Working Paper Series

Stopping Inflations, Big and Small

WP 96-01

Peter N. Ireland
Federal Reserve Bank of Richmond

THE FEDERAL RESERVE BANK OF RICHMOND
RICHMOND ▪ BALTIMORE ▪ CHARLOTTE

WORKING PAPER 96-1

STOPPING INFLATIONS, BIG AND SMALL

Peter N. Ireland*

Research Department
Federal Reserve Bank of Richmond
February 1996

*The opinions expressed herein are my own and do not necessarily represent those of the Federal Reserve Bank of Richmond or the Federal Reserve System.

Abstract

Previous studies of disinflation work with models in which firms use time-dependent strategies, changing nominal prices at intervals of fixed length. These models may be criticized for failing to allow pricing behavior to adjust after a large shift in policy regime. Consequently, this paper develops a model that allows firms to adopt strategies that are partially state-dependent, changing nominal prices whenever they deviate sufficiently from their target values. The paper uses this model to examine how the welfare costs and benefits of disinflation vary with the initial inflation rate and the speed of disinflation.

I. Introduction

Academic economists differ widely in their views on the costs of disinflation, and nowhere are these differences exhibited more strikingly than in two essays, by Robert J. Gordon and Thomas J. Sargent, that appear back-to-back in Robert E. Hall's 1982 volume, *Inflation: Causes and Effects*. In his essay, Gordon considers 14 historical episodes in the United States since 1916 and in eight other countries since the mid-1960s; data from these episodes lead him to conclude that deliberate monetary policy actions taken to reduce the inflation rate are almost always quite costly, resulting in significant short-run declines in aggregate output. Gordon's analysis also suggests that these output costs can be minimized, although still not avoided altogether, by a policy that brings inflation down only gradually.

Sargent, like Gordon, draws on data from historical episodes in which monetary authorities took deliberate actions to reduce inflation; he examines the ends of hyperinflations in four European countries during the 1920s. Sargent finds that in each of these countries, Austria, Germany, Hungary, and Poland, inflation was brought to an abrupt halt with little or no loss in aggregate output. At first glance, therefore, Sargent's findings appear in stark contrast to Gordon's: they suggest that it is possible to end inflation quickly and costlessly.

This paper attempts to reconcile Gordon and Sargent's disparate observations using a single theoretical model of inflation and disinflation. The model's key feature, a fixed cost of price adjustment, makes firms unwilling to immediately reset their nominal prices following a small change in inflation. Most of the disinflationary episodes studied by Gordon began from moderate rates of inflation, and indeed, the end of a small inflation in the model presented here is accompanied by losses in aggregate output, although these losses are minimized when the disinflation is gradual. At the same time, however, firms incur the fixed cost to adjust their prices following a large change in inflation; big inflations, like those studied by Sargent, can be eliminated quickly with no loss in aggregate output.

Parts of this story appear elsewhere. Ball, Mankiw, and Romer (1988), for instance, develop a model of costly price adjustment in which monetary shocks have larger effects on output at lower rates of inflation; the real effects of money tend to vanish as inflation rises. Because of the technical difficulties associated with solving models featuring fixed costs of price adjustment, however, previous studies have been unable to consider the full effects of large changes in policy, such as those required to implement a disinflation. Danziger (1988), for example, addresses the problem of disinflation in a model of costly price adjustment, but confines his analysis to cases in which a small inflation is brought immediately to an end. He therefore stops short of answering the questions considered here: how does the cost of disinflation depend on the initial inflation rate, and how does it depend on the speed of disinflation?

Other studies of disinflation, including Phelps (1979), Taylor (1983), and Fischer (1986), work with models that constrain firms to use time-dependent strategies, changing their nominal prices at intervals of fixed length. While these models generate gradual price-level adjustments that cause monetary shocks to have real effects, they may be criticized for giving firms little flexibility to adjust their pricing behavior after a dramatic shift in policy, such as the end of a hyperinflation. In particular, these models cannot successfully reconcile Sargent's observations with Gordon's.

Models with fixed costs of price adjustment avoid this criticism by allowing firms to adopt state-dependent strategies, changing their nominal prices whenever they deviate sufficiently from their target values. As noted above, however, these models are extremely difficult to solve. Thus, the analysis here combines elements of time-dependent and state-dependent pricing. This combination, borrowed from Ball and Mankiw (1994), gives firms the flexibility to adjust their prices, at some fixed cost, after a large monetary shock, while preserving the tractability of time-dependent specifications.

Section II presents the model, while sections III and IV explore its quantitative implications for the welfare costs and benefits of inflation and disinflation. Section V concludes by bringing the model's implications to bear on the questions first raised by Gordon and Sargent.

II. The Model

The model takes many of its features from those of Ball and Romer (1989) and Blanchard and Fischer (1989, Ch.8). The economy consists of a representative household and a continuum of firms indexed by $i \in [0,1)$. Each firm produces a distinct, perishable consumption good. Hence, goods may also be indexed by $i\in[0,1)$, where firm i produces good i. Time is discrete and indexed by $t=0,1,2,\ldots$. Since there are no sources of uncertainty, all agents have perfect foresight from time t=0 forward.

The representative household trades shares in each firm i, which sell at the nominal price $Q_t(i)$ at the beginning of time t and pay a nominal dividend $D_t(i)$ at the end of time t. By choice of units, each firm has one share outstanding. Thus, if $s_t(i)$ denotes the number of shares in firm i held by the household at the end of time t, market clearing requires $s_t(i)=1$ for all $i\in[0,1)$ and $t=0,1,2,\ldots$.

The household purchases $c_t(i)$ units of each good i from firm i at the nominal price $P_t(i)$ and supplies $n_t(i)$ units of labor to each firm i at the nominal wage W_t during each time $t=0,1,2,\ldots$. The household's preferences are described by the utility function

$$\sum_{t=0}^{\infty}\beta^t(c_t^\alpha/\alpha-n_t), \qquad 1>\beta>0,\ 1>\alpha,$$

where the composites c_t and n_t are defined by

$$c_t = \left[\int_0^1 c_t(i)^{(\theta-1)/\theta}di\right]^{\theta/(\theta-1)}, \qquad \theta>1,$$

and

$$n_t = \int_0^1 n_t(i)di.$$

Its budget constraints are

$$\int_0^1 Q_t(i)s_{t-1}(i)di + \int_0^1 D_t(i)s_t(i)di + W_tn_t$$

$$\geq \int_0^1 P_t(i)c_t(i)di + \int_0^1 Q_t(i)s_t(i)di,$$

for all $t=0,1,2,\ldots$, where for $t=0$, $s_{-1}(i)=1$ for all $i\in[0,1)$ are initial conditions.

