1-\gamma}, \quad (2)$$

where $\gamma > 0$, $\eta > 0$, and $\beta \equiv 1/(1+\rho)$, with $\rho > -1$. We require $c_t(t+j) \geq 0$ and $\ell_t(t+j) \in [0,1]$ for all $j = 0,\ldots,n-1$ and for all $t$. These preferences imply that households' elasticity of intertemporal substitution in consumption, $\sigma$, is the reciprocal of their coefficient of relative risk aversion $v$, where $v \equiv 1-\eta(1-\gamma)$.

Since this is an OLG model, no restrictions need be placed on the value of $\rho$ other than $\rho > -1$. Although $\rho$ is often described as the "(pure) rate of time preference," a more meaningful measure of the households' relative valuation of current versus future consumption is their *effective time preference rate* $\varphi \equiv 1-(1+\rho)^{-1/\gamma}$. If a household with an effective time preference rate of $\varphi$ is faced with a zero net real interest rate, then it will choose an average

[16] An exception is Altig and Carlstrom (1991).

[17] See Haslag (1995) for a discussion of some fairly general conditions under which cash-in-advance and reserve requirement economies are allocationally equivalent.

lifetime consumption growth rate of approximately $-\varphi$.[18]

The parameter $\eta$ is the elasticity of intratemporal substitution of consumption and leisure. It is also the dominant parameter governing the share of households' time that they devote to providing labor. We require households to retire from the labor force at an age $n^* \leq n$; that is, $\ell_t(t+j) = 1$ for $j = n^* - 1, \ldots, n-1$.

## 3.3 Markets

**3.3.1 Inputs Markets.** The firms rent capital and hire effective labor at rental and wage rates equal to these inputs' respective marginal products. The rental rate is $r(t) = \alpha\lambda^{(1-\alpha)(t-1)}k(t)^{\alpha-1}$ and the wage rate is $w(t) = \lambda^{(1-\alpha)(t-1)}(1-\alpha)k(t)^{\alpha}$.

**3.3.2 Asset Markets.** There are four basic types of assets in the model: physical capital, consumption loans, government (consumption) bonds, and fiat currency. For purposes of simplicity, all loans and bonds are assumed to have terms of one period. We assume that households hold any assets other than government bonds indirectly: Their direct holdings consist of deposits issued by perfectly competitive financial intermediaries. The intermediaries use these deposits to make loans to households, to purchase capital goods from households (in order to rent them to firms next period), and to purchase fiat currency. The financial intermediaries incur a constant real cost per unit of goods lent or used to acquire capital. This cost, which is denoted $\xi \in (0,1)$, is assumed to be incurred during the period when the loans are repaid or the capital goods are recovered from storage. Households purchase bonds directly from the government without incurring any transactions costs.[19]

The intermediation cost assumption is intended to act as a crude proxy for the costs associated with risky private lending—including, perhaps, a risk premium. At the calibration stage, the assumption helps us do a better job of mimicking the observed structure of interest rates on private and government liabilities.[20] It has little effect on our estimates of the welfare cost of inflation (see section 6.3.1).

Financial intermediaries are legally required to hold a minimum fraction $\phi \in (0,1)$ of their liabilities in the form of real balances of fiat currency (reserves). We confine ourselves to the study of equilibria in which fiat currency is not return-competitive, so that households do not hold it directly and intermediaries do not hold excess reserves. As we have indicated, however, we use reserve demand as a proxy for base money demand from all sources.

**3.3.3 Government Intervention in Markets.** The government in the model is a consolidated federal, state, and local entity. At date $t$ the government must finance a real expenditure of $G(t) \geq 0$ by a combination of direct taxation and seigniorage. The goods that comprise the expenditure are assumed to leave the economy. We will assume that the level of the expenditure must be constant relative to output: that is, the government chooses $G(t)$ so that $G(t)/Y(t) = g$ for some $g \geq 0$. Thus, the expenditure must grow at the same rate as output, which is $\lambda\psi$ (gross) per period in a steady state.

The government issues two types of liabilities: fiat currency and one-period consumption bonds. The nominal quantity of fiat currency outstanding at the end of date $t$ is denoted $H(t)$. The currency price of a unit of the consumption good at date $t$ is denoted $P(t)$. Aggregate real currency balances at date $t$ are $M(t) = H(t)/P(t)$. The real rate of return on currency balances is $R^h(t) = P(t)/P(t+1)$ and the gross inflation rate is $\Pi(t) = 1/R^h(t) = P(t+1)/P(t)$. We assume that the government issues just enough additional fiat currency, each period, to allow the total nominal stock to grow at a constant rate: $H(t) = \theta H(t-1)$ for some $\theta \geq 1$. Nominal revenue from currency seigniorage at date $t$ is simply $H(t) - H(t-1) = (1 - {}^1/\theta)\,H(t)$. Real revenue from this source is $S^c(t) \equiv M(t) - R^h(t-1)\,M(t-1) = (1 - {}^1/\theta)\,M(t)$.

The aggregate real market value of the consumption bonds issued during period $t$ is denoted $B(t)$. The gross real interest rate that the government pays on these bonds is denoted $R^b(t)$. Real government revenue from bond seigniorage (net extension of real indebtedness) is $S^b(t) \equiv B(t) - R^b(t-1)\,B(t-1)$.[21]

[18] The easiest way to see this is to consider a two-period endowment model with inelastic labor supply and preferences given by $u = \Sigma_{i=0}^{1}(1-\gamma)^{-1}\beta^i c_t(t+i)^{1-\gamma}$. The first-order conditions imply $c_t(t+1) = (\beta R(t))^{1/\gamma}c_t(t)$, where $R(t)$ is the gross real interest rate between dates $t$ and $t+1$. In this case the effective time preference rate $\varphi$ is exactly equal to the arithmetic inverse of the consumption growth rate the household chooses when faced with a zero net real interest rate. Thus, $\varphi = -(c_t(t+1)/c_t(t) - 1)|_{R(t)=1}$, which, from the first-order condition, is $\varphi = 1 - \beta^{1/\gamma}$, or $1 - (1+\rho)^{-1/\gamma}$.

[19] We assume that households who attempt to make private loans directly would face prohibitively high transactions costs.

[20] Recent work on financial intermediation in macroeconomic models includes Boyd and Smith (1998), Bernanke and Gertler (1989), Greenwood and Williamson (1989), Williamson (1987), and Díaz-Giménez et al. (1992). Our approach is most closely related to the latter paper.

[21] In official government budget statistics, bond seigniorage is not regarded as revenue. Thus, the empirical analog of bond seigniorage revenue is the government budget deficit.

The government collects the bulk of its revenue using three direct proportional taxes: a tax on real labor income levied at rate $\tau^w \in [0,1)$, a tax on the nominal interest income of households levied at rate $\tau^i \in [0,1)$, and a "corporate income tax" on firms' nominal returns to capital levied at rate $\tau^c \in [0,1)$. Aggregate real labor income tax revenue at date $t$ is $T^l(t) \equiv \tau^w w(t)L(t)$. The government taxes the net nominal interest that households receive on government bonds or intermediary deposits. Nominal revenue from this source is $\tau^i[R^d(t-1)/R^h(t-1) - 1]$ $A^+(t-1)P(t-1)$, where $A^+(t)$ represents households' aggregate gross real asset holdings at date $t$ and $R^d$ is the gross real rate of return to deposits. Real revenue from interest income taxation is $T^i(t)$ $\equiv \tau^i[R^d(t-1) - R^h(t-1)]A^+(t-1)$. The government also taxes the net nominal returns paid by the firms to the financial intermediaries, after adjustment for depreciation. Thus, the corporate income tax produces nominal revenue of $\tau^c(R^{kn}(t-1)/R^h(t-1)-1)$ $K(t)P(t-1)$, where $R^{kn}$ is the gross real rate of return to capital net of depreciation, $R^{kn}(t) = R^k(t) - \delta$, $R^k$ is the gross real rate of return to capital, $R^k(t) = 1 + r(t+1)$, and $r(t+1)$ is the marginal product of capital. The real revenue from corporate income taxation is

$$T^k(t) \equiv \tau^c\left(R^{kn}(t-1) - R^h(t-1)\right)K(t).$$

Our tax structure is intended to provide a crude but parsimonious representation of the current U.S. tax system. This representation captures two important features of the U.S. system for taxing capital income: double taxation of dividend income and the fact that household income from interest and capital gains is taxed on a nominal basis. By levying the labor tax on real labor income, we are taking a conservative approach so that we do not overstate the impact of inflation on welfare. We will match revenues from taxes on capital to the data on the sources of government revenue, and then we will allow all remaining government revenue to come from labor income taxation.

The government budget constraint can be written

(3) $$gY(t) = T^l(t) + T^i(t) + T^k(t) + S^c(t) + S^b(t).$$

## 3.4 Market Clearing

**3.4.1 The Structure of Real Interest Rates.** Our assumptions about money demand, financial intermediation, and capital income taxation determine the interest rate structure of our economy. In the equilibria that we study, the lowest gross real return rate in our economy is $R^h(t)$, the gross real rate of return on currency. The highest gross real rate of return is the gross pre-depreciation return rate on capital, which is $R^k(t) = 1 + r(t+1)$, where $r(t+1)$ is the marginal product of capital. The gross real rate of return on capital, net of depreciation, is $R^{kn}(t) \equiv R^k(t) - \delta$. Since firms are taxed on their nominal net-of-depreciation returns to capital, the after-tax gross real rate of return they pay to the financial intermediaries is $R^{ka}(t) \equiv (1-\tau^c)R^{kn}(t) + \tau^c R^h(t)$.[22] Arbitrage implies that $R^{cl}(t)$, the gross real rate of return that the intermediary receives on consumption loans, must be equal to $R^{ka}(t)$.

Because financial intermediation is costly, intermediaries are not willing to pay $R^{ka}(t)$ to depositors. The real intermediation cost is $\xi$ per unit of loans or capital intermediated, so the gross real return rate net of this cost is $R^{kc}(t) \equiv R^{ka}(t) - \xi$. In addition, intermediaries must allocate a fraction $\phi$ of their deposits to the acquisition of fiat currency reserves. They pay household depositors a reserve-ratio-weighted average of the real return rate on fiat currency and the real return rate on loans net of intermediation costs: This rate is $R^d(t) \equiv (1-\phi)$ $R^{kc}(t) + \phi R^h(t)$. Since government bonds are not intermediated, arbitrage implies that their gross real interest rate, which we denote $R^b(t)$, must be equal to $R^d(t)$. Finally, households must pay taxes on their nominal interest income at a rate of $\tau^i$, so the gross after-tax real rate of return on deposits is $R^{da}(t) \equiv (1-\tau^i)R^{da}(t) + \tau^i R^h(t)$. Since the government taxes interest income from all sources, $R^{da}(t)$ must be equal to $R^{ba}(t)$, the gross after-tax real interest rate on government bonds.

To summarize, our steady-state asset return structure obeys the following chain of inequalities:

(4)
$$R^h < R^{da} = R^{ba} < R^d = R^b < R^{kc} < R^{cl} = R^{ka} < R^{kn} < R^k.$$

The associated gross nominal interest rates are equal to the gross real rates divided by $R^h$. The monetary authority determines $R^h$ through its conduct of monetary policy. The equilibrium conditions of the model can be thought of as determining the equilib-

[22] The gross nominal return to capital employed at date $t-1$ is $R^{kn}(t-1)/R^h(t-1)$, so at date $t$ the firms have to make a nominal tax payment to the government of $\tau^c((R^{kn}(t-1)/R^h(t-1))-1)K(t-1)P(t-1)$. Firms' total real net-of-depreciation earnings are $R^{kn}(t-1)K(t-1)$. Dividing the former expression by $P(t)$ to put it into real terms and then subtracting it from the latter expression produces the gross rate of return expression given in the text. A similar calculation defines $R^{da}$ below.

rium value of $R^k$, and the capital income tax, intermediation cost, and money demand parameters then determine the remainder of the interest rate structure.[23]

**3.4.2 Household Decisions.** Households maximize (2) subject to a lifetime budget constraint that we will now define. We will let $a_t(t+j)$, $j=0$, …, $n-1$, denote the demand for assets, at date $t+j$, of a household born at date $t$. Households can borrow or lend in any period of life. If they borrow at date $t$, then they pay the gross real rate $R^{ka}(t)$. If they lend by holding deposits with the financial intermediary, then they earn the gross real after-tax return $R^{da}(t)$. The budget constraints of an agent born at date $t$ are

(5)

$$c_t(t)+a_t(t)=(1-\tau^w)w(t)e_1\left(1-\ell_t(t)\right),$$

$$c_t(t+j)+a_t(t+j)=(1-\tau^w)w(t+j)e_{j+1}\left(1-\ell_t(t+j)\right)$$
$$+\hat{R}(t+j-1)a_t(t+j-1)$$

for $j=1,\ldots,n-2$ and

$$c_t(t+n-1)$$
$$=(1-\tau^w)w(t+n-1)e_n\left(1-\ell_t(t+n-1)\right)$$
$$+\hat{R}(t+n-2)a_t(t+n-2),$$

where $w(t)$ is the before-tax real wage at date $t$, and

$$\hat{R}(t+j)=\begin{cases}R^{ka}(t+j) & \text{if} \quad a_t(t+j)<0,\\ R^{da}(t+j) & \text{if} \quad a_t(t+j)\geq 0.\end{cases}$$

Aggregate net household asset holdings at date $t$ are

$$A(t)\equiv\sum_{j=0}^{n-2}\psi^{t-j}a_{t-j}(t),$$

where the population at date $t$ has been normalized to unity. We can decompose $A(t)$ into $A^+(t)$, the net aggregate asset demand of the households whose net asset demand is non-negative (net creditors) at date $t$, and $A^-(t)$, the net aggregate asset demand of the agents whose net asset demand is negative (net debtors) at date $t$. The first group earns a gross rate of return $R^{da}(t)=R^{ba}(t)$, while the second group must pay interest at gross rate $R^{ka}(t)$. We will define $A^-(t)$ as a positive number, so that $A(t)=A^+(t)-A^-(t)$. The liabilities of the financial intermediaries are then $K(t+1)+A^-(t)$. A household's asset demand problem can have a corner solution for a particular life period $j$, in which case the household sets $a_t(t+j)=0$.

