IMF *Staff Papers*
Vol. 45, No. 4 (December 1998)

Warning: Inflation May Be Harmful to Your Growth

ATISH GHOSH and STEVEN PHILLIPS*

While few doubt that very high inflation is bad for growth, there is less agreement about the effects of moderate inflation. Using panel regressions and allowing for a nonlinear specification, this paper finds a statistically and economically significant negative relationship between inflation and growth, which holds robustly at all but the lowest inflation rates. A "decision-tree" technique identifies inflation as one of the most important determinants of growth. Finally, short-run growth costs of disinflation are only relevant for the most severe disinflations, or when the initial inflation rate is well within the single-digit range. [JEL E31, 040]

RAPID OUTPUT GROWTH and low inflation are the most common objectives of macroeconomic policy. It is rather surprising, therefore, that a consensus about the relationship between these two variables is yet to emerge. While early studies by Phillips (1958) suggested an exploitable trade-off between output and price stability, the stagflationary experience of the industrialized countries in the 1970s belied this finding and showed that, beyond the short run, any such trade-off is illusory. More recent cross-country studies, particularly those that include middle- and low-income countries in their samples, suggest a *negative* relationship between growth and inflation.[1] Even

* Atish Ghosh is an Economist and Steven Phillips is a Senior Economist, both in the Policy Development and Review Department. The authors would like to thank Kadima Kalonji for research assistance, and Alan Taylor and Maurice Obstfeld for making available their computer programs. Hugh Bredenkamp, Matthew Canzonieri, Sharmini Coorey, Peter Doyle, Martin Evans, Stanley Fischer, Robert Flood, Manuel Guitian, Javier Hamann, Timothy Lane, Michael Sarel, Susan Schadler, and Tsidi Tsikata provided helpful comments on an earlier draft, as did participants in seminars at Georgetown University and at the Graduate Institute of International Studies (Geneva).

[1] Fischer (1993), Bruno and Easterly (1995), Judson and Orphanides (1996), Sarel (1996), Barro (1995), and IMF (forthcoming).

among these studies, however, there is little agreement on whether the empirical association of lower inflation with faster growth is statistically and economically significant, let alone causal.[2]

If prices exhibit downward rigidity, then very low inflation rates may ossify the structure of relative prices, impeding adjustment to real shocks. A little inflation, therefore, might help to "grease" the economy. On the other hand, high inflation rates, by confounding relative price signals and making efficient resource allocation more difficult, could result in more sluggish economic growth. But whether these or other negative effects begin at single-digit inflation rates, or only at much higher rates, remains a controversial question. Moreover, it is not clear that a rise in inflation causes a *proportional* worsening of the country's growth performance: it might be that, once chaotic inflation rates have been reached, relative prices cease to have much meaning anyway, making further increases in inflation less important.

In a multivariate context, the inflation-growth relationship becomes yet more complicated. Obviously, growth-inflation regressions must include other plausible determinants of growth. Several issues then arise. First, the inflation-growth findings may not be robust once "conditioning" variables are included in a regression analysis. Levine and Zervos (1993), for example, find that inflation does not survive Leamer's extreme bounds tests in growth regressions. Second, the conditioning variables may themselves be functions of the inflation rate. For instance, investment affects GDP growth, but may itself be affected by inflation. To the extent that inflation influences growth through such indirect effects, inclusion of these variables in a growth regression reduces the apparent effect of inflation. Third, there may be rich and important interactions between inflation and the other determinants of growth. For example, the marginal effect of inflation on growth may differ according to the level of physical and human capital in the country. With growth having many possible determinants, it may be difficult to model such interactions, especially since theory provides little guidance on the appropriate specification. Fourth, inflation is not under direct policy control; especially in the short run, it is affected by shocks that can influence both inflation and growth, possibly resulting in spurious correlations. Finally, even if low inflation is generally associated with faster growth, it does not necessarily follow that *dis*inflation is always good for growth. In particular, rapid disinflation may result in lower growth, at least in the short run.

These considerations suggest that, if a relationship between inflation and GDP growth exists, it is not likely to be a simple one. The bivariate relationship will not be monotonic, let alone linear; there may be important interaction effects between inflation and the other determinants of growth; and

[2] Clark (1993).

the correlation between disinflation and growth may be quite different from the steady-state inflation-growth relationship. Perhaps the lack of a consensus about the effects of inflation on growth is not so surprising after all.

In this paper, we try to address these various methodological problems and examine the relationship between inflation and disinflation and output growth. We employ a large panel data set, covering IMF member countries over 1960–96. Our primary analytical tool is a panel regression, in which our main contribution is to combine a nonlinear treatment of the inflation-growth relationship with an extensive examination of robustness. Complementing this analysis is our use of a decision-theoretic ("tree") technique that is more robust to outliers and nonlinearities than is standard regression analysis. Throughout, the emphasis is on examining the still-controversial question of whether there is *any* robust inflation-growth relationship, rather than pinning down the dynamics of such a relationship or identifying specific mechanisms through which inflation (or the policy choices it reflects) might influence growth.

In general, we find a negative relationship between inflation and growth that is statistically significant and of an economically interesting magnitude. This finding survives a battery of robustness checks. While we cannot rule out the possibility that part of this negative relationship stems from effects of growth on inflation, we still find a statistically and economically significant relationship between inflation and GDP growth when we use several sets of instruments to control for such simultaneity. But even if low inflation is associated with more rapid output growth, it is possible that the process of *dis*inflation may—at least in the short run—depress GDP growth. Our results here are striking. Disinflation tends to reduce growth only if the starting level of inflation is already very low, or if the pace of disinflation is severe.

Our more detailed results may be summarized briefly. First, there are two important nonlinearities in the inflation-growth relationship. At very low inflation rates (around 2–3 percent a year, or lower), inflation and growth are positively correlated. Otherwise, inflation and growth are negatively correlated, but the relationship is convex, so that the decline in growth associated with an increase from 10 percent to 20 percent inflation is much *larger* than that associated with moving from 40 percent to 50 percent inflation. Taking both these nonlinearities into account, we find that the negative inflation-growth relationship is evident in both the time and cross-section dimensions of the data, and that it is quite robust. Excluding high inflation observations, time-averaging the data, or using various subsamples (defined according to time period or the degree of inflation) does not alter the basic findings. We also find that inflation is a robust regressor in Leamer's extreme bounds sense, and that allowing for nonlinear relationships between the other regressors and GDP growth does not diminish the inflation-growth association.

To allow for threshold effects and nonlinear interactions, we use a technique known as *binary recursive trees.* The key advantage of this technique is its robustness to alternative specifications and to outliers. Indeed, the results are invariant to any monotone transformation of the variables. Importantly, this decision-theoretic analysis identifies inflation as one of the most important determinants of GDP growth (second only to physical and human capital).

Turning to the short-run consequences of rapid *dis*inflation, we find that starting from inflation rates above 6 percent, only the most drastic disinflations (at least halving the inflation rate in a single year) are associated with any negative impact on growth (which itself is largely offset by the higher growth associated with the new lower level of inflation). Starting from lower inflation rates, however, a rapid disinflation (halving the inflation rate) is associated with a fall in GDP growth.

I. Basic Statistics and Correlations

Our complete data set consists of 3,603 annual observations on real *per capita* GDP growth, and period average consumer price inflation, corresponding to 145 countries, over the period 1960–96.[3]

As a first step in exploring the bivariate relationship between inflation and growth, Figure 1 graphs the joint frequency distribution of inflation and growth. It is noteworthy that, while there are relatively few observations with inflation above 20 percent a year, they occur predominantly with negative per capita GDP growth rates. Indeed, two-fifths of the observations with inflation

[3] Each observation corresponds to an individual country for a given year. Of course, with the wide range of countries in the sample, the quality of the underlying data probably varies enormously. For this reason, the results are presented both in aggregate and with breakdowns by per capita income.

The original data set actually consists of 3,772 observations, but some observations were excluded at the outset. First, observations with GDP growth above 30 percent per year (5 observations) or below –30 percent a year (5 observations) were excluded because such extreme values may be unreliable or occur under exceptional circumstances (e.g., civil war) so that their relevance for economic policymaking is suspect. Their inclusion in the data set does not alter the basic conclusions (see Section III). Second, cases of negative inflation (159 observations) were excluded, not only for being outside the range of interest here, but for the practical reason that the analysis below requires taking the logarithm of the inflation rate. Including such cases by replacing negative inflation rates either by a small positive number (Sarel, 1996) or by their absolute value (IMF, forthcoming) does not alter the basic conclusions of this study.

Figure 1. *Joint Frequency Distribution of Inflation and GDP Growth Rates*
(In percent a year)

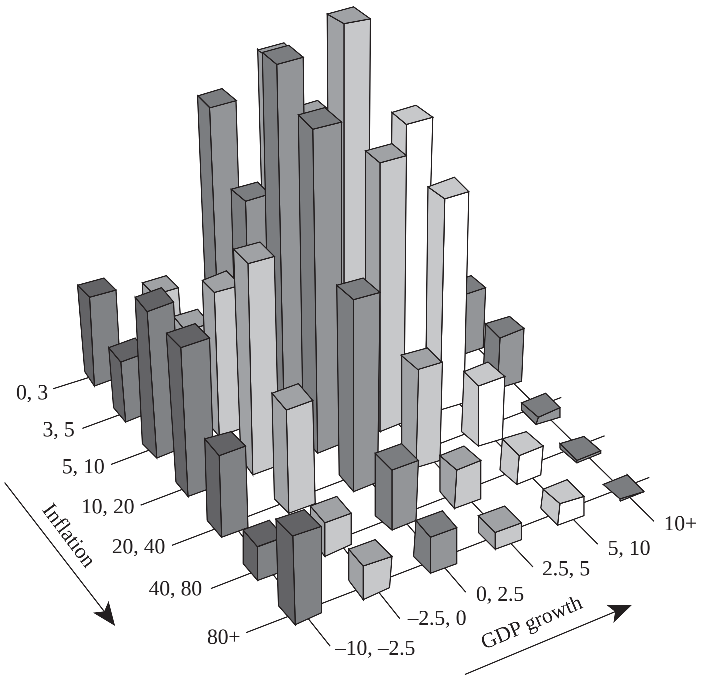

Note: Number of observations shown on vertical scale: total observations: 3,603; maximum: 247; minimum: 1.

above 20 percent show up at negative GDP growth rates, compared to only one-fifth of the cases with inflation below 20 percent. Alternatively, of the observations with positive GDP growth, more than three-fourths occur at inflation rates below 20 percent a year. Thus, grouping the data in this way suggests a negative association between inflation and growth.

Table 1 presents much the same information, but in tabular form, and for several different samples (the "consistent sample" consists of the 2,231 observations, for 103 countries, for which data on all of the conditioning variables used below are available). Again, the bivariate evidence suggests a negative relationship between inflation and growth. This relationship appears to break down, however, somewhere in the very low inflation range.

