

# 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 discussion 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.  
© 2004, The Federal Reserve Bank of St. Louis.

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.<sup>1</sup> 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.<sup>2</sup> 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

<sup>1</sup> 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.

<sup>2</sup> 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.<sup>3</sup>

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

<sup>3</sup> 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.<sup>4</sup>)

<sup>4</sup> 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.

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.<sup>5</sup>

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.<sup>6</sup> 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.<sup>7</sup> They assume that these fiscal policy

<sup>5</sup> 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.

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

<sup>7</sup> 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.<sup>8</sup> 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.

---

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.

<sup>8</sup> 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.

**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.<sup>9</sup> 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.<sup>10</sup>

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.<sup>11</sup> 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.

---

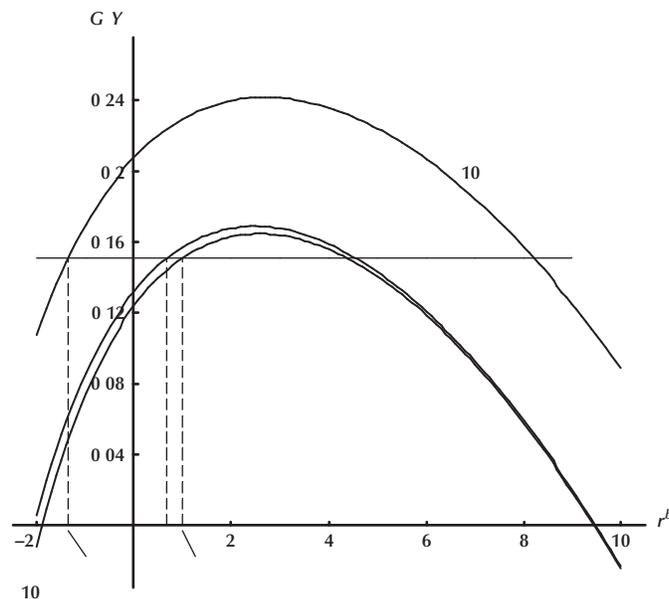
<sup>9</sup> Part of the reason for this is that in Samuelson-case equilibria, capital is overaccumulated: The marginal return to capital is smaller than the marginal cost of maintaining the capital stock via investment. A reduction in the real interest rate increases the capital stock and increases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

<sup>10</sup> 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.

<sup>11</sup> 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.<sup>12</sup> 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.

<sup>12</sup> 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.<sup>13</sup>

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.<sup>14</sup> 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.<sup>15</sup> 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

<sup>13</sup> 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.

<sup>14</sup> 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.

<sup>15</sup> 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.<sup>16</sup> 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.<sup>17</sup>

### 3.2 Primitives

**3.2.1 Demographics.** A generation of identical households is born at each discrete date  $t = \dots, -2, -1, 0, 1, 2, \dots$  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  $e_i$ ,  $i = 1, \dots, n$ . 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

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

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  $\ell_t(t + j)$ ,  $j = 0, \dots, n - 1$ , respectively.

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

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

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  $j = 0, \dots, n - 1$  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  $\nu$ , where  $\nu \equiv 1 - \eta(1 - \gamma)$ .

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

<sup>16</sup> An exception is Altig and Carlstrom (1991).

<sup>17</sup> 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  $-\varphi$ .<sup>18</sup>

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  $j = n^* - 1, \dots, n - 1$ .

### 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.<sup>19</sup>

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.<sup>20</sup> 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  $S^b(t) \equiv B(t) - R^b(t-1)B(t-1)$ .<sup>21</sup>

<sup>18</sup> The easiest way to see this is to consider a two-period endowment model with inelastic labor supply and preferences given by  $u = \sum_{i=0}^1 (1-\gamma)^{-1} \beta^i c_i(t+i)^{1-\gamma}$ . 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  $1 - (1 + \rho)^{-1/\gamma}$ .

<sup>19</sup> We assume that households who attempt to make private loans directly would face prohibitively high transactions costs.

