

Bank of Canada



Banque du Canada

Working Paper 2002-23 / Document de travail 2002-23

A large, stylized white graphic of a classical building facade with a pediment and columns, set against a grey background. The text is centered within this graphic.

How to Improve Inflation Targeting at the Bank of Canada

by

Nicholas Rowe

ISSN 1192-5434

Printed in Canada on recycled paper

Bank of Canada Working Paper 2002-23

September 2002

How to Improve Inflation Targeting at the Bank of Canada

by

Nicholas Rowe

Research Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
nrowe@bankofcanada.ca
and
Department of Economics
Carleton University
Ottawa, Ontario, Canada K1S 5B6
nick_rowe@carleton.ca

The views expressed in this paper are those of the author.
No responsibility for them should be attributed to the Bank of Canada.

Contents

Acknowledgements.....	iv
Abstract/Résumé.....	v
1. Introduction.....	1
2. Targeting Creates a Difficult Data Set.....	2
3. How to Improve Targeting: A Proposed Solution.....	6
4. Simple vs. Multiple Regressions.....	9
5. Serial Correlation.....	11
6. An Application to the Bank of Canada.....	11
6.1 Simple correlations.....	12
6.2 Multivariate regressions.....	16
6.3 Estimating the reaction function.....	21
6.4 Results.....	22
7. Conclusion.....	24
References.....	26
Appendix A: Multicollinearity.....	27
Appendix B: Previous Literature.....	30

Acknowledgements

I thank Allan Crawford, Tiff Macklem, Pierre Duguay, David Longworth, Chuck Freedman, Gabriel Srour, Bob Amano, and workshop participants at the Bank of Canada for their comments, and David Tulk and Marjorie Santos for their good help as research assistants.

Abstract

This paper shows that if the Bank of Canada is optimally adjusting its monetary policy instrument in response to inflation indicators to target 2 per cent inflation at a two-year horizon, then deviations of inflation from 2 per cent represent the Bank's forecast errors, and should be uncorrelated with its information set, which includes two-year lagged values of the instrument and the indicators. Positive or negative correlations are evidence of systematic errors in monetary policy. The econometric evidence suggests that, over the past decade, the Bank has, on average, responded optimally to indicators of inflation, being neither too aggressive nor too timid in raising or cutting the overnight rate. While responding optimally on average, however, the Bank has not responded optimally to each individual indicator. The Bank systematically overreacted to some indicators and underreacted to others. Correcting these errors could improve inflation targeting in Canada.

JEL classification: E5

Bank classification: Inflation targets; Monetary and financial indicators; Monetary policy implementation

Résumé

Cette étude montre que si la Banque du Canada règle de façon optimale son instrument d'intervention à la lumière des indicateurs de l'inflation, afin de viser un taux d'accroissement des prix de 2 % à un horizon de deux ans, les écarts de l'inflation par rapport à cette cible représentent dans ce cas les erreurs de prévision de la Banque et ne devraient pas être corrélés avec les informations à la disposition de celle-ci, y compris les valeurs de l'instrument et des indicateurs retardés de deux ans. Des corrélations positives ou négatives témoignent de la présence d'erreurs systématiques dans la formulation de la politique monétaire. Les résultats économétriques laissent croire que, au cours des dix dernières années, la Banque a généralement réagi de façon optimale à l'ensemble des indicateurs de l'inflation, les majorations ou diminutions du taux cible du financement à un jour auxquelles elle a procédé n'ayant été ni trop fortes, ni trop faibles. Toutefois, on ne peut en dire autant de sa réaction à ces indicateurs considérés isolément. En effet, les mesures prises par l'institution ont été systématiquement excessives face à certains indicateurs, et insuffisantes face à d'autres. En corrigeant ces erreurs, la Banque pourrait améliorer l'efficacité de sa stratégie de maîtrise de l'inflation.

Classification JEL : E5

Classification de la Banque : Cibles en matière d'inflation; Indicateurs monétaires et financiers; Mise en œuvre de la politique monétaire

1. Introduction

. . . observations during the period of the control cannot be used to obtain improved estimates of the parameters, which is a serious drawback.

A.W.H. Phillips (1972), in Leeson (2000, 486)

What can we learn about optimal targeting from observing the correlations between a target variable, the instrument, and the indicators?

The traditional presumption is that indicators should be strongly correlated (either positively or negatively) with the future target variable. The stronger the correlation, the better the indicators are as forecasters, and hence the better they serve as indicators of how the instrument should optimally be set to ensure that the target variable hits the target. Also, the policy-maker's instrument should be strongly correlated (either positively or negatively) with the future target variable. The stronger the correlation, the better can control of the instrument ensure control of the target variable. And, in principle, an econometrician should be able to regress the target variable on the lagged instrument and indicators, and use the resulting parameter estimates to calculate the optimal reaction function, which specifies how the instrument should be set, as a function of the indicators, to try to ensure that the target variable hits the policy-maker's target.

This traditional presumption, however, is wrong. Instead (for a constant target), when the policy-maker is setting the instrument according to the optimal reaction function, the instrument and the indicators will have zero correlation with the future target variable.

To see why the traditional presumption is wrong, assume for concreteness that there is an eight-quarter control lag in the effect of the policy-maker's instrument on the target variable. Note that a policy-maker should set its instrument, as a function of its observed indicators, so that the rational expectation of the target variable, eight quarters ahead, is equal to the target. Deviations of the target variable from the target thus represent the policy-maker's forecast errors. With rational expectations, these forecast errors should be uncorrelated with any variable in the policy-maker's current information set. The indicators, and the instrument, are part of that information set. Deviations of the target variable from the target should therefore be uncorrelated with the policy-maker's instrument and indicators, lagged eight quarters. But if deviations are uncorrelated, a regression of deviations of inflation from target, on an eight-quarter lagged instrument and indicators, would reveal only random errors. With a constant target, it is therefore impossible for an econometrician to directly estimate the reduced-form relationship between the target variable and the instrument and indicators.

Furthermore, even if the policy-maker is not setting its instrument optimally, so that deviations from the target *are* correlated with the lagged instrument and indicators, if the policy-maker is

nevertheless following *some* deterministic reaction function, so that the instrument is *some* function of the indicators, there will be perfect multicollinearity between the instrument and the indicators, making it impossible to estimate separately the coefficients on each.

This paper first explores the nature of this multicollinearity problem, and then shows how to solve it. The proposed solution provides a new and correct way of interpreting correlations between indicators and target variables, and also provides a simple and econometrically efficient way to improve the policy-maker's rule for setting its instrument.

The proposed solution involves running a regression with deviations from the target as the dependent variable, and the lagged indicators as the independent variables, *with the policy-maker's instrument omitted from the regression*. In other words, we deliberately "misspecify" the regression equation by omitting a variable we know (or at least assume) to be causally relevant. If the estimated coefficient on any indicator in the "misspecified" equation is zero, it does not mean that the indicator is useless. Rather, it means that the policy-maker is already reacting to that indicator optimally.¹ If the estimated coefficient on any indicator is of the "right" sign, then the policy-maker is *underreacting* to that indicator. If the coefficient on any instrument is of the "wrong" sign, then the policy-maker is *overreacting* to that indicator.²

I apply this proposed solution to the Bank of Canada. Over the past 10 years, the Bank has been targeting the inflation rate. We want to determine whether the Bank has been choosing its instrument as an optimal function of the available indicators. By estimating a regression of deviations of inflation from the target on various indicators, lagged several quarters (to reflect the Bank's control lag), we can determine whether the Bank has been responding optimally to each indicator.

2. Targeting Creates a Difficult Data Set

Suppose, following Rowe and Yetman (2002), that a policy-maker is targeting some variable, π_t (henceforth called the target variable), to try to keep it as close as possible to some level, π_t^* (henceforth called the target), by adjusting some instrument, R_t , observing some vector of

-
1. Reacting to an indicator optimally might mean that the indicator is indeed useless, and the policy-maker is correctly ignoring it, but it might also mean that the indicator is very useful, and the policy-maker is correctly responding to it.
 2. Strictly, we do not need to have theoretical priors that enable us to distinguish "right" signs from "wrong" signs, but it is easier intuitively to describe our results in this way, to avoid getting muddled when we talk about, say, reducing a negative response to an indicator, especially when the instrument itself may have a negative effect on the target variable.

indicator variables, X_t , where the instrument affects the target variable with a lag of k periods. How can this be done?