The household chooses $c_t(i)$, $n_t(i)$, and $s_t(i)$ for all $i\in[0,1)$ and $t=0,1,2,\ldots$ to maximize its utility subject to its budget constraints. Its first-order conditions are

$$c_t^{\alpha-1+1/\theta}c_t(i)^{-1/\theta} = \lambda_t P_t(i),$$

$$1 = \lambda_t W_t,$$

and

$$\lambda_t Q_t(i) = \lambda_t D_t(i) + \beta\lambda_{t+1}Q_{t+1}(i)$$

for all $i\in[0,1)$ and $t=0,1,2,\ldots$, where $\lambda_t>0$ is the multiplier on the budget constraint for time t.

Following Ball and Romer (1989), a simple quantity-theoretic equation determines the relationship between the nominal money supply M_t and nominal expenditures on goods at time t:

$$M_t = \int_0^1 P_t(i)c_t(i)di. \tag{1}$$

One could derive a similar relationship by subjecting the household to a cash-in-advance constraint, requiring it to make its purchases of goods with money at each time $t=0,1,2,\ldots.$ In that case, however, inflation would act as a distortionary tax; as in Cooley and Hansen (1989), the household would attempt to economize on its money balances by inefficiently substituting out of market activity and into leisure when faced with a positive inflation rate. Equation (1) abstracts from the cost of this inflation tax; however, the conclusion briefly considers the implications of the alternative, cash-in-advance approach.

Combining the household's first-order conditions with (1) yields

(2) $$c_t = x_t,$$

(3) $$c_t(i) = x_t^{1-\theta}[P_t(i)/M_t]^{-\theta},$$

(4) $$W_t = x_t^{1-\alpha}P_t,$$

and

(5) $$Q_t(i) = D_t(i) + (\beta x_t^{1-\alpha}P_t/x_{t+1}^{1-\alpha}P_{t+1})Q_{t+1}(i),$$

where $x_t=M_t/P_t$ denotes the level of real balances and

$$P_t = \left[\int_0^1 P_t(i)^{1-\theta}di\right]^{1/(1-\theta)}$$

defines the aggregate price level P_t at time t. Equation (2) links c_t, a measure of aggregate consumption and output, to the level of real balances at time t, as required by (1). Equation (3) defines the household's demand for good i as a downward-sloping function of firm i's real price $P_t(i)/M_t$ at time t. Equation (4) relates the real wage W_t/P_t to the level of real balances and hence, using (2), the level of aggregate output at time t. Finally, (5) shows that firm i's share price $Q_t(i)$ is determined as the sum of its current-period dividend $D_t(i)$ and the present discounted value of its future share price $Q_{t+1}(i)$.

Each firm i must sell output on demand at its nominal price $P_t(i)$ during time t. It produces this output, denoted $y_t(i)$, with labor $l_t(i)$ according to

$$y_t(i) = l_t(i)$$

for all t=0,1,2,....

The money supply increases at rate μ_t during time t, so that

$$M_t = \mu_t M_{t-1}$$

for all t=0,1,2,.... Since it faces a fixed cost of adjustment, however, firm i does not necessarily change its nominal price every period. Specifically, for t=0,2,4,..., each firm $i\in[0,1/2)$ can freely choose a new price $P_t(i)$ for its output. Each firm $i\in[1/2,1)$, meanwhile, must sell at the same price $P_t(i)=P_{t-1}(i)$ that it set at time t-1 unless it pays the fixed cost $\kappa>0$, measured in terms of labor; if it pays this fixed cost, it can set a new price $P_t(i)\neq P_{t-1}(i)$. For t=1,3,5,..., roles are reversed: firms $i\in[1/2,1)$ freely set new prices, while firms $i\in[0,1/2)$ must sell at prices set during t-1 unless they pay the fixed cost κ. Each firm i makes its price-setting decisions at time t to maximize its share price $Q_t(i)$.

This specification for the cost of price adjustment, taken from Ball and Mankiw (1994), allows pricing strategies to be partially state-dependent, since each firm can change its price whenever it deviates sufficiently from its target value. Thus, firms have some flexibility to adjust their price-setting behavior in response to a large shift in policy, such as an announced disinflation. Nevertheless, the specification retains the analytic tractability of models of time-dependent pricing by allowing each firm to change its price freely after a fixed number of periods. Thus, the specification combines the desirable features of both models of pricing.

In light of the linear production function and the requirement that firms sell output on demand, (3) implies that firm i's dividend at time t is

$$D_t(i) = x_t^{1-\theta}[P_t(i)-W_t][P_t(i)/M_t]^{-\theta} - W_t\chi_t(i)\kappa,$$

where the indicator function $\chi_t(i)=1$ if the firm pays the fixed cost of price adjustment at time t and $\chi_t(i)=0$ otherwise. Hence, using (4) and (5),

$$(6) \qquad Q_t(i) = M_t x_t^{-\alpha} \sum_{j=0}^{\infty} \beta^j x_{t+j}^{\alpha} D_{t+j}(i)/M_{t+j},$$

where

$$(7) \quad x_{t+j}^{\alpha} D_{t+j}(i)/M_{t+j} = x_{t+j}^{1-\theta}\{x_{t+j}^{\alpha}[P_{t+j}(i)/M_{t+j}]^{1-\theta}-[P_{t+j}(i)/M_{t+j}]^{-\theta}\} - \chi_{t+j}(i)\kappa.$$

A firm that can freely choose a new price at time t has two options. It can adopt a single-price strategy, retaining the same price for t and t+1. In this case, it avoids paying the fixed cost κ at t+1, but must sell output at a price that stays constant for two periods. Alternatively, it can adopt a two-price strategy, charging different prices at t and t+1. In this case, it adjusts its price optimally between the two periods, but must pay the fixed cost κ at t+1. Typically, therefore, firms will adopt single-price strategies under moderate rates of inflation, when the cost of retaining the same nominal price for two periods remains small, and switch to two-price strategies under higher rates of inflation, when it becomes worthwhile to pay the fixed cost κ to avoid the larger cost of price rigidity.

