A Framework for Examining the Real Effects of Inflation Volatility

*Christopher Ragan**

Introduction

The Bank of Canada has embarked on an explicit policy of reducing the rate of inflation. Its ultimate goal may even be to reduce inflation to zero. Since Governor Crow's comments in Edmonton in January of 1988 (Crow 1988), there has been active debate in Canada about the benefits and costs associated with achieving this goal. In an address in Moncton in October 1990, Governor Crow said:

> Price stability is not an end in itself. What counts is what price stability can deliver. We have a market economy, and one of its pillars is the institution of money and monetary exchange....The prices of goods and services are set in money terms. We keep economic score and we frame our day-to-day and longer-term decisions in money terms. Damaging that institution, as inflation must, introduces needless uncertainty and thereby makes our economy less efficient and less productive. It also injects inequity, hurting some groups of people to the benefit of others, and for no socially worthwhile purpose (Crow 1990, 3-4).

* *I thank Thor Sigvaldason for research assistance, and Seamus Hogan and Thor Sigvaldason for helpful comments.*

A purist might take exception to a literal interpretation of Crow's comments which would suggest that even a perfectly anticipated and steady rate of inflation would generate significant real effects; most economists would take the view that such a "pure" inflation would, at least in the long run, be approximately neutral. But Governor Crow's statement reflects a belief held at the Bank of Canada (see Freedman 1991) and elsewhere, that such a steady and perfectly anticipated rate of inflation exists only in economic models, and is certainly not something that is possible in the real world.

Crow's statement also reflects a belief that the level of inflation and the uncertainty of inflation, though conceptually distinct, are positively correlated; and, further, that many of the benefits from a reduction in inflation derive from the accompanying reduction in the uncertainty of inflation. The mechanism that has been emphasized by the Bank of Canada (Freedman 1991) is the implied distortion in relative price signals. If inflation (and inflation uncertainty) make it more difficult for firms and individuals to identify changes in relative prices, then many decisions that are made will turn out to be mistakes. Such mistakes represent a misallocation of resources that would be avoided if the inflation – and thus the uncertainty attached to it – were to be eliminated.[1]

1. The Basic Approach

This paper provides a simple theoretical framework for considering the real effects of inflation volatility. In particular, the question that forms the motivation for my paper is:

> ***How does the long-run behaviour of the real economy in an environment of high and unstable inflation differ from that in an environment of low and stable inflation?***

I want to make a number of comments about my interpretation of this question. First, I take very seriously the presence of the term "long-run." This, fortunately, allows me to ignore completely the important and interesting – but

1. This is not a new idea, of course; it is precisely the mechanism at work in the well-known model of monetary non-neutrality by Lucas (1972).

complicated – issues pertaining to the short-run transitional costs of reducing inflation. I focus purely on what I call the medium-run and long-run real effects of inflation volatility. Any costs or benefits I identify with *being in* a low and stable inflation environment must therefore be added to those costs or benefits associated with *getting to* that environment.

Second, my focus on the medium and long runs leads me to concentrate on a model economy in which all prices are perfectly flexible and all markets clear. This means that a pure change in inflation such as the one discussed above would have (approximately) no real effects.[2] In other words, I take the view that an economy in which inflation is 10 per cent with a standard deviation of x per cent behaves identically to an economy in which inflation is 5 per cent with a standard deviation of x per cent. This is particularly convenient since the comparison of a low and stable inflation environment with a high and unstable inflation environment then reduces to an analysis of stability only. Following this logic, the thought-experiment I conduct later, in the more theoretical sections of this paper, is one of a mean-preserving spread in the distribution of inflation.

Third, despite my rigid adherence to what appears to be perfect markets, much of my analysis is based on a particular incompleteness of markets; given the paper's focus on the effects of uncertainty, some such incompleteness is necessary to make the story interesting. My key assumption is that indexed loan contracts are not available in the credit market; all loan contracts therefore specify a given nominal interest rate. We might think of indexed contracts as being prohibitively expensive or as simply being unavailable. In any event, the motivation for this assumption should be obvious; with very few exceptions, indexed contracts in credit markets do not appear to exist, and this non-existence introduces several avenues along which inflation uncertainty – even in the presence of otherwise perfect markets – can have real effects.

2. The only real effects would be on the level of real money demand; such "shoe-leather" costs of inflation, in countries with low or moderate inflation, are surely quite small. In any event, I do not discuss the costs associated with perfectly anticipated inflation (see, for example, Driffill et al. 1990; Fischer 1986b; Foster 1972; and Laidler 1990).

Finally, I wish to emphasize that none of the real effects associated with inflation uncertainty that I discuss in this paper relies on a correlation between inflation uncertainty and relative price uncertainty (except, of course, the *ex post* real intertemporal relative price). Indeed, in some cases I make the explicit assumption that agents know exactly the future path of relative prices, so that uncertainty about the aggregate price level is the only uncertainty at play. Uncertainty about relative prices, generated by uncertainty about inflation, could be added to the analysis, but it is not central to the main arguments.

1.1 Plan of the paper

The paper is organized as follows. Section 2 presents some evidence on the relationship between the level of inflation and the volatility of inflation. My view is that before I can undertake the construction of an answer to the question posed above, the relationship that the question seems to take as given must be shown to exist.

Sections 3 and 4 are the main part of the paper. Section 3 examines the behaviour of the real economy and how that behaviour is affected by a mean-preserving spread in the distribution of inflation. I divide the discussion into the medium run, for which I discuss the determination of the real interest rate and the level of aggregate output, and the long run, for which I discuss the steady-state level of the aggregate capital stock. I also present evidence from G-7 countries that provides some support for my arguments. Section 4 adds financial intermediation to the real economy and examines the relationship between inflation uncertainty and the cost of financial intermediation. My arguments in this section are based largely on the ability of individual borrowers to declare bankruptcy, and on how inflation uncertainty increases the probability of bankruptcy. Using data from Canadian chartered banks for the past 20 years, I show some support for these arguments.

Section 5 contains my summary comments. I discuss the relevance to policymakers of my main arguments. I also discuss several factors that I think are important to the analysis of inflation uncertainty but which I have chosen to ignore here. Tables and figures are given in an Appendix.

2. Does Lower Inflation Mean Less Inflation Volatility?

In this section, I examine the alleged relationship between the average level of inflation and its volatility. In what follows, I provide no behavioural story about why such a correlation might be expected to exist; I simply look at some data – both cross-section and time-series – to see if there is any evidence of such a relationship.

One issue must be addressed at the outset: there is an important difference between the *volatility* of inflation and the *uncertainty* of inflation. I deal with this distinction by examining both cross-section data and time-series data. With the cross-section data, I employ simple measures of volatility based on sample variances. With the time-series data, I attempt to construct simple measures of inflation uncertainty.

2.1 Cross-section measures of inflation volatility

We begin by repeating the exercise conducted by Okun (1971), Logue and Willett (1976) and Foster (1978). Observations on the seasonally unadjusted consumer price index (CPI, all items) for 22 OECD countries are used. The data is monthly and the sample is 1960M1 to 1989M11.[3] The rate of inflation of the CPI for each period is computed as the rate of change over the previous 12 months, so there should be no seasonality in the measured inflation series. For each country, I then compute the sample mean $(\overline{\Pi})$ and the sample standard deviation (SD).

Even if we accept the proposition that a measure of volatility based on the sample variance is indicative of the extent of uncertainty in inflation, there is still the issue of whether the standard deviation or the coefficient of variation is the appropriate measure. To the extent that individuals and firms care about inflation uncertainty because of the costs that it imposes in terms of prediction errors about real rates of return – and I think this is precisely the reason – then the standard deviation of inflation is the appropriate measure: actual inflation of 8 per cent

3. The CPI data are quarterly for Australia, Ireland and New Zealand. Iceland, Turkey and Yugoslavia are omitted from the analysis because they are extreme outliers in the inflation process, and their presence in the sample virtually guarantees a positive correlation between inflation and inflation volatility.

when we are expecting 6 per cent is just as costly as actual inflation of 3 per cent when we are expecting 1 per cent.

Figure 1 (see Appendix) plots *SD* against $\overline{\Pi}$ for the 22 OECD countries. There is a clear positive relationship; the associated regression (with standard errors in parentheses) is shown below the figure. Figure 2 plots a different measure of volatility; the average absolute change in the inflation rate $(AA\Delta)$. Following Foster (1978), I compute for each country the annual inflation rate for each year (my measured rate from December) and then take the absolute value of the change in annual inflation rates. $AA\Delta$ is then the average of these absolute changes. Foster argues that $AA\Delta$ is a good measure of the uncertainty of inflation only in the case of static expectations, in which case the absolute change in the inflation rate from one year to the next is exactly the forecast error. Figure 2 shows a strong positive correlation between this measure of inflation volatility and average inflation. The associated regression is shown below the figure.

2.2 Time-varying measures of inflation uncertainty

The obvious problem with using international cross-section evidence is that such correlations as those in Figures 1 and 2 tell us very little about whether an individual country can expect less inflation volatility as a result of lowering its average inflation rate. One story that is consistent with the above correlations is that all countries are equally able (maybe perfectly) to control their money supply process, but that countries with high inflation just happen to be countries that are politically much less stable than countries with low inflation. This political instability – with its associated uncertainty about future economic policy – might be manifested by, among other things, extreme volatility in money demand. In this setting, a country that lowers its mean rate of inflation is still subject to the same volatility in money-demand shocks and thus may have unchanged inflation volatility.

One would not want to push this particular story too far, but the point should be clear: while the cross-section evidence is suggestive, it is not informative enough about the inflation process (within a country) to form the basis

of an engineered policy to reduce inflation. There may be many excellent reasons to reduce inflation, but the evidence in Figures 1 and 2 is not among them.

Ideally, we would like to know if there is any causal relationship between a country's average rate of inflation and the uncertainty of its inflation. Ball (1990a) offers one reasonable theory, based on the existence of a short-run Phillips curve trade-off.[4] When inflation is low, any costs of inflation are also low and so policymakers are inclined to continue doing what they have been doing. There is no obvious policy dilemma; inflation is low and is expected to be low. But when inflation is high, policymakers are faced with a clear dilemma. They could reduce inflation and suffer the obvious short-run transitional costs or they could avoid these costs but suffer the costs associated with continued high inflation. We can think of policymakers randomizing (between high and low inflation) when inflation is high, whereas they have a dominant policy of low inflation when inflation is low; this implies that inflation is more volatile when average inflation is high.

Whatever the story, it is clear that an improvement over the simple cross-section evidence above involves looking at time-series behaviour within a particular country (or within many countries). Ideally, we would like a time-varying measure of inflation uncertainty rather than simply a measure of inflation volatility. Given such a measure, we could then examine – over time but within a particular country – the correlation between that measure and the rate of inflation. Though this would still leave us without a structural story about the source of any relationship between inflation and uncertainty, it would allow us (implicitly) to control for all of the institutional details unique to a particular country.[5]

---

4. See Devereux (1989) and Ball, Mankiw and Romer (1988) for two others.
5. A simple alternative to the time-varying measure is to compute measures of volatility across different time periods. To do this, the OECD data for each country are divided, rather arbitrarily, into three inflation regimes: 1960-1973; 1974-1982; 1983-1989. I then take the first differences of SD and $\overline{\Pi}$ across regimes and run the simple regression of ΔSD_i against $\Delta\overline{\Pi}_i$. The results are

$$\Delta SD = \underset{(0.210)}{-0.180} + \underset{(0.030)}{0.124} \cdot \Delta\overline{\Pi}\text{; with } N = 44 \text{ and } \bar{R}^2 = 0.27$$

which gives us some reason to believe that a country that lowers its average inflation rate will also experience a reduction in inflation volatility.