**3.4.3 Equilibria.** Households hold aggregate real deposits $D(t)$ with financial intermediaries. The intermediaries use a fraction of these deposits to acquire fiat currency reserves: $\phi D(t)=M(t)$. The remainder of the deposits are lent to firms and households: $(1-\phi)D(t)=A^-(t)+K(t+1)$. The money-market clearing condition is

$$M(t)=\frac{\phi}{1-\phi}\left[A^-(t)+K(t+1)\right]. \tag{6}$$

The credit-market clearing condition is

$$A(t)=M(t)+K(t+1)+B(t), \tag{7}$$

since some household assets may take the form of government debt. These last two conditions can be combined, producing the condition

$$B(t)=A^+(t)-\frac{1}{1-\phi}\left[A^-(t)+K(t+1)\right]. \tag{8}$$

The marginal product condition for the capital rental rate at date $t$ can be used to express the capital-labor ratio $k(t+1)$ as a function of the rental rate $r(t+1)$ and thus of the gross pretax, pre-depreciation capital return rate $R^k(t)$. The value of $k(t)$, which depends on $R^k(t-1)$, determines the equilibrium wage rate $w(t)$ through the marginal product of labor condition. The values of $R^h(t)$ and $R^k(t)$ determine the structure of real return rates, and together with the values of $w(t)$ they provide the data necessary for the households to solve their decision problems. The households' leisure choices imply values for $L(t)$, the aggregate supply of effective labor, and thus for $K(t)=k(t)L(t)$.

Equation (3), the government budget constraint, involves $B(t)$ and $M(t)$. We can substitute equations (6) and (8) into this constraint. As we have seen, the quantities of labor and capital income tax revenue that appear in the government budget constraint also depend on the wage rate, the household labor supply decision, and the household asset demand decisions. After we make these substitutions, the consolidated government budget constraint becomes

[23] In practice, we fix $R^b$ and $R^h$ at their postwar U.S. averages and find a parameterization of the model that has an equilibrium that supports these values, as we will discuss. This parameterization determines the rest of the interest rate structure.

(9)

$$\begin{aligned} gY(t) = {} & \tau^{w} w(t)L(t) + \\ & A^{+}(t-1)\left\{\tau^{i}[R^{d}(t-1) - R^{h}(t-1)] - R^{b}(t)\right\} + \\ & \frac{A^{-}(t-1)}{1-\phi}\left\{R^{b}(t) - \phi R^{h}(t-1)\right\} + \\ & K(t)\left\{ \begin{matrix} \tau^{c}\left(R^{kn}(t-1) - R^{h}(t-1)\right) + \\ \frac{1}{1-\phi}\left(R^{b}(t) - \phi R^{h}(t-1)\right) \end{matrix} \right\} + \\ & A^{+}(t) - A^{-}(t) - K(t+1). \end{aligned}$$

In equilibrium, the level of output at date $t$ depends on real returns prevailing at date $t-1$, which determine the capital-labor ratio $k(t)$, and $L(t)$, which determines the capital stock given $k(t)$. Thus, given the solution to the firms' and households' decision problems, the right-hand side of the resulting equation can be written entirely as a function of $\{R^h(t)\}$, $\{R^k(t)\}$, and the parameters of the model:

$$gY(t) = Z\left[\{R^{h}(t)\}, \{R^{k}(t)\}, F, \Delta\right], \tag{10}$$

where $F$ and $\Delta$ are sets of parameters defined below.

It is readily seen that if $R^h(t)$ and $R^k(t)$ are date-invariant then the value of the function $Z$ will grow at a rate equal to the steady-state output growth rate. Consequently, the steady-state version of the equilibrium budget constraint can be written

$$g = \zeta\left[R^{h}, R^{k}, F, \Delta\right].$$

We will confine ourselves to the study of steady-state equilibria.

As we indicated in Figure 1, if $R^h$ (and thus the inflation rate) and the parameters of the model are held fixed, then the function $\zeta$ takes the form of a downward-opening paraboloid in $R^k$ or (equivalently) $R^b$. Provided the ratio $G/Y$ is low enough, there will consequently be two steady states, one associated with a relatively low real interest rate and the other associated with a relatively high one. In this paper we focus on the steady states that are associated with relatively low real interest rates, which provide a better match for the data.[24]

In steady-state equilibria, the capital-labor ratio and the real wage grow at a gross rate of $\lambda$, while the levels of real aggregates such as output, asset demand, the capital stock, and real balances grow at a gross rate of $\lambda\psi$. Money prices will rise at a gross rate of $\theta/\lambda\psi$, where $\theta$ is the gross growth rate of the stock of fiat currency. Thus, the steady states we study are quantity-theoretic in the sense that the rate of money growth dictates the inflation rate through the standard quantity theory equation.

It is possible, of course, that for given values of the inflation rate, the direct tax rates, and the other parameters of the model, equation (10) will not have real solutions. To fit our model to the data, we need to solve the agents' decision problems and equation (10) repeatedly for a wide variety of parameter values. For this reason, our approach to calculating steady states of calibrated versions of our model is slightly different from the approach implicitly described by the discussion presented above. When we search for plausible values of our model's "deep parameters" we hold $R^h$, $R^b$, and the other parameters of the model fixed at values suggested by the data, and we treat the revenue-output ratio $g$ as endogenous. We then adjust the values of the deep parameters in an attempt to match data-derived targets for many of the model's endogenous variables—including a target for $g$. We now turn to describe this calibration process in detail.

## 4. CONFRONTING THE DATA

### *4.1 Calibration Strategy*

Our goal is to find a specification of the model that is plausible in two senses: in the sense that the values of the parameters are not out of line with published estimates and/or values used elsewhere in the calibration literature, and in the sense that the steady-state values of endogenous variables provide a convincing match for the data. We divide the parameters of our model into two sets. The first set, which we refer to as the *fixed parameter vector F*, consists of parameters whose values are not very controversial because they map into the data in a simple, direct way. We choose values for these parameters, element by element, based on data for the postwar period.

The second set of parameters, the *deep parameter vector* $\Delta$, consists of parameters that do not map into the data directly. Although many of these parameters appear in other calibrated general equilibrium models, the question of their appropriate values is unsettled and controversial, largely because empirical estimates vary widely from study to study. Our

[24] For most of the calibrated specifications we study, the alternative steady state produces unrealistically high values for the rate of return on government debt and the debt-to-GDP ratio.

**Table 1**

**Targets on Observable Quantities**

| Endogenous variable | Target | Range | Source |
|---|---|---|---|
| $K/Y$ | 3.32 | [2.32, 4.32] | Cooley and Prescott (1995) |
| $I/K$ | 0.076 | [0.066, 0.086] | Cooley and Prescott (1995) |
| $B/Y$ | 0.47 | [0.255, 0.686] | U.S. data, 1959-94 |
| $hcg$ | 0.015 | [0.01, 0.03] | Laitner (1992) |
| $alt$ | 0.154 | [0.075, 0.33] | Authors' calculations |
| $I^m/Y$ | 0.06 | [0.05, 0.07] | Díaz-Giménez et al. (1992) |
| $M/Y$ | 0.0592 | [0.041, 0.078] | U.S. data, 1959-94 |
| $g$ | 0.151 | [0.121, 0.18] | U.S. data, 1959-94 |
| $T^k/G$ | 0.119 | [0.036, 0.201] | U.S. data, 1959-94 |

SOURCE: U.S. data were obtained from the 1996 *Economic Report of the President*.

approach to setting these parameters is somewhat novel: We select them jointly, using an iterative nonlinear optimizing procedure (a genetic algorithm) to find a vector of values for the deep parameters that supports a steady state that produces endogenous variables whose values come as close as possible to a vector of targets based on postwar data. The target variables we select are natural ones from the perspective of our model.

## 4.2 The Fixed Parameter Vector

The fixed parameter vector is

$$F = \left[\{e_i\}_{i=1}^{n}, n^*, \lambda, \psi, R^h, R^b, \tau^i\right]. \tag{11}$$

Thus, the fixed parameters are the labor efficiency profile $\{e_i\}_{i=1}^{n}$, the age of retirement $n^*$, the gross productivity growth rate $\lambda$, the gross labor force growth rate $\psi$, the gross real return rate on money $R^h$ (or, equivalently, the gross inflation rate $\Pi = 1/R^h$), the gross before-tax real interest rate on government bonds $R^b$, and the tax rate on interest income $\tau^i$.

We set $n^*$, the retirement age, to household age 44 (figurative age 65). We set the gross productivity growth rate at $\lambda = 1.015$, the gross rate of labor force growth at $\psi = 1.017$, the gross real rate of return on currency at $R^h = 0.9615$ (implying a 4 percent inflation rate), and the before-tax gross rate of return on government bonds at $R^b = 1.01$. All four of these values are based on postwar data. The choices of $R^h$, $\lambda$, and $\psi$ imply a value for $\theta$, the base money supply growth rate, that satisfies $R^h = \lambda\psi/\theta$.

Since our estimate of the postwar-average after-tax gross real return rate on federal government bonds is close to zero, we seek a parameterization that produces $R^{ba} = 1$.[25] The relationship between $R^b$ and $R^{ba}$ then requires us to set the interest income tax rate $\tau^i$ at 0.2. Our life-cycle labor productivity profile is based on estimates constructed by Hansen (1993).[26]

## 4.3 The Deep Parameter Vector

The deep parameter vector is

$$\Delta = \left[\rho, \gamma, \eta, \alpha, \delta, \xi, \phi, \tau^w, \tau^c\right]. \tag{12}$$

Thus, the deep parameters are the pure rate of time preference $\rho$ (or the discount factor $\beta = [1+\rho]^{-1}$), the indifference-curve convexity parameter $\gamma$, which helps determine the intertemporal consumption substitution elasticity $\sigma = [1-\eta(1-\gamma)]^{-1}$, the intratemporal labor-leisure substitution elasticity (or labor share of time) $\eta$, the capital share of output $\alpha$, the net depreciation rate $\delta$, the unit intermediation

25 We constructed this estimate using marginal tax rate data provided by Joseph Peek of Boston College. We thank him for his cooperation.

26 The Hansen data are collected from samples taken in 1979 and 1987. The data separate males from females. We average the data from the two years, and we also average the data across males and females using weights of 0.6 and 0.4. The resulting profile is a step function, because the data are collected for age groupings. We fit a fifth-order polynomial to this step function. This yields the smooth profile $e_{i-20} = m_0 + m_1 i + m_2 i^2 + m_3 i^3 + m_4 i^4 + m_5 i^5$ for $i = 21,\ldots,76$, with the vector of coefficients $m = [-4.34, 0.613, -0.0274, 0.0063, -0.717 \times 10^{-5}, 0.314 \times 10^{-7}]$. This profile peaks at agent age 28 (figurative age 48), when productivity is about 1.6 times its level at agent age 1 (figurative age 21). Productivity in the final year of life is virtually the same as in the first year of life.

cost $\xi$, the reserve ratio $\phi$, the labor income tax rate $\tau^w$, and the corporate income tax rate $\tau^c$.