If the firm adopts a single-price strategy, (6) and (7) imply that its choice of $P_t(i)$ must maximize

$$x_t^{1-\theta}\{x_t^{\alpha}[P_t(i)/M_t]^{1-\theta}-[P_t(i)/M_t]^{-\theta}\} + \beta x_{t+1}^{1-\theta}\{x_{t+1}^{\alpha}[P_t(i)/M_{t+1}]^{1-\theta}-[P_t(i)/M_{t+1}]^{-\theta}\}.$$

In this case, therefore, firm i has

$$(8) \quad P_t(i) = P_{t+1}(i) = [\theta/(\theta-1)](M_t^{\theta}x_t^{1-\theta}+\beta M_{t+1}^{\theta}x_{t+1}^{1-\theta})/(M_t^{\theta-1}x_t^{1-\theta+\alpha}+\beta M_{t+1}^{\theta-1}x_{t+1}^{1-\theta+\alpha}).$$

If, on the other hand, the firm adopts a two-price strategy, (6) and (7) imply that each $P_{t+j}(i)$, j=0 and j=1, must maximize

$$x_{t+j}^{\alpha}[P_{t+j}(i)/M_{t+j}]^{1-\theta} - [P_{t+j}(i)/M_{t+j}]^{-\theta}.$$

In this case, therefore, firm i has

$$P_t(i) = [\theta/(\theta-1)](M_t/x_t^{\alpha}) \tag{9}$$

and

$$P_{t+1}(i) = [\theta/(\theta-1)](M_{t+1}/x_{t+1}^{\alpha}). \tag{10}$$

By (6) and (7), firm i chooses the strategy that maximizes

$$x_t^{1-\theta}\{x_t^{\alpha}[P_t(i)/M_t]^{1-\theta}-[P_t(i)/M_t]^{-\theta}\}$$
$$+ \beta x_{t+1}^{1-\theta}\{x_{t+1}^{\alpha}[P_{t+1}(i)/M_{t+1}]^{1-\theta}-[P_{t+1}(i)/M_{t+1}]^{-\theta}\} - \beta\chi_{t+1}(i)\kappa,$$

where $\chi_{t+1}(i)=0$ if it follows (8) and $\chi_{t+1}(i)=1$ if it follows (9) and (10). When solving this optimization problem, the firm takes the nominal money supply M_t, the aggregate price level P_t, and hence the aggregate level of real balances x_t as given. In equilibrium, however, x_t must be consistent with the individual firms' choices. That is, given x_t, each $P_t(i)$ must be optimal for firm i; and given each optimal $P_t(i)$, x_t must equal M_t/P_t.

III. The Effects of Steady-State Inflation

A. Steady-State Conditions

When money growth is constant, with $\mu_t=\mu$ for all t=0,1,2,..., steady-state equilibria exist in which real balances are also constant, with $x_t=x$ for all t=0,1,2,.... If firm i can freely choose a new price at time t, and if it adopts a single-price strategy in the steady state, (8) implies that it sets

(11) $$P_t(i) = P_{t+1}(i) = [\theta/(\theta-1)][(1+\beta\mu^{\theta})/(1+\beta\mu^{\theta-1})](M_t/x^{\alpha}).$$

If, on the other hand, firm i adopts a two-price strategy in the steady state, (9) and (10) imply that it sets

(12) $$P_t(i) = [\theta/(\theta-1)](M_t/x^{\alpha})$$

and

(13) $$P_{t+1}(i) = [\theta/(\theta-1)](M_{t+1}/x^{\alpha}).$$

Since half of all firms set new prices at $t=0,2,4,\ldots$, and half of all firms set new prices at $t=1,3,5,\ldots$, (11)-(13) imply that steady-state real balances are

(14) $$x = \{[(\theta-1)/\theta][(1+\beta\mu^{\theta-1})/(1+\beta\mu^{\theta})][2/(1+\mu^{\theta-1})]^{1/(1-\theta)}\}^{1/(1-\alpha)}$$

if all firms adopt single-price strategies and

$$x = [(\theta-1)/\theta]^{1/(1-\alpha)}$$

if all firms adopt two-price strategies. Equation (2) then implies that aggregate output is constant in the steady state, with $c_t=x$ for all $t=0,1,2,\ldots.$ Aggregate employment is also constant, with $n_t=n$ for all $t=0,1,2,\ldots$, where

$$n = x^{1-\theta+\alpha\theta}[(\theta-1)/\theta]^{\theta}[(1+\beta\mu^{\theta-1})/(1+\beta\mu^{\theta})]^{\theta}[(1+\mu^{\theta})/2],$$

with x given by (14), if all firms adopt single-price strategies, and

(15) $$n = [(\theta-1)/\theta]^{1/(1-\alpha)} + \kappa/2$$

if all firms adopt two-price strategies. Note that in the latter case, half of all firms incur the fixed cost of price adjustment at each $t=0,1,2,\ldots$, so that employment exceeds output by $\kappa/2$.

B. Model Parameterization

To see how output, employment, and welfare vary with steady-state inflation, one must begin by assigning values to the model's four parameters: α, β, θ, and κ. Equations (9) and (10) show that a firm that resets its price optimally in every period chooses $P_t(i)=P_t^*(i)$ for all $t=0,1,2,\ldots$, where

$$P_t^*(i) \quad = \quad [\theta/(\theta-1)](M_t/x_t^{\alpha}). \tag{16}$$

Dividing both sides of (16) by P_t, taking logs, and using (2) yields

$$\ln[P_t^*(i)] - \ln(P_t) \quad = \quad \ln[\theta/(\theta-1)] + (1-\alpha)\ln(c_t),$$

which is identical to the formula for $P_t^*(i)$ derived by Ball, Mankiw, and Romer (1988). This formula reveals that $1-\alpha$ measures the elasticity of firm i's optimal relative price $P_t^*(i)/P_t$ with respect to aggregate output c_t. Based on their view that this elasticity is quite small, Ball, Mankiw, and Romer set $\alpha=0.9$, a choice that is also used here.

Following Ball and Mankiw (1994), each period in the model is identified as six months. Under moderate rates of inflation, firms adopt single-price strategies, keeping their prices fixed for two periods to avoid the cost of adjustment. Thus, this choice for the period length implies that under moderate rates of inflation, firms adjust their prices annually. The choice is therefore consistent with observations by Carlton (1986), Cecchetti (1986), Blinder (1994), and Kashyap (1995), which suggest that nominal prices in the United States economy are typically adjusted at intervals of one year or more. The choice of period length also dictates a setting of $\beta=0.975$, which implies an annual discount rate of 5 percent.

The linear production function implies that the aggregate markup of price over marginal cost in this economy is simply P_t/W_t, the inverse of the real wage. Thus, (4) and (14) imply that under moderate rates of inflation, where firms adopt single-price strategies, the steady-state markup is

$$x^{\alpha-1} = [\theta/(\theta-1)][(1+\beta\mu^{\theta})/(1+\beta\mu^{\theta-1})][(1+\mu^{\theta-1})/2]^{1/(1-\theta)}.$$

When μ is small, the second and third terms in this expression remain close to unity, so that the markup is approximately $\theta/(\theta-1)$. Based on their review of empirical studies of the markup in the United States economy, Rotemberg and Woodford (1992) select a benchmark value of 1.2 for the steady-state markup; their choice, which implies $\theta=6$, is also used here.