One such time-varying measure of inflation uncertainty uses the ARCH model (Engle 1983), or any of its many variants. Though it is not inherent to the estimation procedure, one problem with the typical (and practical) use of this approach is that the measure of inflation uncertainty at time t requires information from the entire sample, including information that would not yet be available to market participants at time t. Another drawback with this approach is that it typically assumes that there is a single structure to the inflation process over the entire sample.[6]

The time-varying measure of inflation uncertainty used here is the root-mean-square forecast error from a rolling autoregressive inflation prediction equation. Specifically, for each country, I estimate in each period a simple AR(11) equation in the (12-month) inflation rate.[7] The structure of the estimated inflation process is allowed to change by keeping the estimation sample fixed at 60 (monthly) observations; as a new observation is added to the sample, the first observation is removed. In each period, using the estimated coefficients, a forecast is made of the one-year-ahead inflation rate. This gives a series of one-year-ahead inflation forecasts, each one based only on information available at the time the forecast was made.

The one-year-ahead forecast error for period t is denoted by e_t. The measure of inflation uncertainty at time t is then

$$V_t = \left[(1/12) \cdot \sum_{j=t-12}^{t-1} e_j^2 \right]^{1/2},$$

6. There is no reason why the estimation could not take place recursively, estimating the model at each time t, using only information available at that time. But since the model is estimated using maximum likelihood, which often requires considerable time to reach convergence, this estimation method may not be sensible for a sample of 22 countries and 30 years.
7. For those countries with quarterly data, I estimate an AR(3).

so that the uncertainty of inflation at time t is based on the 12 previous forecast errors, all equally weighted.[8]

Figure 3 (see Appendix) shows the behaviour V_t and the 12-month inflation rate (Π_t) for Canada from 1967 to 1989. Table 1 (see Appendix, p. 457) shows the results of the simple OLS regression of V_t against Π_t for three groups of countries. For the entire group of 22 countries as well as for the G-7 countries taken as a group, there is a strong positive relationship between the rate of inflation and the measure of inflation uncertainty. The relationship in Canada is also strongly significant and positive.

Note that the estimated coefficient on Π declines as the sample of countries gets smaller. This suggests that a considerable part of the correlation in the full-country sample is coming from the "between" effect rather than the "within" effect and, as argued above, it is really the within-country correlation that is relevant to the policy debate. When the same regression is estimated for each country separately, however, we still find a statistically significant (and positive) relationship in all but three of the countries (Austria, Luxembourg and Norway are the exceptions).

I conclude that there is considerable empirical evidence of a positive correlation between the rate of inflation and a sensible measure of the uncertainty of inflation. I emphasize that this is only an empirical regularity (albeit one with considerable robustness across countries); I have said nothing about the behaviour that must underlie such a relationship. And until we have solid theoretical reasons to think that such a relationship should exist, a sceptic would be justified in believing that my evidence simply shows that inflation and inflation uncertainty are both strongly correlated to some unobserved third variable.

On this particular relationship, however, I am not that much of a sceptic. The basic story that I outlined above, from Ball (1990a), seems plausible. That story generates a correlation between inflation and inflation uncertainty because high inflation, unlike low inflation, leads to an obvious policy dilemma and thus to

---

8. A slightly different measure of uncertainty weights the forecast error from period j with the weight $(1/2^{j+1})$. This measure performs very similarly to V_t in the regressions in Table 1.

uncertainty about future inflation. Note that the uncertainty about inflation here has its source in the uncertainty about the money supply process. Another reasonable story is based on the existence of small menu costs, and is found in Ball, Mankiw and Romer (1988). The basic idea can be illustrated by a world where the only uncertainty in the money market is on the demand side. With small menu costs, low inflation means that price-setters respond only to the largest of money-demand shocks. With high inflation, however, price-setters respond even to small money-demand shocks. The result is that low inflation is more stable than high inflation.

These two stories provide very different explanations for the predicted correlation between inflation and inflation uncertainty. But both, importantly, give us some solid reasons to believe that a government that undertakes a policy of lowering its average rate of inflation can also expect to observe a reduction in the uncertainty of inflation. With stories like these in mind, combined with supporting evidence such as that in Table 1, I conclude that the central question noted above is indeed a well-posed question. This permits me to go on to an analysis of the real effects of inflation volatility.

3. The Real Economy in the Medium and Long Run

3.1 The basic approach to interest rate determination

An important distinction is made between the short run, the medium run and the long run. For the purposes of this paper, the short run is defined to be the span of time over which prices are not completely flexible. As stated in the Introduction, the approach is to ignore completely the issue of the transition of the economy from high and unstable inflation to low and stable inflation; this amounts to ignoring the short run. I take the medium run to be the span of time over which all prices are fully flexible but the aggregate capital stock can be viewed as more or less fixed. The long run has both fully flexible prices and flexible capital stock; think of this as the length of time it takes the economy to reach its steady state.

In the discussion of both the medium and long run, the economy is made up of many individuals, each of whom has access to perfect capital markets. I assume that loan contracts in the credit market specify a nominal interest rate; indexed bonds are assumed not to exist. Furthermore, borrowers and lenders are

assumed to transact directly; there are no financial intermediaries. I examine firms' investment decisions and individuals' saving and labour-supply decisions when the only uncertainty is over the rate of inflation. I then analyze how a mean-preserving spread (hereafter, MPS) in the distribution of inflation affects behaviour. This tells me about the medium-term equilibrium (output and real interest rates) as well as the long-run steady state of the economy.

3.2 The medium run: investment, saving and labour supply

3.2.1 Firms' investment behaviour

The basic question to be answered in this subsection is: What is the effect on a firm's investment demand of a (mean-preserving) change in the amount of inflation uncertainty? Surprisingly, I have found no evidence that this question has been addressed in the investment literature. This does not mean that the issue of uncertainty has not been addressed, for it has; only that the emphasis on inflation uncertainty has been absent.

3.2.2 Real uncertainty

Hartman (1972), Pindyck (1982) and Abel (1983) have discussed the effects on investment demand of uncertainty in the firm's real product price (p), real wage (w) and real price of capital (q). Hartman's basic result is that an MPS in the distribution of future values of p or w will *increase* the firm's current investment demand. This follows from Jensen's Inequality and the convexity of the firm's profit function.

To see the result, consider a firm that combines capital (k) and labour (n) to produce output (y), according to the production function f:

$$y_t = f(k_{t-1}, n_t) .$$

The firm's capital accumulation is given by

$$k_t = (1-\delta)\, k_{t-1} + i_t ,$$

where δ is capital's rate of physical depreciation and i is gross investment. Now suppose that the firm observes all prices for any period before making any decisions for that period. The firm's optimal choice of labour input is

$$n_t^* = n^*(w_t/p_t;\ k_{t-1}).$$

Then define the firm's period[9] profit function as

$$\Gamma(w_t, p_t;\ k_{t-1}) \equiv p_t f(k_{t-1}, n_t^*) - w_t n_t^*$$

which is convex in p (and w).[10]

Now, make three simplifying assumptions. First, the firm is competitive in both the product and labour markets and thus takes p and w as given. Second, the real price of capital is increasing in k_t, so that $q'(k_t) > 0$.[11] Third, technology displays constant returns to scale, so that the period profit function can be expressed as $k_t \cdot \gamma(w_{t+1}, p_{t+1})$ where γ, which is convex in p and w, is interpreted as the profit function per unit of capital.

A risk-neutral firm is then interested in maximizing the expected discounted stream of cash flow,

$$E\left[\sum_{t=0}^{\infty}(1+R)^t\Big(\Gamma(w_t, p_t;\ k_{t-1}) - q_t(k_t)\cdot(k_t - (1-\delta)k_{t-1})\Big)\right],$$

9. These are profits before any payments to capital are made. But since the level of capital in period t is predetermined, these are the profits that are maximized by the choice of n.
10. The intuition of the convexity of Γ is as follows: a rise in p increases profits by $f(k_{t-1}, n_t^*)$ if the firm does not change behaviour; if the firm re-optimizes, however, profits rise by more.
11. Hartman (1972) simply assumes that the marginal cost of investment is increasing in the amount of investment, reflecting either monopsony power in the capital goods market or the presence of adjustment costs.

where R is the firm's discount rate, which is assumed to be known and constant. The first-order condition for k_t is:

$$q_t + k_t q'(k_t)$$
$$= (1+R)^{-1} \cdot E\left[\gamma(w_{t+1}, p_{t+1}) + q_{t+1}(1-\delta)\right].$$

A mean-preserving spread in the distribution of p increases the right-hand side since γ is convex. Since the left-hand side is increasing in k, an increase in k is the optimal response to the rise in uncertainty. The heightened uncertainty over p therefore increases the flow demand for investment in period t.

It is worth emphasizing the generality of the result that an increase in uncertainty over p *increases* the firm's investment demand. First, it is not important that the firm be a price-taker in the product and labour markets. Even if the firm is a price-setter in those markets, its profit function is still convex in the relevant random variables (such as demand and supply shifters). Second, the result is sensitive to the technology, but not as much as one might think. If f displays constant returns, then Γ is definitely convex; in this case, the positive effect of uncertainty on investment requires only that the marginal cost of capital be increasing in k. This, in turn, only requires some monopsony power in the capital-goods market or, perhaps more reasonably, some costs to adjusting the capital stock. If f does not display constant returns, then the result only requires that the marginal product of capital be convex in the random variable. This is by no means guaranteed, but is certainly possible.[12]

Finally, and perhaps most importantly, even extreme risk aversion on the part of the firm is not sufficient to overturn the result. Risk aversion does imply that an increase in the uncertainty of profits will reduce the firm's expected utility of profits. But this is largely irrelevant to the investment decision. The investment

12. For example, assuming that f does not display constant returns to scale but does display diminishing returns, pf_k will be convex in p if $f_{kn} > 0$. This is not an uncommon assumption on technology.

decision is about the expected payoff to the marginal unit of capital, and crucial to this is how the increase in uncertainty affects the firm's expected *marginal* utility. If we model the firm as maximizing the expected discounted stream of the utility of cash flow, then the firm's first-order condition for k_t becomes (for constant returns to scale and increasing marginal cost of k),

$$q_t + k_t\, q'(k_t)$$
$$= (1+\mathrm{R})^{-1} \cdot E\Big[\mu'(F_{t+1}) \cdot \Big(\gamma(w_{t+1}, p_{t+1}) + q_{t+1}\,(1-\delta)\Big)\Big],$$

where μ is the firm's strictly concave utility function and F_{t+1} is the cash flow in period $t+1$. The MPS in p still increases the expected value of γ, but the effect on the expected marginal benefit of capital depends on the concavity of *marginal* utility, $\mu'(F)$. Only if μ' is sufficiently concave will the MPS have the effect of reducing the current flow demand for investment.

One important problem with this discussion of risk aversion is that there is little economic content to the convexity or concavity of *marginal* utility; it is quite unrelated to the convexity of preferences, which we typically assume for very good reasons. I know of no reason to favour a presumption of convexity over concavity of marginal utility, or vice versa. Given this agnostic position, it is probably not unreasonable to conduct the analysis under the assumption that marginal utility can be taken to be linear, and thus expected marginal utility to be unaffected by changes in uncertainty. The result is that we can carry on the analysis as if the firm is risk-neutral, with the result that an increase in uncertainty leads to an increase in the firm's flow investment demand.

3.2.3 Nominal uncertainty

I have a more significant concern with the analysis of the previous few pages; for the purposes of this paper, the uncertainty discussed above is in the wrong place. In particular, the uncertainty is over real (relative) prices that the firm faces; it is not at all about uncertainty over the aggregate price level. And it is the latter that is the focus of this paper.