Conceptually, the simplest way of deciding how to set the instrument would be to estimate the reduced form of the relationship between the target variable, the instrument, and the indicators. Assuming for simplicity that this relationship is known to be linear, and to be policy-invariant (i.e., not subject to the Lucas critique), the policy-maker would attempt to estimate:

$$\pi_t = A + BR_{t-k} + CX_{t-k} + e_t, \quad (1)$$

and then use the estimated parameters \hat{A} , \hat{B} , and \hat{C} to set the instrument according to the reaction function:

$$R_t = \frac{\pi_{t+k}^* - \hat{A} - \hat{C}X_t}{\hat{B}}. \quad (2)$$

Following this reaction function would ensure that the policy-maker's expectation of the target variable, which we assume is formed rationally—conditional on the policy-maker's information set—is equal to the target:

$$E[\pi_t | \{R_{t-k}, X_{t-k}\}] = \pi_t^*. \quad (3)$$

To see that (3) follows from (1) and (2), rearrange (2) as:

$$\pi_t^* = \hat{A} + \hat{B}R_{t-k} + \hat{C}X_{t-k}. \quad (2a)$$

Then subtract (2a) from (1) to get:

$$\pi_t - \pi_t^* = A - \hat{A} + (B - \hat{B})R_{t-k} + (C - \hat{C})X_{t-k} + e_t. \quad (4)$$

We want to take expectations over the above equation, conditional on R_{t-k} and X_{t-k} . There are two ways to do so: first, we can take an objectivist or “divine” perspective, of someone who knows the true parameters A , B , and C , but sees the estimates \hat{A} , \hat{B} , and \hat{C} , as random variables subject to sampling variation; or, second, we can take the subjectivist or applied econometrician's perspective, of someone who knows their own estimates \hat{A} , \hat{B} , and \hat{C} , but is subjectively uncertain of the true parameters, A , B , and C . In both perspectives, the instrument, R_{t-k} , and indicator, X_{t-k} , are known, and we are taking expectations conditional on this information.

From the objectivist, or divine, perspective, the expectation of the target variable, conditional on R_{t-k} and X_{t-k} , will not generally equal the target, because someone who knew the true parameters

A , B , and C could know whether the instrument had been set too high or too low, given the indicators. But this means only that divine knowledge (even if it excluded knowledge of the residual e_t) could nearly always improve targeting. What matters is not whether divine knowledge could detect mistakes in targeting, but whether an applied econometrician could do so.

Taking expectations over (4) from the subjectivist or applied econometrician's perspective, we get:

$$E[\pi_t | \{R_{t-k}, X_{t-k}\}] - \pi_t^* = E[A] - \hat{A} + (E[B] - \hat{B})R_{t-k} + (E[C] - \hat{C})X_{t-k} + E[e_t] = 0. \quad (5a)$$

Provided we assume that the policy-maker's estimates are unbiased, so that $E(A) = \hat{A}$, $E(B) = \hat{B}$, and $E(C) = \hat{C}$, the expectation of the target variable is equal to the target, which confirms equation (3).

We next examine the data set created by targeting.

Substituting the reaction function (2) into the actual reduced form (1) to eliminate the instrument, the actual process generating the time series for the target variable would then become:

$$\pi_t = \frac{B}{\hat{B}}\pi_t^* + A - \frac{B}{\hat{B}}\hat{A} + \left(C - \frac{B}{\hat{B}}\hat{C}\right)X_{t-k} + e_t. \quad (6)$$

If the reduced form is estimated perfectly, so that the estimated parameters \hat{A} , \hat{B} , and \hat{C} are equal to the actual parameters A , B , and C , then (2) and (6) reduce to:

$$R_t = \frac{\pi_{t+k}^* - A - CX_t}{B}, \quad (2b)$$

$$\pi_t = \pi_t^* + e_t. \quad (6a)$$

If the target π_t^* is varying over time, it will generally³ be possible for the econometrician observing π_t , π_t^* , R_t , and X_t to estimate both the reduced-form equation (1) and the reaction function (2), and thereby to estimate both the true parameters A , B , C and the policy-maker's estimates \hat{A} , \hat{B} , and \hat{C} . But in many real-world examples, the target is instead constant.

Suppose that the target is constant over time, so that $\pi_t^* = \pi^*$. We then face a problem. If the actual processes generating the time series for the instrument and the target variable are (2b) and (6a), it will be impossible for anyone observing that data set to estimate equation (1) and thereby

3. An exception would be if π^* were perfectly correlated with a subset of the indicators.

to estimate the parameters A , B , and C . It will be impossible to estimate these parameters because there will be perfect multicollinearity between the instrument R_t and the indicator X_t .⁴ Furthermore, the right-side explanator in equation (1), namely $(BR_{t-k} + CX_{t-k})$, will have zero variance, leaving only the error term e_t to “explain” the variance in the target variable. It will be possible to estimate the reaction function (2), and thereby to estimate the ratio (\hat{C}/\hat{B}) , but this represents an estimate of the policy-maker’s estimate, not an estimate of the true ratio (C/B) . Of course, if we assume that the policy-maker knows the true ratio (C/B) , then estimating the reaction function would allow the econometrician to estimate (C/B) . But to assume that the policy-maker knows the true ratio does not make sense if the econometrician is trying to improve the policy-maker’s estimate, to improve the policy-maker’s ability to hit the target.

By the definition of a rational expectation, forecast errors of the target variable must be uncorrelated with any and all variables in the policy-maker’s information set, which includes both the instrument and the indicator. And with a target that is constant over time, *all* variance in the target variable is due to the forecast errors of the policy-maker, and so the target variable is uncorrelated with the instrument and uncorrelated with any indicator. The R^2 of the estimated equation must therefore be zero. The estimates of the parameters B and C (whether they are estimated individually in simple regressions or collectively in a multiple regression) should be totally insignificant. The whole estimated equation should look like total rubbish, even though by assumption it is true.⁵

Suppose that attempts to estimate equation (1) do not in practice result in serious multicollinearity. We could perhaps conclude that the policy-maker was making purely random mistakes in setting the instrument, so that the actual reaction function is

$$R_t = \frac{\pi^* - \hat{A}}{\hat{B}} - \frac{\hat{C}}{\hat{B}} X_t + v_t, \quad (2c)$$

where v_t is a random error in policy-making uncorrelated with e_{t+k} . If the variance of v_t is large enough, then multicollinearity will not be a serious problem, and equation (1) can be estimated without difficulty. But an alternative, and surely more plausible, assumption is that the policy-maker is reacting to a larger set of indicators than the econometrician has included in the estimated equation. What appear to the econometrician to be random errors by the policy-maker

-
4. Appendix A explores the multicollinearity problem more fully.
 5. John Chant has suggested to me that it is interesting to reinterpret “Goodhart’s Law” in this light. Even if the underlying structural relationship between a target variable and an instrument or indicator is unaffected by the policy-maker’s use of that relationship, it will appear to the econometrician as if that relationship disappears. In my interpretation, then, it is not the underlying relationship itself, but rather Goodhart’s Law, that is the statistical illusion.

are in fact responses to indicators the policy-maker observes but which the econometrician does not observe. Provided these “hidden” indicators are indeed genuine indicators of the target variable, and therefore represent omitted variables, they will be reflected in the econometrician’s error term, e_t , so the instrument will be correlated with e_t . This will cause the estimated coefficient on the instrument to be biased towards zero, and the estimated coefficients on the remaining indicators will also be biased.

In other words, if the policy-maker is truly trying to hit the target, and not just making random exogenous changes in the instrument, then there *must* be an underlying problem of multicollinearity. Failure to find multicollinearity in practice would merely mean that the econometrician trying to estimate equation (1) has omitted some indicators that the policy-maker believes to be relevant. Unless the econometrician can know a priori that the policy-maker is wrong, so that these omitted variables affect the target variable only via their effect on the policy-maker’s instrument, the estimated coefficients must be biased.

If the policy-maker were simply to follow some random, exogenous process for setting the instrument, it would be possible in principle for an econometrician to estimate equation (1) and thus to estimate an optimal reaction function for the policy-maker. But if the policy-maker is instead already following some deterministic reaction function, which it believes (whether correctly or incorrectly) to be optimal, it will be impossible for the econometrician to estimate equation (1). Targeting thus appears to sow the seeds that prevent its own improvement, because it creates a data set that seems to destroy the possibility of estimating the very parameters that the policy-maker needs to know to calculate the optimal reaction function. If econometricians cannot estimate equation (1), how can they propose a better reaction function for the policy-maker?