Finally, the value of κ determines the inflation rate at which firms abandon single-price strategies, so that the rigidity of individual goods prices disappears. Thus, $\kappa=0.0175$ is selected here, so that all firms adopt two-price strategies when steady-state inflation reaches 198 percent annually. Of course, this choice for κ implies that significant price rigidities remain for annual inflation rates below 198 percent, but evidence presented by Mussa (1981) suggests that this may not be an unreasonable assumption: he finds that in early stages of the German hyperinflation of the 1920s, when inflation averaged 31 percent per month, some individual goods prices remained constant for periods of up to three months.

C. The Effects of Inflation on Output and Welfare

Figure 1 plots the level of steady-state output $c=x$ for annual inflation rates between zero and 250 percent. For inflation rates below 198 percent, all firms adopt single-price strategies, and output varies with inflation. For inflation rates above 198 percent, firms use two-price strategies, so that nominal rigidities, and the effects of inflation on output, disappear.

When firms adopt single-price strategies, (14) reveals that inflation affects output through the second term,

$$(1+\beta\mu^{\theta-1})/(1+\beta\mu^{\theta}),$$

which is decreasing in μ, and the third term,

$$[2/(1+\mu^{\theta-1})]^{1/(1-\theta)},$$

which is increasing in μ. In fact, these two offsetting effects of inflation on output in this general equilibrium model are precisely those identified in partial equilibrium by Benabou and Konieczny (1994).

The second term in (14) works to decrease output as inflation rises since, as shown by (11), this term also works to increase firm i's real price $P_t(i)/M_t$ when it adopts a single-price strategy. As explained by Benabou and Konieczny, this effect results from an asymmetry in the firm's profit function.

Specifically, when firm i adopts a single-price strategy, (11) and (16) show that the real price $P_t(i)/M_t$ initially lies above its optimal value $P_t^*(i)/M_t$; later, inflation erodes the real price, so that $P_{t+1}(i)/M_{t+1}$ lies below $P_{t+1}^*(i)/M_{t+1}$. Equation (7), meanwhile, shows that the real value of firm i's profits at time t, given by $x_t^{\alpha}D_t(i)/M_t$, can be written as a function of $z_t(i)=\ln[P_t(i)/M_t]$, the log of the firm's real price:

$$\pi[z_t(i)] \quad = \quad x_t^{1-\theta}\{x_t^{\alpha}\exp[(1-\theta)z_t(i)]-\exp[-\theta z_t(i)]\}.$$

This function is asymmetric about $z_t^*=\ln[P_t^*(i)/M_t]$, so that $\pi[z_t(i)]$ decreases faster as $z_t(i)$ falls below z_t^* than it does as $z_t(i)$ rises above z_t^*. Thus, to guard against the erosion of its real price brought about by inflation, the firm sets $P_t(i)/M_t$ higher as μ rises. An increase in inflation thereby leads to an increase in the general price level, decreasing aggregate output through what Benabou and Konieczny call the "profit effect."

The third term in (14) works to increase output as inflation rises by capturing the effect of cross-sectional price dispersion on the total demand for goods. When they adopt single-price strategies, (11) implies that firms $i \in [0,1/2)$ and $i \in [1/2,1)$ sell output at different prices at each $t=0,1,2,\ldots$; in fact, the degree of cross-sectional price dispersion increases with the steady-state inflation rate μ. Again in terms of the log real price $z_t(i)$, (3) shows that the household's demand for good i can be written

$$c[z_t(i)] \quad = \quad x_t^{1-\theta}\exp[-\theta z_t(i)].$$

Since this function is convex in $z_t(i)$, the increase in cross-sectional price dispersion brought about by higher inflation works to increase demand, when averaged across all goods. An increase in inflation thereby leads to an increase in aggregate output through what Benabou and Konieczny label the "demand effect."

In this model, Benabou and Konieczny's demand effect dominates at very low rates of inflation, so that output initially rises with μ. This effect is so small, however, that it is nearly imperceptible in figure 1: output peaks at just 0.013 percent above its zero-inflation level when annual inflation reaches 0.85 percent. For annual inflation rates above 0.85 percent, the profit effect dominates, so that output decreases with inflation. At very high rates of inflation, in fact, the profit effect becomes quite strong: output is more than 60 percent below its zero-inflation level when inflation reaches 197 percent annually.

Figure 2 shows that these movements in output are also reflected in welfare, as measured by the household's utility. Utility, like output, initially rises with steady-state of inflation. Since employment must also increase as output rises, however, the utility-maximizing annual rate of inflation, 0.56 percent, falls below the output-maximizing level. Thus, the optimal rate of inflation, while positive, is very close to zero in this model.

Figure 2 also shows that at higher rates of inflation, utility remains below the level achieved under zero inflation. Although output returns to its zero-inflation level when inflation exceeds 198 percent annually, (15) reveals that the costs of price adjustment increase aggregate employment by 5.1 percent; this increase in employment accounts for the welfare cost of inflation after firms switch to two-price strategies in the steady state.

D. Multiple Equilibria

For some inflation rates, the steady states described in figures 1 and 2 fail to be unique. In these cases, there exist three steady states: one in which all firms adopt single-price strategies, one in which all firms adopt two-price strategies, and one in which a fraction of all firms adopt single-price strategies while the remaining firms adopt two-price strategies.

When firm i adopts a single-price strategy, its fixed nominal price deviates from the optimal value $P_t^*(i)$. The loss in profits resulting from this deviation may become larger as aggregate output increases; in this case, the firm's incentive to use a two-price strategy strengthens as x rises. In addition, from (11)-(13), output is given by

$x=[(\theta-1)/\theta]^{1/(1-\alpha)}\{(1-\varphi)[(1+\beta\mu^{\theta})/(1+\beta\mu^{\theta-1})]^{1-\theta}[(1+\mu^{\theta-1})/2]+\varphi\}^{1/[(\theta-1)(1-\alpha)]}$

when the fraction $\varphi \in [0,1]$ of all firms adopt two-price strategies in the steady state. For large enough values of μ, this expression for x may become increasing in φ.

Cases arise, therefore, in which aggregate output increases as additional firms adopt two-price strategies, making it profitable for even more firms to adopt two-price strategies. In these cases, firm i finds it best to adopt a single-price strategy when all other firms use single-price strategies and to adopt a two-price strategy when all other firms use two-price strategies. In addition, a third equilibrium exists, where firm i is indifferent between the two strategies and a fraction $\varphi \in (0,1)$ of all firms use two-price strategies while the remaining firms use single-price strategies.