It is reasonable to believe that any given firm is more certain about its relevant future relative prices (product and input prices) than it is about the future behaviour of the aggregate price level. And this is more true when we consider the longer horizon – the horizon that is presumably relevant to investment decisions. This simply says that the firm knows more about the industry-specific demand and supply shocks that it will face in the future than it does about the future behaviour of the central banker. Another way to say the same thing – using the vocabulary of the 1970s – is that the variance of real shocks in the economy is less than the variance of nominal shocks.

Having made this distinction between real and nominal uncertainty, I am led naturally to the following thought-experiment. Suppose all static relative prices are held constant and we consider a mean-preserving spread in the distribution of the aggregate inflation rate. It is difficult to think how this would have any effect on the firm's investment decision; after all, the firm cares only about its real selling price, its real input prices, and the opportunity cost of investment (which is the rate of return on some interest-bearing asset).

To formalize the intuition about the irrelevance of inflation uncertainty for investment behaviour, I adopt the extreme assumption that the firm knows perfectly the future path of p, w and q, but does not know the future path of the aggregate price level $(\bar{P})$. This means that the firm also does not know the future path of nominal prices $P\,(=p\cdot\bar{P})$, $W\,(=w\cdot\bar{P})$ or $Q\,(=q\cdot\bar{P})$. The firm's (nominal) profit function, $\bar{\Gamma}$, is now given by

$$\begin{aligned}\bar{\Gamma}(w_t, p_t, \bar{P}_t; k_{t-1}) &= P_t f(k_{t-1}, n_t^*) - W_t n_t^* \\ &= \bar{P}_t\left(p_t f(k_{t-1}, n_t^*) - w_t n_t^* \right) \\ &= \bar{P}_t \cdot \Gamma(w_t, p_t; k_{t-1}).\end{aligned}$$

In the case of constant returns technology and an increasing marginal cost of capital, the firm's first-order condition, from maximizing the expected discounted stream of the utility of cash flow, is

$$\mu'(F_t) \cdot (q_t + k_t q'(k_t))$$
$$= E\left[\mu'(F_{t+1})\left(\frac{1+\Pi}{1+R}\right)\left(\gamma(w_{t+1}, p_{t+1}) + q_{t+1}(1-\delta)\right)\right],$$

where Π is the inflation rate of $\bar{P}$ from period t to $t+1$.

Now consider the effect of an MPS in the distribution of inflation; it is clear that the increase in inflation uncertainty has no effect on the two terms in large brackets within the expectations operator. Furthermore, this result is also independent of the amount of risk aversion. As before, unless we have some solid priors concerning the concavity of the firm's *marginal* utility, we can treat the firm as if it were risk-neutral. Then, everything on the right-hand side is unaffected by the MPS in inflation. Under these rather special informational assumptions, the increase in inflation uncertainty has no effect whatsoever on the firm's flow demand for investment.

3.2.4 Summing up

The results concerning inflation uncertainty and investment behaviour can be summarized as follows:

i) If uncertainty is only over the aggregate inflation rate, then there is little reason to believe that a reduction in inflation uncertainty will have any effect on firms' investment demand.

ii) To the extent that inflation uncertainty brings with it uncertainty over future relative prices, then less inflation uncertainty suggests *less* investment demand. Furthermore, the higher the correlation between inflation uncertainty and relative price uncertainty, the more investment will fall when there is a reduction in inflation uncertainty.

iii) Both i) and ii) are independent of the degree to which firms are risk-averse. They do depend, however, on the concavity of the firms' *marginal* utility. For non-random relative prices, if marginal utility is concave (whatever this means, in terms of preferences), then the reduction in inflation uncertainty will raise investment demand. But if the reduction in inflation uncertainty brings with it a reduction in relative price uncertainty, then there is the possibility that investment

demand will fall (in particular, this will happen if marginal utility is not "concave enough" to offset the convexity of the profit function).

iv) Points i, ii and iii are reasonably robust to changes in assumptions about firms' technology and market power.

For those who had strongly held prior beliefs that a decrease in inflation uncertainty would lead to more investment, these four points must either be devastating or completely ridiculous. I must confess that I am surprised at how difficult it is to tell a simple and sensible story in which the firm optimally chooses to undertake less investment when there is an increase in inflation uncertainty. So, what have I missed? What is it about a rise in inflation uncertainty that suggests to many of us (in our hearts) that firms should shy away from new investment opportunities?

Perhaps the answer is found in a slight variant of Ball's (1990a) hypothesis. His view is that when inflation is high, policymakers are faced with a policy dilemma; this dilemma leads to uncertainty about future policy. Maybe the missing piece from the analysis here is that when inflation is uncertain, policy is also necessarily uncertain. Uncertainty about future economic policy generates a wait-and-see attitude on the part of firms. The importance of this "option value" increases when there are significant costs involved in changing the firm's capital stock, a point that I have omitted until now.

I conclude this subsection on investment behaviour by speculating that a reduction in inflation uncertainty probably increases the flow demand for investment. Part of this comes from a belief that firms are probably much more certain of their own future relative prices than they are about the future aggregate price level, especially over longer forecasting horizons. This is enough to render Hartman's analysis largely irrelevant, and leaves my formal analysis at the point where investment behaviour is independent of inflation uncertainty. The rest comes purely from my basic instincts about the option-value argument, and I find myself in the very uncomfortable position of introducing this argument as a *deus ex machina*.

I express my speculations in the following stylized first-order condition for the firm's investment decision:

$$MPK_t = r(\overline{\Pi}) + \delta + \theta_f(\sigma_\pi^2) \tag{1}$$

where $r(\overline{\Pi})$ is what I call the "safe" real interest rate – the real rate that would exist if inflation took its expected value[13] – and θ_f is increasing in the amount of inflation uncertainty (σ_π^2) and captures the extent to which firms want to wait and see when there is an increase in inflation uncertainty. Expressed in the form of equation (1), I think of this uncertainty-generated option value as increasing the user-cost of capital; an increase in inflation uncertainty increases the user-cost of capital and thus lowers the current flow demand for investment.

3.2.5 *Individuals' saving behaviour*

To analyze the behaviour of individuals, I continue with the assumption that there are no indexed bonds in the credit market. I assume that individuals have time-separable preferences over consumption of a composite good, the price of which is the aggregate price level. The first-order condition for the individual's maximization problem is

$$u'(c_t) \cdot (1+\rho) = (1+R_t) \cdot E\left[\frac{u'(c_{t+1})}{1+\Pi_t}\right]$$

where u is the individual's strictly concave utility function and ρ is his/her subjective rate of time preference. The nominal interest rate, R_t, is known at the beginning of the period and so comes out of the expectations operator. The only uncertainty faced by the individual is the inflation rate from period t to $t+1$, and therefore the level of real consumption in period $t+1$.

Now consider a mean-preserving spread in the distribution of inflation. There are potentially two effects. The first involves the convexity of the individual's marginal utility. For the same reasons as above, this is not particularly

13. $1 + r(\overline{\Pi})$ is obviously equal to $(1+R)/(1+\overline{\Pi})$ and so $r(\overline{\Pi}) = (R-\overline{\Pi})/(1+\overline{\Pi})$.

interesting (or meaningful) and so I proceed under the assumption that the MPS has no effect on the individual's expected marginal utility of future consumption.

The second effect is that an MPS in inflation increases the expected real return to saving, since $[1/(1+\Pi)]$ is convex in Π. This is simply Jensen's Inequality again. This generates both income and substitution effects. For any particular individual, the total effect on saving could go either way. But in the aggregate, we expect the substitution effect to dominate the income effect; the income effects for borrowers tend to cancel out those of the lenders, whereas the substitution effect is in the same direction for everybody.

This suggests that a reduction in inflation uncertainty should increase current aggregate consumption demand, as individuals optimally tilt their consumption paths towards the current period. The interpretation that I prefer is that the reduction in inflation uncertainty makes people more prepared to borrow (or less prepared to lend) and, in this sense, makes them less patient. I express this by re-writing the first-order condition as

$$\frac{u'(c_t)}{E[u'(c_{t+1})]} \cdot (1+\rho) = 1 + r(\overline{\Pi}) + \theta_i(\sigma_\pi^2) \tag{2}$$

where the expected real interest rate is the sum of the safe rate, $r(\overline{\Pi})$, and the term, θ_i, which is increasing in the amount of inflation uncertainty.

It is worthwhile noting that θ_i is not a function of preferences; it depends only on the convexity of $[1/(1+\Pi)]$. Thus, the subscript on θ_i does not index individuals; the i stands for (all) individuals. But preferences will affect the extent to which a change in θ_i leads to an intertemporal redistribution of consumption. In particular, the greater the individual's elasticity of intertemporal substitution, the greater the effect on current consumption demand from the MPS in inflation. If the "representative" individual has an elasticity of intertemporal substitution equal to zero, then changes in inflation uncertainty will have no effect on aggregate consumption demand.

3.2.6 Labour supply

Following the same reasoning, it is not surprising that an MPS in inflation leads to an increase in aggregate labour supply. As before, the increase in inflation uncertainty increases the expected real interest rate and thus generates both income and substitution effects. The substitution effect leads individuals to reduce their current demand for leisure and thus increase their current labour supply. Under the assumption that leisure is a normal good, the income effect leads borrowers to increase their labour supply but leads lenders to reduce theirs. For the aggregate economy, the substitution effects should dominate.

This suggests that a reduction in inflation uncertainty should reduce current aggregate labour supply and thus decrease the aggregate supply of output. As with the analysis of the consumption-saving decision, preferences matter because they determine the magnitude of this effect. If the "representative" individual is unwilling to substitute leisure intertemporally in response to changes in the expected real interest rate, then the reduction in inflation uncertainty will have no effect on aggregate supply.

3.2.7 Medium-term equilibrium

I now combine the analysis of the three preceding subsections to determine the effect of inflation uncertainty on the medium-term equilibrium. Figure 4 shows aggregate consumption demand (C^d), aggregate investment demand (I^d), and aggregate output supply (Y^s), all as functions of the real interest rate (r). A reduction in r lowers the user-cost of capital and thus increases the quantity of investment demanded. A reduction in r tilts aggregate consumption demand towards the present. An increase in r leads individuals, in aggregate, to increase their current labour supply and thus to increase the quantity of output supplied. Y^d is the aggregate demand curve and is simply the horizontal summation of I^d and C^d. The intersection of Y^d and Y^s determines the equilibrium real interest rate and the level of aggregate output. This is the standard view of the goods market in intermediate macro textbooks such as Barro (1993).

Intertemporal substitution clearly plays a significant role in this model of the economy, but it is perfectly consistent to use this framework of analysis and

still believe that individuals have very low elasticities of intertemporal substitution. Figure 5 shows my preferred version of the diagram. The aggregate supply curve is vertical, reflecting my belief that the amount of intertemporal substitution of labour supply is approximately zero. The consumption demand curve is drawn downward-sloping but quite steep; there is some intertemporal substitution in consumption, but probably not a lot. The aggregate investment demand schedule is drawn with considerable interest elasticity; since investment in physical capital takes time to generate returns, the opportunity cost of that investment – the interest rate – should be reasonably important.

Now suppose that the central bank undertakes a policy of reducing the uncertainty of inflation – that is, they engineer a mean-preserving compression in the distribution of inflation. The variable on the vertical axis in Figure 5 is now $r(\overline{\Pi})$, the safe real interest rate. I have assumed in drawing Figure 5 that the elasticity of intertemporal substitution of leisure is zero; it follows that the Y^s curve is not affected by the change in inflation uncertainty. The C^d curve shifts a little to the right, reflecting individuals' limited willingness to substitute consumption intertemporally. The I^d schedule also probably shifts to the right, given my speculative arguments in the previous subsection.