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

<sup>21</sup> 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

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

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

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

### 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  $R^{ka}(t) \equiv (1 - \tau^c)R^{kn}(t) + \tau^c R^h(t)$ .<sup>22</sup> 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:

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

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-

<sup>22</sup> 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.<sup>23</sup>

**3.4.2 Household Decisions.** Households maximize (2) subject to a lifetime budget constraint that we will now define. We will let  $a_t(t+j)$ ,  $j = 0, \dots, n-1$ , 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)

$$\begin{aligned} c_t(t) + a_t(t) &= (1 - \tau^w)w(t)e_1(1 - \ell_t(t)), \\ c_t(t+j) + a_t(t+j) &= (1 - \tau^w)w(t+j)e_{j+1}(1 - \ell_t(t+j)) \\ &\quad + \hat{R}(t+j-1)a_t(t+j-1) \end{aligned}$$

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

$$\begin{aligned} c_t(t+n-1) &= (1 - \tau^w)w(t+n-1)e_n(1 - \ell_t(t+n-1)) \\ &\quad + \hat{R}(t+n-2)a_t(t+n-2), \end{aligned}$$

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

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

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

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

The credit-market clearing condition is

$$(7) \quad A(t) = M(t) + K(t+1) + B(t),$$

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

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

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

<sup>23</sup> 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)

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

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:

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

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

$$g = \zeta[R^h, R^k, F, \Delta].$$

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.<sup>24</sup>

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

<sup>24</sup> 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 non-linear 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

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

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  $R^{ba} = 1$ .<sup>25</sup> 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).<sup>26</sup>

## 4.3 The Deep Parameter Vector

The deep parameter vector is

$$(12) \quad \Delta = \left[ \rho, \gamma, \eta, \alpha, \delta, \xi, \phi, \tau^w, \tau^c \right].$$

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 intra-temporal labor-leisure substitution elasticity (or labor share of time)  $\eta$ , the capital share of output  $\alpha$ , the net depreciation rate  $\delta$ , the unit intermediation

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

<sup>26</sup> 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  $i = 21, \dots, 76$ , 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.

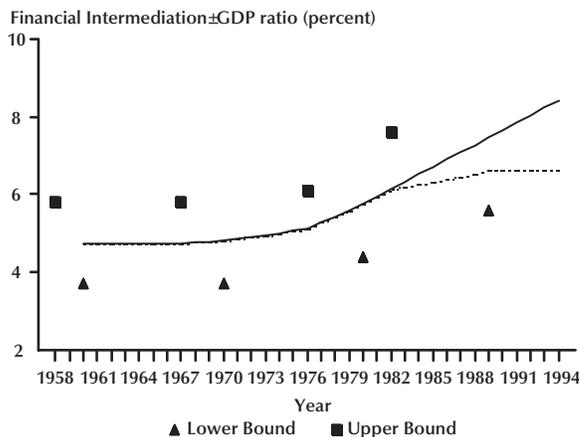
The estimates of the average capital-output ratio  $K/Y$  and the average investment-capital ratio  $I/K$  are due to Cooley and Prescott (1995). These estimates are based on a broad definition of capital that includes consumer durables and government capital. We target the equilibrium bonds-to-output 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.<sup>27</sup> As we have indicated,

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.

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 over-emphasize 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.<sup>28</sup>

<sup>27</sup> 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.

<sup>28</sup> 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  $\phi$ , 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  $\phi$  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  $\phi$  ( $-0.0078$ ) 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).<sup>29</sup>

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.<sup>30</sup> 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

<sup>29</sup> 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.

<sup>30</sup> 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
<b>Aggregate performance</b>			
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
<b>Preferences</b>			
Effective rate of time preference	$\phi$	-0.0067	Open
CRRA	$\nu$	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
<b>Asset holdings</b>			
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
<b>Technology</b>			
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.<sup>31</sup> 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.

<sup>31</sup> 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
<b>Intermediation</b>			
Intermediation-output ratio	$I^m/Y$	0.0599	0.06
<b>Government size</b>			
Government-output ratio	$G/Y$	0.151	0.151
Revenue from firms	$T^k/G$	0.119	0.119
<b>Government revenue sources</b>			
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 overaccumulation. 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.<sup>32</sup> 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.<sup>33,34</sup>

<sup>32</sup> 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).

<sup>33</sup> 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.