3. How to Improve Targeting: A Proposed Solution

How can targeting be improved? The more accurate the policy-maker’s estimates of the parameters A , B , and C , the smaller will be the variance of deviations of the target variable from the target. We therefore need to improve the accuracy of those estimates.

An econometrician trying to estimate equation (1) directly would have difficulty estimating the parameters B and C . If the policy-maker were setting the instrument according to equation (2) (or according to any deterministic linear reaction function), the instrument would be perfectly correlated with the indicator, giving perfect multicollinearity. Even if there were some small amount of noise in the reaction function, so that it is (2c) instead of (2), the high correlation between the instrument and the indicator would mean that the estimates of B and C would have

large standard errors. This section provides an alternative way to improve the relevant estimates and thereby improve targeting.⁶

The key to solving the problem of multicollinearity is to drop the instrument from the estimated equation. Because of the resulting omitted-variable bias, the resulting parameter estimates will, of course, be biased, *but only if they are interpreted as being what they cannot be*: estimates of the effect of changes in the indicator on the target variable, *ceteris paribus*. If, instead, they are correctly interpreted as estimates of how changes in the indicator affect the target variable, *including the feedback effect via policy-induced changes in the instrument*, then they are unbiased estimates. In other words, the estimated coefficients are estimates of the *total* derivative, and not the *partial* derivative, of changes in the target with respect to changes in the indicator.

Furthermore, these estimates can also be interpreted as showing how the policy rule needs to be revised for optimal targeting. The payoff to this approach is that, with multicollinearity eliminated by excluding the instrument from the regression, these estimates should have smaller standard errors than those obtained by estimating equation (1) directly, and therefore will give much more accurate estimates of how the existing policy rule needs to be revised.

Instead of trying to estimate equation (1), the econometrician should run a regression of deviations of the target variable from the target, on the indicators, *with the instrument excluded from the regression equation*:

$$\pi_t - \pi^* = L + MX_{t-k} + e_t. \quad (7)$$

Comparing the terms in equation (7) with terms in equation (6), with the target assumed constant over time,

$$\pi_t - \pi_t^* = \left[A - \pi^* + \frac{B(\pi^* - \hat{A})}{\hat{B}} \right] + \left[C - B \left(\frac{\hat{C}}{\hat{B}} \right) \right] X_{t-k} + e_t, \quad (6)$$

the constant terms are $L = [A - \pi^* + B(\pi^* - \hat{A})/\hat{B}]$ and the slope terms are $M = [C - B(\hat{C}/\hat{B})]$. To see the intuition, note that the slope term, M , can be interpreted as the total derivative of π_t with respect to X_{t-k} , which equals the partial derivative, C , plus the derivative of R_{t-k} with respect to X_{t-k} , which is (\hat{C}/\hat{B}) , times the partial derivative of π_t with respect to R_{t-k} , which is B .

6. In this paper, I claim that direct estimation of the reduced-form relation between the target variable, instrument, and indicators will run afoul of the multicollinearity problem. Estimation of a more structural model, in which the instrument has a lagged effect on some intermediate variable, which in turn has a lagged effect on the target variable, may nevertheless be possible under some circumstances. Full exploration of this question is outside the scope of this paper.

Remember that the reaction function is (rearranging equation (2)):

$$R_t = \frac{\pi^* - \hat{A}}{\hat{B}} - \frac{\hat{C}}{\hat{B}} X_t. \quad (2d)$$

Examining (2d), we see that the policy-maker does not need to know the parameters A , B , and C themselves, but rather the ratios $(\pi^* - A)/B$ and (C/B) , because those ratios define the constant term and the slope in the reaction function.

Having generated estimates \hat{L} and \hat{M} from estimating equation (7), the policy-maker can then revise its reaction function according to the following formulae (where \tilde{A} , \tilde{B} , and \tilde{C} indicate the revised estimates):

$$\frac{\pi^* - \tilde{A}}{\tilde{B}} = \frac{\pi^* - \hat{A}}{\hat{B}} - \frac{\hat{L}}{B}, \quad (8)$$

$$\frac{\tilde{C}}{\tilde{B}} = \frac{\hat{C}}{\hat{B}} + \frac{\hat{M}}{B}. \quad (9)$$

If the estimated coefficient, \hat{M} , is zero, then the deviation of the target variable from the target is orthogonal to the indicator, which means that the existing reaction function already incorporates an optimal response of the instrument to that indicator. Therefore, no revision to the reaction function is required; i.e., $(\tilde{C}/\tilde{B}) = (\hat{C}/\hat{B})$.

If \hat{M} is positive, however (and assuming that B , \hat{B} , and \hat{C} are all positive for simplicity of exposition), then the reaction function needs to be revised to strengthen the response of the instrument to the indicator, so that $(\tilde{C}/\tilde{B}) > (\hat{C}/\hat{B})$. To see this intuitively, suppose that the policy rule initially failed to incorporate any response to this indicator, so that $(\hat{C}/\hat{B}) = 0$. Then the instrument would be uncorrelated with the indicator, and the estimate \hat{M} from estimating the “misspecified” equation would be an unbiased estimate of the parameter C . In other words, the coefficient on the “misspecified” equation would have the theoretically “right” sign and magnitude. As the reaction function changed to respond to the indicator (that is, as (\hat{C}/\hat{B}) approached (C/B)), the coefficient \hat{M} would get progressively smaller, as bigger and bigger changes in the instrument offset more and more of the effects of the indicator, until eventually \hat{M} would become zero when (\hat{C}/\hat{B}) equalled (C/B) .

If the reaction function overreacts to the indicator—if, in other words $(\hat{C}/\hat{B}) > (C/B)$ —then \hat{M} would be negative, because the overreaction from the instrument would cause changes in the indicator to have the “wrong” sign in their effect on the target variable.

Thus, by estimating the sign of the parameter M in equation (7), we can determine whether the policy-maker's reaction function is suboptimal, and if it is suboptimal, know the *direction* in which it must change to improve. But we cannot say *how much* it must change to become optimal, because even if we can estimate M , we need to know (M/B) to tell how much to revise the policy-maker's reaction function. We cannot estimate (M/B) because we cannot estimate B using this proposed solution. Intuitively, we can estimate the effect (M) of a unit change in the indicator on the target variable (in the *total* derivative sense), but unless we know how powerful the instrument is (B), we cannot say how much more the instrument needs to respond to the indicator to eliminate the total effect on the target variable.

4. Simple vs. Multiple Regressions

The proposed solution to improving targeting assumes that X_t represents a vector of indicators, and that the econometrician estimating the forecasting equation (7) includes all the relevant indicators. In practice, the econometrician may not have data on all the relevant indicators, or may for other reasons ignore some indicators. In fact, it seems very likely that a real world econometrician will in practice include only a subset of the set of indicators to which the policy-maker responds, or to which the policy-maker ought optimally to respond. An extreme example of this potential problem is an econometrician estimating (7) using only a single indicator, and thus estimating a simple regression where a multiple regression would appear to be appropriate.

Fortunately, the proposed solution is robust to missing indicators. Since in practice there will almost always be missing indicators, the proposed solution's robustness to this problem makes an even stronger case for its use. I now show that the policy advice from even a simple regression is good advice, though policy advice from the multiple regression, including all the relevant indicators, would be better.