**4.3.1 Targets.** The endogenous variables whose values we target are listed in Table 1. Most of the target values are widely cited estimates of postwar averages, so we will not discuss them in much detail. In some cases, the closeness of a particular variable to a target depends largely on the value of one parameter: When this is the case, we will identify the relevant parameter. It should be emphasized, however, that in each of these cases the parameter in question also plays a role in the determination of other endogenous variables.

The estimates of the average capital-output ratio $K/Y$ and the average investment-capital ratio $I/K$ are due to Cooley and Prescott (1995). These estimates are based on a broad definition of capital that includes consumer durables and government capital. We target the equilibrium bonds-to-output ratio $B/Y$ at the average postwar ratio of gross federal debt to output. The money-output ratio $M/Y$ is targeted at the average postwar ratio of the monetary base to output. Since we are using bank reserve demand as a proxy for base money demand from all sources, the value of $M/Y$ is largely determined by the reserve ratio $\phi$.

The target for intermediation costs relative to output, $I^m/Y$, is based on estimates of the size of the U.S. financial intermediation sector that were constructed by Díaz-Giménez et al. (1992). These estimates are summarized in Figure 2. They indicate that the quantity of resources devoted to financial intermediation is quite large—roughly 5 to 7 percent of GDP in the early 1980s. Our target for intermediation costs relative to output is in the middle of this range. The value of the unit intermediation cost $\xi$ plays a key role in determining whether we hit this target.

Laitner (1992) reports evidence that indicates that the average lifetime consumption growth rate of U.S. households is not very different from the aggregate consumption growth rate. Consequently, we set the target for *hcg*, households' net lifetime consumption growth rate, at 0.015, which is the net rate of technological progress (from *F*) and thus the steady-state net growth rate of aggregate per capita consumption. In OLG models the lifetime consumption growth rate can be very different from the aggregate consumption growth rate, so this target imposes a significant constraint on our parameter choices. Our target for *alt*, the average share of households' time devoted to providing labor, is based on our own calculations.[27] As we have indicated, the choice of $\eta$ essentially determines the value of this variable.

Our target for *g*, the ratio of government expenditures to output, is the average postwar value of consolidated government revenue, net of transfers and government investment, relative to GDP.

That leaves only the sources of government revenue to be determined. We want the revenue coming to the government from the corporate profits tax to look like it does in the data, so that we do not overemphasize this feature of the model economy. Accordingly, we target $T^k/G$, the ratio of corporate profits tax revenues to total government expenditures, to match the average value of this ratio in postwar data.[28]

**Figure 2**

**Financial Intermediation in the United States**

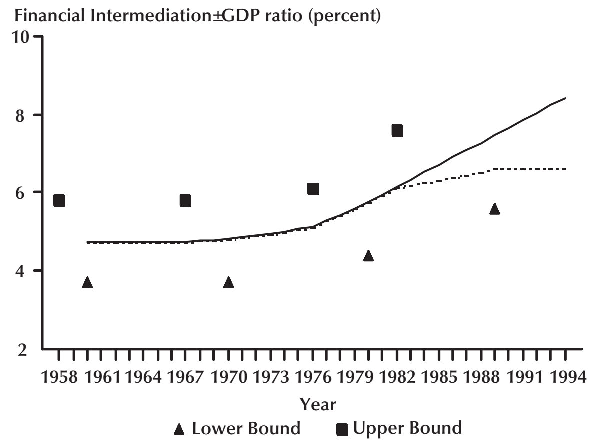


NOTE: The value of financial intermediation services provided in the U.S. economy is large. The boxes and triangles represent upper (total product basis) and lower (value added basis) bounds, respectively, as calculated by Díaz-Giménez et al. (1992). The lines are simple midpoints based on linear interpolation between data points and extrapolation of existing trends (solid line) and no trend (dotted line) at the end of the sample.

[27] The target value is based on a 24-hour day, a 40-hour work week, ten vacation days and ten holidays per year, and a 70 percent labor force participation rate. The 24-hour day assumption seems reasonable because the utility function implies that if leisure hours are zero then the marginal utility of leisure is infinite.

[28] In our model, government revenue comes from five sources: a tax on labor income, a tax on household interest income, a tax on corporate profits, currency seigniorage, and bond seigniorage. The personal interest income tax rate is an element of *F*. The volume of currency and bond seigniorage revenue is determined by the interest rates $R^h$ and $R^b$, which are also part of *F*, and by the ratios of money and bonds to output, both of which we have targeted. By setting a target for $T^k/G$, we leave all remaining government revenue to come from the tax on labor, $\tau^w$, an element of $\Delta$. This tax is levied in real terms and is not affected by inflation.

**Table 2**

**Baseline Parameter Values**

| $\rho$ | $\gamma$ | $\eta$ | $\alpha$ | $\delta$ | $\xi$ | $\phi$ | $\tau^w$ | $\tau^c$ |
|---|---|---|---|---|---|---|---|---|
| –0.223 | 37.4 | 0.154 | 0.26 | 0.0439 | 0.018 | 0.0169 | 0.11 | 0.0742 |

NOTE: Values of deep parameters in the baseline steady state.

Our parameter-selection algorithm also allows us to indicate ranges around the target values that we regard as plausible. These ranges are also displayed in Table 1. Although these ranges have some effect on the operation of the algorithm, the values it selects turn out to be very close to our targets.

**4.3.2 Parameter Choices.** A complete description of our parameter-selection algorithm is presented in the appendix, along with a detailed description of the results it generates. The baseline parameter values we obtain using this algorithm are displayed in Table 2. The characteristics of the associated steady state are described in Tables 3, 4, and 5. We discuss the parameter values first, before discussing the fit to the data in the next section.

Although our baseline value for $\rho$ may seem quite low, household preferences regarding lifetime consumption paths are much more accurately summarized by the value of the effective time preference rate $\varphi$, which is a nonlinear function of $\rho$ and $\gamma$ (see section 3.2.3). Our baseline values for these two parameters imply an effective time preference rate of –0.0067. This value produces very plausible-looking consumption behavior: In particular, it is largely responsible for the fact that households' lifetime consumption growth rate is so close to the aggregate consumption growth rate. Our value for $\varphi$ is quite close to the value (–0.0098) implied by Hurd's (1989) widely cited econometric estimates of $\gamma$ and $\rho$, and it is even closer to a recent direct estimate of $\varphi$ (–0.0078) obtained by Barsky et al. (1997) using experimental methods.

Similarly, while our baseline value for $\gamma$ may seem high, household preferences regarding substituting consumption across periods are much more accurately summarized by $\sigma$, the elasticity of intertemporal substitution in consumption (EISC), which is a nonlinear function of $\gamma$ and $\eta$ (again, see section 3.2.3). Our values for these two parameters produce $\sigma = 0.151$. This EISC value is well within the range of published estimates. It falls particularly closely in line with widely cited econometric estimates due to Hall (1988) and with recent laboratory estimates due to Barsky et al. (1997).[29]

The question of the appropriate value for the EISC is highly controversial. Many empirical studies have produced estimates significantly higher than the value we use: Attanasio and Weber (1995), for example, report a point estimate of 0.56. On the other hand, an argument for EISC values much lower than ours can be based on the fact that under an expected-utility interpretation of the preferences we employ, the coefficient of relative risk aversion is the reciprocal of the EISC.[30] Researchers working with stochastic models typically find that it takes very high degrees of risk aversion to explain the observed differential between risky and risk-free return rates (the equity premium). Kandel and Stambaugh (1991), for example, report that they need a relative risk aversion coefficient of 29 to explain the risk-free rate and the equity premium using standard preferences, while Campbell and Cochrane (1999) use a local relative risk aversion coefficient of 48.4 to accomplish the same task using habit-formation preferences. Later in this paper, after we report estimates of the welfare cost of inflation that are implied by the baseline specification of our model, we will also report alternative cost estimates implied by specifications with EISC values equal to the Attanasio-Weber and Kandel-Stambaugh

[29] In Hall's (1988) introductory summary of his results, he asserts that $\sigma$ is "unlikely to be much above 0.1" (p. 340); later in the paper he says it is "probably not above 0.2" (p. 350). Barsky et al. (1997) report a point estimate of 0.18.

[30] Thus, our estimate of $\sigma$ would be associated with a relative risk aversion coefficient of approximately 6.6. While this value is well within the (very wide) range of published estimates, and satisfies the Mehra-Prescott (1985) plausibility criterion by being below 10, most economists would probably regard it as uncomfortably high. There is no uncertainty in our model, however, and there are good reasons to believe that reluctance to substitute consumption intertemporally (a low value of $\sigma$) is not in fact closely associated with aversion to risk. Indeed, Barsky et al. (1997, p. 568) conclude that for their experimental subjects "there is no significant relationship, either statistically or economically, between risk tolerance and intertemporal substitution."

**Table 3**

**Baseline Steady-State Characteristics**

| | Variable | Model | Target |
|---|---|---|---|
| **Aggregate performance** | | | |
| Real output growth | $\lambda\psi$ | 1.032 | Fixed |
| Inflation | $1/R^h$ | 1.04 | Fixed |
| Technological progress | $\lambda$ | 1.015 | Fixed |
| Labor force growth | $\psi$ | 1.017 | Fixed |
| **Preferences** | | | |
| Effective rate of time preference | $\varphi$ | –0.0067 | Open |
| CRRA | $v$ | 6.6 | Open |
| EISC | $\sigma$ | 0.151 | Open |
| Individual consumption growth | *hcg* | 0.0188 | 0.015 |
| Lifetime average agent time devoted to labor | *alt* | 0.1539 | 0.154 |
| **Asset holdings** | | | |
| Capital-output ratio | $K/Y$ | 3.33 | 3.32 |
| Bonds-output ratio | $B/Y$ | 0.48 | 0.47 |
| Money-output ratio | $H/Y$ | 0.0591 | 0.0592 |
| **Technology** | | | |
| Capital share | $\alpha$ | 0.26 | Open |
| Depreciation rate | $\delta$ | 0.0439 | Open |
| Investment-capital ratio | $I/K$ | 0.0762 | 0.076 |

NOTE: The term "fixed" in a target entry means we set these quantities directly based on U.S. data. The term "open" in a target entry means we did not fix or target these quantities directly. More characteristics are given in Tables 4 and 5.

estimates.[31] These alternative estimates indicate that the value of EISC is qualitatively unimportant to our conclusions, in the sense that our welfare cost estimates are still an order of magnitude larger than estimates in the existing literature.

Our baseline value for $\alpha$, the capital share of output, is quite close to the value of 0.25 that is used by Auerbach and Kotlikoff (1987) and has become standard in the literature on calibrated life-cycle models. Participants in the real business cycle literature typically use higher values for $\alpha$: Cooley and Prescott (1995), for example, use $\alpha = 0.4$. We will estimate the welfare cost of inflation in an alternative specification that uses this capital share value. Again, it turns out that the particular $\alpha$-value we use is not qualitatively important for our results.

[31] Since the Campbell-Cochrane (1999) preferences are not standard (that is, not intertemporally separable CES), the local intertemporal substitution elasticity may not be equal to the reciprocal of the local coefficient of relative risk aversion.

## 5. CHARACTERISTICS OF THE BASELINE ECONOMY

### *5.1 A Quantitative Match for the Data*

The characteristics of our baseline steady state are summarized in Tables 3, 4, and 5. The steady state does a remarkably good job of matching the data along the dimensions we have selected. The only detectable discrepancy between a variable and its target involves households' lifetime consumption growth rate, which is a bit less than 0.4 percentage points higher than the target value, as shown in Table 3. Since our target for this variable was not based on a precise estimate of its value, we do not view a discrepancy of this magnitude as a problem. We included *hcg* on our list of targets to avoid ending up with a baseline economy in which the lifetime consumption growth rate was substantially (multiple percentage points) higher or lower than the aggre-

**Table 4**

**Baseline Steady-State Characteristics**

| | Variable | Model | Target |
|---|---|---|---|
| **Intermediation** | | | |
| Intermediation-output ratio | $I^m/Y$ | 0.0599 | 0.06 |
| **Government size** | | | |
| Government-output ratio | $G/Y$ | 0.151 | 0.151 |
| Revenue from firms | $T^k/G$ | 0.119 | 0.119 |
| **Government revenue sources** | | | |
| Household labor tax | $T^w/Y$ | 0.081 | Open |
| Household interest tax | $T^i/Y$ | 0.037 | Open |
| Corporate profits tax | $T^k/Y$ | 0.018 | Open |
| Bond seigniorage | $B^s/Y$ | 0.010 | Open |
| Currency seigniorage | $C^s/Y$ | 0.004 | Open |

NOTE: This is a continuation of Table 3.

