WORKING PAPERS
RESEARCH DEPARTMENT

WORKING PAPER NO. 10-2
THE IMPLICATIONS OF INFLATION IN AN ESTIMATED NEW-KEYNESIAN MODEL

Pablo A. Guerron-Quintana
Federal Reserve Bank of Philadelphia

December 28, 2009

RESEARCH DEPARTMENT, FEDERAL RESERVE BANK OF PHILADELPHIA
Ten Independence Mall, Philadelphia, PA 19106-1574 • www.philadelphiafed.org/research-and-data/

The Implications of Inflation in an Estimated New-Keynesian Model

Pablo A. Guerron-Quintana*

Research Department, FRB Philadelphia, Philadelphia, PA, USA

December 29, 2009

Abstract

This paper studies the steady state and dynamic consequences of inflation in an estimated dynamic stochastic general equilibrium model of the U.S. economy. It is found that 10 percentage points of inflation entails a steady state welfare cost as high as 13 % of annual consumption. This large cost is mainly driven by staggered price contracts and price indexation. The transition from high to low inflation inflicts a welfare loss equivalent to 0.53%. The role of nominal/real frictions as well as that of parameter uncertainty is also addressed.

*I thank Thomas Grennes, Michel Julliard, John Lapp, Douglas Pearce, Stephanie Schmitt-Grohe, Kei-Mu Yi, an anonymous referee, and seminar participants at the 2008 Midwest Macroeconomics Meetings for comments. A substantial part of this paper was written while I was an assistant professor of economics at NC State University. Financial support from the Gill and FRDP grants is greatly appreciated. All remaining errors are mine. The views expressed here are those of the author and do not necessarily reflect those of the Federal Reserve Bank of Philadelphia or of the Federal Reserve System. This paper is available free of charge at www.philadelphiafed.org/research-and-data/publications/working-papers/. Email: pablo.guerron@phil.frb.org.

1 Introduction

Undoubtedly, the period 1990-2006 will be remembered as one of relative prosperity in the U.S., characterized by high growth and low and stable inflation. Output, for example, averaged an annual growth rate of 3% while the mean annual CPI inflation was 3%, for that period. Yet recent developments in international commodity and financial markets have induced substantial level changes on inflation. Indeed, annualized CPI inflation in the U.S. was close to 6% in July 2008, a level not seeing seen since the early 1990s.[1] Although inflation has subsequently retreated, the low fed funds rate of recent months has sparked renewed interest among economic observers, who argue that inflation may be back sooner than expected. It is unlikely that we will reach the inflation rates of the 1970s, but these variations in prices invite us to revisit some old questions: What are the welfare consequences of inflation? And more important, how much would society willingly sacrifice to avoid, say, 10 percentage points of inflation? This paper tries to answer these questions from the perspective of an estimated New Keynesian model.

Evaluating the (unpleasant) impact of inflation on society has been a recurrent topic in macroeconomics that can be traced back to the seminal contributions of Bailey (1956) and Friedman (1969). This research agenda has typically pursued two distinct approaches. The first line follows Bailey (1956) in measuring the welfare cost of inflation as the area under the money demand curve. Under this tradition, money is a special consumption good while inflation is a direct tax on it. Hence, large inflations are welfare reducing, since they make holding real balances costly. For example, Bailey established that a 10-percentage-point drop in steady state inflation entails a welfare gain equivalent to 1% of annual income. Over the years, authors such as Lucas (1981, and 2000), and Fischer (1981) have persistently found welfare estimates smaller than Bailey's. More recently, Ireland (2008) and Khan, King, and Wolman (2003) estimate the welfare cost to be as low as 0.20% and 0.05%, respectively.

The second strand of the literature has explored the cost of inflation in a general equilibrium context. In these models inflation is costly because households must divert productive time into leisure and financial activities aimed at saving on real cash balances (Ireland, 1997). Cooley and Hansen (1989) and Burdick (1997), for instance, find only modest gains (around 0.5%) of low inflation in a highly stylized real business cycle model. Dotsey and Ireland (1996) study inflation in a richer general equilibrium framework and report that an inflation of 4 percentage points entails a welfare cost as high as 1% of annual output. Furthermore, Ireland (1997) shows that the presence of sticky price contracts only exacerbates the neg-

[1] International data show that rising prices were not only a problem at home but everywhere, with inflation averaging 6% in Belgium, 6.3% in China, 16% in Russia, and 33% in Venezuela during 2008.

ative consequences of inflation. Finally, Schmitt-Grohe and Uribe (2004 and 2005) explore the optimal inflation rate in fully fledged DSGE models.

The contribution of this paper falls within the second line of research. Specifically, I study the welfare implications of inflation by employing a fairly standard DSGE model, entertaining features such as price/wage sluggishness, habit formation, and costly adjustment of investment. The proposed model borrows concepts from Altig, Christiano, Eichenbaum, and Linde (2005, henceforth ACEL) and the important contribution of Christiano, Eichenbaum, and Evans (2005, henceforth CEE). The presence of real and nominal frictions gives the model a more realistic flavor and facilitates comparisons between the predictions of the medium scale New Keynesian models with those from more parsimonious formulations, e.g., Dotsey and Ireland (1996) and Ireland (1997). As will become clear, the predictions from the two setups can be quite different.

When deciding how much in real balances to keep, households confront two tensions in the model. They enjoy utility from directly holding positive money balances as in Sidrausky (1967). However, each dollar kept for utility purposes forgoes a positive return that would be earned if deposited in a financial intermediary. This dual role of money gives rise to a well-defined money demand function as in Khan et al. (2003). This money demand equation has two appealing properties: 1) it is well suited for estimation, and 2) it allows us to evaluate the taxational aspect of inflation as in Bailey (1956). Of course, Sidrausky's method is only one of many ways to justify the presence of money in the economy. For example, Cooley and Hansen (1989) propose a cash-in-advance formulation to analyze the implications of inflation. More recently, Aruoba and Schorfheide (2008) study welfare and prices using a search-based model of money balances.

Studying the cost of inflation in an estimated DSGE model poses some interesting challenges. To begin with, ACEL and CEE estimate a small interest rate semi-elasticity of the demand for money, which seems necessary to properly account for the high frequency properties of the data. However, Lucas (1981 and 2000) argues that the long-run semi-elasticity is the right choice for welfare analysis.[2] Hence I propose a flexible money demand formulation, which can simultaneously capture the short- and long-run properties of money demand. The key ingredient in this formulation is that re-balancing the composition of money balances for utility purposes or for bank deposits is costly. To capture this cost, the model assumes that households use time-dependent rules to re-optimize their money holdings. The presence of those costs in turn implies that households look forward when re-balancing their portfolios between cash and deposits.

[2] For the rest of the paper, I will use the terms semi-elasticity of money demand and semi-elasticity as shorthand for interest rate semi-elasticity of money demand.

Even though the model gives rise to a rich dynamic money demand equation, minimum-distance estimators tend to recover the short-run properties of the data, resulting in the small elasticities reported in the literature (ACEL, and CEE). Therefore, the model is estimated using Bayesian methods similar to those applied in Schorfheide (2000) and Smets and Wouters (2007). This approach has the advantage that priors can be used to simultaneously recover the short- and long-run elasticities of the demand for money. Additionally, the Bayesian methodology allows us to assess the impact of parameter uncertainty in the welfare calculations. It will become clear that this type of uncertainty significantly affects the steady state welfare estimates.

In the benchmark formulation, which includes several nominal/real frictions, an annual inflation of 10 percentage points entails a steady state welfare cost equivalent to around 13% of annual consumption or 6.5% when measured in annual output. This relatively large cost of inflation mainly arises from staggered price contracts and price indexation, which induce significant steady state price dispersion. Habit formation and the interest semi-elasticity of money demand also contribute to making inflation costly, although to a lesser degree. For example, if money demand is relatively inelastic as in CEE, the steady state welfare cost drops to 9.5% of annual consumption. It is also shown that the central bank's response to inflation in the Taylor rule has strong implications for welfare. A mild response, for instance, results into a smaller welfare cost of inflation. Finally, the estimated model implies that Bailey's taxation aspect of inflation imposes a welfare loss of about 1% of consumption, a result consistent with previous studies.

Following Ireland (1997), the estimated model is used to evaluate the transitional costs of moving to a lower inflation state. I find that this transition is welfare reducing, but it is only a small fraction of the benefits from living in a low inflation environment. In fact, the transition amounts for a welfare loss of 0.53% of annual consumption in the benchmark specification. The absence of sticky prices or habit formation makes the transition less costly.

The rest of the paper is organized as follows. Section 2 describes the baseline model, including the money demand formulation. I describe the estimation technique and report estimated parameters in section 3. The welfare analysis is presented in sections 4 and 5. Finally, section 6 contains some concluding remarks.

2 Model

The model builds on ACEL, CEE, and Schmitt-Grohe and Uribe (2005). Since this type of environment has been extensively discussed in the literature, I provide a brief discussion, omitting lengthy derivations. The main features of the model can be summarized as follows:

The economy grows along a stochastic path; prices, wages, and money holdings are assumed to be sticky à la Calvo; preferences display external habit formation; investment adjustment is costly; and finally, there are five sources of uncertainty: neutral and capital embodied technology shocks, preference, government, and monetary shocks.

2.1 Firms

There is a continuum of monopolistically competitive firms indexed by $j \in [0,1]$, each producing a final good out of capital services, k_j, and labor services, $L_{j,t}$. The technology function is given by $k_{j,t}^{\alpha}\left(S_t^L L_{j,t}\right)^{1-\alpha} - S_t^* \psi$; the term ψ makes profits equal to zero in the steady state. S_t^* is the stochastic growth path of the economy (see below for its definition).[3] The neutral technology shock, S_t^L, grows at rate g_t^L, which is assumed to follow the process

$$\ln g_t^L = (1-\rho_{g^L}) \ln g_{ss}^L + \rho_{g^L} \ln g_{t-1}^L + \sigma^{g_L} \varepsilon_{g^L,t},$$

where $\varepsilon_{g^L,t}$ is distributed $\mathbb{N}(0,1)$.

Firms rent capital and labor in perfectly competitive factor markets. I assume that workers must be paid in advance. As a consequence, firms must borrow the wage bill, $W_t L_{j,t}$, from a financial intermediary. The loan plus the interest rate, R_t, must be repaid at the end of the period.

Firms choose prices to maximize the present value of profits; prices are set in a Calvo fashion; that is, each period, firms optimally revise their prices with an exogenous probability $1-\xi_p$. If, instead, a firm does not re-optimize its price, then the price is updated according to the rule: $P_{j,t} = (\pi_{t-1})^{\chi} P_{j,t-1}$, where π_{t-1} is the economy-wide inflation in the previous period and $\chi \in [0,1]$. By allowing partial indexation in the price rule, I follow the common practice in the literature (Schmitt-Grohe and Uribe, 2004, and Fernandez-Villaverde and Rubio-Ramirez, 2008). An optimizing firm at time t sets prices according to the program

$$\max_{P_{j,t}} \mathbb{E}_t \sum_{n=0}^{\infty} \left(\xi_p \beta\right)^n \lambda_{t+n} \left[\frac{P_{j,t} \Pi_{\tau=0}^{n-1} \left(\pi_{t+\tau}\right)^{\chi}}{P_{t+n}} y_{t+n}(j) - mc_{t+n} y_{t+n}(j)\right],$$

Here, P_t is the price index, $y_t(j)$ is the aggregate demand for good type j, mc_t is firm j's marginal cost, β is the discount factor, and λ_t is the marginal utility of consumption at time t.