Suppose, for simplicity, that there are two relevant indicators, X_1 and X_2 . Equation (6) now reads:

$$\pi_t - \pi_t^* = \left[A - \pi^* + \frac{B(\pi^* - \hat{A})}{\hat{B}} \right] + \left[C_1 - B \left(\frac{\hat{C}_1}{\hat{B}} \right) \right] X_{1t-k} + \left[C_2 - B \left(\frac{\hat{C}_2}{\hat{B}} \right) \right] X_{2t-k} + e_t. \quad (6b)$$

Then suppose that the econometrician ignores the second indicator, X_2 , and runs a simple regression on X_1 only:

$$\pi_t - \pi_t^* = L + M X_{1t-k} + e_t. \quad (7a)$$

Suppose also that the two indicators are correlated according to:

$$X_{2t} = DX_{1t} + v_t. \quad (10)$$

Substituting (10) into (6b) gives:

$$\pi_t - \pi_t^* = \left[A - \pi^* + \frac{B(\pi^* - \hat{A})}{\hat{B}} \right] + \left\{ \left[C_1 - B \left(\frac{\hat{C}_1}{\hat{B}} \right) \right] + \left[C_2 - B \left(\frac{\hat{C}_2}{\hat{B}} \right) \right] D \right\} X_{1t-k} + \left[C_2 - B \left(\frac{\hat{C}_2}{\hat{B}} \right) \right] D v_{t-k} + e_t. \quad (11)$$

Note that since v_{t-k} is uncorrelated with X_{1t-k} , a simple regression of the deviation from the target, $\pi_t - \pi^*$, on the first indicator, X_{1t-k} , will yield a slope coefficient, M , that is an unbiased estimate of $\{C_1 - B(\hat{C}_1/\hat{B}) + [C_2 - B(\hat{C}_2/\hat{B})]D\}$. Suppose that M is positive. This means that the policy-maker's reaction function is not optimal. But it does not tell us whether the fault lies in the policy-maker's reaction to the first indicator, the second indicator, or both indicators. For example, the policy-maker might be reacting optimally to the first indicator, that is to say $(\hat{C}_1/\hat{B}) = (C_1/B)$, and underreacting to the second indicator, $(\hat{C}_2/\hat{B}) < (C_2/B)$, and the two indicators are positively correlated, $D > 0$.

On estimating the simple regression, and observing a positive slope coefficient, M , the econometrician would conclude that the policy-maker is underreacting to the first indicator, and advise it to increase (\hat{C}_1/\hat{B}) . On the face of it, this might appear to be bad advice, because by assumption $(\hat{C}_1/\hat{B}) = (C_1/B)$, so the policy-maker is already reacting optimally to the first indicator; the fault lies instead in the policy-maker's reaction to the second indicator.

Nevertheless, taking the econometrician's advice *would* improve inflation targeting. Increasing (\hat{C}_1/\hat{B}) for a given (\hat{C}_2/\hat{B}) would reduce $\{C_1 - B(\hat{C}_1/\hat{B}) + [C_2 - B(\hat{C}_2/\hat{B})]D\}$, and thereby reduce the variance of deviations from the target conditional on X_{1t-k} and v_{t-k} , which is equivalent to reducing the variance of deviations from the target conditional on X_{1t-k} and X_{2t-k} . Of course, increasing (\hat{C}_2/\hat{B}) and holding (\hat{C}_1/\hat{B}) constant (which is the advice that would follow from a multiple regression) would be better advice, because, in terms of equation (11), this would reduce the coefficient on X_{1t-k} and *would also reduce the coefficient on v_{t-k}* , which would further reduce the variance of the deviation from the target.

Intuitively, given that the policy-maker's reaction function is suboptimal with respect to the second indicator, the second-best optimal reaction to the first indicator may be different from the first-best, and will be different if the two indicators are correlated. The simple regression does not tell us in which direction to change the reaction function towards the first-best, but it does tell us in which direction to change the reaction function towards the second-best, holding the reaction to all other indicators constant.

5. Serial Correlation

Suppose that an unforecastable shock hits the target variable at time t , causing the target variable to rise above the target at time t . It is very likely that the effect of that shock will persist, and that the target variable will also be above target in period $t+1$. In other words, the residuals e_t in the forecasting equation will probably be serially correlated. Serial correlation will not bias the estimated coefficients in the forecasting equation, but it will bias the estimated t -statistics and standard errors, so a coefficient that is statistically significantly different from zero may appear to be insignificant, or vice versa.

A standard way to handle serially correlated residuals is to estimate the equation assuming that the errors are an autocorrelated process, like AR(1), for example. But an AR(1) process would violate the requirement that, under optimal targeting, the deviations of the target variable from the target, being forecast errors, must be unforecastable at the control lag. A positive shock in an AR(1) process never dies out completely over time; it approaches zero asymptotically. Therefore, estimating the forecasting equation assuming any AR process would implicitly contradict our method.

The correct alternative way to handle serial correlation is to assume that the errors are a moving-average process. If we are using quarterly data, and the control lag is eight quarters, for example, we could allow the errors to be MA(1), or MA(2), or anything up to MA(7), while maintaining our assumption that the eight-quarter-ahead error is unforecastable and mean zero. In principle, we should allow the errors to be MA(7), but limitations on degrees of freedom could make this impractical.

6. An Application to the Bank of Canada

We apply the above method to the Bank of Canada and its attempt to target inflation over the past 10 years. In February 1991, the Bank announced that the target inflation rate for December 1992 would be 3 per cent, declining to 2 per cent by December 1995, and remaining at 2 per cent thereafter.⁷

7. Strictly, the Bank announced the target for inflation as the midpoint of a range, plus or minus one percentage point around the midpoint. The simple method adopted in this paper requires a point target, and we choose to consider the midpoint of that range to be the point target.

6.1 Simple correlations

Using quarterly data, we calculate simple correlations between deviations of inflation from the target, for the period 1992Q4 to 2001Q3, and lagged values of the Bank's instrument (the overnight rate) and various indicators. We examine two different definitions of the target variable: total CPI inflation, and CPIXFET inflation (the old definition of core inflation, which was used up to May 2001; it excludes food, energy, and the effect of changes in indirect taxes). In each case, inflation was defined as year-over-year. These correlations are reported in Table 1 (for CPIXFET) and Table 2 (for total CPI). Remember that, if monetary policy is adjusting the instrument to respond to these indicators optimally, the instrument, and each indicator, should be uncorrelated with future deviations from the target, at the control lag. If the control lag is eight quarters, for example, then deviations of inflation from the target should be uncorrelated with the instrument and indicators lagged eight or more quarters. Correlation coefficients bigger than 0.25 are statistically significant at the 90 per cent level, and coefficients bigger than 0.3 are significant at the 95 per cent level.

Table 1. Simple correlations with deviations of CPIXFET inflation from target 1992Q4–2001Q3

Indicator	Lagged 4	Lagged 5	Lagged 6	Lagged 7	Lagged 8	Lagged 9	Lagged 10
Total CPI inflation	-0.13	-0.12	-0.07	-0.02	-0.04	-0.14	-0.23
CPIXFET inflation	-0.13	-0.02	0.06	0.03	0.00	-0.14	-0.28
Exchange rate depreciation	-0.16	-0.18	-0.14	-0.07	0.09	0.32	0.43
U.S. inflation	0.12	-0.01	-0.12	-0.18	-0.20	-0.21	-0.22
M1 growth	0.05	-0.01	-0.11	-0.19	-0.27	-0.22	-0.11
Currency growth	-0.01	0.15	0.30	0.36	0.35	0.40	0.39
Output gap	0.45	0.36	0.24	0.13	0.03	-0.06	-0.14
Unemployment	-0.31	-0.24	-0.13	-0.04	0.05	0.10	0.12
Commodity price inflation	0.33	0.32	0.22	0.06	-0.12	-0.23	-0.27
Labour cost inflation	-0.14	-0.16	-0.24	-0.33	-0.27	-0.25	-0.15
U.S. federal funds	0.32	0.15	-0.02	-0.14	-0.22	-0.27	-0.28
Longbond	0.14	0.07	-0.01	-0.10	-0.14	-0.10	-0.10
Overnight rate	0.23	0.13	0.07	-0.00	-0.03	-0.01	-0.03

Table 2. Simple correlations with deviations of total CPI inflation from target 1992Q4–2001Q3

Indicator	Lagged 4	Lagged 5	Lagged 6	Lagged 7	Lagged 8	Lagged 9	Lagged 10
Total CPI inflation	-0.02	-0.06	-0.08	-0.11	-0.18	-0.26	-0.26
CPIXFET inflation	-0.19	-0.15	-0.15	-0.22	-0.26	-0.35	-0.37
Exchange rate depreciation	-0.34	-0.35	-0.25	-0.06	-0.11	-0.30	-0.37
U.S. inflation	-0.10	-0.29	-0.43	-0.49	-0.50	-0.46	-0.43
M1 growth	-0.04	-0.05	-0.14	-0.17	-0.17	-0.07	-0.09
Currency growth	-0.11	-0.01	0.10	0.16	0.20	0.26	0.24
Output gap	0.71	0.65	0.51	0.35	0.19	0.03	-0.06
Unemployment	-0.66	-0.62	-0.49	-0.34	-0.21	-0.08	-0.01
Commodity price inflation	0.37	0.23	0.02	-0.14	-0.33	-0.39	-0.41
Labour cost inflation	0.10	0.04	-0.07	-0.12	-0.09	-0.08	-0.03
U.S. federal funds	0.55	0.44	0.29	0.16	0.06	-0.04	-0.07
Longbond	-0.18	-0.29	-0.36	-0.38	-0.34	-0.28	-0.25
Overnight rate	0.22	0.05	-0.04	-0.08	-0.14	-0.09	-0.11

The first noteworthy result is that the Bank's instrument (the overnight rate) has a fairly high positive correlation with four-quarter-ahead deviations of inflation from the target, but at longer lags the correlation drops quickly towards zero, and indeed becomes negative. This is exactly what we would expect if high inflation causes the Bank to raise the overnight rate, but the higher interest rate takes several quarters to bring inflation back down to the target. For total CPI inflation, the correlation turns negative at six quarters, and for CPIXFET inflation the correlation turns negative at seven quarters. This suggests that the control lag in monetary policy may be slightly shorter than the eight-quarter lag normally assumed. Erring on the side of caution, we shall henceforth assume that the control lag is indeed eight quarters, so that all correlations should be zero at a lag of eight or more quarters, if monetary policy is responding optimally to indicators.