The conclusion about the medium-term economy is that the safe real interest rate will rise when inflation uncertainty falls.[14] Furthermore, it will rise by more, the more willing individuals are to substitute consumption and leisure intertemporally. Also, the more inflation uncertainty leads firms to adopt a wait-and-see attitude, the more the safe real rate will rise. The effect on aggregate output is ambiguous; aggregate supply could fall by more than the increase in aggregate demand, but I think this is very unlikely. To the extent that Figure 5 represents realistic assumptions about individuals' willingness to substitute consumption and leisure intertemporally – and I think that it does – the effect on aggregate supply will be approximately zero. Furthermore, with zero intertemporal

14. Since $\overline{\Pi}$ is unchanged by construction, the nominal interest rate rises.

substitution in leisure, the aggregate supply curve is vertical; this implies that any change in aggregate demand affects the safe real interest rate but has no effect on the level of medium-term output.

I conclude that if the central bank undertakes a policy that successfully reduces the amount of inflation uncertainty, we should expect an increase in the safe real interest rate and, at most, a very small increase in the level of aggregate output. The story is that the reduction in inflation uncertainty makes individuals less patient – thus increasing their current demand for output – and as well makes firms more willing to undertake new investment projects (coming from the *assumed* reduction in policy uncertainty that accompanies the reduction in inflation uncertainty). This increase in demand for current goods pushes up the safe real interest rate. Output rises only to the extent that individuals are prepared to supply more work effort in response to the rise in the safe real interest rate.

3.2.8 Some suggestive evidence

I conclude my discussion of the medium term by looking briefly at some data. The analysis of the previous subsection predicts a negative correlation between the safe real interest rate and the amount of inflation uncertainty. For a pooled sample of G-7 countries (omitting France), I estimate a simple OLS regression of the safe real interest rate on the time-varying measure of inflation uncertainty as well as country and year dummy variables.

Table 2 shows the results for four different specifications. In the first specification, without any year or country controls, the predicted negative relationship is present and is statistically significant. In the second specification, we add country dummies to control for possible differentials in risk premiums. The negative effect remains. One problem with these simple regressions is that they omit all other determinants of real interest rates, which may be correlated with the measure of inflation uncertainty. To capture this, I include a full set of year controls, but I constrain the year effects to be the same across countries. This reflects the view that capital markets are integrated and thus real interest rates in all countries should be influenced by the same factors. The negative relationship between the safe real rate and inflation uncertainty is still

present and significant in this third specification. Finally, the fourth specification includes both country and year dummies, and the negative relationship remains.

Note that the effect of inflation uncertainty on the safe real interest rate is also economically significant; the typical (across the six countries) sample standard deviation of V_{it} is 0.25, so a reduction in inflation uncertainty by one standard deviation has the effect of increasing the equilibrium safe real interest rate by approximately 0.23 of a percentage point $(= (0.9) \cdot (0.25))$.[15]

3.3 The effect on the steady-state capital stock

The steady state to which the economy converges in simple macro models is characterized by the modified golden rule, $MPK = \rho + \delta$, where MPK is the marginal product of capital, ρ is the individuals' subjective rate of time preference, and δ is the rate of depreciation of physical capital.[16] To see the effect of inflation uncertainty on the steady-state capital stock, combine the firm's first-order condition, equation (1), and the individual's first-order condition, equation (2). The modified golden rule, appropriately adjusted for the presence of inflation uncertainty, then becomes

$$\begin{aligned} MPK &= \left(\rho - \theta_i(\sigma_\pi^2)\right) + \left(\delta + \theta_f(\sigma_\pi^2)\right) \\ &= \rho^*(\rho, \sigma_\pi^2) + \delta^*(\delta, \sigma_\pi^2) \end{aligned} \tag{3}$$

which makes clear my association of θ_i with the impatience of individuals and θ_f with the user-cost of capital. I assume that any change in the uncertainty of

15. It is interesting to note that my simple results in Table 2 contrast those of Barnea et al. (1979) who estimate a similar regression for the United States for the period 1959-1974. Their measure of inflation uncertainty, however, is the variance at time t over the cross-section of inflation forecasts found in the Livingston Survey.

16. The intuition of this condition is straightforward. Profit-maximizing firms choose their capital stock in every period so that $MPK = r + \delta$. A steady state of the economy requires that consumption is constant across periods, which in turn requires that $r = \rho$. Thus, the aggregate capital stock accumulates until the marginal product of capital equals $\rho + \delta$. For simplicity, I assume zero population growth.

inflation will have no effect on the rate of physical depreciation (δ) or on the fundamental rate of time preference (ρ). I also assume that it does not affect the marginal product of capital schedule.[17] The effect of any change in inflation uncertainty on the steady-state capital stock therefore depends on the relative magnitudes of the θ terms. The intuition is the following. A reduction in inflation uncertainty (and therefore about future policy) reduces the option value associated with not investing. This effect implies a higher steady-state capital stock. But a reduction in inflation uncertainty also makes individuals less patient, so they prefer consumption to saving. This effect suggests a lower steady-state capital stock.

The conclusion is that, in principle at least, the steady-state capital stock, and with it the level of steady-state output and consumption, could rise or fall as a result of a reduction in the amount of inflation uncertainty. My own intuition is that the effect of a reduction in inflation uncertainty on firms' investment decisions is probably more significant than the effect on individuals' consumption decisions. This intuition comes largely, I think, from my belief that individuals' willingness to substitute consumption intertemporally is quite limited; but this belief is not based on any hard evidence. If my intuition is correct, then a policy that systematically reduces the uncertainty of inflation can reasonably be expected to increase the economy's steady-state capital stock, but I have no idea how large or small this effect might be.

4. An Economy with Financial Intermediation

In contrast to the model examined in the previous section, financial intermediaries do exist, and they exist for many reasons: there are almost certainly fixed costs associated with making credit-market transactions; borrowers and lenders are not all interested in transactions of the same size or for the same amount of time; and borrowers are not all equally likely to default on their loans.

17. To the extent that there is some intertemporal substitution of work effort, a fall in inflation uncertainty would reduce work effort and thus would likely have an effect on the MPK schedule. This effect is probably very small and I ignore it.

As with the output of any industry, there are real costs associated with the provision of financial intermediation, and anything that increases these costs is likely to reduce the amount of financial intermediation in the economy. Such a reduction in financial intermediation (hereafter, FI) is likely to affect real variables such as output or employment. One mechanism through which the amount of FI affects the real economy is on the supply side: credit is an input to the production process and so a reduction in the amount of FI reduces the availability (or increases the price) of credit and thus can be viewed as a negative shock to productivity (King and Plosser 1984). In this interpretation, an increase in the cost of FI is regarded in a similar way as an OPEC-generated increase in the price of oil. Another possible mechanism, emphasized by Bernanke (1983), is on the demand side: a rise in the cost of FI (due to bankruptcies and bank failures) leads to an increase in the lending rate and thus, for a given riskless rate, there is a reduction in demand for current output.

This section examines the relationship between inflation uncertainty and the costs of financial intermediation. The basic hypothesis is that an increase in inflation uncertainty will increase the cost of financial intermediation and, as a result, lead to a reduction in the amount of financial intermediation in the economy.

In what follows, the economy is viewed in two different ways. In Section 4.1, let us suppose that the FI industry is competitive and that all borrowers have the same default risk. In this setting, financial intermediaries exist because there are fixed costs associated with intermediation. In Section 4.2, we continue with the assumption of competitive financial intermediaries, but borrowers differ in their default risk. This allows credit rationing to make an appearance. Though the precise details of the story differ slightly between Sections 4.1 and 4.2, the bottom line is the same: an increase in inflation uncertainty increases the cost of financial intermediation.

4.1 The simplest world with financial intermediation

We continue here with the assumption that indexed loan contracts are not available. All potential borrowers have the same default risk and are known by the

financial intermediaries to have the same default risk. There can therefore be no credit rationing in this world – at least, not of the Stiglitz and Weiss (1981) variety. Furthermore, all credit-market transactions are assumed to be for a maturity of one period. These assumptions imply that there is a single lending rate in the credit market. Finally, I assume that all financial intermediaries are identical and that they offer identical services to all potential depositors. There is, therefore, a single deposit rate in the credit market.

The spread between the lending rate and the deposit rate reflects the profits made by financial intermediaries. Entry to the FI industry then drives profits to zero, at which point the spread reflects the cost of financial intermediation. Anything that increases the cost of FI leads the intermediaries to lower the deposit rate; the lending rate, in contrast, is taken as given by the FI industry and is determined only by demand and supply in the credit market. The fall in the deposit rate leads to a reduction in the amount of deposits as some lenders remove their deposits and choose instead to lend directly to borrowers by buying bonds. The reduction in deposits implies a reduction in the amount of financial intermediation, an effect that is only magnified by the existence of a fractional reserve banking system.

As is often the case in models of the credit market, the existence of bankruptcy plays a central role in my story. Bankruptcy allows a borrower to default on a loan if the return from his investment project falls below some cutoff point that places his wealth below the value required for full repayment of the loan. A mean-preserving spread in the distribution of the project's return therefore increases the probability of default since it pushes more of the distribution below the cut-off point. Since the cost of financial intermediation increases with the probability of default, an MPS in the project's return increases the cost of financial intermediation.

The important point for this paper is that the loan and repayment are specified in nominal terms (owing to the absence of indexed contracts). Thus, it is uncertainty in the borrower's nominal projected return that matters for the probability of default. An MPS in the distribution of inflation therefore has the

effect of increasing the probability of default and therefore increasing the cost of financial intermediation.

In this simple world, what should we observe in the credit market when there is an increase in the amount of inflation uncertainty? First, the rise in the cost of FI should lead to an increase in the spread between the lending and deposit rates.[18] Second, the rise in the spread should lead to a reduction in the amount of deposits and thus to a reduction in the amount of financial intermediation.

4.2 Financial intermediation in a more interesting world

The more interesting world I have in mind changes only one assumption from the simple world: borrowers are allowed to differ in their default risk, in a manner that is not observable by the financial intermediaries. This allows the possibility of equilibrium credit rationing, as in Stiglitz and Weiss (1981).

The basic story of equilibrium credit rationing is a simple one. Borrowers all have projects with risky returns, and their distributions over these returns differ. For simplicity, suppose that all distributions have the same mean, but differ according to the variance, which is known only by the individual borrowers. Because of the borrowers' ability to declare bankruptcy in the event of a low return on their projects, the expected return to an individual borrower – for a given interest rate – increases with the riskiness of the distribution. This means that higher lending rates will encourage potential borrowers with more risky investments. This adverse-selection problem implies that the financial intermediary's expected return may not be an increasing function of the lending rate, a least over some range.[19] Moreover, the lending rate that maximizes expected profits may be one for which the quantity of loans demanded exceeds the quantity of loans supplied – equilibrium credit-rationing.