[3] The growth term is needed to have a well-defined steady state around which we can solve the model.

2.2 Households

The economy is populated by a continuum of households indexed by i. Every period households must decide how much to consume, work, and invest. In addition, they must choose the amount of money to be sent to a financial intermediary. I assume that agents in the economy have access to complete markets; such an assumption is needed to eliminate wealth differentials arising from wage heterogeneity (CEE, and Erceg, Henderson, and Levin, 2000). Households maximize the expected present discounted value of utility

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[S_t^{Uc} \log(C_{i,t} - bC_{t-1}) - \Phi \frac{L_{i,t}^{1+1/\gamma}}{1+1/\gamma} + \psi_m \left(\frac{M_{i,t}}{S_t^* P_t} \right)^{1-\zeta_m} \right] \tag{1}$$

subject to

$$\begin{aligned} P_t C_{i,t} + \frac{P_t}{S_t^K} \left(I_{i,t} + a(x_t) K_{i,t} \right) + \mathcal{M}_{i,t} &= R_t(\mathcal{M}_{i,t-1} - M_{i,t} + T_t) + R_t^K x_t K_{i,t} + W_{i,t} L_{i,t} + M_{i,t} + A_{i,t}, \\ K_{i,t+1} &= (1-\delta) K_{i,t} + I_{i,t} \left(1 - \Gamma\left(\frac{I_{i,t}}{I_{i,t-1}} \right) \right). \end{aligned}$$

Here, S^{Uc} is a preference shock that follows the process $\log S_t^{Uc} = \rho_{Uc} \log S_{t-1}^{Uc} + \sigma^{Uc} \varepsilon_{U_c,t}$ with $\varepsilon_{U_c,t}$ distributed $\mathbb{N}(0,1)$; preferences display external habit formation, measured by $b \in (0,1)$; and Γ is a function reflecting the costs associated with adjusting investment. This function is assumed to be increasing and convex satisfying $\Gamma = \Gamma' = 0$ and $\kappa \equiv \Gamma'' > 0$ in steady state. $\mathcal{M}_{i,t-1}$ is household i's beginning of period t stock of money, whereas T_t is a lump-sum transfer by the government. Households send the amount $\mathcal{M}_{i,t-1} - M_{i,t} + T_t$ to a financial intermediary where it earns the interest rate, R_t. The stochastic trend, $S_t^* = S_t^L \left(S_t^K \right)^{\alpha/(1-\alpha)}$, in the money term is required to have a well-defined steady state. The term S_t^K is an investment-specific shock whose growth rate obeys

$$\log g_t^K = (1-\rho_{g^K}) \log g_{ss}^K + \rho_{g^K} \log g_{t-1}^K + \sigma^{g^K} \varepsilon_{g^K,t},$$

where $\varepsilon_{g^K,t}$ is distributed $\mathbb{N}(0,1)$.

As in ACEL, CEE, and Schmitt-Grohe and Uribe (2004), I assume that physical capital can be used at different intensities. Furthermore, using the capital with intensity x_t entails a cost $a(x_t)$, which satisfies $a(1) = 0; a"(1) > 0; a'(1) > 0$. For future reference, define $\varkappa_a = a"(1)$. The term $A_{i,t}$ captures net payments from complete markets and government bonds, and profits from producers. The individual consumption good is assumed to be a composite made of differentiated goods indexed by j according to the aggregator

$$C_{i,t}=\left(\int_0^1 c_t(i,j)^{\frac{\zeta-1}{\zeta}}dj\right)^{\frac{\zeta}{\zeta-1}},\quad 1\leq\zeta<\infty,$$

where $c(i,j)$ is the demand of household i for good type j. With this type of composite good, the demand for goods of type j is given by $c(i,j)=\left(\frac{P_{j,t}}{P_t}\right)^{-\zeta}C_{i,t}$. Here, the nominal price index is $P_t=\left(\int_0^1 P_{j,t}^{1-\zeta}dj\right)^{\frac{1}{1-\zeta}}$. Similarly, I assume that individual investment obeys $I_{i,t}=\left(\int_0^1 I_t(i,j)^{\frac{\zeta-1}{\zeta}}dj\right)^{\frac{\zeta}{\zeta-1}}$. As with consumption, $I(i,j)$ denotes household i's demand for investment good of type j.

2.3 Wage Setting

Following Erceg, Henderson, and Levin (2000), I assume that each household is a monopolistic supplier of a differentiated labor service, $L_{i,t}$. Households sell these labor services to a competitive firm that aggregates labor and sell it to final firms. The technology used by the aggregator is

$$\widetilde{L}_t=\left[\int_0^1 L_{i,t}^{\frac{\zeta_w-1}{\zeta_w}}dj\right]^{\frac{\zeta_w}{\zeta_w-1}},\quad 1\leq\zeta_w<\infty.$$

It is straightforward to show that the relation between the labor aggregate and the wage aggregate, W_t, is given by $L_{i,t}=\left[\frac{W_t}{W_{i,t}}\right]^{\zeta_w}\widetilde{L}_t$.To induce wage sluggishness, I assume that households set their wages in Calvo fashion. In particular, with exogenous probability ξ_w a household does not re-optimize wages each period. If this is the case, wages are set according to the rule of thumb $W_{i,t}=(\pi_{t-1})^{\chi_w}W_{i,t-1}$. Following Schmitt-Grohe and Uribe (2004) and Fernandez-Villaverde and Rubio-Ramirez (2008), the wage rule allows for partial indexation with parameter $\chi_w\in[0,1]$. Similar to the firms, households set wages according to the program

$$\max_{W_{i,t}}\mathbb{E}_t\sum_{n=0}^{\infty}(\xi_w\beta)^n\left[-\Phi\frac{L_{i,t+n}^{1+1/\gamma}}{1+1/\gamma}+\lambda_{t+n}\frac{W_{i,t}\Pi_{\tau=0}^{n-1}(\pi_{t+\tau})^{\chi_w}}{W_{t+n}}\frac{W_{t+n}}{P_{t+n}}L_{i,t+n}\right].$$

The marginal utility of consumption, λ, is not indexed by i, reflecting our assumption of complete markets.

2.4 Demand for Money

As previously discussed, modeling money demand needs to account for two important regularities found in the literature. On the one hand, authors such as ACEL and CEE argue that in the context of DSGE models the short-run demand for money is what matters. In particular, they estimate an interest rate semi-elasticity of money demand around 1 (this finding is robust across different econometric techniques). On the other hand, studies about the welfare implications of inflation stress the importance of the long-run properties of money demand. For example, Lucas (2000) estimates that the long-run semi-elasticity of money demand lies between 5 and 7. Based on these numbers, he finds the welfare costs of inflation to be on the order of 1% of annual income. Furthermore, an extrapolation of his results implies that for a semi-elasticity of 1 the welfare cost is roughly 0.2%. This evidence raises the following dilemma: Too much elasticity delivers sizable welfare costs but worsens the short-run dynamics of money. Alternatively, low elasticities provide the right high frequency description of money at the expense of predicting too low welfare costs.

A simple yet formal way to solve the money demand dilemma is to assume time-dependent portfolio adjustment; i.e., agents re-optimize their money balances, M, infrequently, similar in spirit to the price- and wage-setting model of Christiano, Eichenbaum, and Evans (2005). Specifically, a fraction, $1-\xi_m$, of randomly chosen households is allowed to re-optimize their balances every period. As far as inactive households, the literature on portfolio choice provides little guidance regarding their behavior (Campbell and Viceira, 2002). Hence, if a household is not allowed to re-optimize today, its money holdings are adjusted according to the rule $M_{i,t} = \pi_{t-1} g^*_{t-1} M_{i.t-1}$, where π_{t-1} represents the last period inflation, and g^* is the growth rate of the aggregate shock S^*.[4] This rule does not allow for partial indexation, since initial estimation attempts clearly showed that the indexation parameter was not identified.

As argued in Guerron-Quintana (2009), the sticky money assumption is likely to capture two important aspects of the economy. First, it proxies the degree of access to financial and banking services enjoyed by households. Prior to the widespread use of ATMs, electronic banking, and the branching liberalization of the 1980s, households spent an important amount of resources managing their accounts. Consequently, households had limited access to such services, which is parsimoniously captured in the model by infrequent portfolio re-balancing.

Second, the time-dependent assumption captures the costs faced by households when assessing the uncertainty surrounding the economy and the financial system. The presence of large costs makes it harder for households to determine the state of the economy and in

[4] The presence of g in the indexation rule implies that there are no distortions from portfolio dispersion along the steady state growth path.

particular the risk exposure of banks. As a consequence, households may opt to limit their participation in financial markets. We can also think of the portfolio friction as indirectly capturing the infrequent participation of trading agents in the equity market reported by Vissing-Jorgensen (2003). As before, I interpret this infrequent re-optimization as the result of costs faced by households. The basic idea is that in the presence of these costs, households fully optimize their portfolio only periodically and follow simple rules for changing their portfolio at other times.

The staggered money setting and the functional forms for the utility function imply that an optimizing (active) household at time t chooses money balances according to the program

$$\max_{M_{i,t}} \mathbb{E}_t \sum_{n=0}^{\infty} (\xi_m \beta)^n \left[\psi_m \frac{\left(\frac{M_{i,t} \Pi_{\tau=0}^{n-1} (g^*_{t+\tau} \pi_{t+\tau})}{S^*_{t+n} P_{t+n}} \right)^{1-\zeta_m}}{1-\zeta_m} - \frac{\lambda_{t+n}}{P_{t+n}} (R_{t+n} - 1) M_{i,t} \Pi_{\tau=0}^{n-1} (g^*_{t+\tau} \pi_{t+\tau}) \right].$$

As shown in the appendix, the solution to the previous program gives rise to a money demand for active households, which requires that the expected marginal benefit of an extra dollar (enjoy additional utility) equals its expected marginal cost (forgone interest rate), i.e.,

$$\underbrace{\psi_m \left(\frac{M^*_t}{S^*_t P_t} \right)^{-\zeta_m} x^1_{m,t}}_{\textit{Marginal Benefit}} = \underbrace{x^2_{m,t}}_{\textit{Marginal Cost}}, \tag{2}$$

where the terms x^1_m and x^2_m are given by

$$x^1_{m,t} = 1 + \beta \xi_m \mathbb{E}_t \left(\frac{g^*_t \pi_t}{g^*_{t+1} \pi_{t+1}} \right)^{1-\zeta_m} x^1_{m,t+1}, \text{ and } x^2_{m,t} = \lambda_t S^*_t (R_t - 1) + \beta \xi_m \mathbb{E}_t \frac{g^*_t \pi_t}{g^*_{t+1} \pi_{t+1}} x^2_{m,t+1}.$$

Here, M^*_t is the money holdings optimally chosen by active households today. Equation (2) implies that the annualized short- and long-run semi-elasticities of money demand are given by

$$E_{SR} = -\frac{(1-\xi_m)}{4(R-1)\zeta_m}, \text{ and } E_{LR} = -\frac{1}{4(R-1)\zeta_m}, \tag{3}$$

respectively; here, R is the steady state quarterly interest rate (see the appendix for details). As long as $\xi_m > 0$, we see that the short-run elasticity is smaller than its long-term counterpart, $|E_{SR}| < |E_{LR}|$. Consequently, the curvature parameter, ζ_m, can be used to describe the money demand in the long run as required by welfare analysis. Furthermore, we can control the short-run dynamics of money via the sluggishness coefficient, ξ_m. This ability to capture the short- and long-run properties of money through separate parameters is exploited in the estimation section.