At lags longer than the control lag, optimal policy would result in a zero correlation between the overnight rate and deviations of inflation from the target. A positive correlation would imply that the Bank had not been sufficiently aggressive in raising interest rates in response to inflation indicators. A negative correlation would imply that the Bank had been too aggressive. The roughly zero correlation we observe means that, on average, the Bank has been responding about optimally to average indicators of inflation, neither overreacting nor underreacting. This is one of the major findings of this paper.

But responding optimally to *average* indicators of inflation does not mean that the Bank was responding optimally to *each* indicator of inflation. It could have been overreacting to some indicators and underreacting to others, causing inflation to fluctuate more than it would if the Bank responded optimally to each indicator. To examine whether the Bank was responding optimally to each indicator, we must examine whether each indicator in turn is correlated or uncorrelated with future inflation.

Canadian inflation (year-over-year), both CPIXFET and total CPI, tends to be negatively correlated with future deviations of inflation from the target (both CPIXFET and total CPI). This is especially true, and the correlations tend to be statistically significant, at longer lags of nine or ten quarters. This result suggests that the Bank was overreacting to current inflation, raising the overnight rate by too much, so that future inflation tends to fall below target.

U.S. CPI inflation (year-over-year) is also negatively correlated with future deviations of Canadian inflation from the target, significantly for total inflation. This suggests that the Bank was overreacting to U.S. inflation.

The (year-over-year) rate of depreciation of the Can\$/US\$ exchange rate gives nearly the same results for both total CPI and CPIXFET inflation targets. At lags of eight quarters and above, the rate of exchange rate depreciation is positively correlated with future deviations of inflation from the target, with the correlations becoming highly significant at nine- and ten-quarter lags. This result implies that the Bank was underreacting to longer-run exchange rate depreciation.

The (year-over-year) M1 growth rate tends to be negatively correlated with deviations of inflation from the target, significant at the 10 per cent level for the CPIXFET inflation target at an eight-quarter lag, and less significant at other lags and for the total inflation target. This suggests that the Bank may have been overreacting to the M1 growth rate.

For CPIXFET inflation, the indicator that has the biggest correlation with future deviations from the target is the (year-over-year) growth rate in the stock of currency outside banks. Indeed, the correlation is highly significant at lags of seven quarters and above. If, instead, we assume that the target variable is total CPI inflation, rather than CPIXFET inflation, the correlations are lower, though still positive, and significant at nine- and ten-quarter lags. This result strongly suggests that the Bank was underreacting to changes in the currency growth rate.⁸

8. I confess that currency was included as an indicator in this study only by sheer accident. One possible explanation of our result is that the growth rate of currency is a measure of inflationary pressures from the underground economy. Perhaps the other indicators the Bank watches happen to exclude the underground economy. If the Bank does not pay much attention to currency, it is understandable that the Bank could be making mistakes in inflation targeting by ignoring this indicator.

The unemployment rate and the output gap (which have a very high (-0.9) negative correlation with each other) have almost identical patterns (with the signs reversed) of correlations with future deviations of inflation from the target. We can therefore best think of them as being one single indicator. Since the unemployment rate is closest to real-time data, let us concentrate on the unemployment rate and ignore the output gap. The unemployment rate is negatively correlated with future deviations of total inflation from the target, but the correlations become weaker and statistically insignificant at longer lags. The unemployment rate has a weaker correlation (sometimes positive, sometimes negative) with future deviations of CPIXFET inflation from the target. Our policy recommendation depends on which of the two measures of inflation we see as being the Bank's target. If total CPI inflation is the target variable, the Bank should react more strongly to the unemployment rate.

Unit labour cost inflation (year-over-year) tends to be negatively correlated with future deviations of inflation from target. The correlation is statistically significant for CPIXFET inflation, but not for total CPI inflation. This suggests that the Bank has overreacted to unit labour cost inflation.

The U.S. federal funds rate is negatively and significantly correlated with future deviations of CPIXFET inflation from the target. This suggests that the Bank has been overreacting to the U.S. federal funds rate, raising the overnight rate by too much when the U.S. Federal Reserve raises its interest rate, and thereby causing future CPIXFET inflation to fall below target. But the U.S. federal funds rate has a much smaller and insignificant correlation (sometimes positive, sometimes negative) with future deviations of total inflation from the target, which suggests that the Bank has been reacting roughly optimally to changes in this indicator, if total CPI inflation is the target variable.

The interest rate on 10-year bonds is negatively and significantly correlated with future deviations of total inflation from the target. It has a smaller, insignificant, though still negative correlation with future CPIXFET inflation. If total CPI inflation is the target, the Bank was overreacting to 10-year bond rates.

Commodity price inflation (year-over-year, measured in U.S. dollars), except for shorter lags, is negatively and significantly correlated with future deviations from target, both for total and for CPIXFET inflation. This result suggests that the Bank was overreacting to commodity price inflation.

The above simple correlations were calculated on the 1992Q4 to 2001Q3 data set. A statistically significant non-zero correlation strongly suggests that the Bank was making systematic mistakes in monetary policy. The Bank's reaction function over those years was either under- or

overreacting to some of the indicators. During the early years of inflation targeting, when the Bank clearly undershot its inflation target, it could be argued that the Bank was learning how to operate under the new regime (or that the rest of the economy was learning how the Bank was operating). Perhaps the reaction function changed after this initial transition period. If so, lessons learned from observing systematic mistakes over the whole period of inflation targeting would be misapplied if directed at the current, improved, reaction function. To test this possibility, I calculated the same simple correlations over the 1995Q1 to 2001Q3 data set (not reported). If the Bank had improved its reaction function from the early years of inflation targeting, we should have expected the correlations to become smaller in the more recent data set. This did not happen. On average, the correlations got bigger (although a bigger correlation is not necessarily more statistically significant when the degrees of freedom are smaller). In particular, the correlations that were statistically significant on the full data set nearly all got bigger in the smaller, more recent data set. The one exception was that the negative correlation between the U.S. federal funds rate and total CPI inflation (at the eight-quarter lag) became much smaller in the more recent data set.

6.2 Multivariate regressions

Tables 1 and 2 report simple rather than multivariate correlations between the indicators and deviations from the target. As discussed in section 5, all of the preceding policy recommendations should therefore be taken as isolated recommendations, and cannot be understood as being made jointly. For example, if exchange rate depreciation is positively correlated with future deviations from the target, and currency growth is also positively correlated with future deviations from the target, we can conclude that the Bank should react more strongly to the exchange rate, *or* that the Bank should react more strongly to currency, but we cannot conclude that the Bank should react more strongly to both indicators jointly. The simple correlation may be positive, but the partial correlation might nevertheless be negative.

To explore the multivariate relationship between indicators and the target variable, I ran ordinary least squares (OLS) regressions, with the dependent variable being the deviation of inflation from the target, and the independent variables being various indicators, lagged eight quarters. The eight-quarter lag reflects the consensus at the Bank that it takes roughly eight quarters for changes in the Bank rate (the Bank's instrument) to have a major impact on the rate of inflation (the Bank's target variable).⁹

9. Strictly, the Bank sees its instrument as having a distributed lag effect on the target, but I have not figured out how to incorporate this assumption properly.