Figure 2, which is central to our results, gives a more direct view of the inflation-growth association by plotting the median GDP growth rate against the median inflation rate (for each of 20 equal-sized subsamples defined

Table 1. *Basic Statistics*
(In percent a year)

	Number of observations	Inflation		GDP growth	
		Mean	Median	Mean	Median
			Large Sample		
All observations	3,603	39.1	8.3	1.8	2.2
$0 < \pi < 3$	628	1.7	1.8	2.6	2.7
$3 < \pi < 5$	525	3.9	4.0	2.8	2.9
$5 < \pi < 10$	913	7.4	7.3	2.4	2.6
$10 < \pi < 20$	843	14.0	13.3	1.8	1.8
$20 < \pi < 40$	394	27.3	26.1	0.4	0.9
$40 < \pi < 80$	142	56.7	54.6	0.9	1.3
$\pi > 80$	158	635.4	166.9	–3.8	–2.9
			Consistent Sample		
All observations	2,231	42.0	9.3	1.9	2.2
$0 < \pi < 3$	321	1.8	1.9	2.6	2.7
$3 < \pi < 5$	303	4.0	4.0	2.8	2.9
$5 < \pi < 10$	570	7.5	7.4	2.6	2.7
$10 < \pi < 20$	568	13.9	13.2	1.6	1.5
$20 < \pi < 40$	272	27.3	26.2	0.6	1.0
$40 < \pi < 80$	104	56.6	55.0	1.0	1.4
$\pi > 80$	93	715.7	163.4	–1.9	–1.0
			Consistent Sample Upper- and upper-middle-income countries		
All observations	937	36.3	6.7	2.7	2.7
$0 < \pi < 3$	180	2.0	2.1	3.6	3.2
$3 < \pi < 5$	183	3.9	3.9	3.5	3.5
$5 < \pi < 10$	244	7.2	7.1	2.8	2.9
$10 < \pi < 20$	177	14.0	13.5	2.0	2.0
$20 < \pi < 40$	66	26.0	25.0	2.1	2.2
$40 < \pi < 80$	37	56.6	56.6	2.5	2.4
$\pi > 80$	50	497.1	168.2	–0.7	0.1
			Consistent Sample Lower- and lower-middle-income countries		
All observations	1,294	46.2	10.8	1.3	1.7
$0 < \pi < 3$	141	1.5	1.7	1.4	1.8
$3 < \pi < 5$	120	4.0	4.1	1.8	2.3
$5 < \pi < 10$	326	7.6	7.8	2.4	2.5
$10 < \pi < 20$	391	13.8	13.0	1.5	1.4
$20 < \pi < 40$	206	27.7	26.6	0.1	0.7
$40 < \pi < 80$	67	56.5	54.5	0.2	0.5
$\pi > 80$	43	969.9	161.0	–3.5	–3.5

Table 1. *(concluded)*

	Number of observations	Inflation		GDP growth	
		Mean	Median	Mean	Median
		Consistent Sample Post-1973 observations			
All observations	1,786	50.1	10.6	1.5	1.9
$0 < \pi < 3$	204	1.7	1.8	2.0	2.3
$3 < \pi < 5$	195	4.0	4.1	2.7	2.7
$5 < \pi < 10$	442	7.6	7.6	2.3	2.6
$10 < \pi < 20$	513	13.9	13.3	1.5	1.4
$20 < \pi < 40$	252	27.4	26.3	0.4	0.9
$40 < \pi < 80$	93	56.5	54.8	1.2	1.4
$\pi > 80$	87	754.6	171.7	–2.1	–1.6

according to degree of inflation).[4] Again, the concentration of inflation observations in the 0–20 percent range is evident, but this data-smoothing technique also makes two key features of the data immediately apparent. First, at the very lowest inflation rates, for which there are quite a few observations, inflation and growth are *positively* associated. Second, at all other inflation rates, the relationship is negative and clearly *convex*—implying, plausibly, that an increase in inflation from 5 percent to 25 percent impairs growth more than an increase from 100 percent to 120 percent. The slope is quite flat over the highest inflation ranges; such observations are a small part of the sample, but as outliers their effective weight in a regression analysis may be considerable. Indeed, Figure 2 suggests that ignoring these nonlinearities and regressing growth linearly on the inflation rate would impart a downward bias to the estimated slope over the range of greatest policy interest. This bias may, at least in part, account for the failure of previous studies to detect a significant and robust negative relationship between inflation and growth.

II. Conditioning Variables: Multivariate Inflation-Growth Regressions

Findings of a negative correlation between inflation and growth suggest—but obviously do not prove—the notion that lower inflation promotes faster growth. Causality aside, it is natural to suspect that part of the correlation may be spurious, reflecting the effects of third factors. This section checks

[4] Using group means rather than medians gives a similar picture, except that the points for the groups with the highest inflation observations shift to the right, making the plotted curve even more convex.

Figure 2. *Inflation and Per Capita GDP Growth*[1]
(In percent a year)

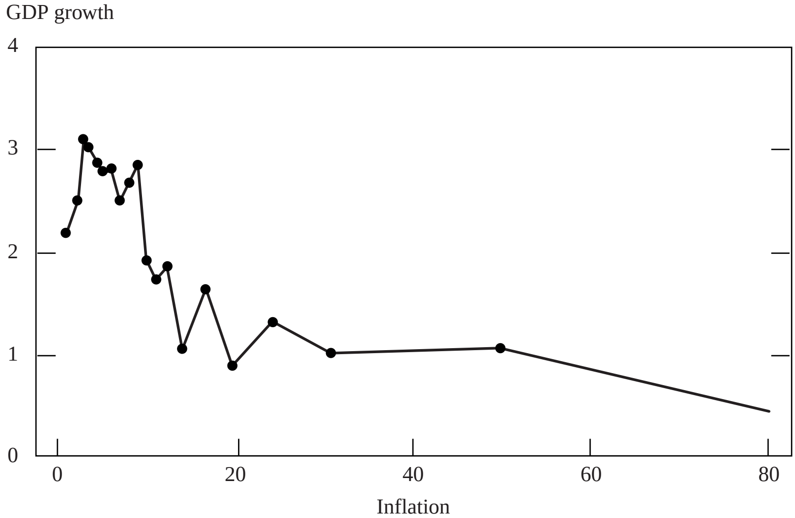

[1] Median inflation and growth rates in equal-sized subsamples, defined according to range of inflation (right-most point not shown).

whether an inflation-growth relationship appears also in multivariate regression analysis. The intent is not to develop an explanatory model of GDP growth, but rather to determine whether the inflation-growth correlation is robust to including a set of *conditioning* variables. The analysis also checks for nonlinearity of the inflation-growth relationship.

A first step is to include annual dummies in regressing per capita GDP growth on inflation.[5] More generally, the regression is augmented with other proposed growth determinants; for this, the empirical literature pullulates with possibilities. Neoclassical theory stresses capital accumulation as the engine of (pre-steady state) growth. More recent growth theories also emphasize the importance of human capital. Various measures of human capital, such as school enrollment rates, average years of primary and secondary education completed, and life expectancy, have been proposed. These tend to be highly correlated, so, rather than include them individually, we use the first principal component of primary and secondary school enrollment rates and life expectancy as a measure of human capital (*HK*).

[5] Here, we do not include country dummies since these are highly correlated with some of the additional regressors. In the robustness section, we show that adding such dummies actually strengthens the inflation-growth findings.

Beyond the physical and human capital variables suggested by theory, a largely ad hoc smorgasbord of factors that might affect "productivity" is usually included in growth regressions. To control for "catch-up" effects, the log of the ratio of U.S. per capita income to country *j*'s per capita income in 1960, and measured in international prices, is used (*GAP*). A large tax burden on the economy or a large share of public consumption might depress economic growth; we include the ratios of revenues to GDP (τ/GDP), public consumption to GDP (*G*/*GDP*), and the fiscal balance (*B*/*GDP*).

A number of studies stress the importance of openness to international trade, both as a means of effecting the transfer of technical progress and as an engine of growth; we use the sample average of the ratio of exports plus imports to GDP, ($(X + M)/GDP$). The (log of the) black market exchange rate premium (*BLK*) provides a measure of the overvaluation of the real exchange rate and, in at least some instances, of economic mismanagement more generally. The terms of trade volatility, σ_{TT}, is used as a measure of the importance of external shocks.[6] Finally, we include indicator variables for cataclysmic events such as drought (*DROUGHT*), or cases where there are war-related deaths (*DEATH*). By controlling for these types of supply shocks, these regressors should reduce the chances of picking up spurious (negative) inflation-growth comovements.

Some of these variables, such as drought or war, are clearly exogenous with respect to inflation. But other variables, most notably the investment rate, are likely to be influenced by inflation. To the extent that inflation affects growth by influencing these conditioning variables (and they, in turn, affect growth), their inclusion in the regression could diminish the *measured* effects of inflation on growth. Since there is no easy way around this problem, we report results both including and excluding the investment ratio. (In the robustness section, we undertake a more systematic analysis of the effects on the inflation-growth relationship of including and excluding the various other regressors.)

The next step is to model the evident nonlinearity of the inflation-growth relationship. From Figure 2, the positively sloped part of this relationship ceases at inflation rates somewhere around 2–3 percent a year. To deal with this nonlinearity, we follow Sarel (1996) and use a spline technique, allowing the relationship to have a "kink"or turning point where $\pi = 2\,^1/_2$ percent.[7] (As shown in Section III, the exact location of this kink turns out not to be important for the questions on which we focus.)

[6] Current and lagged terms of trade changes were found to be insignificant; their inclusion in the regression would change the inflation-growth results only slightly (see Section III).

[7] It turns out that $2\,^1/_2$ percent inflation is also the placement of the kink that yields the best fit of a multivariate growth regression (see Section III).

This leaves the question of how to capture the convexity of the negative inflation-growth relationship. We consider a number of possibilities. Our first model (denoted model (1) in Table 2) simply ignores the convexity, specifying a linear relationship (beyond the kink at 2 $^1/_2$ percent). Model (2) uses the real rate of depreciation of the currency, defined as $\pi/(1+\pi)$, as the measure of inflation.[8] Model (3) uses the log of the inflation rate, and model (4) generalizes this by replacing $\log(\pi)$ with $(1-\gamma)^{-1}\pi^{(1-\gamma)}$, where γ is estimated via nonlinear least squares. This specification collapses to the linear specification as γ approaches zero, and to the logarithmic specification as γ approaches unity.[9]

Estimates for the inflation-related parameters in these four specifications are reported in Table 2, and the implied GDP growth rates at various inflation rates are illustrated in Figure 3. The coefficient on (the various) measures of inflation (when inflation is above 2 $^1/_2$ percent a year), given by β_1, is always negative and statistically significant, with heteroscedastic-consistent *t*-statistics ranging from about 3 to over 10.

In specification (1)—which is linear beyond the kink at 2 $^1/_2$ percent inflation—the inflation coefficient, though statistically significant, is economically paltry. Indeed, the negative slope for this model is barely discernible in Figure 3. The linear model suggests that raising inflation from 10 percent a year to 20 percent a year would be associated with a mere 0.01 percentage point reduction in annual growth. (It is easy to see that even weaker results would appear if the kink at 2 $^1/_2$ percent were not allowed, and complete linearity imposed.)

In contrast, the nonlinear models—the real rate of depreciation (2), the logarithmic (3), and the more general nonlinear variant (4)—are all suggestive of economically important effects over the inflation range of greatest policy interest. According to these models, an increase in annual inflation from 10 percent to 20 percent a year would be associated with a reduction of per capita GDP growth by about 0.3–0.4 percentage points, while an increase in inflation from 10 percent to 40 percent a year would be associated with about 0.8 percentage points slower growth.[10] Figure 3 also shows that these three nonlinear models tend to give relatively similar predictions about the apparent effect of inflation on growth—which are far greater than the predictions implied by the linear model.

[8] See, for example, Cukierman (1992) for use of this specification. A log-based specification has been used by Sarel (1996), IMF (forthcoming), and others.

[9] We do not, however, use $\log(1+\pi)$—a specification suggested by some authors—because this function is close to being linear over the range in which most inflation observations lie.

[10] The predicted changes would be even larger if one assumed that inflation would be reduced in part by raising *B*/*GDP* (cutting the budget deficit). Of course, one wants to be careful about applying causal interpretations to growth regression results; this issue is discussed later.