**Table 5**

**Baseline Steady-State Characteristics**

| | Rates of return | | | | | | |
|---|---|---|---|---|---|---|---|
| | $R^h$ | $R^{da}$,$R^{ba}$ | $R^d$,$R^b$ | $R^{kc}$ | $R^{ka}$,$R^{cl}$ | $R^{kn}$ | $R^k$ |
| Target | Fixed | Fixed | Fixed | Open | Open | Open | Open |
| Real | 0.9615 | 1.0003 | 1.0100 | 1.0108 | 1.0288 | 1.0342 | 1.0781 |
| Nominal | 1.0000 | 1.0403 | 1.0504 | 1.0513 | 1.0700 | 1.0756 | 1.1213 |

NOTE: This is a continuation of Table 3.

gate consumption growth rate—a common occurrence in previous work with calibrated versions of life-cycle models.

## 5.2 Real Rates of Return

Our model produces rate-of-return differentials, as shown in Table 5, that can be compared with average return rates in the data. In our model, the counterpart to the return on a basket of stocks is $R^{ka}$, since corporate profits taxes are deducted from firms' earnings before they pay dividends and the capital gains on firms' stock presumably reflect market adjustments for depreciation. The "equity premium" in our model is then the difference between $R^{ka}$ and $R^b$, the before-tax real interest rate on government bonds. In our baseline steady state this difference is 188 basis points. Eight of these basis points are due to the reserve requirement; the remaining 180 are due to the cost of financial intermediation. Campbell, Lo, and MacKinlay (1997, Table 8.1) report that the equity premium in U.S. data, measured as an average of annual excess returns over a long time horizon, is 418 basis points with an approximate 95 percent confidence band of [64,756]. Thus, although our baseline equity premium accounts for a bit less than half of their point estimate, it is well within their 95 percent confidence band.

As we indicated in section 2, an important element of our approach to identifying the welfare costs of inflation is the assumption that capital is overaccumulated—a situation implied by the fact that in our baseline steady state the real interest rate on government debt is lower than the output growth rate. As we have also indicated, our steady-state values for both these variables are close to estimates of their values based on postwar U.S. data.

In an influential paper, Abel et al. (1989) note that in stochastic models an average risk-free real interest rate lower than the average output growth rate is not a sufficient condition for capital over-accumulation. They derive a sufficient condition for efficiency of a steady state in any model, stochastic or otherwise. The condition is that gross capital income is *always* larger than gross investment. They use data from the national income and product accounts to argue that, in the United States, gross capital income exceeded gross investment in every year from 1929 to 1980. They conclude, on this basis, that capital has not been overaccumulated in the United States.

If we calculate gross capital income using the definition that Abel et al. employed in their empirical analysis, we find that gross capital income exceeds gross investment in our baseline steady state.[32] Thus, our steady state actually passes the Abel et al. test for efficiency. Nevertheless, our baseline steady state has too much capital, and policy-induced increases in the real interest rate increase household welfare. The source of this conundrum is the fact that the theoretical analysis presented by Abel et al. abstracts from capital income taxes, intermediation costs, or other factors that might drive a wedge between income paid by firms and income received by households. In our model, by contrast, we use tax and intermediation-cost assumptions that are based on analysis of postwar data. We suspect that proper accounting of taxes and intermediation costs would reverse the Abel et al. conclusion that gross capital income has always exceeded gross investment.[33,34]

While our baseline steady state matches post-war U.S. data along a number of other dimensions, a more complete description of its features would take us too far afield.

## 6. THE WELFARE COST OF INFLATION

### 6.1 Definition and Measurement

**6.1.1 Assessing Welfare Costs.** Many analyses of the cost of inflation compare a steady state with a given inflation rate with a Pareto optimal Friedman-rule steady state. This approach amounts to assuming that inflation is the only source of distortions in the economy. Our model also has an optimal Friedman-rule steady state. In this steady state, the tax rates on labor and capital income are zero, the inflation rate and the nominal interest rate on government bonds are zero, and the combined real stocks of fiat currency and government debt are just large enough to support a real interest rate equal to the output growth rate (so that currency and bond seigniorage revenue is also zero). If government expenditures are positive then they are financed by lump-sum taxes.

For a practical assessment of the costs of inflation, a steady state of this type does not seem very interesting, because it is so unlike the actual U.S. economy. A procedure for estimating inflation costs that was based on this steady state would ignore the fact that the economy is distorted in many ways that do not involve inflation; it would also ignore the fact that moderate changes in the inflation rate—changes of the sort usually contemplated by policy-makers—would not eliminate distortions caused or aggravated by inflation. Our alternative approach is to model the most important sources of deviation from the Pareto optimum and study the ways in which inflation interacts with them.

As we have indicated, two sources of distortion

---

32 Equivalently, the baseline marginal product of capital, net of depreciation, exceeds the baseline output growth rate. The ratio of gross investment to output is $[K(t+1) - K(t) + \delta K(t)]/Y(t)$, which is $(\lambda\psi - 1 + \delta)(K/Y)$ in a steady state. Our baseline value of this ratio is 0.253—a value consistent with calculations presented by Cooley and Prescott (1995). The baseline ratio of gross capital income to output is $(R^k - 1)K/Y = 0.26$, which exceeds our baseline ratio of investment to output. Abel et al. (1989) report a much lower estimate of the ratio of investment to output. One reason for this is that in performing their calculations they used data on gross investment in private business capital, as opposed to the broader concept of capital used by Cooley and Prescott (1995), which includes government capital and consumer durables. Using this narrower concept, the gross investment figure in the data declines to 0.16. For additional discussion of this question and related questions, see Bullard and Russell (1999).

33 The dynamic efficiency literature does not provide any conclusive test that applies to cases in which gross capital income fluctuates above and below gross investment.

34 In our model, if the real rate of return on capital, net of depreciation and intermediation costs but *not* taxes, is lower than the output growth rate (as in our baseline case) then capital is overaccumulated and our welfare-cost analysis always goes through. However, the case in which the real return rate net of depreciation and intermediation costs is higher than the output growth rate—but the real return rate net of depreciation, intermediation costs, *and taxes* is lower than the output growth rate—is considerably more complicated. In this case capital is not overaccumulated: Maintaining the capital stock at its steady-state level does not reduce aggregate consumption on the margin. However, the tax-induced return rate distortion reduces the welfare of households by artificially increasing the cost of future consumption, and policies that tend to increase the real interest rate (such as a decrease in the inflation rate) may increase household welfare even though they reduce aggregate consumption.

are particularly critical to our analysis. The first is the fact that the government issues bonds at a real interest rate lower than the output growth rate—a policy that amounts to levying a tax on bond holders at a rate equal to the difference between these two rates. Although this tax is conceptually similar to the "inflation tax," it turns out to have much broader welfare implications, partly because the real stock of government bonds is much larger than the real stock of currency and partly because the real interest rate on government bonds plays a much more fundamental role in determining the structure of real rates of return on other assets. The second distortion is the fact that the government taxes nominal capital income—particularly, the fact that the effective tax rate on capital income increases with the inflation rate. This policy means that increases in the money growth and inflation rates would produce, if all else were held constant, large increases in the volume of government revenue. Under our assumptions about fiscal policy behavior, equilibrium is restored by large declines in the volume of government borrowing and large decreases in the real interest rate on government bonds—decreases that have very adverse effects on household welfare.

**6.1.2 Calculating Welfare Costs.** Our welfare cost calculations are based on comparing the steady states of two economies that share exactly the same environment, preferences, and technology, including common values of all parameters except the base money growth rate $\theta$. One economy has a higher money growth rate, and thus a higher inflation rate (a lower value of $R^h$), than the other. A typical household in the high-inflation economy will be worse off, in a welfare sense, than a typical household in the low-inflation economy. We measure the magnitude of the welfare difference by calculating the amount of consumption-good compensation necessary, at each date, to make the agents in the high-inflation economy indifferent between staying in that economy or moving to the low-inflation economy. This measure of the welfare cost of inflation is conceptually similar to measures used by Cooley and Hansen (1989) and others. However, the agent heterogeneity in our model makes our calculations slightly more complicated.

Our procedure compares the welfare levels produced by particular rates of inflation to the levels produced by benchmark inflation rates that exceed the Friedman-rule rate. In one set of calculations, the benchmark inflation rate is our baseline rate and the associated steady state is our baseline steady state. In the other set of calculations, the benchmark inflation rate is zero: The associated steady state is the one generated by this inflation rate when all the other parameters are set at their baseline values. In each case, we begin by calculating the lifetime utility of a representative household from an arbitrary generation $t$ in the steady state associated with the benchmark inflation rate. Next, we solve the model for the new steady state associated with a different inflation rate, and we record the consumption and leisure choices of a representative household from the same generation $t$. We hold these consumption and leisure quantities fixed and imagine giving the household annual compensation, in units of the consumption good, until its augmented lifetime consumption-leisure package gives it the same utility level as in the benchmark steady state. The quantity of consumption-good compensation given to the household during each period of its life is assumed to grow at a gross rate of $\lambda$, the (exogenous) steady-state growth rate of *per capita* consumption. The final steps of our procedure are (i) to calculate the total amount of consumption good compensation necessary, at an arbitrary date $t$, to compensate each household alive at that date in the manner just described; (ii) to calculate the total amount of output produced at the same date in the benchmark steady state; and (iii) to express the former value as a percentage of the latter value.

## 6.2 Results

**6.2.1 Welfare Costs Relative to Baseline Cases.** The triangles in Figure 3 display the welfare costs of a range of alternative inflation rates, relative to a benchmark inflation rate of 4 percent. These cost estimates represent the principal results of our analysis. The triangle at the point with coordinates (4,0) represents our baseline steady state. Lower inflation rates yield welfare benefits to households relative to the benchmark inflation rate, so the welfare cost estimates associated with these rates are negative.

Since increases in the inflation rate reduce the welfare of households, the line formed by the triangles is upward-sloping. The gradient linking the inflation rate to the welfare cost is surprisingly steep. Over the range of inflation rates considered, a 1 percent increase in the annual inflation rate increases the welfare cost of inflation by more than 1 percent of output per year. As we have noted, these cost estimates are an order of magnitude larger than most

## Table 6

**Recent Estimates of the Welfare Cost of Inflation**

| Study | Features | Inflation comparison | Welfare cost |
|---|---|---|---|
| Cooley/Hansen (1989) | RBC model with cash-in-advance | 10% vs. optimal | 0.387% |
| Gomme (1993) | Endogenous growth with cash-in-advance | 8.5% vs. optimal | 0.0273% |
| İmrohoroğlu/Prescott (1991) | Idiosyncratic labor income risk | 10% vs. 0% | 0.9% |
| Dotsey/Ireland (1996) | Endogenous growth with cash-in-advance | 10% vs. 0% | 1.73% |
| Lucas (2000) | Representative agent with shopping time | 10% vs. optimal | 1.3% |
| This paper | Life-cycle economy with financial intermediation | 10% vs. 0% | 12.4% |

**NOTE: Some recent studies of the welfare cost of inflation. Costs are expressed as the compensating consumption necessary to make agents indifferent between the two inflation regimes. The cost calculated in this paper is an order of magnitude larger than those from the earlier literature.**

other estimates from the literature. Table 6 summarizes some of the previous estimates.[35]

There are no triangles in the figure for inflation rates lower than 2.5 percent. The reason for this is that, once the inflation rate falls below this level, the government faces a revenue shortage: No steady-state real interest rate produces enough revenue to meet our target value for $G/Y$. We adopt two different strategies for addressing this situation. The simplest strategy involves modifying our baseline steady state by holding fixed all the parameters of the model except the inflation rate and accepting whatever level of $G/Y$ turns out to be consistent with a benchmark inflation rate of zero. We can then calculate the welfare cost of inflation relative to this zero-inflation steady state.[36] The results of these calculations are indicated by the boxes in Figure 3. Over the range of inflation rates considered, the increase in the welfare cost as the rate of inflation rises is again better than 1 percent of real output per percentage point increase in the inflation rate. These results illustrate the fact that even when the initial inflation rate is low, the marginal distortion produced by increases in the inflation rate can be quite large.

## Figure 3

**The Welfare Cost of Inflation**

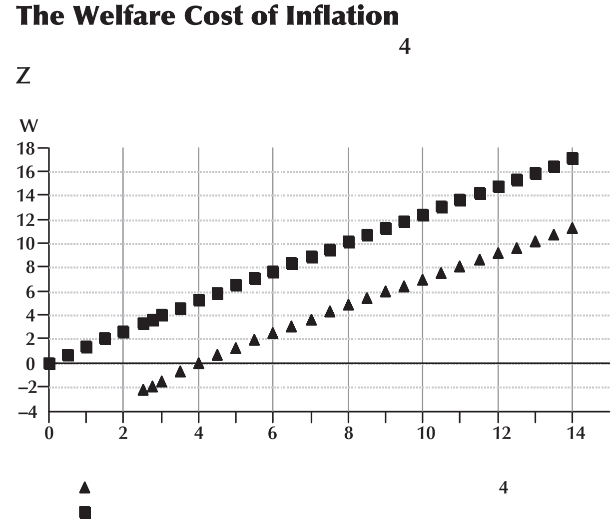


**6.2.2 Welfare Gains from Lower Inflation: An Alternative Approach.** A second approach to confronting the revenue shortage that arises when the inflation rate is allowed to fall below 2.5 percent is to raise direct tax rates to recover the lost revenue.