2.5 Government

The monetary authority sets the quarterly interest rate according to a Taylor rule. In particular, the central bank smooths interest rates and responds to deviations of actual inflation from steady state inflation, π, and deviations of output from its trend level, $(Y/S^{*})_t$.

$$\frac{R_t}{R} = \left(\frac{R_{t-1}}{R}\right)^{\rho_r} \left[\left(\frac{\pi_t}{\pi}\right)^{\phi_\pi} \left(\frac{Y_t/S_t^{*}}{(Y/S^{*})_t}\right)^{\phi_y}\right]^{1-\rho_r} \exp(\sigma_m \varepsilon_{m,t}). \tag{4}$$

The term $\varepsilon_{m,t}$ is a random shock to the systematic component of monetary policy and is assumed to be standard normal; σ_m is the size of the monetary shock. Other authors have implemented similar Taylor rules, e.g., Del Negro et al. (2004) and Justiniano and Primiceri (2006).

As in the related literature (ACEL, and Levin et al., 2005), it is assumed that the government has access to lump-sum taxes and debt. Furthermore, the government consumes a stochastic fraction of output $G_t = S_{g,t} Y_t$ (Justiniano and Primiceri, 2006). The law of motion for S_g is $\log S_{g,t} = (1-\rho_g) \log S_g + \rho_g \log S_{g,t-1} + \sigma_g \varepsilon_{g,t}$, where $\varepsilon_{g,t}$ has a standard normal distribution.

2.6 Financial Intermediaries

Financial intermediaries receive deposits from households in the amount $\int (\mathcal{M}_{i,t-1} - M_{i,t}) di + T_t$, which includes the monetary transfer T_t from the government. All this money is lent to the good firms so they can pay workers at the beginning of each period. Consequently, the clearing condition in the loan market is $\int W_t L_{j,t} dj = T_t + \int (\mathcal{M}_{i,t-1} - M_{i,t}) di$.

3 Estimation

The data come from the Haver Analytics database and span 1984:I to 2004:IV. I opt for this short sample based on two observations. To begin with, Stock and Watson (2007) report that the stochastic process for inflation changed around 1984. Second, Fernandez-Villaverde and Rubio-Ramirez (2007 and 2008) argue that either stochastic volatility or parameter drifting are essential features of any DSGE model to capture the pre- and post-1984 features of the data, i.e., to properly account for the Great Moderation. Since a central point in this paper is the implications of inflation in recent years, introducing those features will only complicate the solution and estimation of the model without adding much substance to the subject.

The model is estimated using eight U.S. variables: the growth rates of output, consumption, investment, real wages, and real money balances $(\Delta \log Y_t, \Delta \log C_t, \Delta \log I_t, \Delta \log (W/P)_t,$

$\Delta \log (M/P)_t$), the level of labor, nominal interest rates, and inflation ($\log L_t, i_t, \pi_t$). The series are built as follows: Real GDP per capita results from dividing nominal GDP by population and the GDP deflator. Real consumption is the sum of personal consumption of non-durables and services. Real investment consists of personal consumption expenditures of durables and gross private domestic investment. Both real consumption and real investment are divided by population to obtain per capita measures. The log of hours of all persons in the non-farm business sector divided by population corresponds to labor in the paper. Real wages result from dividing nominal wage per hour in the non-farm business sector by the GDP deflator. Interest rates correspond to the effective federal funds rate while inflation is the quarterly log difference of the GDP deflator.

CEE interpret the utility from money as capturing the transaction role of money. Furthermore, Feenstra (1986) argues that money in the utility function is equivalent to a formulation where money provides liquidity services. These interpretations point to seasonally adjusted $M1$ as the relevant measure of money for estimation purposes. Specifically, the ratio $M1/P$ will be used as the counterpart of aggregate real balances in the model, $M/P = \int M_{i,t}/P di$. In results not reported here, I find that using $M2$ *minus* and its own opportunity cost as measures of money and interest rates delivers similar implications.

Bayesian Inference

Following Schorfheide (2000), Del Negro et al. (2004), and Smets and Wouters (2007), the linearized version of the model is estimated using Bayesian methods. In particular, the posterior distribution of the structural parameters is characterized using a Markov Chain Monte Carlo (MCMC) approach (for details of this algorithm see the appendix and the excellent surveys of An and Schorfheide, 2007, and Geweke, 1999). Since there are eight observable variables and only five structural shocks, I avoid stochastic singularity by following Sargent (1989) in including measurement errors to the state space representation used to estimate the model.[5] These errors are assumed to be iid and distributed $\mathbb{N}(0, \sigma)$. The scale of these errors can vary across the measurement equations. The results in the next sections are based on a Markov chain of 150,000 draws after discarding 10,000 replications from a burn in phase.

Priors

A subset of the parameter space was fixed: $\alpha = 0.36$, $\delta = 0.025$, $S_g = 0.22$. The steady state fraction S_g was set to match the average share of government expenditure in output in the sample. Since steady state labor, L_{ss}, is estimated, the parameter Φ is endogenously

[5] Since we observe the exact values of interest rates, measurement errors were not included in the equation corresponding to interest rates in the state space representation.

determined.

Based on the discussion in Section 2.4, I set the priors for ξ_m and ζ_m to capture the short- and long-run elasticities of money demand (see Table 1). For the average annualized interest rate of 5.4% in the sample, the implied mean elasticities are roughly 6.25 and 12.5, respectively. The long-run value is consistent with the results in Mankiw and Summers (1986) and Ball (1998). On the other hand, the short-run elasticity is large relative to ACEL and Christiano et al. (1999). I choose to do so to keep symmetry among the sticky contract assumptions in the model. Notice that the price, wage, and money contracts share the same priors. As we will see in the next section, the money demand priors are not very informative in the sense that the inference approach uncovers distinct posteriors. From equation (3), it is clear that the parameters ξ_m and ζ_m completely characterize the dynamics of money demand; i.e., the data are silent about the remaining parameter in the money block, ψ_m. Hence, this parameter is set to the value chosen in CEE: 0.055.

The prior distributions for the remaining parameters are reported in Table 1. These priors are loose and consistent with those typically used in the literature (see Del Negro et al., 2004, Levin et al., 2005, and Justiniano and Primiceri, 2006, Smets and Wouters, 2007). For example, the priors for the dispersion parameters χ and χ_w are beta $B(0.5, 0.2)$. The large standard deviation reflects our relative ignorance about those parameters. The priors for the elasticities of substitution ζ and ζ_w are centered around the values used in Christiano et al. (2005).

Median Estimates

Table 2 reports the median estimates for the structural parameters in my formulation. Numbers in parenthesis correspond to the 5% and 95% percentiles for each parameter (a 90% probability interval). The absence of the price of investment as an observable variable implies that the two trends in the model, S^L and S^K, are not identified separately. Therefore, the steady state growth rate of the investment-specific shock is set to one, $g^K = 1$.

Broadly speaking, the estimates are in line with the results previously found in the literature (CEE, and Smets and Wouters, 2007). For example, the model displays significant habit formation, around 0.93, and adjustment costs of investment on the order of 3.82. The habit formation estimate may seem high relative to that in ACEL; however, the estimate is perfectly in line with the findings in Fernandez-Villaverde and Rubio-Ramirez (2008). The empirical results imply that prices and wages are re-optimized on average every 3 and 1.5 quarters, respectively. It is tempting to contrast the length of the price/wage contracts with the results in Nakamura and Steinnson (2008) and Bils and Klenow (2004). Yet the presence of partial indexation makes such comparison unfeasible. The estimated Taylor rule implies that the central bank actively responds to inflation and smooths interest rates with coefficients similar to those found in Justiniano and Primiceri (2006). In terms of the structural errors, I find that they display significant serial correlation. The inference approach estimates a Frisch elasticity, γ, of 1.61, a value consistent with that reported in Fernandez-Villaverde and Rubio-Ramirez (2007), and Justiniano and Primiceri (2006). When we turn to inflation, we observe that the inference approach places its steady state value around 2%, which is close to the mean inflation in the sample (2.3%).

The median estimates for the money demand coefficients, ζ_m and ξ_m, are 1.85 and 0.86, respectively. The large value for the Calvo lottery in money reflects the estimation's attempt to capture the high frequency properties of money. In fact, its implied short-run elasticity is 1.40; interestingly, CEE report a comparable estimate. On the other hand, the empirical results suggest a long-run elasticity of 10, which is well within the boundaries found in the literature (see Goldfeld and Sichel, 1990).[6] The estimates of ζ_m and ξ_m also indicate that the assumptions outlined in Section 2.4 are flexible enough to simultaneously capture the high and low frequency properties of money demand.

4 Welfare Cost of Deflations

Inflation is potentially welfare reducing in the model due to several factors. To begin with, the presence of money demand (equation 2) makes inflation costly because of its tax implications as in Bailey (1956). Indeed, low inflation implies reduced nominal interest rates (Fisher, 1930), which benefits households because consuming real balances becomes inexpensive. A second source of distortion in the economy is staggered price contracts. To see this point, note that, ignoring growth and capital utilization, aggregate output in the model is

$$\begin{aligned} y_t &= \left[k_t^\alpha (L_t)^{1-\alpha} - \psi\right]/s_t, \qquad (5) \\ s_t &\equiv \int \left(\frac{P_{i,t}}{P_t}\right)^{-\zeta} di. \end{aligned}$$

Schmitt-Grohe and Uribe (2005) establish that s is bounded below by 1 and captures the degree of price dispersion in the economy. Staggered prices force optimizing firms to heavily review their prices to keep up with inflation, which induces dispersion, i.e., $s >> 1$. Large price stickiness (big ξ) or small price indexation (low χ) exacerbates this dispersion, decreases aggregate output, and ultimately reduces welfare (see Section 5.1). Finally, costly investment adjustment and habit formation make consumption and investment decisions rel-

[6] The long-run elasticity is somehow larger than that reported in Ireland (2008). Although we use $M1$ and similar time spans, our approaches differ in two dimensions: 1) while Ireland proposes a static money demand, I propose a dynamic formulation; and 2) Ireland estimates his model using dynamic OLS.

atively inflexible in the short term. Such inflexibility may also amplify the effects of inflation, especially during the transition from high to low inflation.