Table 2. *Alternative Inflation-Growth Regressions*

(1) $\Delta y = \alpha_0 + \beta_{0(\pi < 0.025)} (\pi - 0.025) + \beta_1 \pi$

(2) $\Delta y = \alpha_0 + \beta_{0(\pi < 0.025)} ([\pi/(1 + \pi)] - 0.026) + \beta_1 (\pi/(1 + \pi))$

(3) $\Delta y = \alpha_0 + \beta_{0(\pi < 0.025)} [\log(\pi) - \log(0.025)] + \beta_1 \log(\pi)$

(4) $\Delta y = \alpha_0 + \beta_{0(\pi < 0.025)} (\pi - 0.025) + \beta_1 (1 - \gamma)^{-1}\pi^{(1-\gamma)}$

Inflation measure	α_0	*t*-statistic	β_0	*t*-statistic	β_1	*t*-statistic	γ	*t*-statistic	R^2
Full Sample									
No conditioning variables									
(1) π	0.0198	20.29	0.0767	0.27	–0.0021	–3.55**			0.02
(2) $[\pi/(1 + \pi)]$	0.0294	22.72	0.6106	2.22	–0.0779	–10.59**			0.06
(3) $\log(\pi)$	–0.0050	–2.12	0.0229	5.78	–0.0108	–11.75**			0.07
(4) $(1 - \gamma)^{-1}(\pi^{(1-\gamma)})$	0.0935	2.39	1.5292	4.49	–0.0127	–11.25**	0.8700	15.09	0.06
Annual dummies									
(1) π	0.0200	21.19	0.3461	1.25	–0.0020	–3.59**			0.08
(2) $[\pi/(1 + \pi)]$	0.0287	22.20	0.7771	2.88	–0.0704	–9.39**			0.11
(3) $\log(\pi)$	–0.0026	–1.04	0.0225	5.90	–0.0098	–10.03**			0.11
(4) $(1 - \gamma)^{-1}(\pi^{(1-\gamma)})$	0.0935	2.39	1.5292	4.49	–0.0127	–11.25**	0.8700	15.09	0.10
Conditioning variables except *I*/*GDP*									
(1) π	0.0199	22.72	0.5305	2.20	–0.0013	–2.99**			0.21
(2) $[\pi/(1 + \pi)]$	0.0259	19.34	0.7527	3.18	–0.0487	–5.86**			0.22
(3) $\log(\pi)$	0.0039	1.49	0.0176	5.40	–0.0069	–6.40**			0.22
(4) $(1 - \gamma)^{-1}(\pi^{(1-\gamma)})$	0.0706	2.05	1.3530	4.15	–0.0094	–8.39**	0.8625	11.36	0.22

All conditioning variables									
(1) π	0.0198	22.87	0.3907	1.61	–0.0013	–3.07**			0.23
(2) $[\pi/(1+\pi)]$	0.0252	19.14	0.6020	2.54	–0.0447	–5.46**			0.24
(3) $\log(\pi)$	0.0049	1.91	0.0157	4.74	–0.0064	–6.04**			0.24
(4) $(1-\gamma)^{-1}(\pi^{(1-\gamma)})$	0.0605	2.21	1.0164	3.23	–0.0087	–7.98**	0.8473	10.52	0.24
Upper- and Upper-Middle-Income Countries									
Conditioning variables except *I*/*GDP*									
(1) π	0.0320	11.06	–0.2794	–0.72	–0.0018	–2.94**			0.37
(2) $[\pi/(1+\pi)]$	0.0384	9.19	–0.1572	–0.43	–0.0361	–2.85**			0.38
(3) $\log(\pi)$	0.0204	5.53	0.0063	1.18	–0.0059	–3.70**			0.38
(4) $(1-\gamma)^{-1}(\pi^{(1-\gamma)})$	0.0697	3.10	0.1617	0.45	–0.0081	–5.50**	0.8110	7.15	0.38
All conditioning variables									
(1) π	0.0277	9.54	–0.3571	–0.91	–0.0018	–2.96**			0.41
(2) $[\pi/(1+\pi)]$	0.0323	7.68	–0.2503	–0.68	–0.0279	–2.25*			0.41
(3) $\log(\pi)$	0.0188	5.15	0.0033	0.61	–0.0045	–2.79**			0.41
(4) $(1-\gamma)^{-1}(\pi^{(1-\gamma)})$	0.0492	4.39	–0.2509	–0.77	–0.0071	–5.15**	0.7249	5.68	0.41
Lower- and Lower-Middle-Income Countries									
Conditioning variables except *I*/*GDP*									
(1) π	0.0177	8.78	1.0628	3.49	–0.0012	–2.43*			0.19
(2) $[\pi/(1+\pi)]$	0.0264	10.48	1.4211	4.64	–0.0685	–5.61**			0.21
(3) $\log(\pi)$	–0.0019	–0.47	0.0256	6.41	–0.0088	–5.38**			0.21
(4) $(1-\gamma)^{-1}(\pi^{(1-\gamma)})$	0.0899	1.50	2.1944	4.38	–0.0118	–7.08**	0.8741	10.09	0.19
All conditioning variables									
(1) π	0.0171	8.70	0.9395	3.03	–0.0012	–2.51*			0.21
(2) $[\pi/(1+\pi)]$	0.0255	10.34	1.2894	4.17	–0.0657	–5.44**			0.22
(3) $\log(\pi)$	–0.0019	–0.46	0.0242	5.93	–0.0085	–5.35**			0.22
(4) $(1-\gamma)^{-1}(\pi^{(1-\gamma)})$	0.0856	1.44	1.9737	3.98	–0.0112	–6.80**	0.8744	9.72	0.21

Table 2. *(concluded)*

Inflation measure	α_0	*t*-statistic	β_0	*t*-statistic	β_1	*t*-statistic	γ	*t*-statistic	R^2
				Pre-1973 Observations					
				Conditioning variables except *I/GDP*					
(1) π	0.0386	6.88	–0.4603	–1.03	–0.0260	–3.23**			0.19
(2) $[\pi/(1+\pi)]$	0.0417	6.59	–0.3371	–0.78	–0.0686	–2.50*			0.19
(3) $\log(\pi)$	0.0176	2.15	0.0046	0.79	–0.0071	–2.31*			0.19
(4) $(1-\gamma)^{-1}(\pi^{(1-\gamma)})$	0.0449	5.95	–0.3474	–0.77	–0.0233	–2.24*	0.3373	0.96	0.20
				All conditioning variables					
(1) π	0.0370	6.52	–0.4505	–1.03	–0.0229	–3.00**			0.21
(2) $[\pi/(1+\pi)]$	0.0394	6.08	–0.3565	–0.83	–0.0570	–2.05*			0.21
(3) $\log(\pi)$	0.0201	2.44	0.0027	0.46	–0.0055	–1.75			0.20
(4) $(1-\gamma)^{-1}(\pi^{(1-\gamma)})$	0.0412	7.35	–0.4390	–1.04	–0.0228	–2.31*	0.1999	0.49	0.21
				Post-1973 Observations					
				Conditioning variables except *I/GDP*					
(1) π	0.0153	14.07	0.8676	3.03	–0.0012	–2.91**			0.21
(2) $[\pi/(1+\pi)]$	0.0212	13.72	1.1110	3.99	–0.0477	–5.49**			0.22
(3) $\log(\pi)$	–0.0012	–0.43	0.0232	6.28	–0.0072	–6.20**			0.23
(4) $(1-\gamma)^{-1}(\pi^{(1-\gamma)})$	0.0442	2.34	1.6066	4.49	–0.0082	–7.29**	0.8053	8.59	0.22
				All conditioning variables					
(1) π	0.0161	14.97	0.7095	2.46	–0.0012	–2.96**			0.23
(2) $[\pi/(1+\pi)]$	0.0216	14.20	0.9450	3.37	–0.0446	–5.21**			0.24
(3) $\log(\pi)$	0.0004	0.14	0.0214	5.67	–0.0068	–6.03**			0.24
(4) $(1-\gamma)^{-1}(\pi^{(1-\gamma)})$	0.0477	2.05	1.3407	3.70	–0.0078	–7.05**	0.8259	8.55	0.23

Note: One asterisk and two asterisks indicate statistical significance at the 5- and 1-percent level, respectively. Statistical significance shown only for coefficient β_1 (estimated slope for inflation above 2 $^1/_2$ percent a year).

Figure 3. *Implied Per Capita GDP Growth Rates Under Alternative Models*[1]
(In percent a year)

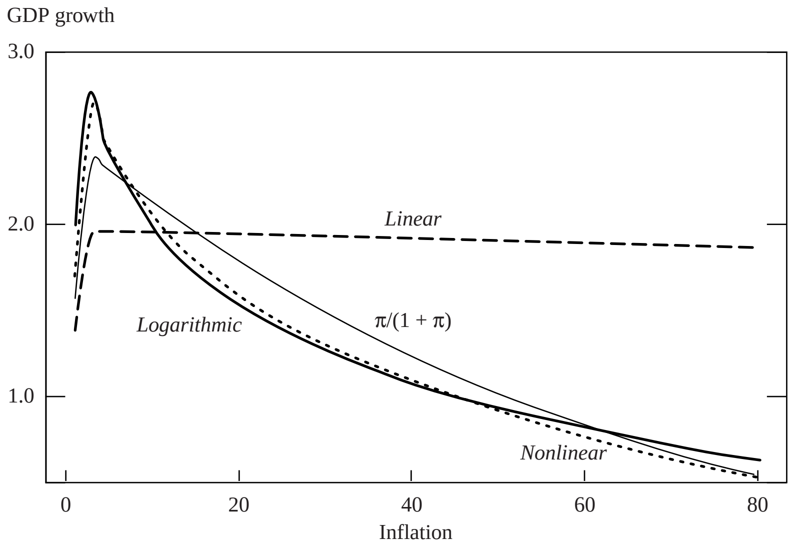

[1] Bivariate inflation-growth relationship, holding constant all other explanatory variables at mean values.

How best to choose between the models? Note that in model (4), the estimated value of γ is significantly different from zero, thus *rejecting the linear model.* The estimates of γ tend to be somewhat smaller than (but not significantly different from) unity, at least in the variant that includes all the conditioning variables. Thus the log model cannot be rejected.

We conclude that the simple logarithmic model—with a low-inflation kink—provides a reasonable characterization of the inflation-growth relationship. Our base model then becomes:

$$\begin{array}{lllll}
\Delta y = & \underset{(1.91)}{0.004} & + \underset{(4.74^{**})}{0.015\, D25\,[\log(\pi) - \log(0.025)]} & & - \underset{(6.04^{**})}{0.0064 \log(\pi)} \\
& + \underset{(2.12^{**})}{0.019\, \sigma_{TT}} & - \underset{(2.66^{**})}{0.027\, \tau/GDP} & + \underset{(0.04)}{0.002\, G/GDP} & - \underset{(4.25^{**})}{0.325\, \Delta POP} \\
& + \underset{(4.46^{**})}{0.008\, HK} & + \underset{(5.39^{**})}{0.109\, B/GDP} & - \underset{(3.24^{**})}{0.009\, BLK} & - \underset{(7.06^{**})}{0.02\, DROUGHT} \\
& - \underset{(2.30^{**})}{0.07\, DEATH} & + \underset{(5.40^{**})}{0.010\, GAP} & + \underset{(1.85^{*})}{0.010\, (X+M)/GDP} & + \underset{(6.04^{**})}{0.086\, I/GDP} ,
\end{array} \quad (1)$$

where *D*25 is a dummy variable equal to unity when inflation is less than $2^1/_2$ percent, and where the coefficients on the annual dummy terms (about half of which are significant) are not reported.[11]

All but one of the coefficient estimates shown above are statistically significant,[12] and all except the one associated with terms of trade volatility have the expected signs. Of greater interest, the coefficient on log inflation is negative and significant by a wide margin. Moreover, the positive coefficient on the spline term is significant, rejecting the hypothesis of no break in the relationship.[13] Despite the dilemma noted above, and somewhat surprisingly, there is little change in the inflation coefficient if the investment term is dropped.[14]

We examine the robustness of these base model results in the next section, but before proceeding, several points on interpretation may be useful. Inflation is of course not under direct policy control; especially in the short run, it is more of an outcome determined by both macroeconomic policy choices and various shocks and is therefore probably best thought of as an *indicator* of those policy choices. We will use several methods to determine whether the inflation-growth correlation found in annual panel data is mainly spurious, being driven for example by short-run shocks or policy responses. On the other hand, we do not attempt to identify the exact mechanisms or channels through which inflation—or the related policy choices it reflects—might hinder output growth. Still, several points can be noted in this connection. The inflation variable in equation (1) evidently picks up the influence of policy choices other than high government consumption, high budget deficits, or high black market exchange rate premiums, since these are also included in the regression. Also, as shown later, the inflation variable captures something other than the effects of the inflation volatility associated with higher inflation. Finally, in whatever way inflation or its correlates influence growth, it does not seem to be mainly an indirect effect through investment.