[35] One interpretation of the results reported by Lacker and Schreft (1996) would place the welfare cost of 10 percent versus 0 percent inflation at 4.27 percent of output. Other interpretations, however, are consistent with the lower cost estimates reported in the papers listed in Table 6. Lacker and Schreft (1996) emphasize resource-costly credit and the impact of inflation on real returns.

[36] The steady state at zero inflation inherits most of the quantitative properties of the steady state at 4 percent inflation, so we do not report these properties here. The principal exception is government revenue as a fraction of real output, which is about 11 percent in the zero inflation case versus 15.1 percent in the baseline case. At the cost of some complications, we could use the Friedman-rule inflation rate—a deflation rate equal to the output growth rate—as a benchmark rate without changing our qualitative conclusions. However, a zero rate of inflation has often been used as a benchmark in the literature, and it has often been proposed as a practical target for monetary policy.

**Figure 4**

**Disinflation with Distorting Taxes**

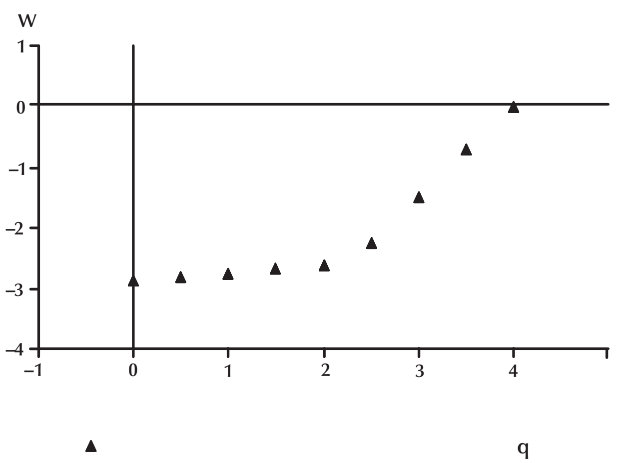


This approach sacrifices comparability with the results reported in Table 6, since the experiments that produce these results often assume either that there are no direct taxes or that any direct tax rates are fixed. However, this alternative approach may give us better insight into the practical problems of disinflation from relatively low initial inflation rates.

Figure 4 illustrates how the need to raise other distortionary taxes affects the welfare benefits of disinflation from low inflation rates. For simplicity, we consider a tax increase scenario in which the lost revenue is made up by increasing all three tax rates ($\tau^w$, $\tau^i$, $\tau^c$) equally in percentage point terms.[37]

The welfare benefit from further reductions in inflation declines dramatically once the initial inflation rate falls below 2.5 percent. The total welfare benefit achieved by moving from 2.5 percent inflation to zero inflation is less than 0.2 percent of the baseline level of real output. While this benefit estimate remains large by the standards of the other papers listed in Table 6, it is quite small compared with the results reported in this paper for disinflation from higher initial inflation rates.

The bottom line of this aspect of our analysis is that when the initial inflation rate is low, the welfare gains from reducing the inflation rate are largely offset by the welfare losses caused by more severe distortions from direct taxes. And while the gains from moving to inflation rates below 2.5 percent are not trivial, achieving them requires a degree of coordination between fiscal policy and monetary policy that is not necessary when disinflating from higher initial inflation rates.

**6.2.3 Decomposing the Welfare Cost of Inflation.** As we have indicated, much of the literature on the welfare cost of inflation has concentrated on what might be called "purely monetary" costs. In our analysis, in contrast, most of the distortions that are aggravated by higher inflation are not directly connected to money demand. From a policy perspective we think our approach is the more useful one, since it allows policymakers to use the total distortion caused by inflation, taking other features of the economy as given, as the yardstick by which judgments are made. In this section of the paper, however, we attempt to apportion our welfare cost estimates by source to allow easier comparisons with other studies and to provide additional intuition for our results.

To repeat, the three basic sources of inflation costs in our model are (i) monetary costs associated with the fact that higher inflation reduces the rate of return on money, (ii) tax-distortion costs associated with the fact that higher inflation increases the effective tax rate on capital income (the "Feldstein effect") and reduces the after-tax real rate of return on capital, and (iii) return-distortion costs associated with the fact that higher inflation produces a decrease in government borrowing that reduces the before-tax real return rate on capital (the "Miller-Sargent effect").

We will begin by attempting to isolate the purely monetary welfare cost, using a procedure as closely analogous to Feldstein's (1997) as we can arrange in the context of our model. In particular, we look for a steady-state equilibrium in which the inflation rate has increased by 10 percentage points, the before-tax real interest rate on government bonds is unchanged from the baseline case, and the ratio of government revenue to output is unchanged from the baseline case. We avoid the Feldstein effect by nominalizing the pretax returns to capital using the baseline inflation rate of 4 percent instead of the new inflation rate of 14 percent. The revenue gain that the government enjoys as a result of the increase in the inflation rate is offset by equal-percentage-point reductions in all three of the direct tax rates. This experiment produces a result similar to Feldstein's: There is actually a small welfare gain from higher inflation—0.2 percent of output—because the inflation tax on bank reserves is slightly

[37] An alternative scenario in which the lost revenue is recovered exclusively by increases in the labor income tax rate produces similar results.

less distortionary than the direct taxes it replaces.

The next step in our decomposition procedure is to identify the welfare costs that are due to the Feldstein effect. Again we increase the inflation rate by 10 percentage points, holding the pretax real government bond rate at its baseline level. This time, however, we nominalize pretax capital income at the new inflation rate. Again we offset the resulting increase in government revenue by reducing the three direct tax rates. The before-tax real return rate on capital is virtually unchanged, but the real interest rate facing household savers—the after-tax real return rate on government bonds or intermediary deposits—falls by almost exactly 1 percentage point. The welfare cost of inflation is approximately 1.3 percent of output. Given the small monetary benefit, the total Feldstein-effect cost of inflation is roughly 1.5 percent of output.

Finally, we move to the new steady state associated with the much-larger welfare cost estimate we reported in Figure 3. We begin by restoring the direct tax rates to their baseline levels: The increase in inflation-tax and capital-income-tax revenue that was previously offset by a reduction in these tax rates is now offset by a decline in pretax real interest rates that reduces government revenue from bond seigniorage. The pretax real return rate on capital falls by 1.7 percentage points, the pretax real interest rate on government debt falls by 2.4 percentage points, and the after-tax real deposit rate falls by 3.6 percentage points.[38] As we have indicated, the welfare cost of inflation is roughly 11.2 percent of output.

On the basis of this decomposition, we can conclude that about 85 percent of the inflation cost we describe (9.7 percentage points of the 11.2 percent of output) is caused by the decline in the before-tax real interest rate, while the Feldstein effect accounts for about 15 percent of the cost and the purely monetary cost is negligible. However, the small size and perverse sign of the monetary cost hinges partly on the way in which that cost is defined. If we give back the additional currency seigniorage revenue produced by an increase in the inflation rate by allowing the pretax real interest rate to fall, rather than by increasing direct tax rates, then the pretax real capital and real bond return rates fall by 20 and 30 basis points, respectively, and the after-tax real deposit rate falls by 25 basis points. The monetary cost of inflation rises to 0.6 percent of output, which amounts to approximately 5 percent of the total welfare cost.[39] The Miller-Sargent effect now accounts for about 80 percent of the total cost of inflation.

**6.2.4 Moderate Changes in the Inflation Rate.** As we have indicated, reporting the results of experiments in which the inflation rate is permanently increased by 10 percentage points or more is a standard practice in the welfare cost of inflation literature. One reason for this is that, in much of the literature, the marginal welfare effects of changes in the inflation rate are so small that it takes larger changes to generate effects of any practical significance. In our model, however, the effects of smaller changes in the inflation rate on welfare and other endogenous variables are quite substantial. Moreover, from a practical point of view, reporting the effects of moderate changes in the inflation rate seems more interesting than reporting the effects of changes that are very large relative to modern U.S. historical experience.

Table 7 reports the effects on a variety of endogenous variables of a permanent, 1 percent increase in the inflation rate, starting from the baseline inflation rate of 4 percent. The before-tax real rate of return on capital and the before-tax real interest rate on government bonds fall by 38 and 43 basis points, respectively, while the real interest rate on bank deposits (which is also the after-tax real rate on bonds) falls by 53 basis points. The level of output rises by 3.9 percent. Most of the increase in output is due to a large (9.3 percent) rise in investment: Aggregate consumption increases by only 0.8 percent. Government spending increases in proportion to the increase in output; the capital stock increases in proportion to the increase in investment. The real wage increases by 1.8 percent and aggregate labor hours rise by 2.1 percent.[40]

[38] Real interest rate declines of this magnitude are quite moderate by comparison with historical data from high-inflation periods. During 1973-79, for example, the average inflation rate was 9.1 percent (CPI-U) and the average real government bond rate was –2.0 percent (ex post real yield on 3-month T-bills). In our baseline specification, a steady-state inflation rate of 14 percent is associated with a real government bond rate of –1.4 percent, while an inflation rate of 9.1 percent produces a real government bond rate of –0.5 percent.

[39] One can think of this alternative monetary cost estimate as an estimate of the component of the welfare cost of inflation that has nothing to do with the fact that inflation changes the effective tax rate on capital income. Thus, our results suggest that 95 percent of the total welfare cost of inflation is associated, directly or indirectly, with this feature of the tax system.

[40] It should be emphasized that these are level effects: In the new steady state, these variables continue to grow at exogenously determined rates that are invariant to changes in the inflation rate. Output, consumption, and investment grow at the exogenous output growth rate, while labor hours grow at the exogenous labor force growth rate and the real wage grows at the exogenous productivity growth rate. Also, these are, of course, steady-state comparisons, and so it would take time for these impacts to occur.

## Table 7

**Real Effects of a Permanent 1-Percentage-Point Increase in the Inflation Rate**

| Variable | % Change | Variable | % Change |
|---|---|---|---|
| Before-tax real MPK | –0.38 | Real wage | +1.8 |
| Before-tax real bond rate | –0.43 | Labor hours | +2.1 |
| After-tax real deposit rate | –0.53 | | |
| | | Consumption growth | –0.08 |
| Output | +3.9 | First 5 periods | +1.5 |
| Consumption | +0.8 | Last 5 periods | –3.7 |
| Investment | +9.3 | Welfare | –1.3 |

NOTE: The figures for return rates are changes in percentage points. The figures for output, consumption, investment, the real wage, labor hours, and first- and last-five-periods consumption are percent changes in levels. The figure for consumption growth is the percentage point change in the lifetime growth rate. The welfare cost figure is the required consumption compensation as a percent of the initial level of output.

The decline in the real interest rate makes it more difficult for households to consume during the later years of their lives. Although households' lifetime consumption growth rate falls by only 0.1 percent, this change has a big effect on their consumption levels at the beginning and end of their lives. Households' average consumption during the first five years of life rises by 1.5 percent, but their consumption during their last five years falls by 3.7 percent. The welfare loss from the 1 percent increase in the inflation rate amounts to 1.3 percent of output. The small increase in aggregate consumption does not suffice to compensate households for the increase in their labor hours and the decrease in the slope of their lifetime consumption path. Households work harder, but most of their increased work effort goes to support accumulation of additional capital that produces low marginal returns.

The relationship between the size of the welfare cost of inflation and the magnitude of the increase in the inflation rate is a little less than linear: The welfare cost of a 10-percentage-point increase in the inflation rate is only 8.7 times larger than the cost of a 1-percentage-point increase in the inflation rate. For other variables the departure from linearity is more pronounced. The percentage point decline in the real interest rate produced by a 10-percentage-point increase in the inflation rate is only six times larger than the decline produced by a 1-percentage-point increase in inflation. Roughly the same ratio holds for the percentage increases in the levels of output, investment, and work effort. For the level of aggregate consumption, the ratio is only 3.7.

### *6.3 Alternative Parameterizations*

In this subsection, we report the results of welfare-cost experiments that are conducted using alternative parameterizations of our baseline economy. These experiments are intended to explore the robustness of our results to changes in the values of parameters whose baseline values are potentially controversial.