Under the parameter settings used in figures 1 and 2, multiple equilibria arise when the annual inflation rate lies between 100 and 198 percent; for each of these cases, the figures show values from the equilibrium in which all firms adopt single-price strategies. The other two equilibria feature higher levels of output and welfare. In particular, levels of output and welfare in the equilibrium in which all firms adopt two-price strategies coincide with those shown in the figures for inflation rates above 198 percent. Levels of output and utility in the equilibrium in which some firms adopt single-price strategies and others adopt two-price strategies lie between those in the other two equilibria.

IV. The Effects of Disinflation on Output and Welfare

The effects of disinflation can be analyzed in this perfect foresight model, following Phelps (1979), Taylor (1983), Fischer (1986), and Danziger (1988), by considering examples in which the economy begins time $t=0$ in steady state. Prior to $t=0$, therefore, all agents expect money growth to remain constant forever. After firms $i\in[0,1/2)$ set their prices $P_0(i)$ at $t=0$, however, the monetary authority announces that it will implement an alternative policy, summarized by a new sequence of money growth rates μ_t, $t=0,1,2,\ldots$, instead.

Given this alternative policy, all agents behave optimally from $t=0$ forward. At $t=0$ and $t=1$, however, each firm $i\in[0,1/2)$ must sell output at the price $P_0(i)$ it set before the change in policy unless it pays the fixed cost κ. Likewise, each firm $i\in[1/2,1)$ must sell output at the price $P_{-1}(i)$ it set at $t=-1$ unless it pays the fixed cost κ. Thus, if the change in policy calls for a reduction in money growth, it may be accompanied by short-run losses in aggregate output.

The algorithm used here to compute the effects of disinflation finds the equilibrium in which the fewest costly price changes occur. Thus, in cases where multiple equilibria exist, the ones discussed below are those that feature the slowest rate of price-level adjustment and hence the largest output effects; this approach focuses on equilibria in which the problem of disinflation is most severe.

Figure 3 shows the output effects of a policy that brings money growth to an immediate halt when the initial annual steady-state inflation rate is 3 percent, approximately equal to the rate of core consumer price inflation in the United States from 1993 through 1995. This change in policy is too small to induce firms to immediately change their prices; hence, output falls nearly 1.5 percent at t=0. Furthermore, as in Danziger (1988), the first firms to respond to the disinflation actually raise their nominal prices; each firm $i \in [1/2, 1)$ adjusts its price upwards at t=1, so that output falls still further. Thereafter, the staggered price-setting structure causes the decline in output to persist, as in Ball and Romer (1989) and Blanchard and Fischer (1989, Ch.8). Output does not return to its initial steady-state level until 3 1/2 years after the change in policy.

Eventually, however, output rises to its new steady-state level under zero inflation, which is 0.070 percent higher than under 3 percent inflation. This long-run output gain, although small, suffices to offset the costs of the short-run output loss, so that the immediate disinflation is welfare-improving relative to the initial policy of continuing 3 percent inflation. For the parameter values used here, in fact, an immediate disinflation yields a welfare gain starting from any annual inflation rate above 2.7 percent.

To illustrate the effects of a more gradual approach to disinflation, figure 4 displays the effects of policies that reduce money growth linearly over T periods. These policies set

$$\mu_t = \mu_{t-1} - \mu/T$$

for $t=0,1,\ldots,T-1$, where $\mu_{-1}=\mu$ is the initial rate of inflation, and

$$\mu_t = 0$$

for $t=T,T+1,T+2,\ldots$. Thus, T=1 corresponds to an immediate disinflation, while larger values of T make disinflation more gradual.

Figure 4, like figure 3, traces out the output effects of ending a 3 percent annual inflation. A more gradual approach reduces the initial output losses. When T=6, so that money growth is reduced to zero over a three-year period, the initial output decline is only 0.24 percent. When T=17, the initial output decline is just 0.086 percent.

In fact, both of the gradual disinflations considered in figure 4 actually allow output to rise above its new steady-state level for a period of time. Ball (1994) shows that in models of staggered price setting, the aggregate price level at time t can be approximated by a weighted average of past and future money supplies. When disinflation proceeds slowly, the gradual decline in money growth makes the money supply a concave function of time. Hence, the price level, as an average of past and future money supplies, falls relative to the current money supply. Real balances, and hence aggregate output, rise.

Thus, figures 3 and 4 corroborate Gordon's (1982) conclusions: starting from a moderate inflation rate, a rapid disinflation yields short-run output losses, but these losses can be minimized by ending inflation more gradually. Starting from higher inflation rates, however, Sargent's (1982) views apply. Figure 5 displays the effects of an immediate disinflation, with T=1, when the economy begins with a 200 percent annual rate of inflation. This dramatic change in policy induces firms to promptly adjust their prices. Hence, the rapid disinflation yields no loss in aggregate output.

Figure 5 also shows that the gradual approach that works well in ending moderate inflations becomes quite costly when the initial inflation rate is high. Starting from a 200 percent annual inflation, the changes in money growth under the policy with T=6 are not large enough to trigger immediate adjustments to each firm's price. Hence, output falls sharply; by t=1, it is nearly 28 percent below its initial level.

Comparing figures 3-5, therefore, suggests that radically different prescriptions apply when stopping inflations of different sizes: small inflations should be ended gradually, while big inflations should be ended immediately. Figure 6 confirms this suggestion by plotting the welfare-maximizing value of T for initial inflation rates between 1 and 250 percent.

Overall, the curve in figure 6 slopes downward. Starting from an annual inflation rate of 1 percent, the best linear disinflation takes 15 1/2 years to complete, so that T=31 is optimal. Starting from 3 percent inflation, the best disinflation takes 8 1/2 years, so that T=17 as in figure 4. All inflations exceeding 136 percent annually are best ended at once. The curve is not entirely monotonic, however; the gain from a more gradual approach works to increase the optimal length of disinflation as the initial annual inflation rate rises from 79 to 135 percent.

Figure 6 shows, therefore, that while small inflations are best ended gradually and big inflations are best ended immediately, the optimal speed of disinflation is not always increasing in the initial inflation rate, at least within the class of linear disinflations considered here. For every initial inflation rate, however, the optimal disinflation shown in figure 6 yields a welfare gain relative to the policy of continuing inflation: stopping inflation is always worthwhile.

V. Conclusion

Gordon (1982) and Sargent (1982) arrive at very different conclusions regarding the costs of disinflation. Gordon's evidence suggests that efforts to reduce inflation quickly are almost always quite costly; he therefore recommends a gradual approach to disinflation. Sargent, however, provides historical examples in which inflation was brought to an immediate end with little or no loss in aggregate output.