III. Robustness

The question of robustness is of particular interest in the empirical analysis of growth, since economic theory provides little guidance on the "true"

[11] Even with annual dummies, the R^2 of the regression is only about 0.25. With time-averaged data, however, the R^2 rises to as much as 0.70 (see Section III).

[12] The exception is *G*/*GDP*; we choose to retain this regressor because it is statistically significant in the variant excluding *I*/*GDP*.

[13] The coefficient on the spline term implies an estimated slope of +0.086 in the low inflation range (i.e., 0.015 greater than the slope elsewhere).

[14] The main effect of dropping the investment term is to increase the coefficient on human capital (with which investment is highly correlated).

specification. Here we examine the robustness of the negative inflation-growth association.

Robustness of the Negative Inflation-Growth Association

The robustness of this association has already received some attention in the literature. Perhaps best known are the results of Levine and Zervos (1993), suggesting that inflation-growth findings can depend on a very few countries with high inflation. Similarly, Bruno and Easterly (1995) show that excluding from a growth regression all countries with inflation above 40 percent can make inflation lose statistical significance. In a more general study, Sala-i-Martin (1997) finds a number of variables to be robustly associated with growth, but not inflation; however, he notes that his methodology presumes a linear relationship.

Indeed, we suspect that failure to allow a nonlinear association between inflation and growth is responsible for such negative results.[15] As noted, the likely consequence of imposing linearity is a very large downward bias in measuring the inflation-growth slope. Accordingly, we use as our base model the growth regression reported in equation (1) above (with all of the conditioning variables, including annual dummies).

In examining robustness, we consider questions in four categories. First, is the estimated coefficient on (log) inflation stable across various alternate samples? One specific concern here is the role of high inflation outliers. Second, does the relationship found in annual panel data derive from both the time-series and cross-sectional dimensions? Here, the concern is that the panel results might be spurious, driven either by fixed country-specific factors or reflecting mainly shocks that induce short-run correlations. Third, does the coefficient on log inflation remain stable when the specification of the conditioning variables is changed in various ways? Finally, we also check whether the results are sensitive to the exact placement of the low inflation kink allowed in the inflation-growth relationship.

Stability Across Samples

There is some suspicion in the literature that the apparent negative effect of inflation on growth arises mainly from a small number of outlying cases, that is, countries with unusually high inflation and weak growth. For example, Levine and Zervos (1993) demonstrate that merely

[15] While the Bruno and Easterly specification is not strictly linear, their $\log(1 + \pi)$ regressor is essentially linear over the range in which most inflation observations lie. Moreover, they allow no kink in the low inflation range.

Table 3. *OLS Estimates from Alternative Specifications and Samples*

Regression variant	Coefficient on log inflation: Estimate	*t*-statistic	Sample size
(0) Base model, full sample	–0.00639	–6.04**	2231
(1) Excluding inflation > 40 percent	–0.00572	–3.79**	2034
(2) In differenced form			
5-year changes	–0.0108	–6.93**	1685
10-year changes	–0.00926	–6.85**	1250
15-year changes	–0.00837	–5.32**	871
(3) In differenced form, excluding inflation > 40 percent			
5-year changes	–0.00996	–5.07**	1459
10-year changes	–0.00768	–4.16**	1067
15-year changes	–0.00822	–4.12**	744
(4) Adding country dummies	–0.0083	–5.78**	2231
(5) Adding change in log inflation			
current change	–0.00629	–5.76**	2218
current change and 2 lags	–0.00597	–5.27**	2182
(6) Base model with annual data,			
1967–96 only	–0.00664	–6.19**	2139
excluding inflation > 40 percent	–0.00635	–4.18**	1947
(7) Pre-averaged data, 1967–96			
5-year averaging	–0.00501	–3.69**	360
10-year averaging	–0.00303	–2.37*	150
15-year averaging	–0.00423	–2.59*	74
(8) Pre-averaged data, 1967–96, excluding inflation > 40 percent			
5-year averaging	–0.00659	–3.20**	323
10-year averaging	–0.00229	–0.80	132
15-year averaging	–0.00606	–2.00	63

Note: One asterisk and two asterisks indicate statistical significance at the 5- and 1-percent level, respectively.

dropping Nicaragua and Uganda from a large cross-section regression can cause the observed relationship to break down. A similar, but somewhat broader, suspicion is that the apparent effect of inflation only becomes serious at rates above some fairly high threshold, perhaps 40 percent a year.[16] Thus, while no one is likely to argue that hyperinflation is good for growth, there is much less agreement on whether inflation in the 10–40 percent range has any deleterious effects on growth. In contrast,

[16] Such an interpretation is sometimes given to the Bruno and Easterly (1995) study.

Figure 4. *Coefficient Estimates from Regressions on Inflation Subsamples*

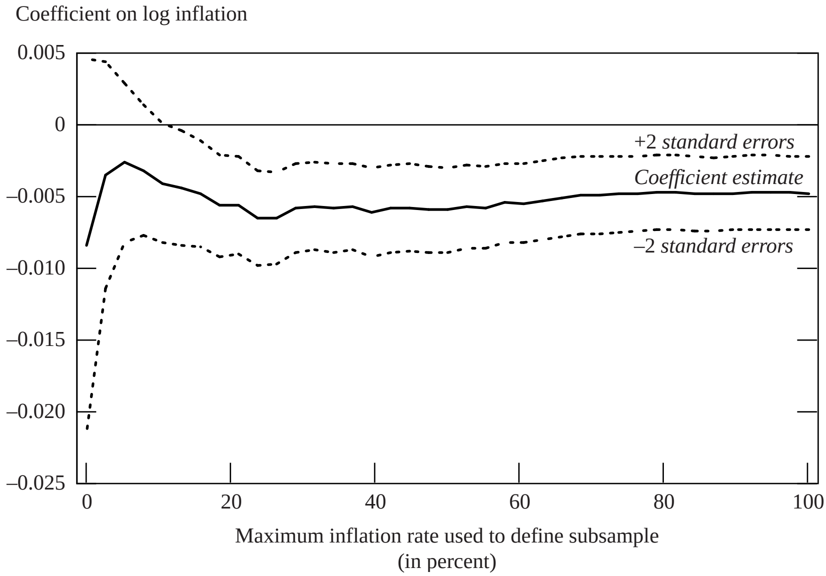

the results presented above suggest that the slope of the inflation-growth relationship is *steeper* in the 10–40 percent range than in the range above 40 percent.

The issue is disposed of readily, by reestimating the base model in a restricted sample. Rather than excluding only select high-inflation outliers, we subject the base model to a more comprehensive and severe test, excluding *all* observations with inflation greater than 40 percent.[17] The result is reported in Table 3 (regression (1)). The estimated coefficient changes only from –0.00639 to –0.00572, and it remains statistically significant (despite a much-reduced variation of inflation in this truncated sample).

Turning from the role of high inflation to more general questions of stability across samples, we examine what happens to the estimated coefficient as the range of inflation rates allowed in the estimation is varied systematically. We start with a very restricted sample, consisting of inflation rates in the 0–5 percent range, and then gradually expand the upper bound of the sample in small increments (2 $^1/_2$ percentage points of inflation). Figure 4 shows the coefficient estimate (±2 standard errors) in each sample. The point estimates are always negative, and the sample need include only the 0–17 $^1/_2$

[17] Results excluding just Nicaragua and Uganda were even closer (essentially identical) to the full sample results.

percent inflation range before statistical significance is found. Of particular interest, as the sample is extended to include inflation of greater than 40 percent, the estimated coefficient does not grow in absolute value. More generally, the estimated coefficient appears fairly stable across all but the smallest samples, and the width of the standard error bands never flares but instead slowly tapers as larger, more diverse samples are considered. These are signs of a well-specified model.

Similarly, we examine stability over time by segmenting the data into time periods. Figure 5a shows the coefficient estimates, starting with a subsample consisting of observations through 1966 only, and then adding one year of data at a time. Of the 31 point estimates, all but one is negative (the exception occurs in the smallest subsample). While these results seem to imply that almost 20 years of data are required before a finding of statistical significance can be assured, this could reflect the lesser information in the earlier years (missing observations and lower variation of inflation). When the same procedure is run “backwards” over time (i.e., starting with 1996 data only and progressively expanding the sample to include earlier years), it takes only two years of data to find statistical significance (Figure 5b). Note that neither Figure 5a nor Figure 5b suggests a structural break occurring at any time during the sample period. As in Figure 4, the standard error bands narrow as the sample is increased.

The Roles of the Cross-Sectional and Time Dimensions

The above results all suggest that the log-based model (with a kink for the low inflation observations) is well specified and the negative inflation-growth relationship robust. But do these panel data results arise from both the sample’s cross-sectional and time dimensions? Results coming from only one of these dimensions would be suspect.

The panel results imply that, comparing two countries with different inflation rates, the country with lower inflation may be expected to have higher growth. But this is not necessarily the same as saying that an individual country that achieves lower inflation is likely to achieve faster growth (even ignoring any possible short-run contractionary effects of disinflation). It could be that the panel data results are driven entirely by cross-country variation in inflation or in unmeasured country-specific factors associated with inflation.