**6.3.1 Intermediation Costs.** As we indicated in the introduction, we introduce intermediation costs primarily because they allow our baseline steady state to do a better job of matching postwar data. These costs allow the steady-state values of endogenous variables from our nonstochastic model economy to more closely resemble long-run averages of variables from an actual economy that is presumably stochastic.[41] Since the magnitude of the unit intermediation cost is not affected by changes in the inflation rate, there is no reason to expect changes in the value of the intermediation cost to have substantial effects on our welfare cost estimate.

As a simple test to confirm the robustness of our results to dropping the cost of intermediation, we calculate the effects of a permanent increase in

[41] Ríos-Rull (1994, 1995, 1996) studies a general equilibrium life-cycle model with aggregate production and return risk but no cost of intermediation. Since the model does not generate a significant equity premium, it does not solve the problem of reconciling the high real average rate of return on equity with the relatively low real interest rate on government debt. See Bullard and Russell (1999) for a more detailed discussion of the impact of intermediation costs in a nonstochastic model.

the inflation rate starting from the steady state that our model produces when the intermediation cost is set at zero but none of the other baseline values are changed. (This new steady state is not a very good match for postwar data along some of the dimensions we have described.) The cost of a 1-percentage-point increase in the inflation rate is 1.1 percent of output, while the cost of a 10-percentage-point increase is 9.8 percent of output. Although the differences between these estimates and our original estimates are not entirely inconsiderable, they would not affect our conclusions in any important way.

**6.3.2 Intertemporal Substitution in Consumption.** Our parameter selection procedure produces relatively low values for the intertemporal elasticity of substitution in consumption and the pure rate of time preference. Some economists might argue that households with the preferences we describe are implausibly resistant to intertemporal substitution and/or place an implausibly high value on future consumption relative to current consumption. It turns out that by reducing the value of the indifference curvature parameter $\gamma$—which increases the value of $\sigma$ associated with a fixed value of $\eta$—and increasing the pure time preference rate $\rho$ so as to keep the ratio of government expenditures to output constant, we can create a family of alternative steady states that match postwar data nearly as well as our baseline steady state. To investigate the robustness of our results to different substitution elasticities and time preference rates, we examined the characteristics of a member of this family of specifications with an IESC of 0.56—a point estimate due to Attanasio and Weber (1995) that is substantially larger than our baseline value for $\sigma$. The corresponding value of the pure time preference rate is –0.048.[42] The welfare cost of a 1-percentage-point increase in the inflation rate is now 1.1 percent of output, while the cost of a 10-percentage-point increase in inflation becomes 8.3 percent of output. Although these estimates are somewhat smaller than our original estimates (particularly in the latter case), the differences remain too small to affect our conclusions.[43]

An interesting feature of these alternative experiments is that the changes in the values of other endogenous variables that result from increases in the inflation rate are smaller, and in some cases different in sign, than the changes produced by our original experiments. For example, when the inflation rate is increased by 1 percentage point, the before-tax real government bond rate falls by only 15 basis points (compared with 43 basis points) and the level of output rises by only 0.3 percent (compared with 3.9 percent). In the 10-percentage-point–increase case, moreover, the real bond rate falls by only 0.7 percentage points (compared with 2.4 percentage points) and the level of output *falls* by 3.4 percent (compared with rising by 24.1 percent).

The small size of the changes in the pretax real bond rate are attributable to the fact that when $\sigma$ is high, aggregate saving is very sensitive to changes in the real interest rate, so that a small reduction in the rate produces a large decline in the real stock of bonds and a large decrease in the revenue from bond seigniorage. The changes in output are small because small declines in the pretax real rate produce small increases in the capital stock—increases that are largely offset, or more than offset, by the decrease in the capital stock that is caused by the increase in the effective tax rate on capital income. The changes in welfare remain large because households with large values of $\sigma$ are very sensitive to small changes in the slope of their lifetime consumption paths. However, the Feldstein effect becomes a much larger component of the cost of inflation, accounting for a bit more than half the total cost: 4.3 percentage points of 8.3 percent of output.

The results presented in this subsection demonstrate that it is possible to believe that inflation has very high welfare costs without believing that permanent increases in the inflation rate produce large permanent increases in the level of output. In addition, these results make it possible to believe that the Feldstein effect is nearly as large as Feldstein's (1997) estimate, while at the same time believing that the total cost of inflation is substantially larger.

**6.3.3 Capital Share of Output.** Our best-fit value for $\alpha$, the capital share of output, is somewhat lower than values used in the literature on real business cycles. Cooley and Prescott (1995), for example, recommend a value of 0.4 for this param-

---

[42] Hurd (1989) estimates the difference between the risk-free real interest rate facing households and the pure time preference rate at 4.1 percent. This estimate is widely used in the literature on calibrated OLG models. Since our estimate of the risk-free rate facing households is zero, the associated estimate of the pure time preference would be –0.041.

[43] We also conducted welfare-cost experiments from a baseline steady state with $\sigma = 1/29$, an estimate based on Kandel and Stambaugh (1991). The welfare costs of 1-percentage-point and 10-percentage-point increases in the inflation rate are 1.4 percent and 13.8 percent of output, respectively.

eter. We will investigate the robustness of our results to changes in the capital share by examining an alternative steady state in which $\alpha$ is increased to 0.4 but the other parameters retain their baseline values. (Again, this alternative steady state is not a good match for postwar data along some dimensions.) The welfare costs of 1-percentage-point and 10-percentage-point increases in the inflation rate fall to 0.7 percent and 7.0 percent of the initial steady-state level of output, respectively. Although these cost estimates are substantially smaller than the estimate produced by our baseline steady state, they remain large enough, relative to other values that have appeared in the literature, to support our qualitative conclusions.[44]

## 7. QUALIFICATIONS

### 7.1 The Magnitude of Our Cost Estimates

The continuing interest in academic research on the cost of inflation grows out of the fact that most of the cost estimates the literature has produced are quite small—too small to satisfy the intuition of many economists and far too small to explain the abhorrence with which inflation is regarded by many members of the business and policy communities. Our cost estimates, on the other hand, may be large enough to produce the opposite problem: If inflation is really this costly, it seems hard to understand why monetary policymakers would ever allow the inflation rate to rise above very low levels.

Our model provides one explanation for the persistence of moderately high inflation rates that seems natural to us: Higher inflation rates produce higher levels of output and employment. Policymakers typically assume that policies that increase output and hours worked also increase public welfare. Economic theory does not necessarily concur, and we think our model provides an empirically plausible counterexample.

It must be noted, however, that our cost estimates are founded on some basic assumptions that may not be very reasonable in practice, and that more realistic assumptions could reduce the size of the estimates substantially. Our principal defense for these assumptions is that analogous ones have been adopted by most other participants in the inflation-cost literature, so that adopting them facilitates comparisons with the literature. In addition, these assumptions have not been considered very controversial.

One of our most important assumptions is that changes in monetary policy are not accompanied by changes in the active elements of fiscal policy (expenditure or tax policy). Most of the previous literature makes the somewhat weaker assumption that changes in monetary policy do not result in changes in government expenditures. In the "monetary cost" literature, the increased seigniorage revenue produced by higher inflation rates is usually returned to the public via lump-sum taxes; Feldstein (1997) and Abel (1997) assume that inflation-induced increases in government revenue are offset by decreases in direct tax rates. We follow our predecessors by assuming that changes in the inflation rate do not result in changes in government expenditures (relative to output), but we also assume that there are no changes in direct tax rates. We are able to make this assumption because in our model the increased government revenue generated by higher inflation can be offset by reductions in government borrowing that decrease the amount of revenue from bond seigniorage.

We think the assumption that the government does not respond to changes in monetary policy by changing direct tax rates is quite plausible for estimating the welfare effects of small or moderate changes in the inflation rate—much more plausible, in fact, than the alternatives just described. However, this assumption begins to seem much less plausible when applied to large changes. One reason for this is that in our model, large changes in the inflation rate produce huge changes in the volume of government borrowing. When the inflation rate rises by 10 percent, for example, a steady-state government budget deficit amounting to 1 percent of output becomes a steady-state budget surplus of 5.5 percent of output. (The government dissipates the surplus by lending to the public.[45])

If we assume that large increases in government revenue are partly offset by cuts in direct tax rates, then the pretax real interest rate becomes less sensitive to increases in the inflation rate and the welfare

[44] One partial explanation for the decline in the cost estimates here, and also in the zero-intermediation-cost case, is that the initial $G/Y$ ratio is substantially smaller than in the baseline steady state. Since the inflation-cost-generating mechanism we describe works through the government budget constraint, lower values of $G/Y$ tend to produce lower cost estimates. A more thoroughgoing respecification that kept this ratio constant would produce substantially higher estimates.

[45] The government budget deficit is equal to the amount of revenue from bond seigniorage. In our model, a permanent 1 percent increase in the inflation rate cuts the steady-state budget deficit from 1 percent of output to 0.5 percent of output.

costs of such increases become smaller. Suppose we assume, for example, that when the monetary authority increases the inflation rate by 10 percentage points the fiscal policy authority responds by allowing the deficit-output ratio to fall by only half the amount just described—a decline that still leaves the government with a very substantial surplus—and offsets the rest of the revenue increase by equal-percentage-point reductions in the three direct tax rates. In this case, the pretax real bond rate falls by only 1.5 percentage points instead of 2.4 percentage points, and the welfare cost of the 10-percentage-point increase in inflation is only 7.0 percent of output. Smaller deficit reductions produce smaller cost estimates.[46]

We have also followed most of the previous literature by confining our analysis to comparisons of steady states. Thus, we are estimating the costs of permanent changes in the inflation rate produced by permanent and perfectly credible changes in monetary policy. A more complete analysis of the welfare implications of changes in inflation would study costs (or benefits) incurred during the transition path from one steady state to another, and it would also incorporate the fact that there is inevitably uncertainty about exactly what the new inflation target is and whether the monetary authority will persist in trying to reach it. While uncertainty about policy implementation is beyond the scope of this paper, we have done preliminary research on the properties of transition paths. Our findings indicate that in our model, a complete or nearly complete transition from one steady state to another is likely to take quite a long time. Thus, our results should not be interpreted as suggesting that the welfare gains from lower inflation can be realized quickly, and it is not entirely inconceivable that a transition analysis might reveal that a permanent decrease in the inflation rate imposes large costs on households whose lives overlap the date of the policy change or who are born in the years immediately following the change.

We can cite one interesting piece of evidence which suggests that transition analysis is unlikely to have any dramatic effect on our conclusions. Suppose we imagine that the shift from the old steady state to the new steady state occurs instantly at some date $T$, so that the households who are alive at date $T$—the members of generations $T$, $T-1$, $T-2,\ldots,T-54$, who have 55, 54,...,1 years left to live, respectively—switch immediately from the consumption-leisure bundles associated with the old steady state to the bundles associated with the new one. We can then use our utility function to conduct across-steady-state comparisons of the welfare of these 55 cohorts of households during the remaining years of their lives. In our baseline economy, the "remaining welfare" of the members of each cohort is higher in a steady state with a 3 percent inflation rate than in the baseline steady state with its 4 percent inflation rate. Thus, if the transition path between the steady states is monotonic, as our preliminary results also suggest, then the benefits of a disinflation undertaken at a date $T$ should start accruing immediately to all members of the society.

### 7.2 Modeling the Tax System

Our assumptions about the tax system are at best a crude approximation of the complex and nonlinear array of taxes imposed by U.S. federal, state, and local governments. We have adopted the conservative approach of allowing a large fraction of government expenditures to be financed by a tax on real labor income—a tax whose effective rate does not depend on the rate of inflation. This decision probably causes us to understate the historical welfare cost of inflation: Actual income taxes are levied on nominal income in a progressive manner and, prior to the 1980s, "bracket creep" allowed increases in inflation to increase both effective labor income tax rates and government labor income tax revenues. We also ignore the historical effect of "bracket creep" on income from interest and capital gains. (In the case of capital gains, the tax reforms of the 1980s reduced this effect but did not eliminate it entirely.) On the other hand, our assumption that the corporate profits tax acts analogously to the interest income tax as a tax on nominal returns to capital is at least partly counterfactual: Under the U.S. tax system, the effective corporate profits tax rate is not directly increased by inflation. However, we think this assumption is reasonable, as a first approximation, for two reasons. First, our corporate profits tax is intended partly as a proxy for a tax on capital gains, which is absent from our model: The effective tax rate on capital gains does increase with higher inflation. Second, the fact that the U.S. tax system uses historic cost depreciation allows inflation to increase the effective tax rate on corporate

[46] A policy of holding the deficit-output ratio fixed produces an unchanged pretax real bond rate and thus duplicates the Feldstein effect, yielding a welfare cost of 1.2 percent of output.

profits indirectly, by reducing the real value of depreciation allowances.[47]

Our model also abstracts from another important feature of the U.S. tax system, which is that households may deduct mortgage interest payments (and before the 1980s, other interest payments) from their taxable incomes. But our tax assumptions account for this effect by taxing households on their *net* interest income. In our general equilibrium environment, the net asset position of the households is essentially the capital stock of the economy. Similarly, our model does not distinguish returns paid by firms as dividends from returns paid as interest: Under the U.S. tax system, firms are taxed on the former but not the latter. As a result, it may seem that we are overstating the extent of double taxation of capital. We address this problem by choosing a relatively low corporate profits tax rate—a rate that allows us to duplicate the observed ratio of corporate profits tax revenue to government expenditures in our baseline equilibrium.