How does inflation uncertainty enter this story? As in the simple world above, an increase in inflation uncertainty increases the (nominal) riskiness of all borrowers' distributions and, therefore, because of the possibility of bankruptcy,

---

18. The two rates, individually, may rise or fall, since the inflation uncertainty may have an effect on the lending rate (see Section 3), but the spread between the two should clearly rise.
19. See also Jaffee and Modigliani (1969) for a discussion of this relationship.

increases the expected return for every potential borrower. Hence, at a given lending rate, more potential borrowers will demand credit. But the increased expected return for a borrower simply reflects the increased probability of default, and thus a decrease in expected return for the financial intermediary. The supply of credit therefore falls. Not only is there a reduction in the amount of credit that is actually transacted, but the amount of credit-rationing (the gap between demand and supply) also increases.[20]

What should we observe in this more complicated world when there is an increase in inflation uncertainty? As in the simple world, we should observe both an increase in the spread between the lending and deposit rates, and a reduction in the amount of credit extended by financial intermediaries. The benefit from considering a world in which credit-rationing takes place is that it allows us to make use of the fact that all borrowers are not the same. In particular, since governments typically cannot (or do not) declare bankruptcy, we should observe financial intermediaries expressing a preference for the safer government assets during times of heightened inflation uncertainty. This preference might be reflected in an increase in the rate differential on loans to the private sector relative to government. Or it might be reflected in a substitution in the financial intermediaries' assets away from private-sector loans and towards government loans or bonds.

4.3 A look at some Canadian data

My basic premise is that an increase in inflation uncertainty pushes up the cost of financial intermediation; this, in turn, leads to a reduction in the amount of financial intermediation and a reduction in the level of aggregate economic activity. In this subsection, I present some evidence of the first effect. In particular, I show evidence that is consistent with the view that inflation uncertainty has

20. This effect is made worse if the increase in inflation uncertainty brings with it added uncertainty about relative prices which is different across potential borrowers. The potential borrowers who experience the largest increase in relative price uncertainty have the largest increase in expected return and thus the largest increase in the probability of default. In this case, the adverse-selection problem is exacerbated, generating even more credit-rationing than in the case discussed in the text.

significant effects on the actions of financial intermediaries; I interpret these actions as reflecting the increased cost of financial intermediation.[21]

4.3.1 Inflation uncertainty and relative prices in the credit market

I estimate the lending-deposit rate spread as the difference between the prime lending rate and the rate paid on 90-day deposit receipts. I estimate the private/government risk premium as the difference between the prime rate and the rate on short-term treasury bills.[22] Figures 6 and 7 (see Appendix) plot the spread and the risk premium, respectively, for Canada from 1961 to 1989. Table 3 shows the results of estimating simple OLS regressions of the spread and risk premium against expected inflation and the measure of inflation uncertainty.

The top half of Table 3 shows the regression for the lending-deposit rate spread. The first equation shows a significant positive relationship between inflation uncertainty and the spread. When expected inflation is included as an additional regressor, the estimated coefficient increases considerably.[23] The relationship is also economically significant. The sample mean of inflation uncertainty is 0.47, indicating that the "typical" forecast error for inflation is about one-half of a percentage point. The standard deviation of these errors is about 0.18. Thus, if inflation uncertainty in Canada were reduced by one standard deviation, the lending/deposit rate spread would be expected to fall by about 25 basis points.

The bottom half of Table 3 shows the regressions for the measure of the private/government risk premium. The rate of expected inflation does not appear to be a determinant of the risk premium; but there is a positive and statistically significant relationship between inflation uncertainty and the risk premium. The estimated coefficient of 1.86 suggests that a reduction in inflation uncertainty by

---

21. I present no evidence on the relationship between the cost of financial intermediation and the level of aggregate economic activity; for this paper, I take the existence of this relationship as an article of faith. Bernanke (1983) presents some evidence supporting the existence of such a relationship for the United States during the Great Depression.
22. An alternative measure of the private/government risk premium is the difference between the rate on prime corporate paper and the treasury bill rate. My results are not sensitive to the choice of measure.
23. The results are the same if actual inflation rather than expected inflation is used as a regressor.

one standard deviation would reduce the risk premium by about 37 basis points; clearly, this is economically significant.

I take the evidence in Table 3 as support for a relationship between inflation uncertainty and the cost of financial intermediation. The rise in inflation uncertainty leads to a rise in the cost of financial intermediation, as reflected by the behaviour of the spread, as well as to a change in the perception of risk associated with private versus government borrowers. My argument is that both effects reflect the increase in the probability of default (of private borrowers) which is caused by the increase in inflation uncertainty. Finally, note that these effects, besides being statistically significant, are not small; anyone who has shopped around for mortgages knows that 1/2-1/3 of a percentage point is not unimportant.[24]

4.3.2 *Inflation uncertainty and relative quantities in the credit market*

The final piece of evidence that I present is the relationship between inflation uncertainty and the composition of the assets of financial intermediaries. In a world where borrowers differ in terms of default risk, a rise in inflation uncertainty should lead financial intermediaries to substitute away from more risky borrowers. The most significant difference across borrowers in terms of default risk is the difference between private-sector borrowers and governments; I assume that governments cannot (or do not) declare bankruptcy, making their probability of default equal to zero. Such a portfolio switch by the financial intermediaries implies a reduction in the availability of credit to the private sector and thus, through the channels discussed by Bernanke (1983) and King and Plosser (1984), a likely reduction in aggregate economic activity.

---

24. There is one important caveat that should be mentioned: if either deposit rates or lending rates are regulated (e.g., Regulation Q in the United States until 1986), then the spread will not exactly equal the cost of financial intermediation, even if the FI industry is perfectly competitive. However, it is still true that an increase in the spread, other things equal, can be expected to reduce the amount of deposits and thus reduce the amount of financial intermediation in the economy. Thus, the results in Table 3 still suggest that an increase in inflation uncertainty, through its effect on the spread, reduces the amount of financial intermediation.

I use data on the chartered banks' non-reserve assets from 1969 to 1989. Non-reserve assets are divided into: treasury bills; (long-term) government bonds; loans; and corporate securities. My hypothesis is that the share of assets that represents lending to the private sector,

$$\Phi = \frac{loans + corporate\ securities}{total\ non\text{-}reserve\ assets},$$

should be negatively related to the amount of inflation uncertainty. Table 4 shows the regressions of Φ on expected inflation and inflation uncertainty.

The top two rows in Table 4 show a negative relationship between inflation uncertainty and Φ. Note, however, that when only inflation uncertainty is used as a regressor, the relationship is not statistically significant. When the measure of expected inflation is included in the regression, however, the effect of inflation uncertainty becomes strongly significant. My interpretation is that an increase in inflation – for a given amount of inflation uncertainty – causes financial intermediaries to reduce their holdings of government bonds and treasury bills, which have fixed nominal interest rates, and substitute towards loans where nominal interest rates can rise in step with inflation. The positive correlation between (expected) inflation and inflation uncertainty therefore biases towards zero the estimated coefficient on inflation uncertainty in the first equation.

The bottom part of Table 4 shows the same basic regression separately for each of the four components of the chartered banks' assets. For both treasury bills and government bonds, an increase in expected inflation reduces their asset share, while an increase in inflation uncertainty increases their asset share. For loans and corporate securities, on the other hand, increases in expected inflation increase their asset share while an increase in inflation uncertainty decreases their asset share. As must be the case, the sum of the estimated coefficients in each column (other than for the constant term) is zero, since the four asset shares sum to one.

5. Conclusions

I have tried to present a consistent framework of analysis for thinking about the real effects of inflation uncertainty. I have perhaps gone too far in removing some interesting and complicating factors. But the result has been the discussion of some real effects of inflation uncertainty in its purest form. In particular, I examine a world in which uncertainty only exists over the rate of inflation; static relative prices are known perfectly. I have even ruled out changes in the expected rate of inflation by conducting the thought-experiment of a mean-preserving spread in the distribution of inflation.

Given the extreme thought-experiment that I conduct, it is perhaps surprising that I find any real effects at all, especially since my model economy is entirely a market-clearing one. But it is important to emphasize that my model contains within it a very important market incompleteness – the absence of indexed contracts in the credit market – and this incompleteness of markets is the source of all my analytical results.

5.1 Policy relevance

Especially at conferences like this one, but also more generally, we should ask ourselves about the policy relevance of our economic modelling. There are two broad types of results that I have presented here, and I would like to give my brief assessment of their potential policy importance.

The first result suggests that even in a world of perfectly flexible prices, a (mean-preserving) reduction in inflation uncertainty will have real effects. Specifically, it will lead to an increase in aggregate demand which pushes up real interest rates; aggregate output might rise, but only to the extent that individuals are prepared to substitute leisure intertemporally in response to changes in the expected real interest rate. The driving force behind this result is Jensen's Inequality; contrary to the views held by a whole decade (or more) of macroeconomists who are used to violating Jensen's Inequality with abandon, a mean-preserving compression in the distribution of inflation (for a given nominal rate) does reduce the expected real interest rate.

But do I really expect central bankers across the world to feel enlightened now that they have been reminded of Jensen's Inequality? Not likely. Table 5 shows the practical importance of this effect in three hypothetical worlds: low, moderate and high (expected) inflation. In each, the nominal interest rate is three percentage points above the expected rate of inflation. The safe real interest rate, though, is different in the three worlds; this point has recently been emphasized by Patinkin (1993) in his discussion of the Israeli stabilization of 1985.[25] Each horizontal row in the table shows a particular standard deviation of inflation; inflation is always distributed binomially, with equal probability for each outcome. In each world, a rise in the standard deviation of inflation (holding the mean constant) generates an increase in the expected real interest rate. But the magnitude of this effect is very small. For example, in a world of 5 per cent expected inflation, a standard deviation of inflation of 5 per cent – which constitutes enormous volatility in such a world – pushes up the expected real rate by less than 25 basis points. Furthermore, for higher-inflation worlds, the same increase in standard deviation has a smaller effect on the expected real rate.

The policy conclusion from Table 5 has to be that the real effects that I have discussed in Section 3 are far too small to worry about. This is even more true if the typical individual has low elasticities of intertemporal substitution of consumption and leisure, for in this case, even large changes in the expected real rate brought on by large changes in inflation uncertainty would have only trivial effects on aggregate demand and supply.

I have mixed feelings about this conclusion. On the one hand, it is never very satisfying to reach the end of a piece of analysis only to conclude that the effects you have uncovered are irrelevant from a practical point of view. On the other hand, my results suggest that if we are to care about the level of inflation uncertainty, then it *must* be because inflation uncertainty brings with it something else; uncertainty about future policy and uncertainty about relative prices are two

25. The failure to realize that the safe real interest rate is $(R-\overline{\Pi}) / (1+\overline{\Pi})$ rather than simply $(R-\overline{\Pi})$ is, of course, due to our eagerness to immediately approximate $ln\,(1+x)$ by x, even when x is not very small.

effects that I have mentioned briefly, but there are certainly others that might be considered.

My second general result comes from thinking about the behaviour of financial intermediaries. A reduction in inflation uncertainty leads to a decrease in any individual borrower's probability of defaulting on a loan and thus to a decrease in the cost of financial intermediation. Through mechanisms like those discussed by Bernanke (1983) and King and Plosser (1984), this is likely to increase the level of aggregate economic activity.

It seems to me that this argument is worth thinking about seriously. My first efforts at looking at the data seem to confirm the view that changes in inflation uncertainty do change the composition of assets held by financial intermediaries; in particular, more inflation uncertainty seems to lead FIs to substitute away from private borrowers and towards government bonds. This process of financial disintermediation is damaging. In modern economies, where the needs of individual borrowers and lenders differ along so many dimensions, financial intermediaries are extremely important; indeed, financial intermediaries have a special role as "market makers." Central bankers should design their monetary policy with an eye to ensuring that their actions do not unnecessarily increase the cost of financial intermediation. Such a consideration may be one strong argument in favour of reducing the volatility of inflation.

5.2 What I have ignored

Let me close by briefly outlining several things that I have ignored in this paper, but that I think are important to the relationship between inflation uncertainty and the real economy.