A simple way to capture the cost of inflation is to measure households' dislike for high-inflation environments. Following Cooley and Hansen (1989), Ireland (1997), and Lucas (2000), let us define the welfare cost of a high-inflation regime, Λ, as the fraction of consumption in the low-inflation steady state that households are willing to give up to be indifferent between the low- and high-inflation regimes.[7] To simplify the calculations below, the growth rate of the economy, $g_t^* = S_t^*/S_{t-1}^*$, is set to 1. This assumption is inconsequential for the rest of the analysis as I am solely interested in measuring the welfare costs under perfect foresight. Define the social utility function by

$$\begin{aligned} V &\equiv \int \sum_{t=0}^{\infty} \beta^t \left[S_t^{Uc} \log(C_{i,t} - bC_{t-1}) - \Phi \frac{L_{i,t}^{1+1/\gamma}}{1+1/\gamma} + \psi_m m_{i,t}^{1-\zeta_m} \right] di \qquad (6) \\ &= \sum_{t=0}^{\infty} \beta^t \left[S_t^{Uc} \log(C_t - bC_{t-1}) - \Phi \frac{L_t^{1+1/\gamma}}{1+1/\gamma} + \psi_m \left((1-\xi_m)(m_t^*)^{1-\zeta_m} + \xi_m m_{t-1}^{1-\zeta_m} \right) \right], \end{aligned}$$

where, C, and L correspond to both aggregate and individual consumption, and labor. This is a direct consequence of the complete market assumption. For the money demand block, m^* and m are real balances chosen by active and inactive households, respectively (see Appendix A). The dynamic nature of the model allows us to distinguish two types of welfare costs: in steady state and during the transition. In the absence of uncertainty and using the functional forms given in Section 3, the social utility function in steady state collapses to

$$V^i \left(C^i, L^i, m^i \right) \equiv (1-\beta)^{-1} \left[\log(1-b)C^i - \Phi \frac{(L^i)^{1+1/\gamma}}{1+1/\gamma} + \psi_m \left(m^i \right)^{1-\zeta_m} \right].$$

Here, the index i indicates whether we refer to the high-inflation steady state ($i = H$), or the low-inflation steady state ($i = L$). In addition, C^i, L^i, and m^i correspond to the steady state consumption, labor, and real balances on regime i. The rule of thumb for money choices implies that households choose the same steady state money balances, i.e., $m = m^*$. With these definitions in place, the steady state welfare gain, Λ_{ss}, is given by

$$\begin{aligned} V^H &= V^L \left((1-\Lambda_{ss})C^L, L^L, m^L \right), \\ V^H &= \log(1-\Lambda_{ss}) + V^L. \qquad (7) \end{aligned}$$

The second line is a consequence of the functional forms used in this paper. A positive

[7] Schmitt-Grohe and Uribe (2005) use a related measure to analyze the implications of alternative monetary rules.

Λ_{ss} indicates that households prefer the low-inflation regime, i.e., they willingly give up consumption to avoid the high inflation equilibrium. In other words, inflation entails a steady state welfare loss if $\Lambda_{ss} > 0$.

As previously argued, nominal and real frictions can make painful the transition from high to low inflation. To quantify their effect on welfare, suppose as in Taylor (1983) and Ireland (1997) that the monetary authority fully commits to a new low-inflation policy at time $t = 0$. In the model, such an exercise requires moving the target inflation in the Taylor rule (4) from a high rate, π^H, to a new low inflation π^L. Households and firms observe this change and conclude that the interest rate in the old inflationary regime is large relative to the new steady state.[8] This high interest rate in turn discourages economic activity, since it makes the wage bill ($R_t W_t L_t$) more expensive and consumption less attractive (it is more rewarding to send money to the bank). Therefore, from the point of view of the new low-inflation regime, the old inflationary steady state resembles the initial response of a contractionary monetary shock. It is precisely this contractionary aspect that makes the deflationary path costly, i.e., households who live in a low-inflation scenario are willing to sacrifice consumption to avoid undertaking the transition. How painful this transition is depends, among other things, on the length of the sticky contracts, habit formation, and the shape of the Taylor rule.

Let $\{C_t^+\}_{t=0}^{\infty}$, $\{L_t^+\}_{t=0}^{\infty}$, and $\{m_t^+, m_t^{*+}\}_{t=0}^{\infty}$ denote the sequence of consumption, labor, and real balances associated with the transitional path from the high to the low inflation steady states. These sequences in turn define the transitional social utility function immediately following the adoption of the new policy

$$V^+ \equiv \sum_{t=0}^{\infty} \beta^t \left[\log(C_t^+ - bC_{t-1}^+) - \Phi \frac{(L_t^+)^{1+1/\gamma}}{1+1/\gamma} + \psi_m \left((1-\xi_m) \left(m_t^{*+}\right)^{1-\zeta_m} + \xi_m \left(m_{t-1}^+\right)^{1-\zeta_m} \right) \right],$$

As with the steady state welfare case, define the transitional cost of the lower inflation policy as the fraction, Λ_+, of the low-inflation regime's consumption that consumers surrender to avoid the transition. That is,

$$V^+ = V^L \left((1-\Lambda_+) C^L, L^L, m^L \right). \tag{8}$$

In the current setup, reducing inflation is potentially harmful because it requires lowering real economic activity, which is achieved through an initial surge in interest rates with an unpleasant decline in consumption. Hence from the perspective of an economy with low

[8] By the Fisher equation, the steady state nominal interest rate equals the real interest rate, given by the discount factor, plus inflation. Other things equal, interest rates in the low regime, R_L, are lower than those in the high regime, R_H, if and only if $\pi^L < \pi^H$.

inflation a positive value of Λ_+ indicates that households are better off by not facing the transitional trajectory from high to low inflation, i.e., they must be compensated to face the deflationary path. Using equations (7) and (8), we conclude that the total cost of high inflation is $\Lambda = \Lambda_{ss} + \Lambda^+$. Under the convention previously discussed, positive values of Λ indicate that inflation is indeed costly.[9]

Before fleshing out the results, we must decide the values for the high and low steady state inflation rates. Two factors are decisive in selecting the low inflation rate. First of all, note that steady state inflation is an estimated parameter in the model. Furthermore, welfare will be computed using the low inflation regime as the reference point. Therefore, I set π^L to 2%, the value reported in Table 2, in an attempt to keep consistency between the estimation and welfare parts of the model. The high inflation rate is 12%, a number that will make the results comparable to those in the related literature (Ireland, 1997, Lucas, 2000, and Cooley and Hansen, 1989).

5 Results

To estimate the effects of inflation in the model, suppose that the economy is initially in a steady state with an annual inflation of 12 percentage points. At time $t = 0$ the central bank fully commits to bringing inflation down to 2%. Table 3 reports the steady state, transitional, and total annualized costs from the deflationary exercise. The first row presents the results when welfare is computed using the median estimates reported in Table 2. The welfare estimates indicate that 10 percentage points of inflation entail a steady state cost, Λ_{ss}, equivalent to 13.3% of annual consumption. Using the ratio of consumption to output in the steady state, we find that the cost of inflation represents 6.8% of annual income.[10] This result is substantially larger than that reported in Lucas (2000). As we shall see in the next section, frictions such as habit formation and price stickiness explain the difference between the results here and in Lucas.

When we turn to the transitional path, we note that the change from the high- to the low-inflation environment imposes a significant burden on households, Λ_+, which roughly amounts to 0.53% of annual consumption. The positive sign indicates that households give up consumption to avoid the deflationary path (see previous section). To understand this finding, recall that reducing inflation requires lowering real economic activity, which

[9] To get a description of the transitional dynamics, I use a second-order perturbation algorithm to evaluate V^+ (see Schmitt-Grohe and Uribe, 2004, and Judd, 1998). The approximation is done about the low-inflation steady state. I compute V^+ using the difference between the high- and low-inflation states as the initial condition for the transitional path.

[10] The ratio of consumption to output in steady state equals 0.51 in the model.

is achieved through an initial surge in interest rates. Because of the presence of real and nominal frictions, this spike in turn induces a persistent decline in economic activity, in particular in consumption. The transitional welfare loss results from the (unpleasantly) large and persistent contraction in consumption associated with the recessionary monetary policy. Unlike as in Ireland (1997), the transitional cost here is only a modest fraction of the steady state welfare gain. It may be tempting to contrast our results, but we must note that such a comparison is not straightforward, since our models differ along several dimensions. To name a few: 1) my formulation has capital accumulation, while Ireland's does not; and 2) his price setting mechanism is a mixture of time- and state-dependent formulation, whereas in my work it is solely time-dependent.

The impulse responses (solid lines) in Figure 1 confirm the contractional effects of pushing the economy to the low-inflation state. These impulse responses are computed using the median of the posterior distributions.[11] The new policy successfully brings inflation down to 2%; at the same time, the interest rate initially rises to near 15%, but, as inflation retreats, interest rates converge to its new steady state of 5%. The surge in interest rates makes working capital ($W_tL_tR_t$) expensive, which discourages production. Note, however, that as inflation declines, so does price dispersion. Eventually, this second force takes over and contributes to the recovery of output (equation 5), which ends up in a higher steady state. Consumption reaches its lowest level, -2.7%, about 2 quarters after the adoption of the new policy. The contraction in output reduces the demand for labor and hence increases leisure along the deflationary path. Its surge helps to make the transition less costly because leisure is part of the welfare criterion (equation 6). Finally, the model predicts that it takes less than 20 quarters for most variables to converge to the new steady state (notable exceptions are consumption and real wages). This convergence seems consistent with the evidence from Volcker's deflationary era. (It took roughly from 1981 to 1985 for the U.S. inflation to fall from 11% to a value below 4%.)

The steady state and transitional results indicate that an inflation of 12%, relative to an equilibrium with 2%, entails a welfare cost of 13.9% of annual consumption or 7% of annual income. A way to understand this total welfare effect is as follows. Imagine that households initially live in an economy with an inflation of 2% per year. Suddenly, they find themselves in a new situation in which inflation is 10 percentage points higher. Now households are worse off for two reasons. First, steady state consumption is lower than before, and holding real balances is more expensive due to higher interest rates. Second, if household would like to return to their initial situation with low inflation, they have to endure a transitional path,

[11] The impulse responses for inflation and interest rates are expressed as percentage points. For all other variables, the impulse responses are percentage deviations from the steady state with 2% inflation.

which from the perspective of the low-inflation steady state looks like a recession. The first effect is captured by Λ_{ss} while Λ_{+} measures the impact of the second force.

Although our steady state welfare estimate (13.3%) looks out of touch with the results in Cooley and Hansen (1989) and Lucas (2000), it is consistent with the findings in a recent paper by Aruoba and Schorfheide (2008). Indeed, these later authors find that 10 percentage points of inflation can represent as much as 16% of annual consumption. Yet our welfare estimates are still large relative to those in the sticky-price formulation pursued in Ireland (1997). Hence it seems necessary to assess whether the additional nominal/real frictions in the model drive the different welfare estimates.

5.1 Role of Frictions

In the experiments to follow, one of the estimated parameters in the benchmark formulation will be fixed at a time while the remaining ones are re-estimated. The welfare numbers are based on the median of the new posterior distributions.

Schmitt-Grohe and Uribe (2005) argue that partial price indexation induces significant price dispersion in DSGE models. With partial indexation, inactive firms cannot fully incorporate changes in past inflation, since prices are adjusted according to $P_{i,t} = (\pi_{t-1})^{\chi} P_{i,t-1}$. Hence, once a firm happens to re-optimize, it does so aggressively to keep up with future inflation. High inflation and low price indexation (low χ) induce stronger price revisions by active firms, leading to substantial price dispersion. But equation (5) shows that as price dispersion increase, output declines, which is potentially welfare reducing. Figure 3 illustrates this interaction between inflation and indexation and their effects on steady state consumption and output.