We tackle this problem first by focusing on the effect of changes in the inflation rate on the change in the growth rate (the regression includes changes of all the independent variables that vary over time). This allows us to examine whether a country changing its inflation rate can expect a shift in its growth rate, while still pooling observations. In taking changes of the

Figure 5a. *Coefficient Estimates from Regressions on Time Subsamples*
(Samples start in 1963 and end in final year of regression subsample)

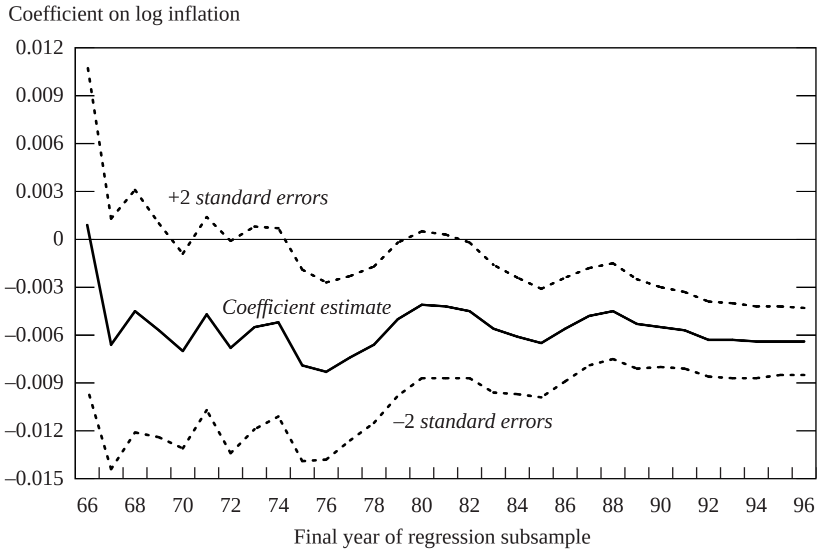

Figure 5b. *Coefficient Estimates from Regressions on Time Subsamples*
(Samples start in first year of regression subsample and end in 1996)

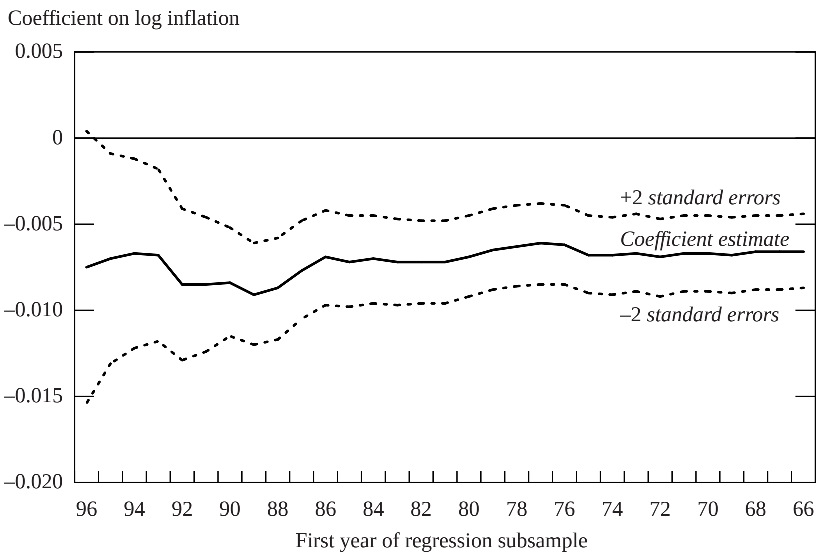

variables, we want a fairly long horizon, since results for short horizons—say, one-year changes—might be influenced by spurious short-run correlations induced by supply (or demand) shocks. Returning to Table 3, regressions (2)–(3) report the results for 5-, 10-, and 15-year changes, including and excluding high-inflation observations.

The inflation coefficient is not even slightly diminished by this transformation of the base model; indeed, it becomes considerably larger in absolute value. Moreover, this result cannot be attributed to outliers (i.e., a few cases of countries moving out of, or into, very high inflation) since it holds even when observations for which either current or (long-) lagged inflation exceeds 40 percent are excluded.

The fact that the negative inflation coefficient becomes steeper when the base model is specified in terms of changes may reflect fixed, country-specific effects: by purging the data of such effects, differencing the data could be correcting for the base model's omission of country dummy variables.[18] Indeed, adding country dummies also makes the inflation-growth slope steeper (Table 3, regression (4)).

These fixed-effects results are striking, since it is all too easy to imagine that unmeasured characteristics of some countries (e.g., weak institutions, or political polarization) somehow drive them to have both low growth and high inflation, inducing at least some degree of spurious correlation in cross-sectional or panel analysis. Were this the case, however, we would expect that adding country dummies to a panel regression would diminish—rather than steepen—the inflation coefficient.

Having confirmed that the negative inflation-growth relationship is apparent in the time dimension of the data, we now address the possibility that the panel data results might reflect *only or primarily* this dimension. For example, supply shocks and policy responses of various kinds could induce negative short-run correlations between inflation and growth. The question is how much our panel estimate might be spuriously influenced by such short-run comovements (keeping in mind that demand shocks may also be at work).

As a first check, we add to the base model the change in (log) inflation. This augmentation, however, as well as one also adding lagged changes, fails to move the estimated coefficient on log inflation more than slightly (Table 3, regression (5)). Moreover, none of the inflation *changes* terms are statistically significant. While a comprehensive analysis of inflation-

[18] Country dummies were not included in the base model because some of the conditioning variables are constant or nearly constant over the sample period, causing colinearity problems. It is, however, possible to estimate the inflation coefficient despite this colinearity. (The standard errors for certain other regressors jump, but this need not concern us here.)

growth dynamics is beyond the scope of this paper, some basic results can be noted briefly. When the base model is augmented with three lagged *levels* of log inflation, or with only two such lags, none of these terms is statistically significant. When just one lagged level is added, however, it is significant (at the 5 percent level). Interestingly, the coefficient on this lagged level is positive, but the coefficient on the current level of inflation becomes more negative. It turns out that the sum of the (current and lagged) coefficients on log inflation is –0.610, quite similar to the base model's single coefficient of –0.639. This result suggests that it may be reasonable to think of the base model results as indicating the long-run relationship between inflation and growth.

Pre-averaging the data over time is a traditional way to reduce the potential influence of any spurious short-term correlations, although it does have its own difficulties. One of these is that the choice of data-averaging horizon is arbitrary.[19] Another problem is that time averaging may be inappropriate when the relationship is nonlinear. Indeed, in the case of the inflation-growth association—which seems in the annual data to be negative and convex for moderate and higher inflations, but concave within the low inflation range—the potential for a bias toward zero is clear. With these caveats in mind, we report in Table 3 the coefficient on inflation using the time-averaged data. For comparability, we first repeat the annual panel data regression over 1967–96, the data set's most recent 30 years (Table 3, regression (6)); we then run analogous regressions within this same period after averaging the data at 5-, 10-, and 15- year horizons, both with and without inflation values exceeding 40 percent (Table 3, regressions (7) and (8)).

Our main interest is the magnitude of the estimated coefficient on log inflation: in particular, does this drop off substantially as the data averaging period is extended beyond one year (i.e., annual data)? We see no clear drop off in moving from annual to 5-year data. Note also that the 15-year estimates lie between the 5- and 10-year estimates. Still, some of the estimates (especially those at the 10-year horizon) are considerably smaller than their corresponding estimates in annual panel data. As regards statistical significance, the results are mixed, as the estimated standard errors tend to be much larger with time-averaged data. As with the annual data, the results using 5-year averaged data are significant at the 1 percent level, in both samples. At the two longer horizons, however, only the results in the unrestricted samples are significant (at the 5 percent level).

[19] Choosing the longest available horizon for data averaging is one response to this problem, but this makes it impossible to control for country-specific "fixed effects."

Thus we find that researchers' arbitrary choice of data-averaging horizon can have nontrivial implications for inflation-growth regressions.[20] However, the finding of a significant inflation-growth relationship is *not limited to annual panel data,* and it does not appear that the results in annual panel data are driven mainly by spurious short-term correlations. (Further support for this conclusion can be found in Section V, where we discuss the simultaneity issue in depth and apply instrumental variables to the problem.)

In sum, the negative inflation-growth association is evident in both the cross-section and time-series dimensions of the data.

Other Augmentations, and an Extreme Bounds Test

Our base model includes a wide variety of conditioning variables; we now extend the robustness analysis by considering alternative specifications of such variables. Such analysis is important because some researchers (e.g., Levine and Zervos, 1993; and Sala-i-Martin, 1996) have found that inflation-growth correlations are not robust to changes in the conditioning variables, at least not when a linear relationship is imposed. We consider a number of augmentations to the base model, nonlinear specifications of the conditioning variables, and a Leamer extreme bounds test.

Fischer (1993) argues that the inclusion of changes in the terms of trade as a regressor goes a long way toward dealing with the problem of a spurious inflation-growth correlation. In our base model we do not include ΔTT because it is not significant (and considerably reduces the sample size). Nonetheless, Table 4, regression (1), reports the inflation coefficient when the change of the terms of trade is included in the regression. (Since many of the ΔTT data are missing, the sample size shrinks; to aid comparison, the base model estimated over this smaller sample is also reported.) We see that adding the change in the terms of trade (and its lag) barely changes the coefficient on inflation.

Nevertheless, it is possible that oil price shocks are responsible for the negative inflation-growth association (although the base model's annual dummies serve to control for common global shocks in any given year). Table 4, regression (2), shows the effect of adding the change in real oil prices (the average spot price of crude oil in dollars deflated by the U.S. wholesale price index) to the base model; the inflation coefficient is virtually unchanged. Since oil price changes affect countries differently, however, their effects would not be perfectly captured in this specification. To get around this complication, we also try restricting the regression sample to

[20] Note that the base model is at some disadvantage in the time-averaged regressions. As noted, averaging the data seems likely to bias the inflation coefficient toward zero. Also, the regression's kink is still imposed at $2^1/_2$ percent inflation, rather than letting the data determine a possibly better fit.

Table 4. *Estimates from Augmented Regressions and Other Regression Variants*

Regression variant	Coefficient on log inflation: Estimate	*t*-statistic	Sample size
(0) Base model, full sample	–0.00639	–6.04**	2,231
(1) Adding change in terms of trade			
current change	–0.00541	–3.66**	1,511
current and lagged change	–0.00534	–3.59**	1,511
(base model, in identical sample)	–0.00541	–3.66**	1,511
(2) Adding change in real oil prices			
current and lagged change[a]	–0.00742	–7.39**	2,231
(base model, w/o annual dummies)	–0.00743	–7.60**	2,231
(3) Excluding 1973–75, 1979–81, and 1990–92	–0.00589	–4.69**	1,570
(4) Adding institutions index (BERI)	–0.00628	–5.89**	2,167
(5) Adding inflation volatility	–0.00625	–5.76**	2,055
(6) Nonlinear conditioning variables			
adding squared terms	–0.00683	–6.43**	2,231
adding log terms	–0.00694	–6.42**	2,175
log terms replacing linear terms	–0.00666	–6.20**	2,175
(7) Extreme bounds test (4,096 regressions)			
weakest estimate	–0.00571	–6.25**	2,231
strongest estimate	–0.01020	–11.12**	2,231

Note: One asterisk and two asterisks indicate statistical significance at the 5- and 1-percent level, respectively.

[a] Because of colinearity between the oil price index and the annual dummies, it is necessary to exclude the latter in this case.

exclude all observations from 1973–75, 1979–81, and 1990–92. The resulting inflation coefficient remains statistically significant at the 1 percent level, with its magnitude only slightly diminished (Table 4, regression (3)).

We take account of political and economic institutions using an index created by the Business Environment Research Institute (BERI).[21] When this regressor is added to the base model, the inflation coefficient and its t-statistic are virtually unchanged (Table 4, regression (4)). Indeed, the BERI index does not enter the growth equation significantly.

Another possibility is to add a term for the *volatility* of inflation. By making it more difficult for economic agents to discern and respond to shifts in

[21] We use a composite index measuring (1) the degree of bureaucratic delays; (2) the enforceability of contracts; (3) the risk of nationalization or expropriation; and (4) the quality of communication and transportation infrastructure. The data were kindly provided by the IRIS Center, University of Maryland, with the permission of Ted Haner, President of BERI.

relative prices, inflation volatility might be expected to negatively affect growth. Since the average level of inflation and its standard deviation tend to be positively correlated,[22] it is possible that the base model's inflation term is picking up this channel. However, adding the volatility of inflation (measured as a 3-year moving standard deviation of log inflation) to the base model reduces the estimated coefficient on log inflation only slightly (Table 4, regression (5)). Still, it is interesting to note that the coefficient on the inflation volatility term is negative (–0.0048, with *t*-statistic of –2.03).