First, my analysis has implicitly assumed that any non-neutralities that might exist in the tax system are not affected by a change in the amount of inflation uncertainty. This assumption is almost certainly (and significantly) false; I have ignored the issue of non-neutral taxes partly because such non-neutralities were discussed in another paper at this conference (Black, Macklem and Poloz), but mainly because it makes my life easier; this omission, however, in no way reflects a belief that it is an unimportant issue. To the extent that corporate and personal

income tax structures are not indexed to inflation – so that a reduction in (even perfectly anticipated) inflation leads to reductions in real income-tax revenue – then a mean-preserving "compression" in the distribution of inflation implies a decrease in expected marginal tax rates and thus a decrease in the expected value of real income-tax revenue.[26] Through both income and substitution effects, this will clearly have real consequences, that go beyond the ones I have discussed in this paper, on both investment and consumption behaviour.[27]

Second, I have drawn a clear distinction between uncertainty about the inflation rate and uncertainty about relative prices; in most of my analysis I assumed that relative prices were known with certainty because the points that I wanted to make did not require uncertainty about (static) relative prices. But there is considerable evidence that inflation and inflation uncertainty are positively correlated with relative price uncertainty (e.g., Fischer 1986a). Such uncertainty generates an inference problem for firms and individuals with obvious implications for the misallocation of resources. It seems easy to make the argument, as John Crow (1990) does in the quote in the Introduction, that a reduction in inflation uncertainty in Canada can be expected to lead to a reduction in such inefficiencies.

I have also completely ignored the labour market and how behaviour there might be affected by a reduction in inflation uncertainty. There are two types of things I have ignored here. The first is recognition of the fact that a reduction in inflation uncertainty probably lengthens contracts and decreases indexation provisions contained in those contracts. Both changes represent a saving of real resources. But, real though these gains may be, they are probably not much larger in magnitude – especially in a country with moderate inflation – than the reduction in the shoe-leather costs that accompanies a reduction in (anticipated) inflation.

---

26. Such "bracket-creep" occurs because of incomplete indexation and a progressive tax structure (increasing marginal income tax rates). My claim about inflation volatility is then just another application of Jensen's Inequality applied to a convex tax function.
27. See Feldstein (1982) for the relationship between inflation and investment in a world of distortionary taxation.

I think my more serious omission regarding the labour market relates to the argument put forward by Friedman (1977), and explored further by Mullineaux (1980), that the natural rate of unemployment is positively correlated with inflation uncertainty. Friedman's argument is that the transition to an environment of more volatile inflation introduces additional noise that increases frictional unemployment, but that eventually institutions develop to deal with this heightened inflation uncertainty, thus leaving the long-run natural rate unaffected by changes in inflation uncertainty.

My own view is that there are considerable fixed costs associated with the development of institutions that could make inflation uncertainty irrelevant, and that inflation and inflation uncertainty have probably never been significant enough in Canada to justify their evolution. Given this belief, I think it is reasonable to expect that a reduction in inflation uncertainty in Canada would reduce the natural rate of unemployment. It is relatively easy to construct a sensible search model of the labour market in which a reduction in inflation uncertainty reduces the option value (for firms and workers) associated with not making a worker-firm match, and thus reducing the "natural" rates of unemployment and vacancies. Such a reduction in the natural rates simply reflects an increase in the efficiency of the process that matches workers and firms in the labour market.[28]

Finally, I have completely ignored any relationship that might exist between inflation uncertainty and the rate of productivity growth. Indeed, if such a relationship exists, this might be by far the most important of all the effects discussed here. The problem is that until we have a clear framework in which to think about the sources of productivity growth, we are unable to discuss the effects of other factors on that growth. If productivity "shocks" are purely exogenous, then there is no reason to expect any relationship between growth and inflation uncertainty. If productivity growth is endogenous, however, as much recent literature suggest (e.g., P. Romer 1986, 1989), the relationship between investment

28. See Hogan and Ragan (1993), for example, for a simple equilibrium search model of the labour market which could be modified to generate this result.

(in all types of capital) and inflation uncertainty will be central to the relationship between economic growth and inflation uncertainty.

There is considerable evidence of a strong positive relationship between the level of inflation and the rate of productivity growth.[29] This body of research, however, is focussed on the relationship between inflation and growth rather than on the relationship between inflation *uncertainty* and growth. Given the strong positive correlation that exists between the rate of inflation and the uncertainty of inflation, it is entirely possible that these studies are detecting the true negative relationship between inflation uncertainty and economic growth, but are interpreting it as evidence of a relationship between inflation and growth when in fact no such relationship exists. Another problem is that there has not been extensive theoretical modelling of the relationship between inflation uncertainty and economic growth. This would seem to be a fruitful area for future research.

* * *

Postscript

The real effects of inflation volatility discussed in Section 3 should be viewed as resulting from distortions introduced to the economy through the volatility of inflation. Suppose, for the sake of argument, that one takes the rather extreme view that the monetary authority can choose precisely the degree of inflation volatility; in this respect, I view inflation volatility like a distortionary tax. The choice of more variable inflation (for a given mean) increases the expected real interest rate (for a given nominal interest rate) and thus distorts the consumption and leisure decisions of individuals. An MPS in inflation reduces aggregate demand and increases aggregate supply with the result that the equilibrium nominal interest rate (as well as the safe real interest rate) falls

29. Jarrett and Selody (1982), using Canadian data from 1963-1979, conclude that a one percentage point reduction in inflation would increase the rate of productivity growth by three-tenths of a percentage point. Selody (1990a) has recently revised this estimate to two-tenths of a percentage point. Lebow et al. (1992) find a similar negative correlation for the United States using data from 1950 to 1989. Fischer (1991) uses cross-section data from a sample of 73 countries from 1970-85 and finds a strong negative correlation between inflation and real economic growth.

unambiguously; the effect on aggregate output depends on the relative shifts of aggregate demand and aggregate supply.

In Section 5, I argued that the practical importance of these Jensen Inequality effects are almost certainly too small to worry about. Table 5 shows that enormous amounts of inflation uncertainty are required to bring about even moderate changes in the expected real interest rate (for a given nominal interest rate). Table 5 also shows that the effect of any given increase in inflation uncertainty on the expected real interest rate declines as expected inflation rises. This result is immediately clear from Figure 8, which shows that the function $\omega(\Pi) = (1+\Pi)^{-1}$ is "less convex" when Π is large than when Π is small. The practical relevance of this point is that the distortionary costs associated with a given amount of inflation volatility, though probably small, *are larger in low-inflation countries than in high-inflation countries.*

This has one immediate implication for inflation policy, and it concerns the optimal target level of inflation. Suppose that inflation volatility is simply unavoidable but that the monetary authority can still choose some target level of inflation. More precisely, suppose that inflation is distributed with a mean of $\overline{\Pi}$ and a standard deviation of σ; the monetary authority can choose $\overline{\Pi}$ but can do nothing to change σ. In this situation, what is the optimal target value of $\overline{\Pi}$?

As is clear from Figure 8, the function $\omega(\Pi) = (1+\Pi)^{-1}$ is the "most convex" around the point $\Pi = 0$. And the distortionary costs of inflation volatility – the Jensen Inequality effects – are maximized when inflation is variable over this "most convex" region of ω. Thus, the conclusion is that, when the monetary authority is saddled with some unavoidable amount of inflation volatility, the inflation target that *maximizes* the distortionary costs of inflation volatility is zero. Considering these costs only, therefore, the worst possible policy is one of zero (mean) inflation!

While the argument in the preceding paragraph is intriguing in its own right, let me immediately distance myself from the use of this argument in guiding inflation policy. As I mentioned in the body of the paper, my instincts are that the typical individual is probably not prepared to make significant intertemporal substitutions of consumption or leisure in response to even moderately large

changes in the expected real interest rate. If my instincts are correct, then the distortionary costs of inflation volatility are probably trivial, and so these distortionary costs associated with a zero-inflation target can be ignored.

Finally, let me close by saying that none of the discussion in this postscript in any way diminishes the arguments made in Section 4 about inflation volatility and the costs of financial intermediation. Those arguments were based on the effects of heightened nominal uncertainty in a world where individual borrowers had the ability to declare bankruptcy. Nothing in those arguments hinged on Jensen's Inequality or on the willingness of individuals to substitute consumption or leisure intertemporally.

Appendix

Table 1

Inflation and Inflation Uncertainty

$$V_{it} = \alpha + \beta \cdot \Pi_{it} + \varepsilon_{it}$$

	Constant	Inflation	N	$\bar{R}^2$
22 OECD countries	0.282 (0.020)	0.085 (0.002)	5,717	0.23
G-7 countries	0.296 (0.020)	0.049 (0.002)	1,930	0.19
Canada	0.317 (0.024)	0.024 (0.003)	276	0.15

Note: Standard errors are in parentheses.

Table 2

Interest Rates and Inflation Uncertainty

$$\left(\frac{R_{it} - \Pi^e_{it}}{1 + \Pi^e_{it}}\right) = \alpha + \beta \cdot V_{it} + \gamma\ D_{it} + \varepsilon_{it}$$

where R_{it} = rate on short-term treasury bills

Π^e_{it} = one-year-ahead inflation forecast

Countries: Canada, Germany, Great Britain, Italy, Japan, United States

	Constant	Inflation uncertainty	N	$\bar{R}^2$
No controls	2.834 (0.122)	–2.753 (0.131)	1,449	0.23
Country controls	---	–2.943 (0.147)	1,449	0.27
Year controls	---	–1.475 (0.098)	1,449	0.67
Country and year controls	---	–0.906 (0.116)	1,449	0.69

Note: Standard errors are in parentheses.

Table 3

Spread, Risk Premium and Uncertainty in Canada

$$(R_t^p - D_t^{90}) = \alpha + \beta \cdot V_t + \gamma \cdot \Pi_t^e + \varepsilon_t$$

where R_t^p = prime interest rate

D_t^{90} = rate paid on 90-day deposit receipts

Sample period: 1966M12 - 1989M11

Constant	Inflation uncertainty	Expected inflation	N	$\bar{R}^2$
0.910 (0.124)	0.953 (0.244)	---	276	0.05
1.057 (0.135)	1.208 (0.261)	–0.041 (0.016)	276	0.07

$$(R_t^p - R_t) = \alpha + \beta \cdot V_t + \gamma \cdot \Pi_t^e + \varepsilon_t$$

where R_t^p = prime rate

R_t = rate on short-term treasury bills

Sample period: 1966M12 - 1989M11

Constant	Inflation uncertainty	Expected inflation	N	$\bar{R}^2$
0.742 (0.112)	1.816 (0.222)	---	276	0.19
0.770 (0.124)	1.863 (0.240)	–0.008 (0.015)	276	0.19

Note: Standard errors are in parentheses.

Table 4

Inflation Uncertainty and Canadian Chartered Bank Asset Shares

$$\Phi_t = \alpha + \beta \cdot V_t + \gamma \cdot \Pi_t^e + \varepsilon_t$$

where Φ_t = share of chartered bank assets that represents lending to private sector

Sample period: 1969M4 - 1989M11

Constant	Inflation uncertainty	Expected inflation	N	$\bar{R}^2$
0.863 (0.011)	–0.020 (0.022)	---	248	0.00
0.825 (0.012)	–0.071 (0.021)	0.009 (0.001)	248	0.17

$$S_{jt} = \alpha + \beta \cdot V_t + \gamma \cdot \Pi_t^e + \varepsilon_t$$

where S_{jt} = share of chartered bank assets in form j

Sample period: 1969M4 - 1989M11

Asset form	Constant	Inflation uncertainty	Expected inflation	$\bar{R}^2$
T-bills	0.099 (0.002)	0.023 (0.004)	–0.004 (0.000)	0.49
Govt. bonds	0.076 (0.010)	0.048 (0.018)	–0.005 (0.001)	0.07
Loans	0.765 (0.009)	–0.051 (0.017)	0.008 (0.001)	0.20
Corporate securities	0.060 (0.003)	–0.019 (0.005)	0.001 (0.000)	0.05

Note: Standard errors are in parentheses.