While our baseline steady state matches postwar 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 policymakers—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

<sup>34</sup> 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%
İmrohorođlu/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.<sup>35</sup>

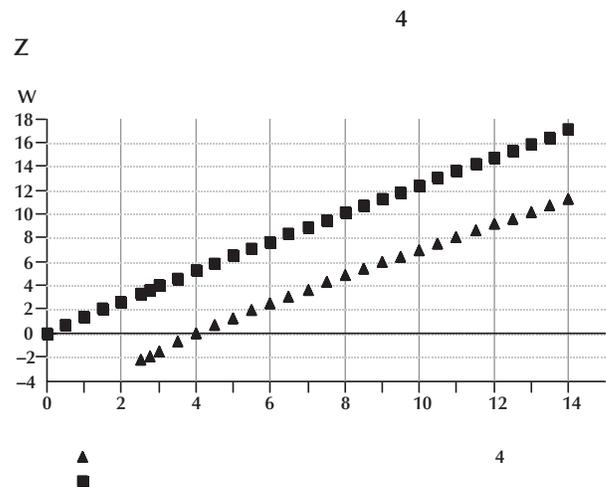
There are no triangles in the figure for inflation rates lower than 2.5 percent. The reason for this is that, once the inflation rate falls below this level, the government faces a revenue shortage: No steady-state real interest rate produces enough revenue to meet our target value for *G/Y*. We adopt 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.<sup>36</sup> 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 per-

<sup>35</sup> 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.

<sup>36</sup> 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 3**

**The Welfare Cost of Inflation**

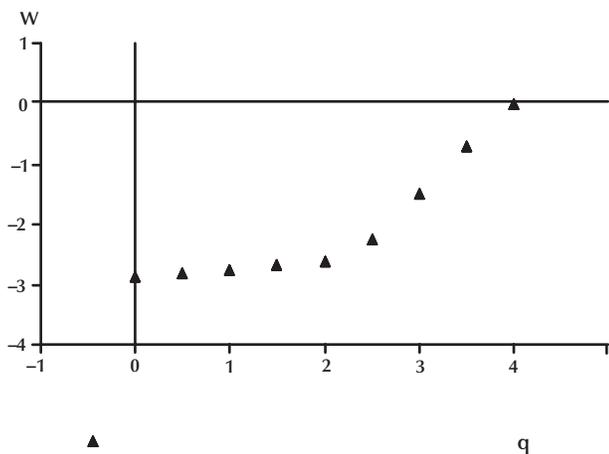


centage point increase in the inflation rate. These results illustrate the fact that even when the initial inflation rate is low, the marginal distortion produced by increases in the inflation rate can be quite large.

**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.

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.<sup>37</sup>

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

<sup>37</sup> 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.<sup>38</sup> 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.<sup>39</sup> 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.<sup>40</sup>

<sup>38</sup> 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.

<sup>39</sup> 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.

<sup>40</sup> 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.<sup>41</sup> 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

<sup>41</sup> 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 non-stochastic 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  $-0.048$ .<sup>42</sup> 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.<sup>43</sup>

<sup>42</sup> 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  $-0.041$ .

<sup>43</sup> We also conducted welfare-cost experiments from a baseline steady state with  $\sigma = 1/29$ , an estimate based on Kandel and Stambaugh

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-

(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.<sup>44</sup>

## 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.<sup>45</sup>)

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

<sup>44</sup> 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.

<sup>45</sup> 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.<sup>46</sup>

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  $T$ ,  $T-1$ ,  $T-2$ , ...,  $T-54$ , 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

<sup>46</sup> 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.<sup>47</sup>

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.<sup>48</sup>

## 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.

<sup>47</sup> 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.

<sup>48</sup> 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, Guglielmo. "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-Superneutrality 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, Schmuël 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.
- Kousta, 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.<sup>49</sup> 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  $\Delta_j^*$ . We conduct ten such searches and report the best-fit vectors.

<sup>49</sup> 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  $\underline{\theta}_i$  and  $\bar{\theta}_i$ , respectively, where  $i = 1, \dots, 9$ , then the fitness of the candidate vector,  $F[\Delta_{js}]$ , is given by

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

where

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

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,  $\Delta_{j,s+1}$ , 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  $i_{cross} \in [1, \dots, 9]$  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 programming 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

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

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  $\Delta_7^*$  according to our set of targets defined in Table 1.<sup>50</sup> The results are reported in Table A1.

<sup>50</sup> We set the parameters in the genetic algorithm, {population,  $p^c, p^m, T, b$ } as {30, 0.95, 0.11, 1000, 2} based on standards in the evolutionary programming 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  $\delta_i, i = 1, \dots, 9$ , 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	<i>K/Y</i>	<i>I/K</i>	<i>B/Y</i>	<i>hcg</i>	<i>alt</i>	<i>I<sup>m</sup>/Y</i>	<i>H/Y</i>	<i>G/Y</i>	<i>T<sup>K</sup>/G</i>	<i>F</i>
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.