To fully characterize the effects of indexation, the second row in Table 3 reports the welfare results when χ and χ_w are set to 1. Full price/wage indexation drives down the steady state welfare cost by a factor of 4. Indeed, the welfare cost is equivalent to 1.70% (3.32×0.51) of annual output, which is surprisingly close to the welfare cost reported in the sticky-price model of Ireland (1997). This finding stresses the importance of a better understanding of the mechanisms behind price setting at the micro level. In terms of the transitional path, we observe that full indexation amplifies the dynamic welfare effect, although by a small margin. As shown in Figure 1, this increase results from the strong decline in real balances following the deflationary shock.

Lucas (2000) emphasizes the crucial role of the elasticity of money demand for welfare analysis. Indeed, he finds that a small elasticity is typically associated with negligible welfare costs of inflation. To assess the implications of Lucas' observation in the context of DSGE models, I re-estimate the model setting the parameter ζ_m such that the implied long-run semi-elasticity matches that of CEE, 1 (a value 10 times smaller than in the baseline case). By shrinking ζ_m we are effectively reducing the area underneath the demand for money curve. Bailey's (1956) theory in turn suggests that inflation should become less costly with the reduced elasticity. Accordingly, the results in Table 3 show that the welfare cost in the steady state is roughly two-thirds of that under the benchmark formulation. This finding concurs with those of Lucas but it also highlights the fact that other frictions, such as price indexation or price stickiness, play an even more important role in the welfare calculations.

The results in Table 3 indicate that if habit formation vanishes ($b = 0$), the steady state cost is smaller than in the benchmark scenario. To understand this finding, note that steady state real balances, m, and marginal utility of consumption, λ, are given by

$$\begin{aligned} m &= \left(\frac{\psi_m}{(\pi/\beta - 1)} \frac{1}{\lambda} \right)^{1/\zeta_m}. \\ \lambda &= \frac{1}{C(1-b)} \end{aligned} \tag{9}$$

Clearly as habit formation declines, so does the marginal utility of consumption. To compensate for the lost utility, households substitute consumption with real money balances. Since inflation acts as a tax on real balances, households have a stronger desire for real balances in the low inflation environment. Therefore, the substitution effect is larger in the low-inflation scenario, i.e., $\partial m^H/\partial b < \partial m^L/\partial b$. These arguments in turn indicate that the difference between utility from money in the low and high states, $\frac{(m^L)^{1-\zeta_m} - (m^H)^{1-\zeta_m}}{1-\zeta_m}$, is increasing in habit formation. But this difference is precisely what matters for steady state welfare comparisons (equations 6 and 7). Hence the lower habit formation is, the lower the steady state welfare costs of inflation.

Notice that the transitional welfare cost almost disappears in the absence of habit formation. This finding is the product of two forces. First, smaller habit formation makes consumption more flexible, allowing it to quickly adjust to the new steady state. Second, households heavily substitute consumption with leisure when habit formation is low (an explanation along the lines of the previous paragraph applies). This intuition is readily confirmed by the impulse responses reported in Figure 1. The strong substitution toward leisure is apparent from the large decline in labor. Furthermore, consumption converges relatively fast to its pre-shock steady state. In the benchmark scenario, consumption is still far away from its steady state even five years after the shock.

Table 3 shows that price flexibility makes the welfare cost of inflation decline in the steady state as well as (in absolute value) during the transition. As previously argued, the absence of sticky price contracts completely eliminates price distortions; firms can freely adjust prices every period, which increases output and hence decreases the benefits of a low-inflation environment. In fact, the welfare cost without sticky prices is almost identical to that computed with full price indexation. Figure 2 depicts the transitional responses in the absence of sticky prices (dashed line). Relative to the baseline formulation, we note two main features: 1) consumption, output, and investment are less responsive, and 2) all variables converge more quickly to the new steady state. For example, inflation falls below 4 percent 5 quarters after the monetary shock, which is 2 quarters faster than in the benchmark case.

When we eliminate all real/nominal frictions in the model, the steady state inflation entails a modest welfare loss of 0.22% of annual consumption (0.1% of annual income). Recall that in a frictionless economy, inflation is solely costly due to its tax implications, as in Bailey (1956). Therefore, it is not surprising that the welfare estimates are more in line with the findings of Lucas (2000), who argues that inflation is welfare reducing solely due to its effects on money demand. Figure 2 (starred lines) show that inflation adjusts to its new steady state without any real effect on the economy, resulting in the nil welfare cost during the transition reported in Table 3. Put differently, in the absence of frictions, the model displays Modigliani's (1963) real dichotomy, which explains the costless deflationary path.

In DSGE models, the Taylor rule is a parsimonious characterization of the central bank's views about the short-run dynamics of inflation and output. Hence the shape of the Taylor rule may 1) influence the model's high frequency properties; and 2) potentially affect the estimation of the model, leading to distinct welfare orderings. Therefore, a natural exercise is to investigate the effects (if any) that alternative Taylor rules have on the welfare implications of inflation. Let us consider the case in which the central bank's systematic response to inflation is fixed to a counterfactually low number, say, $\phi_\pi = 1.25$ (as with the previous experiments, the rest of the parameters were re-estimated). The results in Table 3 indicate that the welfare cost of inflation in steady state declines, while that during the transition jumps up. The overall effect is a decline in the welfare estimate relative to that under the benchmark formulation. To understand this finding, note that fixing the Taylor parameter ϕ_π results in a lower estimated Calvo probability for prices, $\xi = 0.6$; i.e., price contracts are reviewed more frequently.[12] As previously argued, less price stickiness reduces price dispersion, which ultimately lowers the benefits of living in low-inflation environments.

If one repeats the previous exercise but with a large response to inflation, $\phi_\pi = 2.5$, the welfare cost of 10 percentage points of inflation substantially increases in the steady state. Indeed, it now entails a whopping cost equivalent to almost 41% of annual consumption!

[12] The value $\xi = 0.6$ corresponds to the median of the posterior distribution obtained by re-estimating the model when the response to inflation is fixed at 1.25.

A careful analysis of the parameter estimates indicates that this surge results entirely from larger estimated price stickiness, $\xi = 0.76$, and lower price and wage indexation, $\chi = 0.35$ and $\chi_w = 0.54$ (see previous footnote). These new estimates exacerbate price dispersion, which makes inflation more costly. Indeed, if we use the Calvo and indexation parameters from the benchmark estimation, the steady state welfare cost declines to 16%. For completeness, Figure 2 presents the impulse responses following the contractionary monetary policy for the low and high response to inflation in the Taylor rule.

The last two paragraphs highlight the crucial role of the Taylor rule in determining the welfare implications of inflation. Note how a misspecified Taylor rule can easily lead to misleading welfare comparisons. As a consequence, we need to reconsider whether for welfare analysis the use of a Taylor rule is an appropriate choice for the pre-Greenspan era. This reconsideration is further granted by the findings in Sims and Zha (2006), who suggest that using a money growth rule may be better capture the thinking of policymakers in the 1960s and 1970s.

To conclude this section, it is worth briefly mentioning the role of parameter uncertainty. The 90% probability intervals reported in Table 3 indicate that this type of uncertainty primarily affects the steady state welfare estimates. In fact, the upper bound of the welfare cost can be as high as 49 percentage points of annual consumption in the baseline model. The results in Tables 2 and 3 suggest that uncertainty about the sticky price parameter, ξ, and price indexation, χ, drive the large probability intervals associated with welfare. Note that when those frictions vanish, the 90% probability intervals for welfare shrink by a significant amount. This situation does not happen if we suppress, for example, habit formation.

6 Conclusion

What is the welfare cost of inflation? How much should society give up to live in a low-inflation environment? Answering these questions has been a major endeavor in economics for the past half century. This paper has revisited those questions but from the perspective of an estimated New Keynesian model. According to the benchmark model, which entertains real and nominal frictions, 10 percentage points of inflation entail a total welfare cost of 13.9% of annual consumption (7% in annual output). From the point of view of a policymaker, this result suggests that environments with low inflation are desirable. The result is even more appealing for it comes from the type of models now being used for policy analysis around the world.

A second important contribution of this paper is the analysis of the different frictions on both the static and dynamic welfare costs of inflation. The results indicate that price inflexibility increases the transitional welfare cost by roughly 0.30 percentage point. Furthermore, the shape of the Taylor rule and price indexation are crucial components for welfare analysis in the steady state as well as during the transition. These findings call for a better understanding of the microfoundations behind price setting and its interaction with monetary policy as captured by the Taylor rule. At stake is whether inflation imposes a large burden on society.

As in Levin et al. (2005), I find that parameter uncertainty does indeed influence our inference of the welfare costs, in particular those of inflation. As a consequence, the results in this paper suggest that parameter uncertainty must be incorporated into policy analysis to have a better understanding of the welfare costs of deflations.

References

[1] Altig, David, Lawrence Christiano, Martin Eichenbaum, and Jesper Linde (2005) Firm-Specific Capital, Nominal Rigidities and the Business Cycles. NBER Working Paper 11034.

[2] An, Sungbae and Frank Schorfheide (2007) Bayesian Analysis of DSGE Models. Econometric Reviews 26, 113-172.

[3] Aruoba, Boragan, and Frank Schorfheide (2008) Insights from an Estimated Search-Based Monetary Model with Nominal Frictions. Mimeo University of Pennsylvania.

[4] Bailey, Martin (1956) The Welfare Cost of Inflationary Finance. Journal of Political Economy 64, 93-110.

[5] Ball, Lawrence (1998) Another Look at Long-Run Money Demand, Johns Hopkins University Working Paper.

[6] Bils, Mark, and Peter J. Klenow (2004) Some Evidence on the Importance of Sticky Prices. Journal of Political Economy 102, 947–985.

[7] Burdick Clark (1997) A Transitional Analysis of the Welfare Cost of Inflation. Federal Reserve Bank of Atlanta, working paper 97-15.

[8] Campbell, John, and Luis Viceira (2002) Strategic Asset Allocation: Portfolio Choice for Long-Term Investors, Oxford University Press.

[9] Casella, George and Christian Roberts (2004) Monte Carlo Statistical Methods. New York, NY: Springer.

[10] Christiano, Lawrence, Martin Eichenbaum, and Charles Evans (1999) Monetary Policy Shocks: What Have We Learned and to What End? Handbook of Macroeconomics, vol. 1.A., edited by Michael Woodford and John Taylor, North-Holland.

[11] (2005) Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy. Journal of Political Economy 113, 1-45.

[12] Cooley, Thomas and Gary Hansen (1989) The Inflation Tax in a Real Business Cycle Model. American Economic Review 79, 773-748.

[13] Del Negro Marco, Frank Schorfheide, Frank Smets and Ralf Wouters (2004) On the Fit and Performance of New Keynesian Models. Federal Reserve Bank of Atlanta, Working Paper 2004-37.

[14] Dotsey, Michael, and Peter Ireland (1996) The Welfare Cost of Inflation in General Equilibrium. Journal of Monetary Economics 37, 29 - 47.

[15] Erceg, Christopher, Dale Henderson and Andrew Levin (2000) Optimal Monetary Policy with Staggered Wage and Price Contracts. Journal of Monetary Economics 46, 281-313.

[16] Feenstra, Robert (1986) Functional Equivalence Between Liquidity Costs and the Utility of Money. Journal of Monetary Economics 17, 271-291.

[17] Fernández-Villaverde, Jesus and Juan Rubio-Ramírez (2007) Estimating Macroeconomic Models: A Likelihood Approach. Review of Economic Studies 74, 1059-1087.

[18] (2008) How Structural Are Structural Parameters? NBER Macroeconomics Annual 2007, 83-137.