Table 5

The Policy Irrelevance of Jensen's Inequality

World 1 (low inflation)	World 2 (moderate inflation)	World 3 (high inflation)
$E(\Pi) = 5\%$	$E(\Pi) = 10\%$	$E(\Pi) = 30\%$
$R = 8\%$	$R = 13\%$	$R = 33\%$
$r(\overline{\Pi}) = 2.857\%$	$r(\overline{\Pi}) = 2.727\%$	$r(\overline{\Pi}) = 2.308\%$
$\Pi \varepsilon \{4.5,5.5\}$	$\Pi \varepsilon \{9.5,10.5\}$	$\Pi \varepsilon \{29.5,30.5\}$
$\sigma = 0.5$	$\sigma = 0.5$	$\sigma = 0.5$
$E[r] - r(\overline{\Pi}) = 0.002\%$	$E[r] - r(\overline{\Pi}) = 0.002\%$	$E[r] - r(\overline{\Pi}) = 0.001\%$
$\Pi \varepsilon \{4,6\}$	$\Pi \varepsilon \{9,11\}$	$\Pi \varepsilon \{29,31\}$
$\sigma = 1.0$	$\sigma = 1.0$	$\sigma = 1.0$
$E[r] - r(\overline{\Pi}) = 0.009\%$	$E[r] - r(\overline{\Pi}) = 0.009\%$	$E[r] - r(\overline{\Pi}) = 0.006\%$
$\Pi \varepsilon \{2,8\}$	$\Pi \varepsilon \{7,13\}$	$\Pi \varepsilon \{27,33\}$
$\sigma = 3.0$	$\sigma = 3.0$	$\sigma = 3.0$
$E[r] - r(\overline{\Pi}) = 0.084\%$	$E[r] - r(\overline{\Pi}) = 0.077\%$	$E[r] - r(\overline{\Pi}) = 0.054\%$
$\Pi \varepsilon \{0,10\}$	$\Pi \varepsilon \{5,15\}$	$\Pi \varepsilon \{25,35\}$
$\sigma = 5.0$	$\sigma = 5.0$	$\sigma = 5.0$
$E[r] - r(\overline{\Pi}) = 0.234\%$	$E[r] - r(\overline{\Pi}) = 0.213\%$	$E[r] - r(\overline{\Pi}) = 0.152\%$
$\Pi \varepsilon \{-5,15\}$	$\Pi \varepsilon \{0,20\}$	$\Pi \varepsilon \{20,40\}$
$\sigma = 10.0$	$\sigma = 10.0$	$\sigma = 10.0$
$E[r] - r(\overline{\Pi}) = 0.942\%$	$E[r] - r(\overline{\Pi}) = 0.856\%$	$E[r] - r(\overline{\Pi}) = 0.609\%$

Note: R is the nominal interest rate, $r(\overline{\Pi})$ is the safe real interest rate, and Π is the actual inflation rate. $E(\Pi)$ and $E(r)$ denote, respectively, the expected rate of inflation and the expected real interest rate.

Figure 1
Standard Deviation and Mean of Inflation Rates
22 OECD Countries: 1960-1989

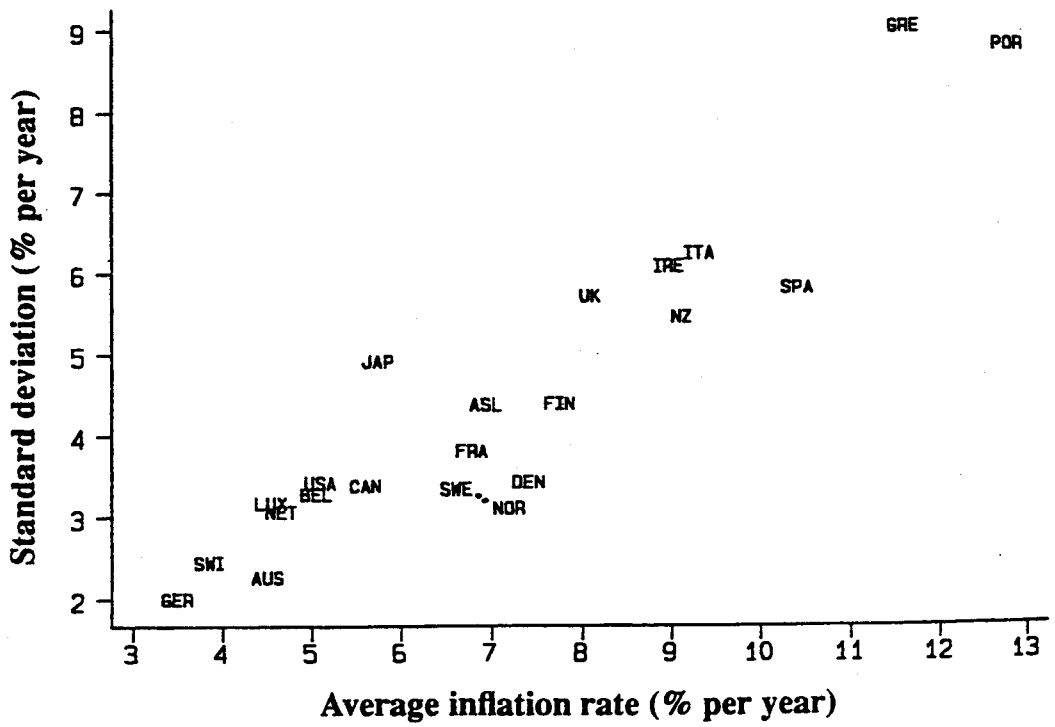

Associated regression for Figure 1

$$SD = -0.634 + 0.705 \cdot \overline{\Pi} \text{ ; with } N = 22 \text{ and } \bar{R}^2 = 0.85 \text{ .}$$
$$(0.476) \quad (0.063)$$

Figure 2

Average Absolute Change and Mean of Inflation Rates

22 OECD Countries: 1960-1989

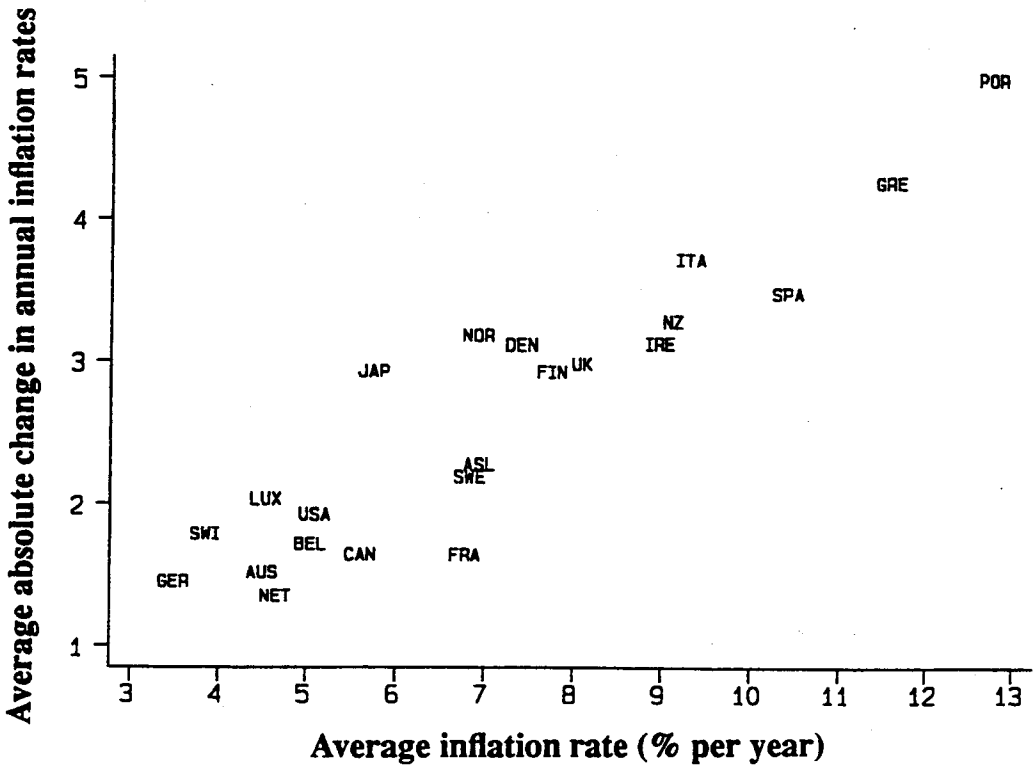

Associated regression for Figure 2

$$AA\Delta = \underset{(0.257)}{0.003} + \underset{(0.034)}{0.363} \cdot \overline{\Pi} \text{ ; with } N = 22 \text{ and } \bar{R}^2 = 0.84 \text{ .}$$