It is not obvious a priori which indicators to include in the regression equation. Ideally, we would want to include all potential indicators of future inflation, to examine whether the Bank was reacting optimally to each of them. Initially, therefore, I included the full set of potential indicators in the estimated regression equations. But if a set of two or more indicators were highly correlated with each other, multicollinearity could make each indicator statistically insignificant, even if the set of indicators as a whole were useful in forecasting future inflation. Usually, economic theory tells us which variables should and should not be included in a regression equation. But in this case, even if economic theory can perhaps tell us which variables are indicators of inflation, to which the Bank should respond, economic theory is silent about whether the Bank will underreact, overreact, or react optimally to that indicator. A priori, for someone who knows nothing about the actual conduct of monetary policy, a reasonable point expectation would be that the Bank would react optimally, with an equal 50-50 chance of overreacting or underreacting. Wishing to have some objective procedure for deciding which variables to include and exclude from the regression equation, I allowed the RATS program's stepwise regression procedure to do this for me, having set the cut-in and cut-out significance levels at 50 per cent.

Two indicators that perhaps should have been included were deliberately excluded. The first is the output gap. This variable seemed promising as an indicator in the simple correlations, but I excluded it because it is doubtful whether it represents real-time data—i.e., the data available to the Bank at the time it set its instrument, because both GDP data and estimates of potential output can be heavily revised. Moreover, the unemployment rate, on which data are available with a short lag, seemed a very good substitute for the output gap empirically. The second deliberately excluded indicator is the yield spread. This variable also seemed very promising in preliminary estimations, but it made the results very hard to interpret, because the yield spread is a composite indicator that reflects public expectations of future inflation, the phase of the business cycle, and the difference between current and future settings of the Bank's instrument. Excluding these two possible indicators avoids these difficulties. Future work, however, should address these issues more directly.

The "forecasting" regressions can tell us whether the Bank was underreacting, overreacting, or reacting optimally to those various indicators. In addition, I estimated an actual reaction function for the Bank. From the estimated actual reaction function, together with the results of the forecasting regressions, we can get some idea of what the optimal reaction function should look like.

Table 3 shows the estimated forecasting regression, with deviations of CPIXFET inflation from the target as the dependent variable, using all indicators, lagged eight quarters.

Table 3. Forecasting deviations of CPIXFET inflation from the target, eight-quarter lag

Indicator	Coefficient	Std. error	Significance
Constant	-1.13	1.38	0.42
Total CPI inflation	0.01	0.13	0.92
CPIXFET inflation	0.19	0.23	0.43
Exchange rate depreciation	-0.02	0.02	0.25
U.S. inflation	-0.15	0.11	0.19
M1 growth	-0.04	0.02	0.04
Currency growth	0.10	0.05	0.07
Unemployment	0.19	0.13	0.16
Commodity price inflation	0.01	0.01	0.12
Labour cost inflation	-0.10	0.04	0.02
U.S. federal funds	0.21	0.13	0.12
Longbond	-0.32	0.10	0.00
$r\bar{a}r^2$	0.76	D.W.	1.72

Table 4 shows the estimated regression with deviations of total CPI inflation from the target as the dependent variable, using all indicators, lagged eight quarters.

Table 4. Forecasting deviations of total CPI inflation from the target, eight-quarter lag

Indicator	Coefficient	Std. error	Significance
Constant	2.42	4.13	0.56
Total CPI inflation	0.16	0.40	0.69
CPIXFET inflation	0.62	0.70	0.39
Exchange rate depreciation	0.01	0.05	0.79
U.S. inflation	-0.15	0.34	0.65
M1 growth	-0.03	0.05	0.50
Currency growth	0.05	0.16	0.76
Unemployment	-0.19	0.39	0.64
Commodity price inflation	-0.02	0.02	0.37
Labour cost inflation	-0.31	0.13	0.02
U.S. federal funds	0.36	0.38	0.36
Longbond	-0.48	0.29	0.11
$r\bar{a}r^2$	0.62	D.W.	1.41

Note that very few indicators are statistically significant, at conventional levels. One exception is unit labour cost inflation, which is negative and highly significant in both equations. (This negative correlation also appeared in the simple correlations of Tables 1 and 2.) This result suggests that the Bank overreacted to unit labour cost inflation. A second exception is the interest rate on long bonds, which is negative and significant in the equation for deviations of CPIXFET inflation from the target, suggesting that the Bank also overreacted to this indicator.

I repeated the above regressions using the stepwise regression procedure to eliminate statistically insignificant indicators.

Respectable econometricians normally frown upon using the stepwise regression procedure. They argue that economic theory should dictate which variables are included or excluded from a regression equation, and argue that excluding a variable that theory says should be included, even if it is statistically insignificant, will bias the estimated coefficients on the other variables. But respectable econometricians would be wrong in this particular case. Before running these regressions, I certainly knew less about the Bank's optimal reaction function than did the Bank itself. My prior point expectation is that the Bank's reaction function is optimal, with a 50 per cent probability that the Bank overreacts to any particular indicator, and a 50 per cent probability that it underreacts. In other words, my prior theory tells me absolutely nothing about whether and in what direction the Bank is making mistakes in responding to indicators, so it cannot tell me which variables to include or exclude from the forecasting equation, and does not predict what the sign on any variable should be. Moreover, as I showed in section 3, even if some indicators are excluded that ought to be included, the policy advice based on the smaller set of indicators is still valid, in the second-best sense.

Table 5 shows the estimated regression with deviations of core inflation from the target as the dependent variable, using a stepwise regression procedure with a 50 per cent significance level cut-in and cut-out.

Table 5. Forecasting deviations of CPIXFET inflation, eight-quarter lag, 50 per cent stepwise

Indicator	Coefficient	Std. error	Significance
Constant	-1.17	1.28	0.37
M1 growth	-0.04	0.01	0.01
Currency growth	0.10	0.05	0.05
Labour cost inflation	-0.10	0.03	0.01
U.S. federal funds	0.21	0.12	0.09
U.S. inflation	-0.15	0.11	0.18
Longbond	-0.33	0.07	0.00
Commodity price inflation	0.01	0.01	0.11
Unemployment	0.20	0.11	0.09
CPIXFET inflation	0.20	0.15	0.18
Exchange rate depreciation	-0.02	0.01	0.14
$r\bar{a}r^2$	0.77	D.W.	1.71

Table 6 shows the estimated regression with deviations of total CPI inflation from the target as the dependent variable, using a stepwise regression procedure with a 50 per cent significance level cut-in and cut-out.

Table 6. Forecasting deviations of total CPI inflation, eight-quarter lag, 50 per cent stepwise

Indicator	Coefficient	Std. error	Significance
Constant	5.78	0.90	0.00
Longbond	-0.30	0.17	0.09
Commodity price inflation	-0.03	0.01	0.03
Total CPI inflation	0.21	0.27	0.45
CPIXFET inflation	0.49	0.59	0.41
Unemployment	-0.52	0.12	0.00
Labour cost inflation	-0.32	0.08	0.00
$r\bar{a}r^2$	0.67	D.W.	1.34

6.3 Estimating the reaction function

The forecasting equations, which tell us whether the deviations of inflation from the target are correlated with the various indicators, tell us whether the Bank's reaction function overreacts, underreacts, or reacts optimally to those indicators. But it is also useful to estimate the Bank's actual reaction function. First, having an estimate of the actual reaction function, together with information from the forecasting equation that tells us in which direction the actual reaction function ought to be adjusted, can give us some idea of what the optimal reaction function might look like. Second, the estimated actual reaction function can serve as a sort of check on the plausibility of the results of our forecasting equations. For example, suppose that there is some indicator that theory suggests should be a positive indicator of future inflation (i.e., a rise in that indicator suggests that future inflation should be expected to increase, unless the Bank takes offsetting action). Suppose further that the forecasting equation finds that the Bank has overreacted to that indicator. And suppose that the estimated reaction function finds that the Bank has reacted negatively to that indicator, by reducing the overnight rate when that indicator rises. This suggests that the optimal reaction function would involve reducing the overnight rate by even more when that positive inflation indicator increased, which is a counterintuitive result, and would raise some doubt about the plausibility of our policy advice.¹⁰

Table 7 shows the estimated regression for the Bank's reaction function, with the actual overnight rate as the dependent variable and the indicators as the independent variables. To prevent simultaneity problems, the reaction function assumes that the Bank sets the current overnight rate as a function of one-quarter lagged values of the indicators.