The model developed here reconciles Gordon and Sargent's conclusions, assigning a key role to the fixed cost of price adjustment. When faced with this fixed cost, firms do not immediately reset their nominal prices when a moderate inflation is brought to an end; these disinflations, like those studied by Gordon, are accompanied by short-run losses in output. Firms incur the fixed cost to adjust their prices following larger changes in policy, however; hyperinflations, like those studied by Sargent, can be eliminated quickly and costlessly.

The analysis shows that, indeed, the best disinflations are gradual when the initial inflation is small and more rapid when the initial inflation is big. But while this prescription applies fairly generally, there are exceptions: over some ranges, the optimal speed of disinflation decreases with the initial inflation rate. Across all of the examples considered, however, stopping an inflation, big or small, turns out to be welfare-improving.

The model developed here is detailed enough to provide quantitative answers to the two questions raised by Gordon and Sargent's studies: how does the cost of disinflation depend on the initial inflation rate, and how does it depend on the speed of disinflation? Nevertheless, the model remains stylized along some dimensions, leaving room for extensions in future work. The model abstracts from the process of capital accumulation, for instance, and although they are given more flexibility here than in previous work on disinflation, firms are constrained to set prices either once or twice per year. In pursuing these extensions, however, the basic specification used here may prove useful. This specification, borrowed from Ball and Mankiw (1994), combines elements of time-dependent and state-dependent models to allow firms to change their pricing behavior after a large shock while still retaining the tractability needed to consider a wide range of alternative monetary policies.

There are, in addition, two other dimensions along which the model presented here can be modified or extended, both of which would yield important changes in the results. The first concerns the way in which money is introduced into the model. Here, money and nominal output are linked by the simple quantity-theoretic equation (1). If, instead, money were introduced through a cash-in-advance constraint, as in Cooley and Hansen (1989), inflation would act as a distortionary tax; the representative household would be given an incentive to economize on its money balances by inefficiently substituting out of market activity and into leisure when faced with a positive inflation rate. Eliminating these effects of the inflation tax would yield additional benefits of disinflation. In fact, Ireland (1995) shows that these additional benefits are large enough to outweigh any short-run output costs, so that even starting from a low rate of inflation, a very rapid disinflation becomes optimal in the cash-in-advance case.

Second, Sargent (1986) criticizes the gradual approach to disinflation advocated by Gordon (1982) and prescribed by this model for low initial inflation rates, arguing that such policies are likely to be regarded with skepticism by the public. Here, any announced disinflationary policy, however gradual, is taken as fully credible by private agents, who immediately begin acting without concern that the change in policy will be reversed. If gradual policies lack credibility, however, the effects of uncertainty about future reversals will be incorporated into firms' pricing behavior, making disinflation more costly. Thus, Ireland (1995) also considers the effects of imperfect credibility in a model of disinflation with time-dependent pricing and finds that, in the absence of full credibility, the optimal policy calls for a more rapid decrease in money growth.

To the extent that the inflation-tax effects emphasized by Cooley and Hansen (1989) and the credibility effects emphasized by Sargent (1986) play important roles, therefore, the optimal disinflationary policies will be less gradual than those prescribed here.

References

Ball, Laurence. "Credible Disinflation with Staggered Price-Setting." *American Economic Review* 84 (March 1994): 282-289.

Ball, Laurence and N. Gregory Mankiw. "Asymmetric Price Adjustment and Economic Fluctuations." *Economic Journal* 104 (March 1994): 247-261.

Ball, Laurence, N. Gregory Mankiw, and David Romer. "The New Keynesian Economics and the Output-Inflation Trade-off." *Brookings Papers on Economic Activity* (1:1988): 1-65.

Ball, Laurence and David Romer. "The Equilibrium and Optimal Timing of Price Changes." *Review of Economic Studies* 56 (April 1989): 179-198.

Benabou, Roland and Jerzy D. Konieczny. "On Inflation and Output with Costly Price Changes: A Simple Unifying Result." *American Economic Review* 84 (March 1994): 290-297.

Blanchard, Olivier Jean and Stanley Fischer. *Lectures on Macroeconomics.* Cambridge: MIT Press, 1989.

Blinder, Alan S. "On Sticky Prices: Academic Theories Meet the Real World." In N. Gregory Mankiw, ed. *Monetary Policy*. Chicago: University of Chicago Press, 1994.

Carlton, Dennis W. "The Rigidity of Prices." *American Economic Review* 76 (September 1986): 637-658.

Cecchetti, Stephen G. "The Frequency of Price Adjustment: A Study of the Newsstand Prices of Magazines." *Journal of Econometrics* 31 (April 1986): 255-274.

Cooley, Thomas F. and Gary D. Hansen. "The Inflation Tax in a Real Business Cycle Model." *American Economic Review* 79 (September 1989): 733-748.

Danziger, Leif. "Costs of Price Adjustment and the Welfare Economics of Inflation and Disinflation." *American Economic Review* 78 (September 1988): 633-646.

Fischer, Stanley. "Contracts, Credibility, and Disinflation." In *Indexing, Inflation, and Economic Policy.* Cambridge: MIT Press, 1986.

Gordon, Robert J. "Why Stopping Inflation May Be Costly: Evidence from Fourteen Historical Episodes." In Robert E. Hall, ed. *Inflation: Causes and Effects*. Chicago: University of Chicago Press, 1982.

Ireland, Peter N. "Optimal Disinflationary Paths." *Journal of Economic Dynamics and Control* 19 (November 1995): 1429-1448.

Kashyap, Anil K. "Sticky Prices: New Evidence from Retail Catalogs." *Quarterly Journal of Economics* 110 (February 1995): 245-274.

Mussa, Michael. "Sticky Individual Prices and the Dynamics of the General Price Level." *Carnegie-Rochester Conference Series on Public Policy* 15 (Autumn 1981): 261-296.

Phelps, Edmund S. "Disinflation Without Recession: Adaptive Guideposts and Monetary Policy." In *Studies in Macroeconomic Theory.* New York: Academic Press, 1979.

Rotemberg, Julio J. and Michael Woodford. "Oligopolistic Pricing and the Effects of Aggregate Demand on Economic Activity." *Journal of Political Economy* 100 (December 1992): 1153-1207.

Sargent, Thomas J. "The Ends of Four Big Inflations." In Robert E. Hall, ed. *Inflation: Causes and Effects*. Chicago: University of Chicago Press, 1982.

________. "Stopping Moderate Inflations: The Methods of Poincare and Thatcher." In *Rational Expectations and Inflation*. New York: Harper and Row, 1986.