In sum, we think our tax assumptions provide an approximation of the U.S. tax system that is adequate for our purposes. However, further research on the nature of the interaction between inflation and the tax code in general equilibrium models is certainly warranted.[48]

## 8. CONCLUDING REMARKS

In this paper, we use a dynamic general equilibrium model to estimate the welfare cost of inflation in the U.S. economy. According to our estimates, inflation is far more costly than most of the literature to date has indicated. However, our estimates of the purely monetary component of the cost of inflation—the component studied in most previous work on this topic—are of the same order of magnitude as previous estimates. Our much-larger total cost estimates grow out of the fact that in our model, inflation has a substantial impact on the real rates of return on nonmonetary assets. Thus, our results provide a formal interpretation of a view about the source of inflation costs that is often expressed in business and policy circles as well as academic discussions.

In our model, most of the welfare cost of inflation is attributable to the fact that higher inflation rates increase the effective tax rate on capital income. However, the portion of the cost that is driven by the tendency of increases in the capital tax rate to widen the spread between the before-tax and after-tax real return rates on capital—a distortion whose impact has been studied by Feldstein (1997) and Abel (1997)—is relatively small. Instead, the lion's share of the welfare cost of higher inflation is attributable to its tendency to produce a downward shift in the entire structure of real interest rates, both before and after taxes. This general decline in real interest rates is a consequence of our assumption that fiscal policymakers respond to the increase in tax revenue that higher inflation produces by reducing government borrowing rather than by cutting tax rates.

Our results have at least two important implications for further study of inflation costs and related issues. First, they indicate that abstracting from general equilibrium considerations may lead to serious underestimates of the welfare cost of inflation—a conclusion we share with Dotsey and Ireland (1996). Second, they demonstrate that plausible alternatives to the standard calibrated dynamic general equilibrium model may produce very different estimates of the cost of inflation—and, by extension, very different answers to other outstanding questions in macroeconomics.

## REFERENCES

Abel, Andrew B. "Comment on 'The Costs and Benefits of Going from Low Inflation to Price Stability' by Martin Feldstein," in Christian D. Romer and David H. Romer, eds., *Reducing Inflation: Motivation and Strategy*. Chicago: University of Chicago Press, 1997.

Abel, Andrew B.; Mankiw, N. Gregory; Summers, Lawrence H. and Zeckhauser, Richard J. "Assessing Dynamic Efficiency: Theory and Evidence." *Review of Economic Studies*, January 1989, *56*(1), pp. 1-20.

Altig, David and Carlstrom, Charles T. "Inflation, Personal Taxes, and Real Output: A Dynamic Analysis." *Journal of Money, Credit, and, Banking*, August 1991, *23*(3, Part 2), pp. 547-71.

[47] Both these effects are discussed in Feldstein (1997), who concludes that inflation does indeed increase the effective tax rate on corporate profits and that, overall, the effect of inflation on the effective tax rate on income produced by firms in the form of dividends and capital gains is actually somewhat larger than its effect on the effective tax rate on interest paid by firms.

[48] Black et al. (1994) also use a calibrated general equilibrium model to study welfare costs of inflation that are driven largely by inflation's interaction with the nominal tax system. The version of their model most comparable to ours produces costs that are less than one-fifth the size of our estimates. However, they find that introducing endogenous growth and/or an open economy increases the cost of inflation significantly. We think these are important directions for future research in this area.

Ahmed, Shagil and Rogers, John H. "Long-Term Evidence on the Tobin and Fisher Effects: A New Approach." Working Paper No. 566, International Finance Discussion Papers, Board of Governors of the Federal Reserve System, September 1996.

Attanasio, Orazio P. and Weber, Gugliemo. "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey." *Journal of Political Economy*, December 1995, *103*(6), pp. 1121-57.

Auerbach, Alan J. and Kotlikoff, Laurence J. *Dynamic Fiscal Policy*. Cambridge: Cambridge University Press, 1987.

Barsky, Robert B.; Juster, F. Thomas; Kimball, Miles S. and Shapiro, Matthew D. "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study." *Quarterly Journal of Economics*, May 1997, *112*(2), pp. 537-79.

Bernanke, Ben and Gertler, Mark. "Agency Costs, Net Worth, and Business Fluctuations." *American Economic Review*, March 1989, *79*(1), pp. 14-31.

Black, Richard; Macklem, Tiff and Poloz, Stephen. "Non-Superneutralities and Some Benefits of Disinflation: A Quantitative General-Equilibrium Analysis," in Bank of Canada, ed., *Economic Behaviour and Policy Choice Under Price Stability, Proceedings of a Conference Held at the Bank of Canada*, October 1993. Ottawa: Bank of Canada, 1994, pp. 477-516.

Boyd, John H.; Levine, Ross and Smith, Bruce D. "The Impact of Inflation on Financial Sector Performance." *Journal of Monetary Economics*, April 2001, *47*(2), pp. 221-48.

Boyd, John H. and Smith, Bruce D. "Capital Market Imperfections in a Monetary Growth Model." *Economic Theory*, March 1998, *11*(2), pp. 241-73.

Bullard, James and Keating, John W. "The Long-Run Relationship Between Inflation and Output in Postwar Economies." *Journal of Monetary Economics*, December 1995, *36*(3), pp. 477-96.

Bullard, James and Russell, Steven. "An Empirically Plausible Model of Low Real Interest Rates and Unbacked Government Debt." *Journal of Monetary Economics*, December 1999, *44*(3), pp. 477-508.

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

Campbell, John Y. and Cochrane, John H. "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior." *Journal of Political Economy*, April 1999, *107*(2), pp. 205-51.

Cooley, Thomas F. and Hansen, Gary D. "The Inflation Tax in a Real Business Cycle Model." *American Economic Review*, September 1989, *79*(4), pp. 733-48.

Cooley, Thomas F. and Prescott, Edward C. "Economic Growth and Business Cycles," in Thomas F. Cooley, ed., *Frontiers of Business Cycle Research*. Princeton: Princeton University Press, 1995, pp. 1-38.

Darby, Michael R. "Some Pleasant Monetarist Arithmetic." Federal Reserve Bank of Minneapolis *Quarterly Review*, Spring 1984, *8*(2), pp. 15-20.

Díaz-Giménez, Javier; Prescott, Edward C.; Fitzgerald, Terry and Alvarez, Fernando. "Banking in a Computable General Equilibrium Economics." *Journal of Economic Dynamics and Control*, July-October 1992, *16*(3-4), pp. 533-59.

Diamond, Peter A. "National Debt in a Neoclassical Growth Model." *American Economic Review*, December 1965, *55*(5), pp. 1126-50.

Dotsey, Michael and Ireland, Peter. "The Welfare Cost of Inflation in General Equilibrium." *Journal of Monetary Economics*, February 1996, *37*(1), pp. 29-47.

Espinosa-Vega, Marco A. and Russell, Steven. "Can Higher Inflation Reduce Real Interest Rates in the Long Run?" *Canadian Journal of Economics*, February 1998a, *31*(1), pp. 92-103.

Espinosa-Vega, Marco A. and Russell, Steven. "The Long-Run Real Effects of Monetary Policy: Keynesian Predictions from a Neoclassical Model." Working Paper 98-6, Federal Reserve Bank of Atlanta, April 1998b.

Feldstein, Martin. "The Costs and Benefits of Going from Low Inflation to Price Stability," in Christina D. Romer and David H. Romer, eds., *Reducing Inflation: Motivation and Strategy*. Chicago: University of Chicago Press, 1997, pp. 123-56.

Gale, David. "Pure Exchange Equilibrium of Dynamic Economic Models." *Journal of Economic Theory*, February 1973, *6*(1), pp. 12-36.

Gomme, Paul. "Money and Growth Revisited: Measuring the Costs of Inflation in an Endogenous Growth Model." *Journal of Monetary Economics*, August 1993, *32*(1), pp. 51-77.

Greenwood, Jeremy and Williamson, Stephen D. "International Financial Intermediation and Aggregate Fluctuations Under Alternative Exchange Rate Regimes." *Journal of Monetary Economics*, May 1989, *23*(3), pp. 401-31.

Hall, Robert E. "Intertemporal Substitution in Consumption." *Journal of Political Economy*, April 1988, *96*(2), pp. 339-57.

Hansen, Gary D. "The Cyclical and Secular Behaviour of the Labour Input: Comparing Efficiency Units and Hours Worked." *Journal of Applied Econometrics*, January-March 1993, *8*(1), pp. 71-80.

Haslag, Joseph H. "The Effects of Monetary Policy in a Model with Reserve Requirements." Working Paper 94-15, Federal Reserve Bank of Dallas, 1994.

Haslag, Joseph H. "A Comparison of Alternative Monetary Environments." Working Paper 95-11, Federal Reserve Bank of Dallas, 1995.

Huizinga, John and Mishkin, Frederic S. "Inflation and Real Interest Rates on Assets with Different Risk Characteristics." *Journal of Finance*, July 1984, *39*(3), pp. 699-712.

Hurd, Michael D. "Mortality Risk and Bequests." *Econometrica*, July 1989, *57*(4), pp. 779-813.

İmrohoroğlu, Ayse and Prescott, Edward C. "Evaluating the Welfare Effects of Alternative Monetary Arrangements." Federal Reserve Bank of Minneapolis *Quarterly Review*, Summer 1991, *15*(3), pp. 3-10.

Jones, Larry E. and Manuelli, Rodolfo E. "Growth and the Effects of Inflation." *Journal of Economic Dynamics and Control*, November 1995, *19*(8), pp. 1405-28.

Kandel, Schmuel and Stambaugh, Robert F. "Asset Returns and Intertemporal Preferences." *Journal of Monetary Economics*, February 1991, *27*(1), pp. 39-71.

King, Robert G. and Watson, Mark W. "Testing Long-Run Neutrality." Working Paper No. 4156, September 1992.

King, Robert G. and Watson, Mark W. "Testing Long-Run Neutrality." Federal Reserve Bank of Richmond *Economic Quarterly*, Summer 1997, *83*(3), pp. 69-101.

Koustas, Zisimos. "Canadian Evidence on Long-Run Neutrality Propositions." *Journal of Macroeconomics*, Spring 1998, *20*(2), pp. 397-411.

Lacker, Jeffrey and Schreft, Stacey L. "Money and Credit as Means of Payment." *Journal of Monetary Economics*, August 1996, *38*(1), pp. 3-23.

Laitner, John. "Random Earnings Differences, Lifetime Liquidity Constraints, and Altruistic Intergenerational Transfers." *Journal of Economic Theory*, December 1992, *58*(2), pp. 135-70.

Lucas, Robert E. Jr. "Inflation and Welfare." *Econometrica*, March 2000, *68*(2), pp. 247-74.

Marshall, David A. "Inflation and Asset Returns in a Monetary Economy." *Journal of Finance*, September 1992, *47*(4), pp. 1315-42.

Mehra, Rajnish and Prescott, Edward C. "The Equity Premium: A Puzzle." *Journal of Monetary Economics*, March 1985, *15*(2), pp. 145-61.

Michalewicz, Zbigniew. *Genetic Algorithms + Data Structures = Evolution Programs*. New York: Springer-Verlag, 1994.

Miller, Preston J. and Sargent, Thomas J. "A Reply to Darby." Federal Reserve Bank of Minneapolis *Quarterly Review*, Spring 1984, *8*(2), pp. 21-26.

Rapach, David E. "International Evidence on the Long-Run Impact of Inflation." *Journal of Money, Credit, and Banking*, February 2003, *35*(1), pp. 23-48.

Ríos-Rull, José-Víctor. "On the Quantitative Importance of Market Completeness." *Journal of Monetary Economics*, December 1994, *34*(3), pp. 463-96.

Ríos-Rull, José-Víctor. "Models with Heterogeneous Agents," in Thomas F. Cooley, ed., *Frontiers of Business Cycle Research*. Princeton: Princeton University Press, 1995, pp. 98-125.

Ríos-Rull, José-Víctor. "Life Cycle Economies and Aggregate Fluctuations." *Review of Economic Studies*, July 1996, *63*(3), pp. 465-89.