While we have allowed inflation to enter the growth regression in a nonlinear manner, we have not given the conditioning variables the same degree of attention, specifying all but one in a linear fashion.[23] If other nonlinear relationships were allowed, would the coefficient on inflation be much affected? To check, we add squared terms for each regressor (other than inflation and the dummies). Interestingly, 6 of these 10 new regressors turn out to be statistically significant, but their inclusion fails to diminish the inflation-growth relationship, with the coefficient on log inflation actually growing slightly (Table 4, regression (6)). We also try using the logarithms of the independent variables (for those variables that do not have numerous negative values), both adding these to the base regression and using them to replace the linear terms. Again, the effect is to slightly increase the absolute value of the inflation coefficient, which remains statistically significant by a wide margin.

Going beyond such augmentations of the base model, we also perform a Leamer extreme bounds test on the inflation term. An extreme bounds test determines whether the inflation term is *always* significant regardless of which combination from a (finite) set of conditioning variables is included as regressors. Thus we run all possible regressions based on the 12 conditioning variables in the base model (1). All regressions include the annual dummies, the log of inflation, and the low-inflation kink term. This gives 2^{12} possible combinations (ranging from no additional variables to all 12 variables). In contrast to others' results, we find that inflation does enter robustly: in over 4,000 regressions, the inflation coefficient is significantly negative in all cases.[24] Indeed, the coefficient estimates range from –0.0057 to –0.0102, with the associated *t*-statistics ranging from –6.25 to –11.12 (Table 4). Moreover, limiting the data set to observations with inflation below 40 percent does not alter this finding.

[22] Ball (1992) discusses why the average level of inflation and its volatility tend to be correlated. Judson and Orphanides (1996) use intrayear volatility of inflation and find a significant effect. A problem with using intrayear volatility is that much of it may be seasonal.

[23] An exception is the exchange rate premium, for which the logarithm is used.

[24] As Sala-i-Martin emphasizes, his test is less "extreme," requiring only that a weighted average of the *t*-statistics be significant.

What accounts for the difference between our results and those of others? The negative extreme bounds results of both Levine and Zervos and Sala-i-Martin are based on a strictly linear specification of the inflation-growth association. As we have seen, the linear model is misspecified and subject to a severe downward bias.

The Role, and Interpretation, of the Kink at 2¹/₂ Percent Inflation

The need to allow *some* kink in the low-inflation range was first emphasized by Sarel (1996). Our placement of this kink at $2\frac{1}{2}$ percent inflation is suggested first by visual inspection of the (bivariate) inflation-growth relationship (recall Figure 2); it also happens to be the placement that yields the best fit of the multivariate regression.[25] The basic results are not sensitive to this placement, however. For example, specifying the kink at possible arbitrary definitions of "low inflation," such as 5 or even 10 percent, yields similar results for the estimated slope to the right of the kink.[26]

We would not interpret the results of this study as indicating precisely $2\frac{1}{2}$ percent as an optimal or growth-maximizing rate of inflation. Rather, our interest is in whether a robust negative inflation-growth relationship is limited only to the high inflation range—say, above 40 percent—or whether it extends down much further, say to the single-digit range. Since all our findings point to the latter, it is natural to wonder *exactly* how far down the negative relationship extends, but we leave this more precise and therefore more difficult question to other researchers. For the record, in a likelihood ratio test, we cannot reject the alternative specification of a kink at 3 percent, but we can reject the alternative of a kink at 5 percent inflation.[27] Against this apparent precision, however, one should consider others' recent results, based on somewhat different samples and regression specifications: Sarel (1996) found that a kink at about 8 percent inflation gave the best fit, while IMF (forthcoming) found a best fit at about 5 percent inflation.

IV. Thresholds and Interactions: A Decision-Tree Technique

Just as there are threshold effects of inflation on growth, there may be threshold effects of the other determinants of GDP growth. For instance,

[25] That is, the R^2 of the multivariate regression has a maximum when the kink is at $2\frac{1}{2}$ percent (searching in $\frac{1}{2}$-point steps between $\frac{1}{2}$ and 20 percent inflation).

[26] The negative coefficient becomes somewhat larger in absolute value, and the associated *t*-statistic remains clearly significant at about –6.

[27] The χ^2 (1) statistic for the likelihood ratio test is 1.62 for the alternative of a kink at 3 percent, and 6.92 for the alternative of a kink at 5 percent.

even if most marginal increases in school enrollment rates have only small effects on growth, there may be some threshold level below which growth suffers greatly because of a lack of sufficient human capital. Moreover, the *interaction* between inflation and other growth determinants may be nonlinear and complex. For example, perhaps having "low" human capital essentially determines slow growth for some countries, almost regardless of their inflation rate, while countries with "high" human capital have a potential for either average or very high growth, with their inflation rate largely determining their position within this range.

In principle, a regression analysis could deal with such complications by including enough interactive dummy variables. Thus, there would be a dummy variable for low human capital *and* high inflation *and* severe terms of trade shocks, another dummy variable for high human capital *and* high inflation *and* severe terms of trade shocks, and so on. In practice, this is quite infeasible, since theory provides little guidance and the number of potential interaction specifications is vast. At best, a few arbitrarily chosen dummy variables could be included.

Fortunately, more systematic methods are available. Recently, Ghosh and Wolf (1998) have proposed the use of *binary recursive trees* as means to identify the most important determinants of economic growth.[28] This technique, while less familiar than standard regression analysis, is actually much simpler and perhaps more intuitive. A binary recursive tree begins from observations being classified as either "high growth" or "low growth."[29] After a researcher proposes a set of possible determinants of growth performance, a search algorithm creates a hierarchal decision "tree" by sequentially splitting the sample observations into (predicted) high and low growth groups, based on the values of the explanatory variables. Thus, at each branch of the decision tree, the algorithm finds the explanatory variable (and the associated threshold point of that variable) that best separates the high-growth observations from the low-growth observations.

For example, suppose that human capital is positively correlated with high growth. Of course the correlation will not be perfect, and there will be some countries that have plenty of human capital but low growth (a type I error), or that have little human capital but high growth (a type II error). The algorithm would search over all observed values of the human capital variable until it finds the threshold value at which the number of such errors is minimized.

[28] A nonlinear discriminator technique, recursive trees are often used in the medical sciences, for example, to analyze the determinants of patient mortality.

[29] For ease of interpretation, this type of analysis is usually done on binary variables. Here, we define "high"-growth observations as those in the top third of the data set, and "low"-growth observations as those in the bottom third. The middle third is excluded from the analysis.

The algorithm then repeats this process for each of the proposed determinants of growth. The variable (and its associated threshold) that minimizes the number of errors is chosen to form the first branch of the tree (with the sample now split into two). The process then continues, generating from each branch further subbranches until a terminal point is reached. To restrict the tree to a sensible (and interpretable) size, a stopping rule—somewhat like an adjusted R^2 statistic—eventually stops the tree from splitting into further subbranches.[30]

Such an exercise has several advantages over standard growth regression analysis. First, it allows for general complementarities between the different regressors—thus, the effects of inflation on growth, for instance, are allowed to vary according to the value of the other variables. Second, the branch level at which an explanatory variable appears provides an intuitive measure of its importance in determining growth. Third, and perhaps most important for our purposes, the results tend to be robust to outliers and are invariant to any monotone transformation of the variables.[31]

Figure 6, following Ghosh and Wolf (1998), illustrates such a decision tree applied to our data set, with all of the regressors in equation (1) offered to the search algorithm as potential explanatory variables. The first branch turns out to be based on the investment ratio: countries with investment ratios below 22 percent have only a 0.36 probability of high growth, whereas those with investment rates above 22 percent have a 0.65 probability of high growth.

For the countries with low investment, the second branching depends upon the level of human capital. Countries with low human capital have a 0.32 probability of high growth (conditional on being at that node, i.e., being a low-investment country) versus 0.51 probability for countries with high human capital. The third branch depends upon the inflation rate, with countries that have less than 15 percent inflation a year almost doubling their chances of high growth, from 0.37 to 0.65.

On the right-hand side of the tree in Figure 6 (countries with high investment), the second branching depends on the inflation rate: countries with inflation below 14 percent have a 0.75 (conditional) probability of high growth, while countries with higher inflation have only a 0.45 conditional probability of high growth.

[30] It would, of course, be possible to continue subdividing the tree until every observation is in its own branch. This would be akin to including as many variables as observations in a regression and thus attaining a perfect fit.

[31] The procedure's focus on classifying cases according to their position above or below threshold levels is similar to analysis focusing on variables' medians rather than means.

Figure 6. *Binary Recursive Tree of High- Versus Low-GDP Growth*

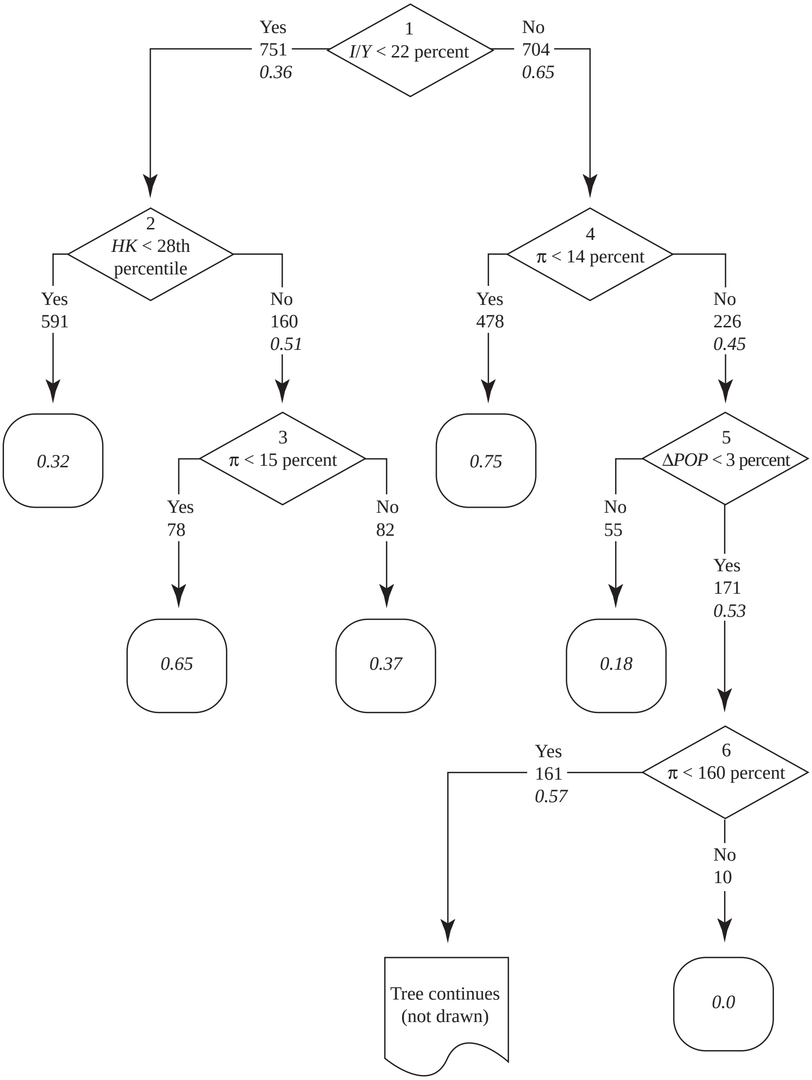

Note: All decision criteria in percent, unless otherwise noted. Figures in italics are the probability of high growth, conditional on being at the current node.

Source: Ghosh and Wolf (1997).

The tree continues with further subbranches (the tree reported in Ghosh and Wolf, for instance, has a total of 8 nodes).[32] For our purposes, however, it suffices to note that—of the various explanatory variables—only physical and (to a lesser extent) human capital are better able to discriminate between low- and high-growth countries than the country's inflation rate.