Taylor, John B. "Union Wage Settlements During a Disinflation." *American Economic Review* 73 (December 1983): 981-993.

Fig. 1. Steady State Output

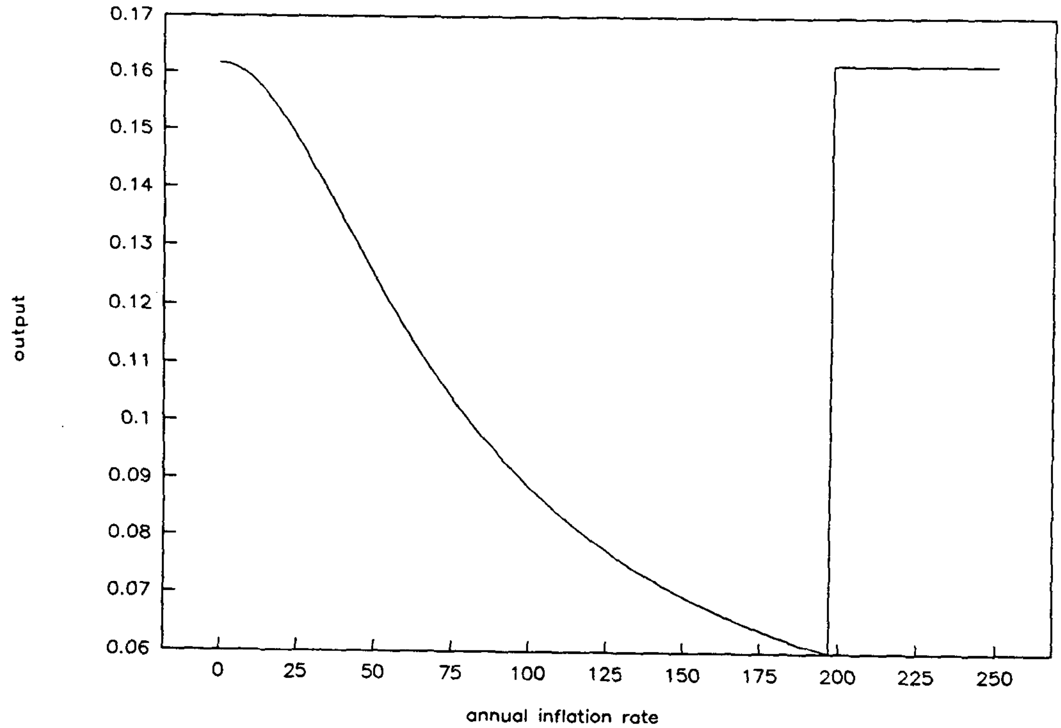

Fig. 2. Steady State Utility

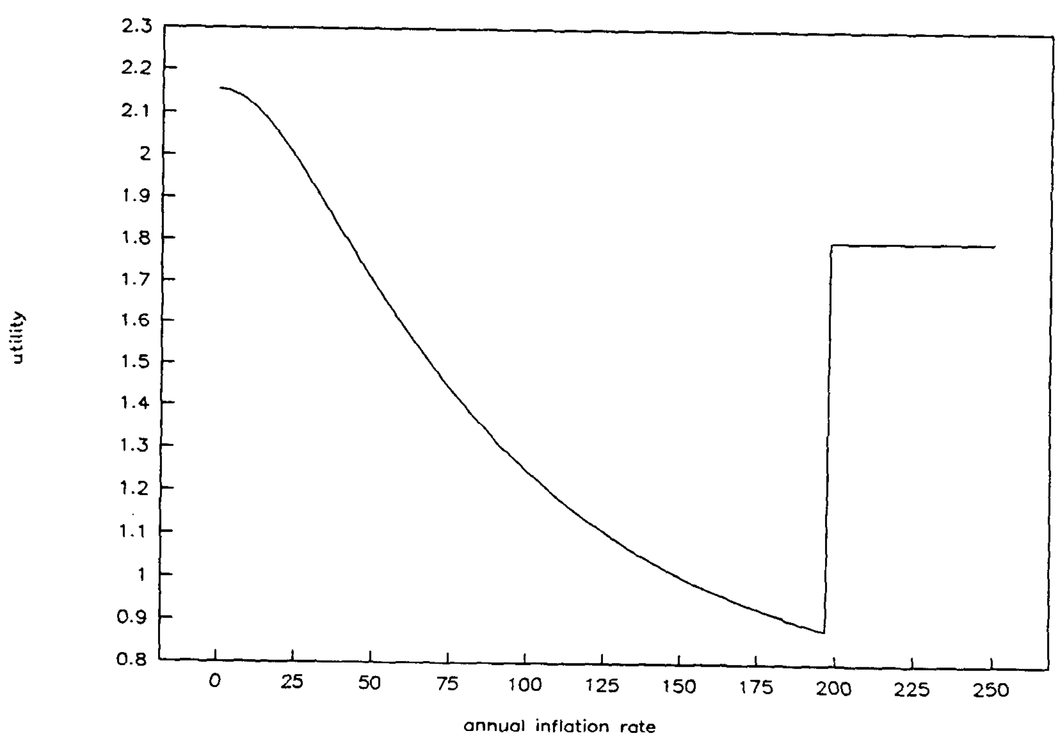