Samuelson, Paul A. "An Exact Consumption-Loan Model of Interest With or Without the Social Contrivance of Money." *Journal of Political Economy*, December 1958, *66*(6), pp. 467-82.

Sargent, Thomas J. and Wallace, Neil. "Some Unpleasant Monetarist Arithmetic." Federal Reserve Bank of Minneapolis *Quarterly Review*, Fall 1981, *5*(3), pp. 1-17.

Sargent, Thomas J. and Wallace, Neil. "The Real-Bills Doctrine Versus the Quantity Theory: A Reconsideration." *Journal of Political Economy*, December 1982, *90*(6), pp. 1212-36.

Sargent, Thomas J. and Wallace, Neil. "Interest on Reserves." *Journal of Monetary Economics*, May 1985, *15*(3), pp. 279-90.

Serletis, Apostolos and Koustas, Zisimos. "International Evidence on the Neutrality of Money." *Journal of Money, Credit, and Banking*, February 1998, *30*(1), pp. 1-25.

Weber, Axel A. "Testing Long-Run Neutrality: Empirical Evidence for G7-Countries with Special Emphasis on Germany." *Carnegie-Rochester Conference Series on Public Policy*, December 1994, *41*(0), pp. 67-117.

Wallace, Neil. "A Legal Restrictions Theory of the Demand for 'Money' and the Role of Monetary Policy." Federal Reserve Bank of Minneapolis *Quarterly Review*, Winter 1983, *7*(1), pp. 1-7.

Wallace, Neil. "Some of the Choices for Monetary Policy." Federal Reserve Bank of Minneapolis *Quarterly Review*, Winter 1984, *8*(1), pp. 15-24.

Wallace, Neil. "A Suggestion for Oversimplifying the Theory of Money." *Economic Journal, Supplement: Conference Papers*, 1988, *98*(390), pp. 25-36.

Williamson, Stephen D. "Financial Intermediation, Business Failures, and Real Business Cycles." *Journal of Political Economy*, December 1987, *95*(6), pp. 1196-1216.

## Appendix

## CHOOSING PARAMETERS BY MEANS OF NONLINEAR OPTIMIZATION

As we have indicated in the text, we use a genetic algorithm to learn about the irregular nonlinear map between the vector $\Delta$ of "deep parameters" and the endogenous variables whose values we target.[49] Given a candidate vector $j$ at algorithm time $s$, $\Delta_{js}$, we can calculate the solution to the agents' decision problem, and based on that information, we can find the implied steady-state equilibrium values for the targets associated with candidate vector $j$. We define a fitness criterion for a candidate vector $\Delta_{js}$ based on deviations of these implied values from targets. We use a genetic algorithm with real-valued coding, and operators providing tournament reproduction, three types of crossover, and non-uniform mutation, as explained below. Because our non-uniform mutation procedure slowly reduces the mutation rate to zero by time $T$, separate genetic algorithm searches can yield different best-fit candidate vectors $\Delta_T^*$. We conduct ten such searches and report the best-fit vectors.

[49] A more complete discussion of the principles of genetic algorithms would take us too far afield. For an introduction, as well as detailed description of the real-valued approach we use and the associated genetic operators, see Michalewicz (1994).

We begin by defining a fitness criterion across the nine targets of our system. We want to consider a criterion on the order of sum of squared deviations from target, but we also want the genetic algorithm to consider the fact that some targets are tighter than others in that the plausible deviation from them is smaller. Accordingly, we think of the target ranges as defining the space of plausible outcomes, and we design our fitness criterion to penalize candidate vectors $\Delta_{js}$ more severely if they deliver values outside the target range. This will prevent the genetic algorithm from spending a lot of time searching areas of the parameter space that are good on many dimensions but bad on a few dimensions. We assign penalty points based on deviations from target on each dimension. The penalty points are assigned linearly up to the boundary of the target, such that a candidate vector is penalized one point if a particular value is at the boundary of the target range. Outside the target range, an appropriately scaled quadratic penalty in the difference between the value and target boundary is added to the linear penalty. If we denote implied values of a candidate vector $\Delta_{js}$ by $\theta_{ijs}$, target values by $\theta_i^*$, and upper and lower target bounds by $\overline{\theta}_i$ and $\underline{\theta}_i$, respectively, where $i = 1, \ldots, 9$, then the fitness of the candidate vector, $F[\Delta_{js}]$, is given by

(A.1) $$F\left[\Delta_{js}\right]=\sum_{i=1}^{9}P_{ijs},$$

where

(A.2) $$P_{ijs}=\begin{cases}\left(\theta_{ijs}-\theta_i^*\right)/\left(\bar{\theta}_i-\theta_i^*\right)+\left(\theta_{ijs}-\bar{\theta}_i\right)^2/\bar{\theta}_i\left(\bar{\theta}_i-\theta_i^*\right) & \text{if } \theta_{ijs}>\bar{\theta}_i\\ \left(\theta_{ijs}-\theta_i^*\right)/\left(\bar{\theta}_i-\theta_i^*\right) & \text{if } \theta_{ijs}\in\left[\theta_i^*,\bar{\theta}_i\right]\\ \left(\theta_i^*-\theta_{ijs}\right)/\left(\theta_i^*-\underline{\theta}_i\right) & \text{if } \theta_{ijs}\in\left[\underline{\theta}_i,\theta_i^*\right]\\ \left(\theta_i^*-\theta_{ijs}\right)/\left(\theta_i^*-\underline{\theta}_i\right)+\left(\underline{\theta}_i-\theta_{ijs}\right)^2/\underline{\theta}_i\left(\underline{\theta}_i-\theta_i^*\right) & \text{if } \theta_{ijs}<\underline{\theta}_i\end{cases}.$$

This definition means that the better-fit vectors will have lower fitness values, and a vector that delivers an exact fit on all targets will have a fitness of zero.

The genetic algorithm is an iterative, directed search procedure acting on a population of $j$ candidate vectors at algorithm time $s$. At time $s$, the fitness of all candidate vectors in the population is calculated. To obtain the next set of candidate vectors, $\Delta_{js+1}$, we apply genetic operators. The first operator is tournament reproduction. We select two vectors at random with replacement from the time $s$ population. The vector with the better fitness value is copied into the time $s+1$ population. This operator is repeated enough times to produce a time $s+1$ population equal in size to the time $s$ population. Reproduction provides most of the evolutionary pressure in the search algorithm, but we need other operators to allow the system to experiment with new, untried candidate vectors. Crossover and mutation provide the experimentation and operate on the time $s+1$ population before the fitness values for that population are calculated.

To implement our crossover operators, we consider the time $s+1$ population two vectors, $j$ and $j+1$, at a time, and we implement the crossover operator with probability $p^c$. If crossover is to be performed on the two vectors, we use one of three methods with equal probability. In single-point crossover, we choose a random integer $i_{cross}\in[1,\ldots,9]$ and swap the elements of $\Delta_{j,s+1}$ and $\Delta_{j+1,s+1}$ where $i\geq i_{cross}$. In arithmetic crossover, we choose a random real $a\in[0,1]$ and create post-crossover vectors $a\Delta_{j,s+1}+(1-a)\Delta_{j+1,s+1}$ and $(1-a)\Delta_{j,s+1}+a\Delta_{j+1,s+1}$. In shuffle crossover, we exchange elements of $\Delta_{j,s}$ and $\Delta_{j+1,s}$ based on draws from a binomial distribution, such that if the $i$th draw is unity, the $i$th elements are swapped, otherwise the $i$th elements are not swapped. Each of these operators has been shown to have strengths in the evolutionary programing literature in certain types of difficult search problems, and we use them all here to improve the prospects for success.

We implement a non-uniform mutation operator that makes use of upper and lower bounds, $\bar{\delta}_i$ and $\underline{\delta}_i$, respectively, on the elements of a candidate vector $\Delta_{j+1,s}$. This operator is implemented with probability $p^m$ on element $\delta_{ij,s+1}$. If mutation is to be performed on the element, we choose a pair of random reals $r_1,r_2\sim U[0,1]$. The new, perturbed value of the element is then set according to

(A.3)
$$\delta_{ij,s+1}^{new}=\begin{cases}\delta_{ij,s+1}+\left(\bar{\delta}_i-\delta_{ij,s+1}\right)\left(1-r_2^{\left(\left(1-\frac{s}{T}\right)^b\right)}\right) & \text{if } r_1>0.5,\\ \delta_{ij,s+1}-\left(\delta_{ij,s+1}-\underline{\delta}_i\right)\left(1-r_2^{\left(\left(1-\frac{s}{T}\right)^b\right)}\right) & \text{if } r_1<0.5,\end{cases}$$

where $b$ is a parameter. With this mutation operator, the probability of choosing a new element far from the existing element diminishes as algorithm time $s\rightarrow T$, where $T$ is the maximum algorithm time. This operator is especially useful in allowing the genetic algorithm to more intensively sample in the neighborhood of the algorithm time-$s$ estimate of the best-fit vector in the latter stages of the search.

We conducted ten genetic algorithm searches to identify a best-fit deep parameter vector $\Delta_T^*$ according to our set of targets defined in Table 1.[50] The results are reported in Table A1.

[50] We set the parameters in the genetic algorithm, $\{\text{population}, p^c, p^m, T, b\}$ as $\{30, 0.95, 0.11, 1000, 2\}$ based on standards in the evolutionary programing literature. In our final search, we set $T=2500$, but we did not observe a commensurate improvement in performance, and so we did not pursue higher values of $T$ any further. We set the bounds on elements $\delta_i, i=1,\ldots,9$, according to [–0.3, 0.1], [1.1,40], [0.075, 0.33], [0.25, 0.4], [0.025, 0.075], [0.01, 0.04], [0.01, 0.08], [0.01, 0.4], [0.01,0.25]. This amounts to a set of constraints on the search to values that are typically viewed as economically plausible. We initialize the system by choosing elements in an initial population of vectors $\Delta$ randomly from uniform intervals defined by these bounds.

## Table A1

### Results of Nonlinear Optimization

| Search | *K/Y* | *I/K* | *B/Y* | *hcg* | *alt* | $I^m/Y$ | *H/Y* | *G/Y* | $T^K/G$ | *F* |
|---|---|---|---|---|---|---|---|---|---|---|
| 1 | 0.02 | 0.14 | 0.00 | 0.24 | 0.32 | 0.01 | 0.00 | 0.00 | 0.05 | 0.78 |
| 2 | 0.05 | 0.04 | 0.00 | 0.31 | 0.29 | 0.00 | 0.00 | 0.00 | 0.00 | 0.69 |
| 3 | 0.01 | 0.00 | 0.00 | 0.29 | 0.27 | 0.00 | 0.00 | 0.00 | 0.00 | 0.57 |
| 4 | 0.00 | 0.06 | 0.00 | 0.29 | 0.28 | 0.01 | 0.00 | 0.00 | 0.01 | 0.65 |
| 5 | 0.01 | 0.02 | 0.00 | 0.25 | 0.02 | 0.00 | 0.00 | 0.00 | 0.00 | 0.30 |
| 6 | 0.02 | 0.08 | 0.00 | 0.31 | 0.33 | 0.00 | 0.00 | 0.00 | 0.02 | 0.76 |
| 7 | 0.01 | 0.01 | 0.00 | 0.29 | 0.31 | 0.00 | 0.00 | 0.00 | 0.00 | 0.61 |
| 8 | 0.08 | 0.06 | 0.00 | 0.33 | 0.31 | 0.00 | 0.00 | 0.00 | 0.00 | 0.78 |
| 9 | 0.04 | 0.04 | 0.00 | 0.33 | 0.46 | 0.00 | 0.00 | 0.00 | 0.03 | 0.91 |
| 10 | 0.06 | 0.02 | 0.00 | 0.30 | 0.23 | 0.00 | 0.00 | 0.00 | 0.01 | 0.62 |

NOTE: The fit to the data. The entries are deviations from target, by target and total, in penalty points, for each of the ten searches we conducted. Columns 2 through 10 are averages across algorithm time $s = T$ populations. Within algorithm time $s = T$ populations, we found little or no variation across fitness components.

We find that the algorithm time-$T$ population of parameter vectors $\Delta_{jT}^{*}$ provide a close fit on our target data. The only quantitatively significant discrepancies from targets occur on individual consumption growth and individual time devoted to market, and then the implied values are typically only 0.2 to 0.35 of a penalty point from target, meaning that implied values on these dimensions lie away from the target only 20 to 35 percent of the distance between the target and a target bound. We found little or no variation among individual parameters within algorithm time $s = T$ populations. Across searches, we found some variance, almost all of it in the preference parameters. The estimates of the elasticity of intertemporal substitution, for instance, ranged from a low of 0.114 to a high of 0.185. Search number 5 provided the best overall fit, so we use the parameters from this search in the baseline specification of our model.