[19] Fischer, Stanley (1981) Towards an Understanding of the Costs of Inflation II. Carnegie-Rochester Conference Series on Public Policy 15, 5 - 41.

[20] Fisher, Irving (1930) The Theory of Interest. Macmillan, New York.

[21] Friedman, Milton, (1969) The Optimum Quantity of Money and Other Essays. Chicago: Aldine.

[22] Geweke, John (1999) Using Simulation Methods for Bayesian Econometric Models: Inference, Develpment, and Communication. Econometric Reviews 18, 1 - 126.

[23] Goldfeld, Stephen., and Daniel Sichel (1990) The Demand for Money, in B. Friedman and F. Hahn (eds.), Handbook of Monetary Economics, Elsevier.

[24] Guerron-Quintana, Pablo (2009) Money Demand Heterogeneity and the Great Moderation. Journal of Monetary Economics 56, 255-266.

[25] Ireland, Peter (1997) Stopping Inflation, Big and Small. Journal of Money, Credit, and Banking 29, 759 - 775.

[26] (2008) On the Welfare Cost of Inflation and the Recent Behavior of Money Demand. Mimeo, Department of Economics, Boston College.

[27] Judd, Kenneth (1998) Numerical Methods in Economics. Cambridge, MA: MIT Press.

[28] Justiniano, Alejandro and Giorgio Primiceri (2006) The Time Varying Volatility of Macroeconomic Fluctuations. NBER, working paper 12022.

[29] Khan, Aubik, Robert King, and Alexander Wolman (2003) Optimal Monetary Policy. Review of Economic Studies 70, 825-860.

[30] Levin, Andrew, Alexei Onatski, John Williams and Noah Williams (2005) Monetary Policy Under Uncertainty in Micro-Founded Macroeconometric Models. NBER Macroeconomics Annual.

[31] Lucas, Robert (1981) Discussion of Towards an Understanding of the Costs of Inflation II. Carnegie-Rochester Conference Series on Public Policy 21, 53 - 55.

[32] (2000) Inflation and Welfare. Econometrica 68, 247 - 274.

[33] Mankiw, Gregory and Lawrence Summers (1986) Money Demand and the Effects of Fiscal Policies. Journal of Money Credit and Banking 18, 415-429.

[34] Modigliani, Franco, 1963. The Monetary Mechanism and Its Implications with Real Phenomena. Review of Economics and Statistics 45, 79-107.

[35] Nakamura, Emi, and Jon Steinsson (2008). Five Facts about Prices: A Reevaluation of Menu Cost Models." Forthcoming, Quarterly Journal of Economics.

[36] Sargent, Thomas (1989) Two Models of Measurements and the Investment Accelerator. Journal of Political Economy 97, 251 - 287.

[37] Schmitt-Grohé, Stephanie and Martin Uribe (2004) Optimal Operational Monetary Policy in the Christiano-Eichenbaum-Evans Model of the U.S. Business Cycle. Mimeo, Duke University.

[38] (2005) Optimal Fiscal and Monetary Policy in a Medium Scale Macroeconomic Model. NBER Macroeconomic Annual 2005, 382-425

[39] Schorfheide, Frank (2000) Loss Function-Based Evaluation of DSGE Models. Journal of Applied Econometrics 15, 645 - 670.

[40] Sidrausky, Miguel (1967) Rational Choice and Patterns of Growth in a Monetary Economy. American Economic Review 57, 534-535.

[41] Sims, Christopher and Tao Zha (2006) Were There Regime Switches in U.S. Monetary Policy? American Economic Review 96, 54-81.

[42] Smets, Frank and Ralf Wouters (2007) Shocks and Frictions in U.S. Business Cycles: A Bayesian Approach. American Economic Review 97, 586-606.

[43] Stock, James and Mark Watson (2007) Why Has U.S. Inflation Become Harder to Forecast? Journal of Money, Credit and Banking 39, 3-33.

[44] Taylor, John (1983). Union Wage Settlements during a Disinflation. American Economic Review 73, 981 - 993.

[45] Vissing-Jorgensen, Annette (2003) Perspectives on Behavioral Finance: Does 'Irrationality' Disapper with Wealth? Evidence from Expectations and Actions. NBER Macroeconomics Annual.

7 Appendix A: Money Demand Equation

Active households in the money market set their real balances according to the program

$$\max_{M_{i,t}} \mathbb{E}_t \sum_{n=0}^{\infty} (\xi_m \beta)^n \left[\psi_m \frac{\left(\frac{M_{i,t} \Pi_{\tau=0}^{n-1} (g_{t+\tau}^* \pi_{t+\tau})}{S_{t+n}^* P_{t+n}} \right)^{1-\zeta_m}}{1-\zeta_m} - \frac{\lambda_{t+n}}{P_{t+n}} (R_{t+n} - 1) M_{i,t} \Pi_{\tau=0}^{n-1} (g_{t+\tau}^* \pi_{t+\tau}) \right].$$

The presence of the terms $g_{t+\tau}^* \pi_{t+\tau}$ is a consequence of the rule of thumb followed by inactive households (non-optimizing households). This is so because a household given the option to re-optimize money balances must take into account that with positive probability it may not re-optimize ever again. Taking derivatives with respect to $M_{i,t}$ we arrive at the following first-order condition

$$\mathbb{E}_t \sum_{n=0}^{\infty} (\xi_m \beta)^n \left[\psi_m \left(\frac{M_{i,t} \Pi_{\tau=0}^{n-1} (g_{t+\tau}^* \pi_{t+\tau})}{S_{t+n}^* P_{t+n}} \right)^{-\zeta_m} \frac{\Pi_{\tau=0}^{n-1} \pi_{t+\tau}}{S_{t+n}^* P_{t+n}} - \frac{\lambda_{t+n}}{P_{t+n}} (R_{t+n} - 1) \Pi_{\tau=0}^{n-1} (g_{t+\tau}^* \pi_{t+\tau}) \right] = 0.$$

Next, re-write the first-order condition as a restriction on money holdings, M_t^*, which requires that the expected marginal benefit of an extra dollar equals its expected marginal cost, i.e.,

$$\underbrace{\psi_m \left(\frac{M_t^*}{S_t^* P_t} \right)^{-\zeta_m} x_{m,t}^1}_{\textit{Marginal Benefit}} = \underbrace{x_{m,t}^2}_{\textit{Marginal Cost}} . \tag{10}$$

where the terms x_m^1 and x_m^2 are given by

$$\begin{aligned} x_{m,t}^1 &\equiv \mathbb{E}_t \sum_{n=0}^{\infty} (\xi_m \beta)^n \left(\frac{\Pi_{\tau=0}^{n-1} (g_{t+\tau}^* \pi_{t+\tau})}{\Pi_{\tau=1}^{n} (g_{t+\tau}^* \pi_{t+\tau})} \right)^{1-\zeta_m}, \\ x_{m,t}^2 &\equiv \mathbb{E}_t \sum_{n=0}^{\infty} (\xi_m \beta)^n \lambda_{t+n} S_{t+n}^* (R_{t+n} - 1) \frac{\Pi_{\tau=0}^{n-1} (g_{t+\tau}^* \pi_{t+\tau})}{\Pi_{\tau=1}^{n} (g_{t+\tau}^* \pi_{t+\tau})}. \end{aligned}$$

After some tedious algebraic manipulations these last equations collapse to

$$\begin{aligned} x_{m,t}^1 &= 1 + \beta \xi_m \mathbb{E}_t \left(\frac{g_t^* \pi_t}{g_{t+1}^* \pi_{t+1}} \right)^{1-\zeta_m} x_{m,t+1}^1, \\ \text{and } x_{m,t}^2 &= \lambda_t S_t^* (R_t - 1) + \beta \xi_m \mathbb{E}_t \frac{g_t^* \pi_t}{g_{t+1}^* \pi_{t+1}} x_{m,t+1}^2. \end{aligned}$$

Given the complete market assumption, the fraction $(1 - \xi_m)$ of households re-optimizing money holdings today chooses the same level, M_t^*, governed by equation (10) . Moreover, the random nature of the Calvo lottery implies that inactive households in this period have on average the same money levels as yesterday, M_{t-1}, adjusted by technology growth and inflation. Hence, aggregate money balances, M_t, result from the combination of both active and inactive money:

$$M_t = \int M_{i,t} di = (1 - \xi_m) M_t^* + \xi_m \pi_{t-1} g_{t-1}^* M_{t-1}. \tag{11}$$

Let $\widetilde{\lambda}_t = \lambda_t S_t^*$ and $m_t = \frac{M_t}{S_t^* P_t}$, then the aggregate equation (11) becomes $m_t = (1 - \xi_m) m_t^* + \xi_m \frac{g_{t-1}^* \pi_{t-1}}{g_t^* \pi_t} m_{t-1}$. In the steady state this condition implies $m = m^*$. Furthermore, the optimal condition for money in the steady state collapses to $\psi_m (m)^{-\zeta_m} = \widetilde{\lambda} (R - 1)$. From this equation it is clear that the annualized long-run interest semi-elasticity of money demand is

$$E_{LR} = -\frac{\partial \log m}{\partial R} = \frac{1}{4(R-1)\zeta_m}.$$

In the short term, contemporaneous changes in interest rates, R_t, affect aggregate money holdings, M_t, only through its influence on active households in the financial market; i.e., those who re-optimize money balances today, M_t^*. This observation in turn implies that the short-term elasticity of money demand is $E_{SR} = -\frac{\partial \log m_t}{\partial R_t} = -(1 - \xi_m) \frac{\partial \log m_t^*}{\partial R_t}$. Finally, optimizing households understand that 1) the change in interest rates is permanent, and 2) they may not re-optimize real balances ever again. Hence, households who re-optimize at time t adjust their real balances to the new steady state, i.e., $-\frac{\partial \log m_t^*}{\partial R_t} = \frac{1}{4(R-1)\zeta_m}$. The combined effect is that the short-run semi-elasticity is given by

$$E_{SR} = (1 - \xi_m) \frac{1}{4(R-1)\zeta_m}.$$

8 Appendix B: MCMC Algorithm

Let $p(\varphi)$ and $p(Y_T|\varphi)$ be the prior distribution of the parameter vector φ and the likelihood of the data conditional on the parameter vector, respectively. I use the data, a state-space representation of the model, and the Kalman Filter to evaluate the posterior distribution $p(\varphi|Y_T)$. A random walk Metropolis-Hastings algorithm is applied to generate 150000 draws $\varphi_{(n)}$ from the posterior distribution. At each iteration n, a candidate parameter vector $\widetilde{\varphi}$ is drawn from the distribution $\mathbb{N}(\varphi_{(n-1)}, c_o^2 \sum)$ and the acceptance ratio, r, is computed $r = \frac{p(Y_T|\widetilde{\varphi})p(\widetilde{\varphi})}{p(Y_T|\widetilde{\varphi}_{(n-1)})p(\widetilde{\varphi}_{(n-1)})}$.The new draw $\widetilde{\varphi}$ is kept with probability $\min(r, 1)$ and rejected otherwise. To characterize the variance of the jumping distribution, $c_o^2 \sum$, I proceed as follows. First, I apply Christopher Sims' *csminwel* code to compute the mode of the posterior distribution. To that end, 1000 draws from the prior distribution were used to calculate the posterior. The 10 draws achieving the highest posteriors are the initial points for Sims' minimization algorithm. Second, after checking that the algorithm delivers similar modes, at least for three different initial conditions, I compute the inverse Hessian at the mode and use it as variance of the jumping distribution. Third, the constant c_o is set to achieve an acceptance rate close to 0.35, a value typically suggested in the literature (Casella and Roberts, 2004).