Fig. 3. Output Effects of Disinflation

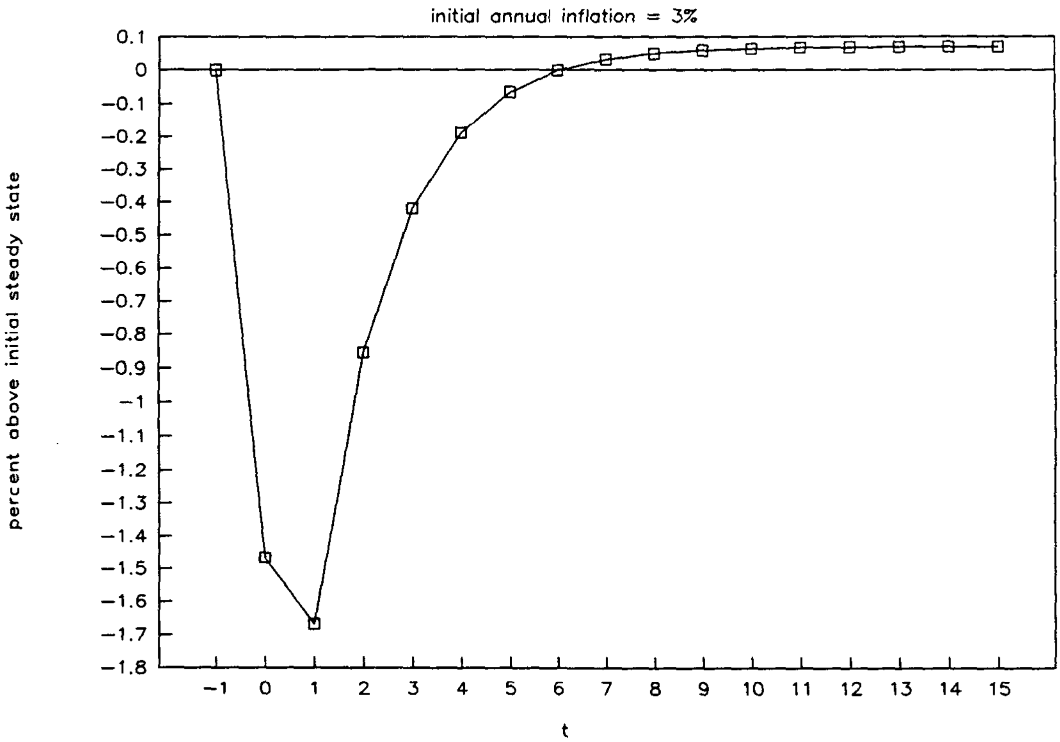

Fig. 4. Output Effects of Disinflation

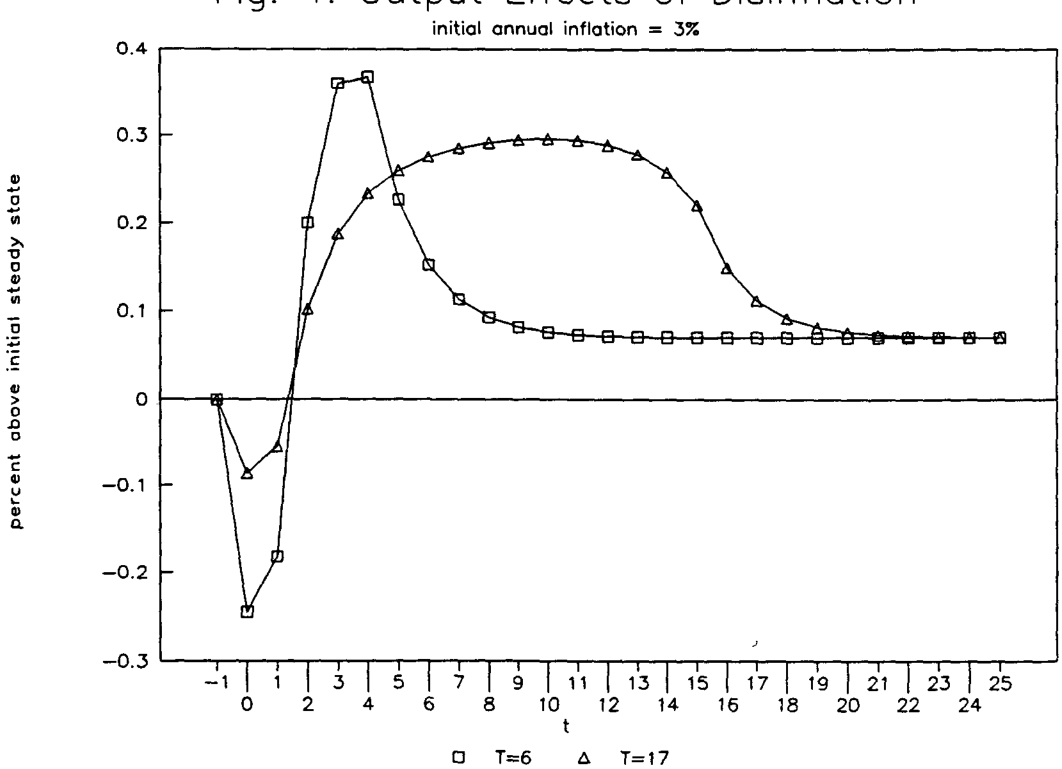

Fig. 5. Output Effects of Disinflation

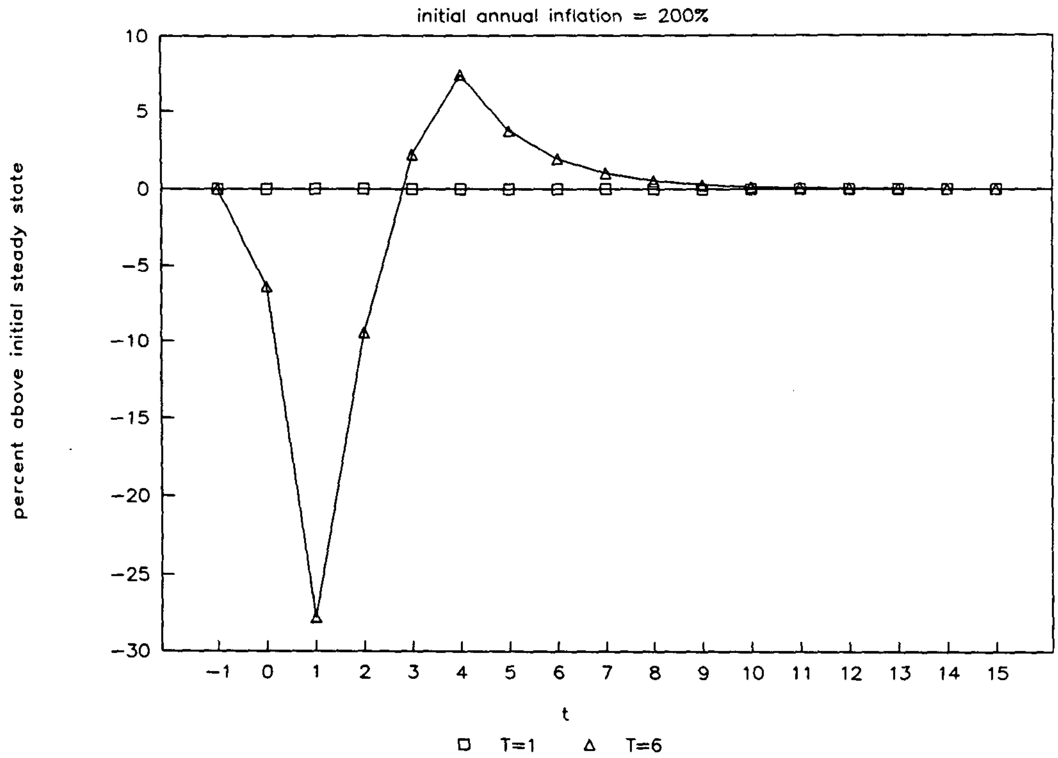

Fig. 6. Optimal Length of Disinflation