Table 7. Reaction function, all indicators, one-quarter lag

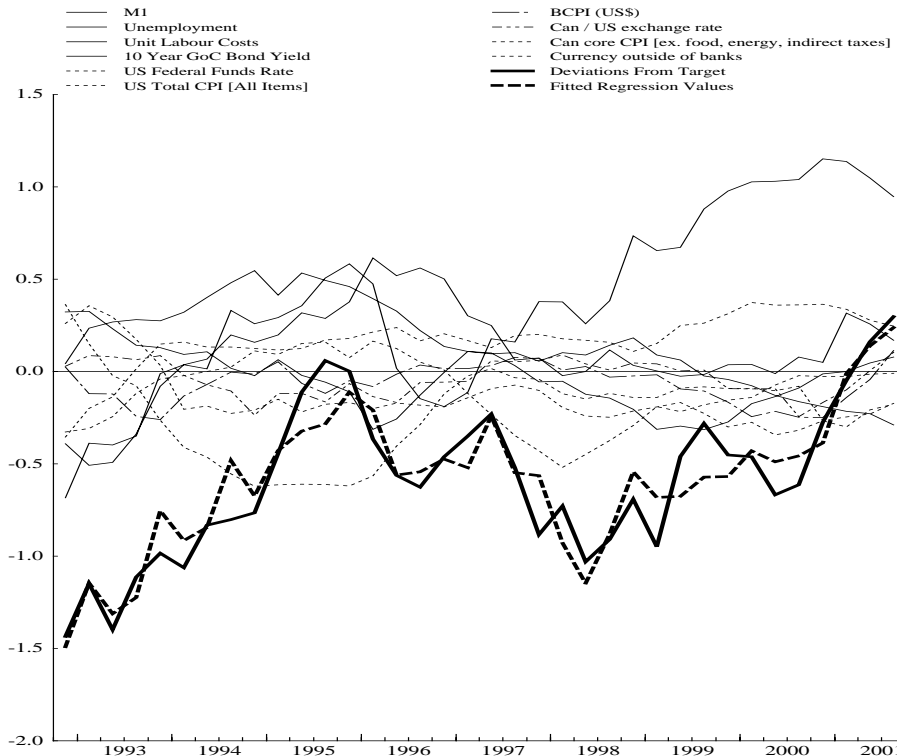
Indicator	Coefficient	Std. error	Significance
Constant	-2.46	3.10	0.43
Total CPI inflation	0.14	0.38	0.71
CPIXFET inflation	0.16	0.75	0.83
Exchange rate depreciation	0.13	0.05	0.02
U.S. inflation	0.09	0.44	0.83
M1 growth	-0.12	0.05	0.02
Currency growth	0.23	0.11	0.05
Unemployment	-0.12	0.26	0.65
Commodity price inflation	0.00	0.03	0.85
Labour cost inflation	-0.02	0.13	0.87
U.S. federal funds	0.77	0.31	0.02
Longbond	0.53	0.31	0.09
r_{bar}^2	0.70	D.W.	1.68

10. This is just what we find for unit labour cost inflation and the M1 growth rate.

6.4 Results

Figure 1 illustrates the results of the forecasting regression reported in Table 5. This figure uses CPIXFET inflation as the target variable, and assumes an eight-quarter lag in targeting. The solid line (labelled “deviations from target”) shows the actual deviation of CPIXFET inflation from the target. The line labelled “fitted regression values” shows the forecast deviation from target from the regression equation. The other lines show the contribution of each of the indicators individually towards that forecast.

**Figure 1: Forecasting Deviations of CPIXFET Inflation
(eight-quarter lag, 50 per cent stepwise)**



The results for each possible indicator are summarized below.

The most surprising result is that unit labour cost inflation appears with a statistically significant negative coefficient in all forecasting equations. This suggests that the Bank was overreacting to this indicator, increasing the overnight rate by too much whenever unit labour cost inflation increased, thereby causing inflation to fall below target eight quarters later. But the estimated

reaction function shows almost no reaction to this indicator. Taken together, these results suggest that the Bank should have *reduced* the overnight rate in response to increases in unit labour cost inflation, other things equal! Clearly, this result is paradoxical.

The M1 growth rate has a small but negative significant coefficient in the CPIXFET inflation forecasting equation, insignificant in the total CPI inflation forecasting equation. But the reaction function shows a negative highly significant coefficient. This suggests that the Bank *reduced* the overnight rate in response to a faster M1 growth rate, and should have reduced the overnight rate by an even greater amount! This result is also paradoxical.

The interest rate on the long (10-year) bond is also negative and significant in all forecasting equations, which suggests that the Bank overreacts to this indicator. The estimated reaction function shows a positive, strong, and highly significant response by the Bank to this indicator. Presumably the Bank should weaken its response.

Commodity price inflation has a positive and nearly statistically significant coefficient in the CPIXFET inflation forecasting equation, but a negative coefficient, significant in the stepwise regression, in the total CPI inflation forecasting equation. The estimated reaction function shows no response by the overnight rate to this indicator. The positive coefficient in the CPIXFET inflation forecasting equation suggests that this was too weak a response for targeting CPIXFET inflation. The negative coefficient in the total CPI inflation forecasting equation suggests that this was too strong a response for targeting total inflation.

The currency growth rate is positive and statistically significant in the CPIXFET inflation forecasting equation, but has a smaller and statistically insignificant coefficient in the total CPI inflation forecasting equation. The estimated reaction functions show that the Bank does respond to higher currency growth by raising the overnight rate, but the forecasting equations, plus the findings from the simple correlations, at longer lags, suggest that this response is not as strong as it should be to target CPIXFET inflation optimally.

The unemployment rate appears with a statistically significant coefficient in both stepwise forecasting equations, but with opposite signs! The coefficient is positive in the CPIXFET, and negative in the total CPI, inflation forecasting equations. The estimated reaction function shows that the Bank does cut the overnight rate in response to higher unemployment, but this result is not statistically significant. The forecasting equations, plus the findings from the simple correlations, suggest that the Bank does not respond strongly enough to unemployment, if total CPI inflation is the target, and responds too strongly, if CPIXFET inflation is the target.

The U.S. federal funds rate is positive and significant in the stepwise CPIXFET inflation forecasting equation. This is especially noteworthy given the large (and significant) positive coefficient (0.77) on this indicator in the reaction function. This suggests that the Bank should respond even more strongly to the Fed's moves, if CPIXFET inflation is the target.

No other indicators are statistically significant in the forecasting equations, which suggests that either the Bank was responding optimally to these indicators, or else the test is not powerful enough to prove otherwise. But one indicator deserves special consideration, because it appears significant in the estimated reaction function but insignificant in the forecasting equations. It is the rate of exchange rate depreciation, which has a positive (0.13) highly statistically significant coefficient in the estimated reaction function. The coefficients are near zero and statistically insignificant in the forecasting equations. But what is most interesting is that the standard errors of those estimates are also very small. This suggests that the test in this case is powerful enough to discover whether the Bank's response to this indicator is non-optimal, and therefore the near-zero coefficients mean that the Bank has been responding very close to optimally to changes in the exchange rate. But the positive and statistically significant simple correlations at the longer nine- and ten-quarter lags might suggest, to the contrary, that the Bank has responded too weakly to longer-run exchange rate depreciation.

7. Conclusion

In this paper, I have proposed a new way of interpreting correlations between target variables and indicators. Traditionally, for example, econometricians would expect to observe a positive correlation between the growth rate of some monetary aggregate and future inflation. They would interpret an observed positive correlation as evidence to support their view that money is a good indicator of future inflation, and they would recommend that the Bank should therefore closely monitor and react to this indicator. They would interpret an observed zero correlation as evidence that money is not a good indicator of future inflation, and recommend that the Bank should ignore this indicator. And they would interpret an observed negative correlation as evidence that more research was needed to resolve this puzzling finding. Instead, I have shown that a positive correlation means that the Bank should react more strongly to this indicator, that a negative correlation means that the Bank should react less strongly, and that a zero correlation means that the Bank has been reacting exactly right.

I have used this insight to review the Bank's policy over the past decade of inflation targeting. I have looked for consistent mistakes in inflation targeting, and found some. But with the benefit of new theory and econometric hindsight, it is not surprising that I can spot the Bank's past mistakes.

I have a new method of interpreting correlations, and I have 10 years' more data than the Bank did when it started inflation targeting 10 years ago. Also, I have used final revised data rather than using the real time data available to the Bank.