These trees draw a fairly complex picture of the interaction between the various determinants of output growth, thus highlighting the limitations of regression analysis. Of course, this technique has its own limitations, but it provides an interesting complement to the regression analysis discussed earlier. As before, the basic finding is that lower inflation is associated with faster growth. Moreover, we again see no sign that the negative effects of inflation only begin, or begin to pick up, after inflation has become rather high.

V. Simultaneity

The results reported above suggest no reason for skepticism about the existence of a robust negative inflation-growth relationship, but this correlation should be interpreted with some caution. A particular concern is that if growth somehow negatively influences inflation, then the inflation-growth findings presented above could, at least in part, reflect simultaneity bias. In the absence of a methodology to tell us whether inflation *causes* lower growth, here we pursue a more modest goal: to check that the negative inflation-growth correlation does not disappear once an effort is made to remove simultaneity bias using instrumental variables. In fact, several authors have already shown that this correlation survives this kind of test (see, e.g., Barro, 1995, and Cukierman and others, 1993).

In choosing instruments, it is helpful to consider how growth might negatively affect inflation. One potential channel can be seen by considering a simple money demand function, with real money demand as a function of real income. Taking logs and first differences:

$$\Delta m - \Delta p = \alpha \Delta y,$$

where α is income elasticity of money demand. If Δm is not immediately adjusted to growth shocks and if α is fixed, this money demand function would imply a negative correlation between inflation and growth. As Barro

[32] Based on the algorithm, the full tree chooses only *I*/*GDP*, *HK*, π, Δ*POP*, *G*/*Y*, and *GAP* as explanatory variables.

(1996) argues, however, the simultaneity bias arising from this channel is probably not very important, since for plausible values of α, shocks to output growth are too small to account for much of the observed variation in inflation rates.[33] In turn, if there is a large variation in inflation rates, then the component of the shock to inflation that is correlated with the shock to GDP growth must be small. Moreover, outside the short run, one might expect policymakers to adjust Δm in response to, and in the same direction as, changes in trend Δy; this also suggests that simultaneity bias might not be much of a concern. To the extent that Δm were adjusted *negatively* in response to short-term changes in real growth, however, simultaneity would become more of an issue.[34]

Note that the above discussion centers on within-country variation, primarily around short-term responses to output shocks; in this context, the potential for simultaneity bias seems clear. On the other hand, it is difficult to see why moving from one steady-state growth rate to another might itself lead policymakers to pursue a different steady-state inflation rate.[35] It is therefore not clear what instruments would adequately deal with such "long-term" channels of simultaneity.

Our method is two-stage least squares (2SLS): we first regress log inflation on a set of instruments, each entered as both linear and squared terms. We then use the fitted log inflation values in a growth regression (again, the base model used in Sections II and III).[36] As always, the validity of potential instruments is an issue. Thus, such variables as the ratio of the fiscal deficit to GDP, lagged money growth, and lagged inflation might be expected to be correlated with inflation, but their validity as instruments is suspect.[37] We use instead instruments in several other categories. The first is the nominal exchange rate regime; for example, Ghosh, Gulde, Ostry, and Wolf (1996) show that pegged exchange rate regimes are associated with lower inflation. Second, we consider three measures of legal central bank independence, as well as the central bank governor turnover rate (a proxy for independence); these are

[33] For example, α equal to 0.5 in Cagan's formulation and ranging up to the unit elastic case.

[34] This might be the case if tax revenues were countercyclical and policymakers used seigniorage (and the inflation tax) to complement conventional tax receipts.

[35] Possibly, considerations of optimal seigniorage would lead governments with small tax bases to compensate by choosing a higher "inflation tax." However, while it seems plausible that weak tax bases could be produced by a low *level* of per capita output, it is not clear why they would be correlated with low *growth.* (Recall that our growth regression controls for countries' initial level of output.)

[36] Thus, we do not instrument for the model's regressors other than inflation. As usual, we view these merely as conditioning variables; we are therefore not concerned with any simultaneity that might affect their associated coefficients.

[37] In principle, (sufficiently long-) lagged versions of these variables might be valid instruments.

Table 5. *Two-Stage Least Squares Results*

Instrument set[a]	Coefficient on log inflation: Estimate	Coefficient on log inflation: *t*-statistic	R^2 of first stage	Sample size
1. Exchange regime	–0.00622	–2.19*	0.17	2,130
2. Legal central bank independence	–0.00781	–2.35*	0.13	1,906
3. Central bank governor turnover	0.00621	1.63	0.20	1,955
4. 1 + 2 + 3	–0.00001	0.00	0.30	1,602
5. 1 + 2	–0.00704	–2.69**	0.18	1,816
Memorandum item: OLS regression	–0.00639	–6.04**	. . .	2,231

Note: One asterisk and two asterisks indicate statistical significance at the 5- and 1-percent level, respectively.
[a] All 2SLS variants include annual dummies in the first-stage regression.

reported by Cukierman (1992). Finally, we also use the base model's time dummies as instruments. Reflecting our interest in assessing robustness, we use a number of different combinations of these instruments.

The 2SLS findings, shown in Table 5, turn out to be sensitive to the choice of instruments. Thus, the results based on the exchange regime indicators, or on the measurements of legal central bank independence, or on both sets together, would suggest that essentially no part of the correlation between inflation and growth reflects a growth-to-inflation channel. That is, the magnitudes of the inflation coefficient estimates are very nearly as great as, or are greater than, the negative ordinary least squares (OLS) estimate. Moreover, although their standard errors are much larger than in the OLS case, these three estimates are statistically significant, at least at the 5 percent level. In contrast, when the central bank governor turnover rate is used as an instrument, the estimated coefficient on log inflation is positive, albeit not statistically significantly different from zero. Alternatively, using the turnover rate together with all the other instruments, the inflation coefficient estimate is negative but extremely small.

Thus, the 2SLS estimates paint a somewhat mixed picture. Results using several sets of instruments suggest that the strong OLS results in Sections II and III are not even slightly influenced by simultaneity bias. On the other hand, using the central bank governor turnover rate as an instrument upsets this result. However, a key shortcoming of this instrument is that it is available only as an average rate over 1950–89 (Cukierman, 1992); without any time variation, it is probably a poor instrument for a panel regression.

VI. Disinflation and Growth

If inflation is bad for growth, is *dis*inflation good? Not necessarily. In particular, the process of disinflation may lower GDP growth, at least in the short run. To the extent that high inflation is bad for growth, of course, any such negative effects of disinflation may be offset, at least partly, by the benefits of lower inflation. In this section, we focus on the simple contemporaneous association between growth and changes in inflation rate.[38]

Figure 7 provides a first pass at this issue. Plotted along the *x*-axis is the current inflation rate (π), along the *y*-axis the percent (not percentage point) change in the inflation rate since the last period ($\Delta\pi/\pi_{-1}$), and along the vertical axis is the GDP growth rate (Δy). Along the π dimension, the response surface is downward sloping: higher inflation is again seen to be associated with lower GDP growth. Along the ($\Delta\pi/\pi_{-1}$) dimension, things are more complicated. For low current inflation rates ($\pi < 10$), growth is decreasing with disinflation. At higher inflation rates, however, the surface flattens, until at 20 percent inflation growth is increasing with disinflation (the surface slopes downward for $\Delta\pi/\pi_{-1} > 0$).

Table 6 reports the results of a simple regression intended to capture the impact effect of severe disinflation. The base model is augmented to include four dummy variables: ($\pi_{-1} < 0.10$, and $\Delta\pi/\pi_{-1} < -0.5$), ($\pi_{-1} < 0.10$, and $-0.5 < \Delta\pi/\pi_{-1} < -0.2$), ($\pi_{-1} > 0.10$, and $\Delta\pi/\pi_{-1} < -0.5$) and ($\pi_{-1} > 0.10$, and $-0.5 < \Delta\pi/\pi_{-1} < -0.2$), in addition to the usual explanatory variables. Results including and excluding *I*/*GDP* are given.

The results suggest that, when the initial level of inflation is above 10 percent a year, even severe disinflations (at least halving the inflation rate) do not have a negative impact on output growth. More moderate disinflations, indeed, are associated with 0.8–0.9 percentage points *higher* GDP growth (*t*-statistic: 2.62 and 2.41).

On the other hand, when the initial level of inflation is below 10 percent a year, severe disinflations are associated with a fall in GDP growth of about 1 percentage point (with the effect statistically significant at the 5 or 1 percent level). More moderate disinflations are also associated with lower GDP growth, by about 0.5 percentage points, except for the upper-income and upper-middle-income countries, where growth picks up with moderate disinflation.

[38] Thus we do not attempt to disentangle short-run and long-run effects of moving from one inflation rate to another, nor do we consider whether possible contractionary effects of disinflation on the level of output might be permanent or temporary.

Figure 7. *Per Capita GDP Growth Versus Inflation and Disinflation*
(In percent a year)

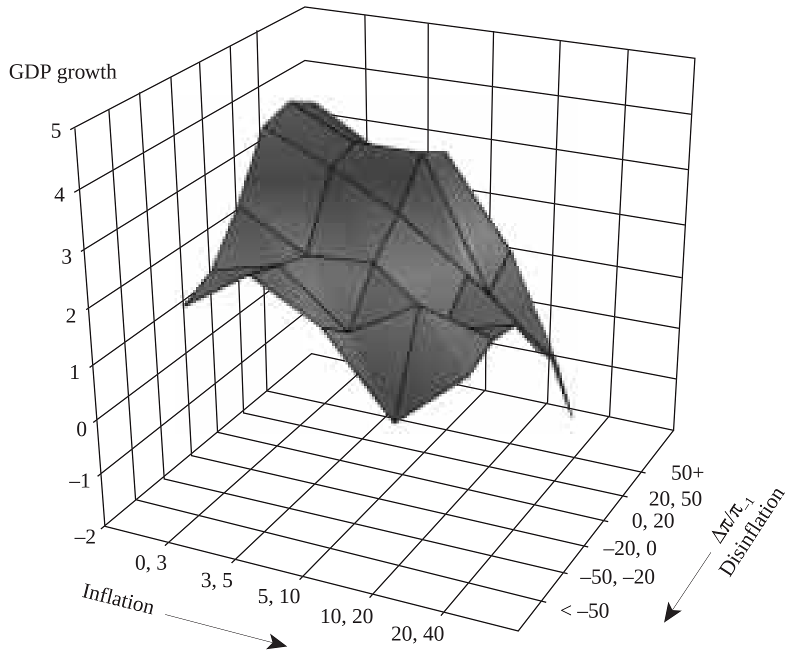

Of course, the 10 percent inflation cutoff, and the definition of "severe" and "moderate" disinflations are chosen arbitrarily on the basis of Figure 7. A more methodical approach is to maximize the likelihood function, where the two dimensions are the initial level of inflation and the degree of disinflation:

$$\log(L) = \sum_{i=1}^{k_i}\sum_{j=1}^{k_j} z_{ij} - \frac{1}{2}\left[\log(2\pi) + \log(\sigma_{ij}^2) + \epsilon_{ij}^2\right] \quad (2)$$

$$z_{ij} = 1 \text{ if } \begin{array}{c} \underline{\pi}_j < \pi < \overline{\pi}_j \\ \underline{\Delta\pi/\pi}_i < \Delta\pi/\pi < \overline{\Delta\pi/\pi}_i \end{array}.$$

In these regressions, we control for current inflation, but we exclude the investment ratio because some of the adverse growth effects associated with