To check convergence of the resulting algorithm, I run three separate chains, each starting from a different random draw of the jumping distribution centered at the mode. The medians of the resulting chains lie within 2% of each other. To further confirm convergence, I compute the potential scale reduction factors, which were less than 1.005.

Table 1: Priors Densities for Structural Parameters

σ_m	σ^{g_L}	σ^{g_K}	σ^{U_C}	b	ξ_w	ξ_p	γ	ρ_R	ϕ_π
IG [2,2]	IG [2,2]	IG [2,2]	IG [2,2]	B [0.5,0.1]	B [0.5,0.1]	B [0.5,0.1]	N [1,0.15]	B [0.75,0.1]	N [1.70,0.3]
ϕ_y	κ	ξ_m	ζ_m	$100(g_L{-}1)$	$100(g_K{-}1)$	$100(\pi-1)$	L	$\varkappa_a$	β
G [0.12,0.1]	N [3,1]	B [0.5,0.1]	N [1.5,0.15]	N [0.5,0.1]	N [0.5,0.1]	N [0.5,0.1]	N [52.89,5]	N [0.17,0.1]	B [0.99,0.002]
ρ_{g_L}	ρ_{g_K}	ρ_{Uc}	g_K	ζ		ζ_w	χ	χ_w	σ_{error}
B [0.5,0.15]	B [0.5,0.15]	B [0.5,0.15]	N [1.01,0.003]	N [6,0.5]		N [21,1]	B [0.5,0.2]	B [0.5,0.2]	IG [0.05,0.03]

Notes: IG˜Inverse Gamma, B˜Beta, N˜Normal, G˜Gamma

Mean and Standard Deviation in square brackets

All measurement errors have priors as described by σ_{error}

Table 2: Estimated Parameters Benchmark Case

σ_m	σ_{g_L}	σ_{g_K}	σ_{U_C}	σ_{gov}	b	ξ_w	ξ_p	γ	ρ_R	ϕ_π	ϕ_y
0.18 [0.14,0.21]	0.36 [0.29,0.45]	0.31 [0.21,0.42]	3.08 [2.07,4.64]	0.53 [0.37,0.72]	0.93 [0.87,0.96]	0.34 [0.25,0.43]	0.68 [0.59,0.78]	1.61 [0.85,2.59]	0.82 [0.72,0.86]	1.81 [1.67,1.94]	0.034 [0.021,0.049]
κ	ξ_m	ζ_m	$\widetilde{g}_L$	$\widetilde{\pi}$	L	χ	χ_w	$\varkappa_a$	ζ	ζ_w	β
3.82 [2.68,5.08]	0.86 [0.81,0.89]	1.85 [1.42,2.16]	0.39 [0.29,0.50]	0.52 [0.38,0.68]	58.17 [51.7,63.2]	0.40 [0.19,0.64]	0.66 [0.41,0.85]	0.24 [0.11,0.43]	6.15 [5.34,6.93]	20.77 [19.1,22.5]	0.994 [0.992,0.995]
ρ_{g_L}	ρ_{g_K}	ρ_{Uc}	ρ_g		σ_{out}	σ_{cons}	σ_{invest}	σ_{wage}	σ_{inflat}	σ_{money}	σ_{labor}
0.67 [0.56,0.75]	0.83 [0.73,0.97]	0.98 [0.97,0.99]	0.55 [0.29,0.79]		0.37 [0.32,0.42]	0.32 [0.25,0.40]	1.66 [1.45,1.93]	0.15 [0.12,0.18]	0.48 [0.41,0.56]	1.11 [0.89,1.31]	0.053 [0.03,0.11]

Notes: $\widetilde{g}_L = 100(g_L{-}1)$, $\widetilde{\pi} = 100(\pi - 1)$

L: steady state labor; π: steady state inflation; g_L: growth rate of neutral technology

Table 3: Welfare Cost Estimates[a]

	Steady State (Λ_{ss})	Transition (Λ_{+})	Total (Λ)
Benchmark	13.30 [5.27,48.92]	0.53 [0.14,1.87]	13.90 [5.81,49.64]
Full Indexation (χ=1,χ_w=1)	3.32 [2.12,5.06]	0.68 [0.36,1.21]	4.04 [3.0,5.72]
Small Elasticity: $E_{LR} = 1$	9.56 [2.21,42.81]	0.60 [0.23,1.94]	10.19 [2.75,44.73]
No Habit Formation	8.27 [2.88,22.83]	0.12 [0.04,0.36]	8.40 [2.93,23.16]
No Price Stickiness	3.82 [2.38,7.20]	0.20 [0.04,0.43]	4.05 [2.33,7.44]
No Frictions	0.22 [0.17,0.30]	0.00	0.22 [0.17,0.30]
Low Taylor Rule: $\phi_\pi = 1.25$	7.55 [4.42,14.14]	0.68 [0.38,1.22]	8.18 [5.05,15.23]
High Taylor Rule: $\phi_\pi = 2.50$	40.88 [13.76,178.8]	1.36 [1.10,5.50]	42.88 [14.82,184.2]

[a]Estimates expressed as percentage of annual consumption.

90% probability interval in square brackets.