Figure 3
Inflation and Inflation Uncertainty
Canada: 1967-1989

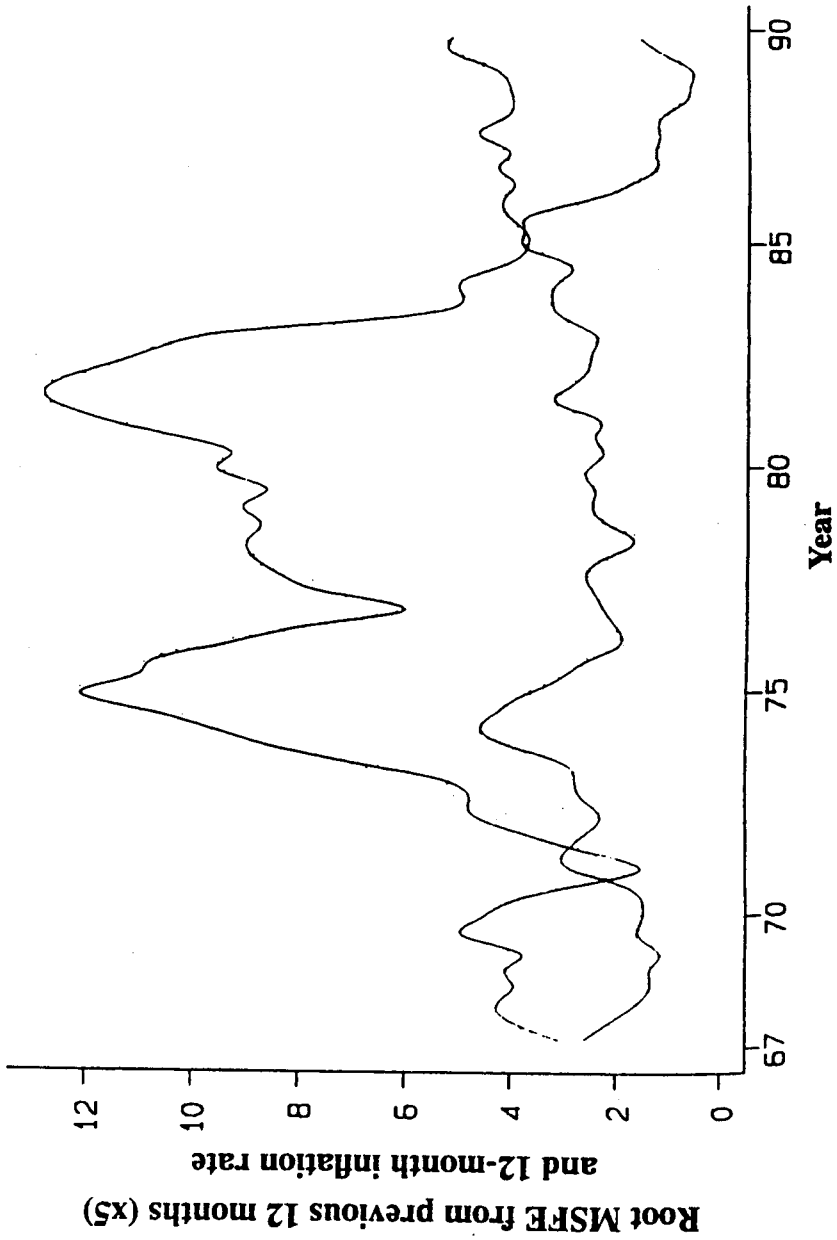

Figure 4

Determination of r and Y

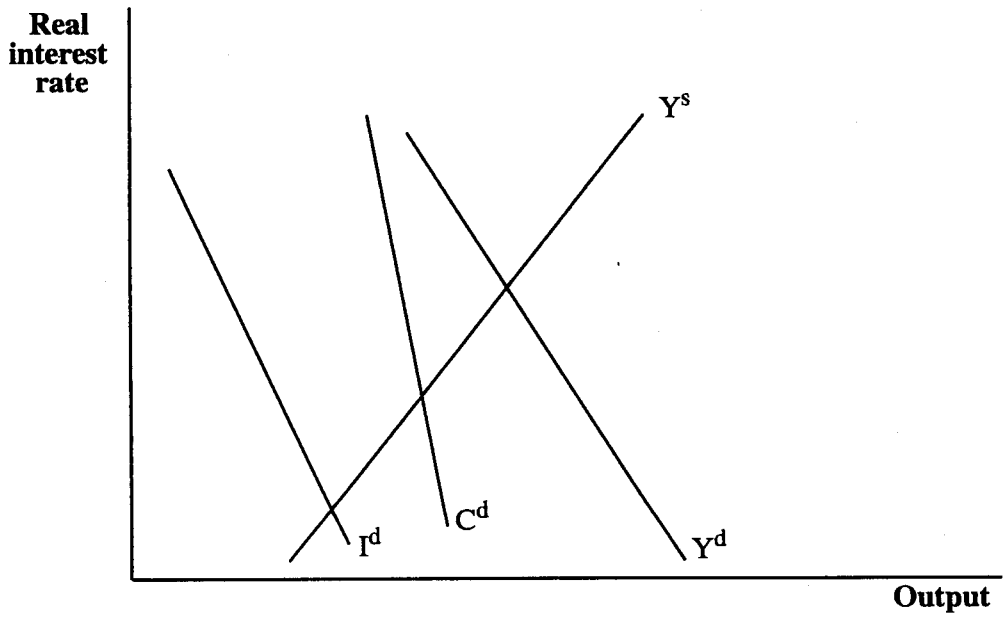

Figure 5

Determination of $r(\overline{\Pi})$ and Y

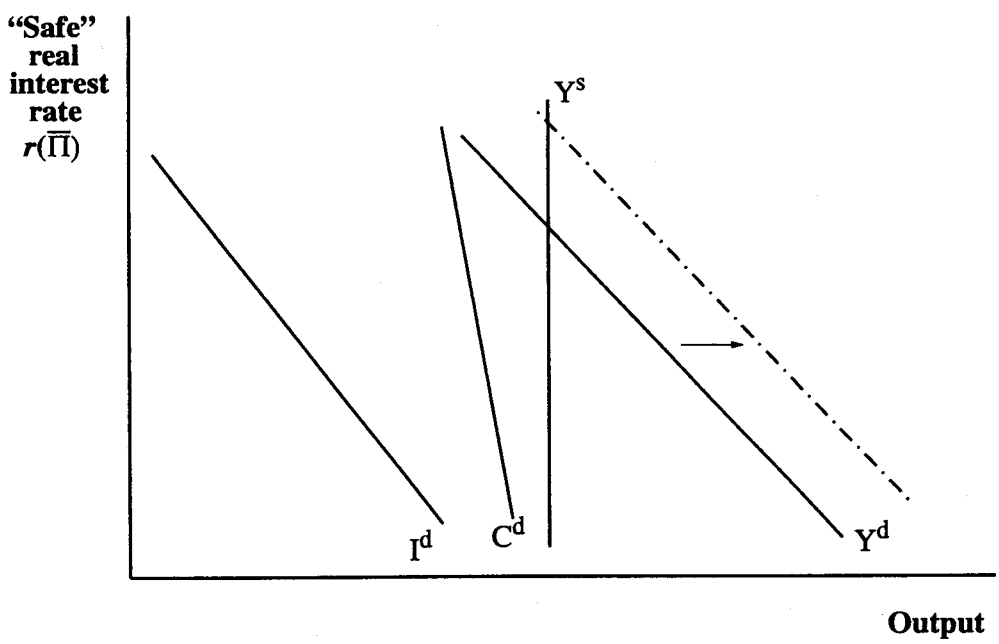

Figure 6
Spread (= Prime Rate - 90-Day Deposit Rate)
Canada: 1961-1989

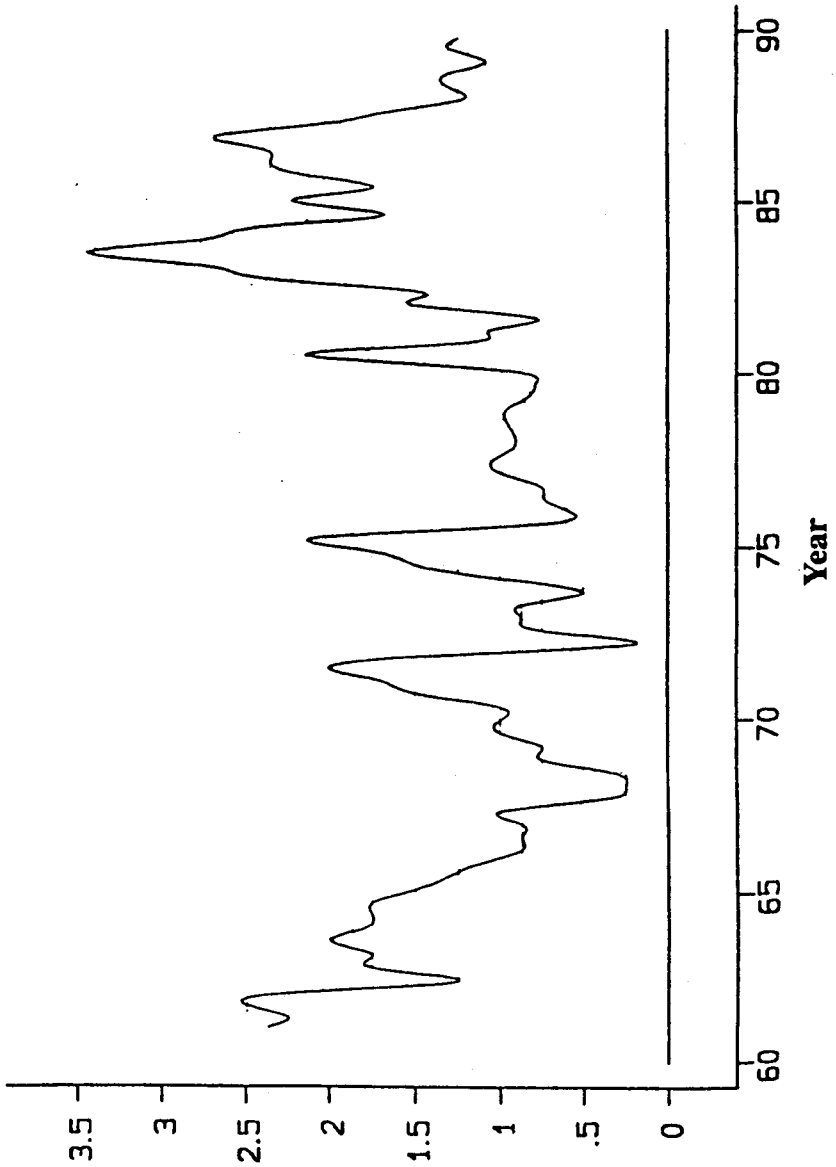

Figure 7
Private-Government Risk Premium (= Prime Rate - Treasury Bill Rate)
Canada: 1961-1989

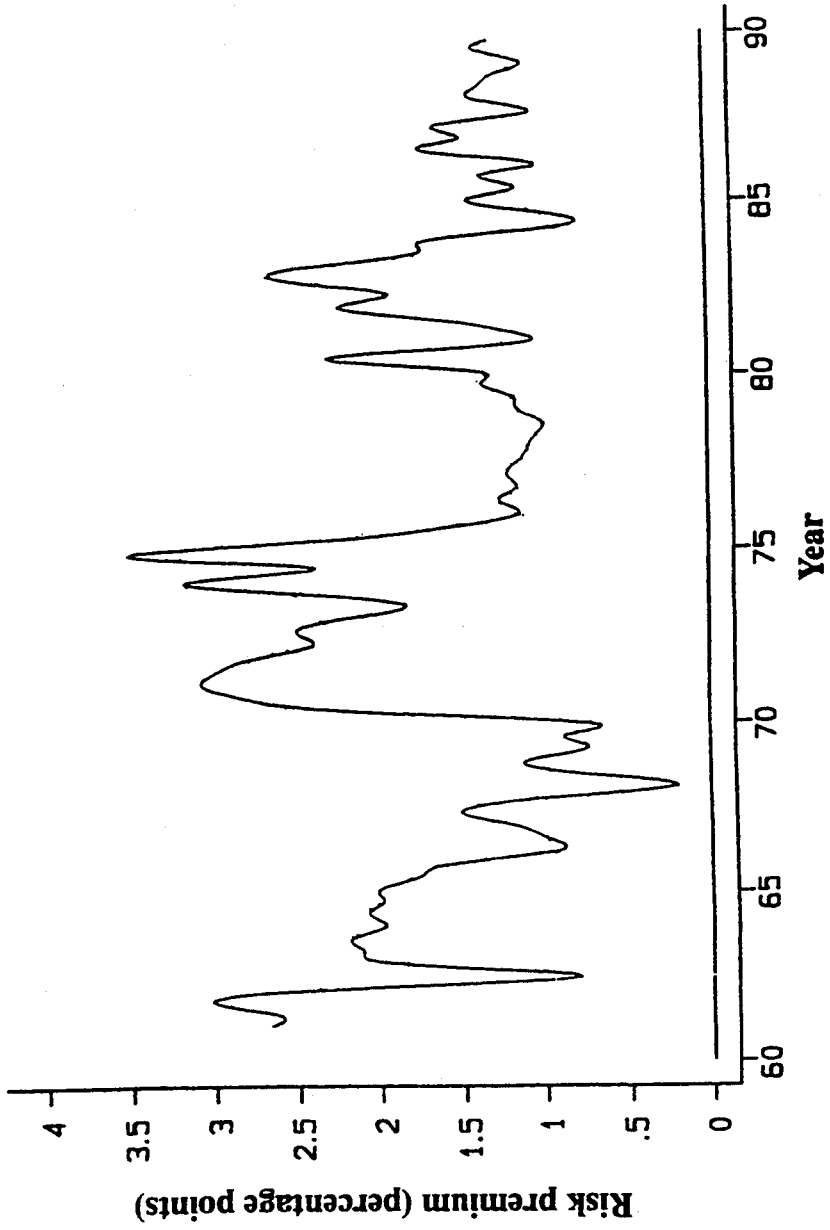

Figure 8