Table 6. *Effects of Disinflation on GDP Growth*

	$\pi_{-1} < 0.10$				$\pi_{-1} > 0.10$				
	$\Delta\pi/\pi < -0.5$		$-0.5 < \Delta\pi/\pi < -0.2$		$\Delta\pi/\pi < -0.5$		$-0.5 < \Delta\pi/\pi < -0.2$		
	Coefficient	*t*-statistic	Coefficient	*t*-statistic	Coefficient	*t*-statistic	Coefficient	*t*-statistic	R^2
Full Sample									
Including *I/GDP*	–0.0133	–2.67**	–0.0086	–3.05**	0.0051	1.04	0.0090	2.62**	0.24
Excluding *I/GDP*	–0.0168	–3.32**	–0.0097	–3.34**	0.0043	0.87	0.0083	2.41*	0.22
Upper-Income and Upper-Middle-Income Countries									
Including *I/GDP*	–0.0058	–2.20*	0.0229	2.45*	0.0122	2.45*	–0.0003	–0.09	0.43
Excluding *I/GDP*	–0.0070	–2.57*	0.0219	2.30*	0.0105	2.01*	0.0029	0.81	0.40
Lower- and Lower-Middle-Income Countries									
Including *I/GDP*	–0.0209	–2.98**	–0.0123	–2.08*	–0.0018	–0.32	0.0067	1.51	0.21
Excluding *I/GDP*	–0.0240	–3.41**	–0.0136	–2.30*	–0.0024	–0.42	0.0061	1.38	0.20
Pre-1973 Observations									
Including *I/GDP*	–0.0107	–1.08	–0.0159	–2.09*	–0.0092	–0.61	0.0115	0.98	0.23
Excluding *I/GDP*	–0.0124	–1.26	–0.0166	–2.10*	–0.0072	–0.49	0.0110	0.95	0.21
Post-1973 Observations									
Including *I/GDP*	–0.0174	–2.86**	–0.0078	–2.62**	0.0062	1.17	0.0087	2.44*	0.24
Excluding *I/GDP*	–0.0211	–3.38**	–0.0087	–2.89**	0.0053	0.98	0.0079	2.22*	0.22
Pegged Rates									
Including *I/GDP*	–0.0119	–1.97*	–0.0113	–3.12**	0.0065	0.95	0.0100	1.97*	0.22
Excluding *I/GDP*	–0.0152	–2.50*	–0.0124	–3.38**	0.0061	0.89	0.0094	1.86	0.20
Floating Rates									
Including *I/GDP*	–0.0156	–1.40	–0.0003	–0.07	0.0028	0.43	0.0051	1.11	0.35
Excluding *I/GDP*	–0.0178	–1.61	–0.0001	–0.02	0.0009	0.13	0.0039	0.81	0.33

Note: One asterisk and two asterisks indicate statistical significance at the 5- and 1-percent level, respectively.

disinflation may arise from contractionary effects on investment.[39] The corresponding estimates are as follows:

For $\pi_{-1} < 0.063$

$$\Delta y = \begin{cases} \underset{(2.25)}{0.023} + \underset{(2.16)}{0.014}\log(\pi/\pi_{-1}) & \text{if } \pi/\pi_{-1} < -0.48 \\ \underset{(16.89)}{0.021} + \underset{(4.02)}{0.017}\log(\pi/\pi_{-1}) & \text{if } -0.48 < \pi/\pi_{-1} < 0.70 \\ \underset{(5.08)}{0.027} - \underset{(1.58)}{0.005}\log(\pi/\pi_{-1}) & \text{if } \pi/\pi_{-1} > 0.70. \end{cases}$$

For $\pi_{-1} > 0.063$

$$\Delta y = \begin{cases} \underset{(5.52)}{0.046} + \underset{(3.93)}{0.015}\log(\pi/\pi_{-1}) & \text{if } \pi/\pi_{-1} < -0.63 \\ \underset{(17.0)}{0.019} - \underset{(3.65)}{0.011}\log(\pi/\pi_{-1}) & \text{if } -0.63 < \pi/\pi_{-1} < 1.28 \\ \underset{(0.49)}{0.029} - \underset{(0.54)}{0.031}\log(\pi/\pi_{-1}) & \text{if } \pi/\pi_{-1} > 1.28. \end{cases}$$

Thus, the procedure segments the data according to whether the initial level of inflation is above or below about 6 percent (in Table 6, the cutoff was chosen at an inflation rate of 10 percent a year). When the initial inflation rate is above 6.3 percent a year, only the most severe disinflations—cutting the inflation rate by more than 63 percent (not percentage points)—is associated with lower growth.[40] Except for these severe disinflations, however, an increase in the rate of inflation is associated with lower GDP growth.

When the initial inflation rate is below 6 percent a year, severe disinflations are again associated with lower GDP growth—as are increases in inflation by more than 70 percent (not percentage points).

It bears emphasizing that these effects are conditional on the current inflation rate (the current inflation rate is included among the regressors).

[39] Controlling for investment actually makes little difference to the results. Starting from low inflation ($\pi_{-1} < 0.063$), increases in inflation are associated with higher growth (and disinflations with lower growth). If the inflation rate rises by more than 70 percent, however, there is again a negative impact on growth. When the starting inflation rate is above 6 percent, however, only disinflations of more than 70 percent ($\Delta\pi/\pi_{-1} < -0.70$) are associated with lower growth. Again, these regressions are conditional on the current inflation rate.

[40] Notice that the coefficient on $\log(\Delta\pi/\pi)$ is positive here, meaning that—over this range—an *increase* in the inflation rate would raise growth.

Thus, even the effects of severe disinflation will be partly offset by the (positive) effects of lower current inflation.

While one should not take these results too literally, especially without an examination of their robustness, they are at least consistent with the idea that, starting from even moderately high inflation rates, all but the most severe disinflations are beneficial for growth, even in the short run.[41] When the starting inflation rate is already low, however, greater caution may be required. In all likelihood, it is not the fact of disinflation itself that matters for short-run growth—rather, that rapid disinflation will generally be associated with tightening monetary conditions. Thus the disinflation variable can be replaced by, say, the change in real money or real credit growth, with broadly similar results.[42]

VII. Conclusions

There are several reasons why governments might want to achieve low inflation, perhaps the most compelling being the potential for faster output growth. Indeed, of the various factors that might affect growth, perhaps none is as readily changed in the short run as the inflation rate. Few would doubt the negative growth effects of high inflation—say, above 40 percent a year—but there has been much less consensus on the effect of less severe inflation. Yet from a policy perspective it is the moderate or intermediate inflation range—perhaps 5 to 30 percent a year—that is of greatest interest.

The results presented here suggest a negative relationship between inflation and growth that is both statistically and economically significant. The relationship is nonlinear, in two senses: first, at very low inflation rates, the relationship is positive; second, at all other inflation rates, the apparent *marginal* effect of inflation on growth becomes less important as higher inflation rates are considered. Failure to take account of both these nonlinearities can seriously bias results toward finding only a small effect, giving the misleading impression that inflation must become quite high before its *cumulative* effect becomes important.

We cannot of course claim to have shown that inflation *causes* lower growth; indeed, it is difficult to conceive of any methodology that would

[41] Focusing on cases with substantially higher initial inflation, Bruno and Easterly (1995) find that growth resumes almost immediately after disinflation.

[42] Starting from inflation below 8 percent, contractions of the real money supply of greater than 7 percent are associated with lower growth—although the effect is not statistically significant (t-statistic: 1.01); when the starting level of inflation is above 8 percent, real contractions of the money supply of greater than 12 percent are associated with lower GDP growth, but again the effect is not statistically significant.

decisively prove causality from inflation to growth. Rather, this study's more modest contribution is its failure, despite a battery of tests, to find any evidence that casts doubt on the idea that inflation (or the policy choices it reflects) reduces growth. Of course, inflation is not under direct policy control; especially in the short run, it is an outcome of both macroeconomic policy choices and exogenous shocks. Inflation is therefore probably best thought of as an *indicator* of those policy choices. Still, we find no sign that the inflation-growth association found in annual panel data is spurious, arising only from short-run correlations induced by shocks. Moreover, while we have not sought to identify the particular mechanisms or channels through which inflation (or its associated policy choices) might hinder growth, it is interesting that a statistically and economically significant inflation-growth association is found even controlling for such likely policy correlates as government consumption, fiscal deficits, and black market exchange rate premiums.

Finally, it bears emphasizing that this study does not claim to precisely locate a "growth-maximizing" rate of inflation (any such rate might be expected to differ, at least somewhat, across countries). Rather, our focus is on the more basic question of whether the negative inflation-growth relationship occurs only at very high inflation rates, or whether it extends down much further, perhaps to the single-digit range. All our findings suggest the latter. *Exactly* how far this negative relationship extends, however, remains an open and difficult question—and one worthy of future research.

REFERENCES

Ball, Laurence, 1992, "Why Does High Inflation Raise Inflation Uncertainty?" *Journal of Monetary Economics*, Vol. 29 (June), pp. 371–88.

Barro, Robert J., 1995, "Inflation and Economic Growth," *Bank of England Quarterly Bulletin*, Vol. 35 (May), pp. 166–76.

———, 1996, "Determinants of Economic Growth: A Cross-Country Empirical Study," NBER Working Paper No. 5698 (Cambridge, Massachusetts: National Bureau of Economic Research).

Bruno, Michael, and William Easterly, 1995, "Inflation Crises and Long-Run Growth," NBER Working Paper No. 5209 (Cambridge, Massachusetts: National Bureau of Economic Research).

Clark, Todd E., 1993, "Cross-Country Evidence on Long Run Growth and Inflation," Federal Reserve Bank of Kansas City, Research Working Paper No. 93-05; also published in *Economic Inquiry*, Vol. 35 (January), pp. 70–81.

Cukierman, Alex, 1992, *Central Bank Strategy, Credibility, and Independence: Theory and Evidence* (Cambridge, Massachusetts: MIT Press).

———, and others, 1993, “Central Bank Independence, Growth, Investment, and Real Rates,” *Carnegie-Rochester Conference Series on Public Policy*, Vol. 39 (December), pp. 95–140.

Easterly, William, 1996, “When Is Stabilization Expansionary? Evidence from High Inflation,” *Economic Policy* (April), pp. 67–107.

Fischer, Stanley, 1993, “The Role of Macroeconomic Factors in Growth,” *Journal of Monetary Economics*, Vol. 32 (December), pp. 485–512.

Ghosh, Atish, Anne-Marie Gulde, Jonathan Ostry, and Holger Wolf, 1996, “Does the Exchange Rate Regime Matter?” NBER Working Paper No. 5874 (Cambridge, Massachusetts: National Bureau of Economic Research).

Ghosh, Atish, and Holger Wolf, 1998, “Thresholds and Context Dependence in Growth,” NBER Working Paper No. 6480 (Cambridge, Massachusetts: National Bureau of Economic Research).

International Monetary Fund, forthcoming, *Economic Adjustment and Reform in Low-Income Countries: Studies by the Staff of the IMF*, ed. by Hugh Bredenkamp and Susan Schadler (Washington: IMF).

Judson, Ruth, and Athanasios Orphanides, 1996, “Inflation, Volatility and Growth,” Finance and Economics Discussion Paper No. 96-19 (Washington: Board of Governors of the Federal Reserve System).

Levine, Ross, and Sara Zervos, 1993, “Looking at the Facts: What We Know About Policy and Growth from Cross-Country Analysis,” World Bank Policy Research Working Paper No. 1115 (Washington: World Bank).

Phillips, A., 1958, “The Relation Between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861–1957,” *Economica*, Vol. 25 (November), pp. 283–99.

Sala-i-Martin, Xavier, 1997, “I Just Ran Two Million Regressions,” *American Economic Review, Papers and Proceedings*, Vol. 87 (May), pp. 178–83.

Sarel, Michael, 1996, “Nonlinear Effects of Inflation on Economic Growth,” *Staff Papers*, International Monetary Fund, Vol. 43 (March), pp. 199–215.