Figure 1: Impulse Responses Deflationary Shock

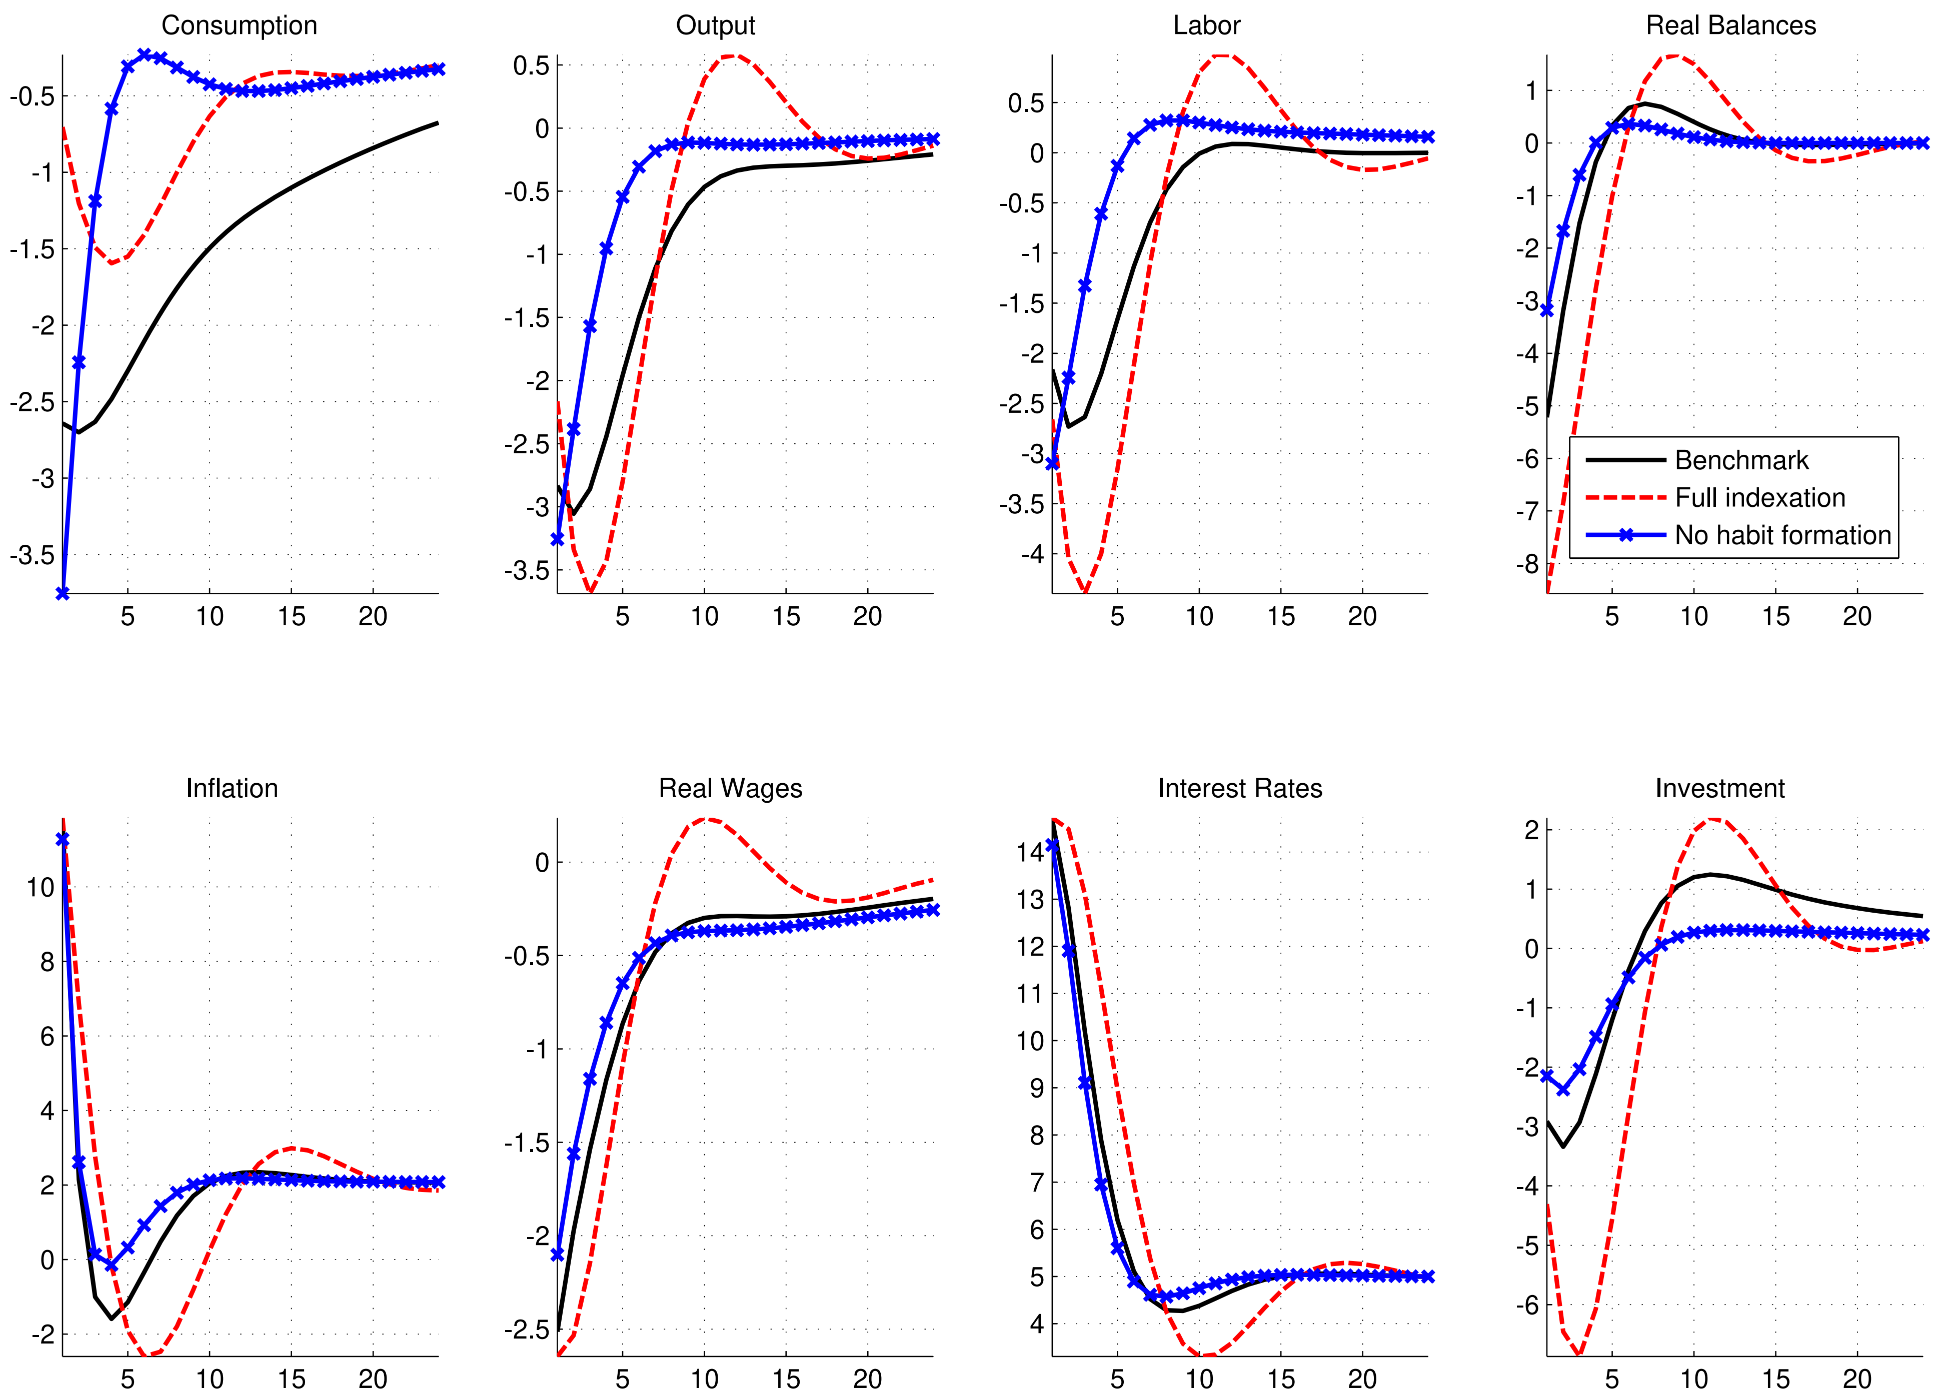

Figure 2: Impulse Responses Deflationary Shock

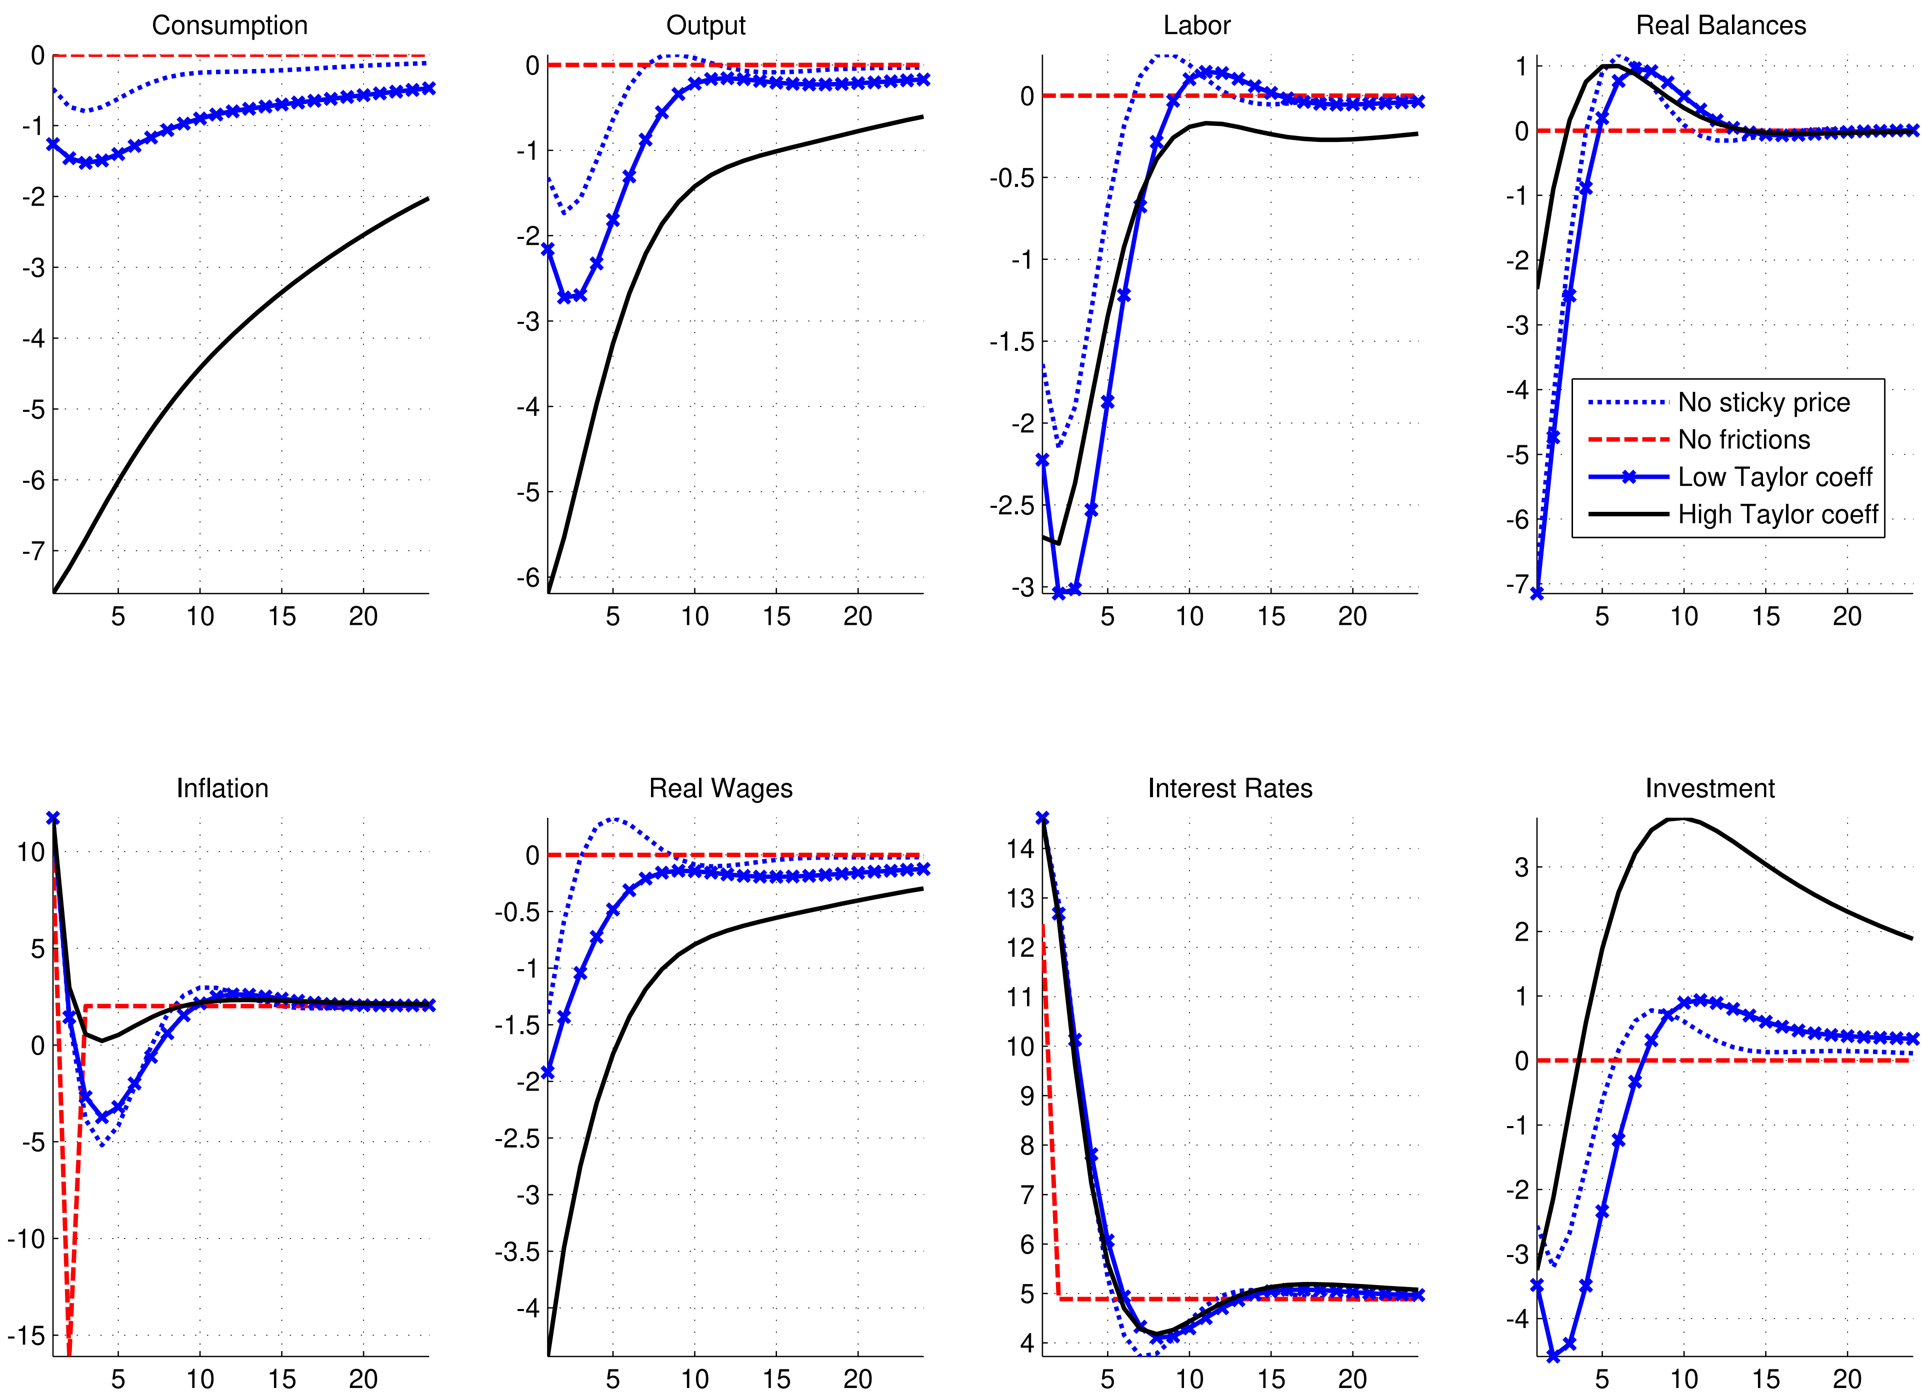

Figure 3: Effects of Price Indexation in Steady State