I have used this review to recommend how the Bank's monetary policy should be changed. These recommendations are qualitative rather than quantitative. I can recommend whether the Bank should respond more or less strongly to various indicators than it has done in the past 10 years, but I cannot say how much more strongly or less strongly. Furthermore, the recommended policy changes refer only to how policy should be changed from what it was *on average over the past 10 years*. If the Bank's monetary policy reaction function has changed during the past 10 years, I cannot tell whether the current reaction function is optimal or not.

References

- Goldfeld, S.M. and A.S. Blinder. 1972. "Some Implications of Endogenous Stabilization Policy." *Brookings Papers on Economic Activity* 3: 585–640.
- Goodhart, C.A.E. 1975. *Money, Information and Uncertainty*. London: MacMillan.
- Henry, S.G.B. and M. Desai. 1975. "Fiscal-Policy Simulations and Stabilization Policy." *Review of Economic Studies* 42(3): 347–59.
- Kareken, J. and R.M. Solow. 1963. "Lags in Monetary Policy." In *Stabilization Policies: A series of research studies prepared for the Commission on Money and Credit*. Englewood Cliffs, N.J.: Prentice-Hall.
- Leeson, R. (editor). 2000. *A.W.H. Phillips: Collected Works in Contemporary Perspective*. Cambridge: Cambridge University Press.
- Peston, M.H. 1972. "The Correlation between Targets and Instruments." *Economica* 29(156): 427–31.
- Poole, W. 1994. "Monetary Aggregates Targeting in a Low-Inflation Economy." In *Goals, Guidelines, and Constraints Facing Monetary Policymakers*, edited by J.C. Fuhrer. Federal Reserve Bank of Boston.
- Razzak, W.A. 2001. "Money in the Era of Inflation Targeting." Reserve Bank of New Zealand Discussion Paper, July.
- Rowe, N. and J. Yetman. 2002. "Identifying a Policymaker's Target: An Application to the Bank of Canada." *Canadian Journal of Economics* (May): 239–56.
- Taylor, J. 2001. "An Interview with Milton Friedman." *Macroeconomic Dynamics* 5(1): 101–31.
- Worswick, G.D.N. 1971. "Fiscal Policy and Stabilization in Britain." In *Britain's Economic Prospects Reconsidered*, edited by A.K. Cairncross. London: George Allen and Unwin.

Appendix A: Multicollinearity

When the policy-maker follows the optimal reaction function exactly, it will be impossible for an econometrician to estimate the reduced form.¹ This appendix shows that the above problem does not depend on knife-edge assumptions. Even if the policy-maker is close to the optimal reaction function, but not following it exactly, an econometrician with a finite sample of data will be unable to generate precise estimates of the parameters B and C by estimating equation (1).

To understand this econometric problem better, let us temporarily change the assumptions slightly, and introduce for pedagogical purposes a random error into the setting of the policy-maker's instrument, so that the reaction function now becomes:

$$R_t = \frac{\pi^* - \hat{A}}{\hat{B}} - \frac{\hat{C}}{\hat{B}} X_t + v_t, \quad (2c)$$

where v_t is white noise and is uncorrelated with X_t and e_{t+k} .

The original problem can now be seen as the limiting case, where (\hat{C}/\hat{B}) approaches (C/B) and the variance of v_t approaches zero, and so the stochastic, suboptimal policy rule (2c) approaches the deterministic, optimal policy rule (2b).

Suppose initially that $(\hat{C}/\hat{B}) = 0$, so that the instrument and the indicator are uncorrelated, and that the variance of v_t is roughly the same size as the variance of the indicator.² In this case, an econometrician estimating equation (1) would face no difficulties. R_t and X_t are now uncorrelated random variables, with roughly the same variance, and provided that the errors have roughly symmetric distributions, the confidence region for the estimates of the parameters B and C would

-
1. In practice, any econometric attempt to estimate the policy-maker's reaction function may not generate an equation that perfectly explains all the variance in the instrument. The R^2 may be less than 1.0, so there will be less than perfect multicollinearity in practice. But an R^2 of less than 1.0 should not be interpreted as evidence of the sort of exogenous policy shocks that are needed to estimate equation (1) directly. Rather, there may be indicators, like the occurrence of special events, to which the policy-maker reacts, but on which data are not available to the econometrician, and so the list of indicators to which the policy-maker reacts may in practice be larger than the list of indicators available to the econometrician. If these "hidden" indicators are important enough, an econometrician attempting to estimate equation (1) directly will fail to observe a significant problem of multicollinearity. But that apparent immunity from the multicollinearity problem really hides a deeper danger—that of omitted variable bias. Since the policy-maker is, by assumption, responding to those "hidden" indicators, the policy-maker's instrument will be correlated with those omitted variables, and the estimate of the coefficient on the instrument will be biased accordingly.
 2. This would be the case if the policy-maker were making no attempt to hit the target, and were ignoring the indicator and just making random changes in the instrument. We are considering this counterfactual case only to help us understand the more realistic case.

be vertically and horizontally symmetric around the point $\{B,C\}$, and would be roughly circular. Indeed, with the instrument uncorrelated with the indicator, so that the confidence region is symmetric, it would make no difference whether the econometrician estimated the parameters B and C jointly in a multiple regression, or separately in two simple regressions, because the omitted variable would just be added to the error term, which would remain uncorrelated with the other independent variable, so the OLS estimators in the simple regressions would remain unbiased.

This means that the ideal conditions under which the econometrician can estimate the reduced form occur when the policy-maker totally ignores any useful information provided by the indicator, and merely makes random fluctuations in the instrument that have no purpose and only worsen fluctuations in the target variable!

If the variance of v_t got smaller, the variance of X_t would stay the same, and so the confidence interval around C would stay the same, but the variance of R_t would shrink, and so the confidence interval around B would get wider, making the confidence region elliptical, but still symmetric around the point $\{B,C\}$. In the limit, as the variance of v_t approaches zero, the ellipse stretches until the confidence interval becomes two parallel lines (vertical, if B is on the vertical axis), meaning that it becomes impossible to estimate the parameter B . The reason is that the instrument never varies in this case, so it is impossible for the econometrician to discover what effect changes in the instrument would have on the target variable.

Suppose instead that (\hat{C}/\hat{B}) has the opposite sign to (C/B) . This would mean that the policy-maker is changing the instrument in a way that would be exactly opposite of what was needed to hit the target.³ For example, if the true parameters B and C were both positive, but (\hat{B}/\hat{C}) were negative, then the reaction function (2c) would make the instrument positively correlated with the indicator. Here, we have the standard multicollinearity problem, where the econometrician knows that R and X together are affecting P , but cannot disentangle the separate effects of each. The confidence region is no longer symmetric, but becomes an ellipse with a negatively sloped major axis. (I think that the slope of the major axis would equal (\hat{B}/\hat{C})). Estimating B and C separately in simple regressions would now give upward-biased estimates, since the omitted variable, and hence the errors in the simple regressions, would be positively correlated with the included independent variable. And if the variance of v_t approached zero, the instrument would become more perfectly correlated with the indicator, the ellipse would lengthen into two negatively sloped parallel lines, and it would be possible only to estimate a linear combination of B and C , not the two parameters themselves.

3. Don't ask why the policy-maker would do this; we are making this counterfactual assumption temporarily merely to help us understand multicollinearity better.

Finally, suppose realistically that (\hat{C}/\hat{B}) has the same sign as (C/B) . The policy-maker is moving the instrument in the right direction to help hit the target. If B and C are both positive, the reaction function will ensure that the instrument is negatively correlated with the indicator. The confidence interval becomes an ellipse with a positively sloped major axis. And as the policy-maker's estimates approach the true parameters, so that (\hat{C}/\hat{B}) approaches (C/B) , the major axis of the confidence region will become part of a ray from the origin to the point $\{B,C\}$. As the variance of v_t approached zero, the ellipse would lengthen to include the origin, and the econometrician estimating equation (1) would be unable to determine whether both parameters were zero, or indeed any scalar, positive or negative, multiple of the true values $\{B,C\}$. That is the problem that this paper seeks to explore and solve.

Appendix B: Previous Literature

Kareken and Solow (1963, 16), in criticism of Friedman, state:

Suppose that by heroic [means] . . . the Federal Reserve manages deftly to counter all disturbing impulses and to stabilize the level of economic activity absolutely. Then an observer following the Friedman methodology would see peaks and troughs in monetary change accompanied by a steady level of aggregate activity. He would presumably conclude that monetary policy has no effects at all, which would be precisely the opposite of the truth . . . Such conclusions [about the effects of monetary policy] are *ceteris paribus* statements, partial derivatives not total derivatives.” (italics original, square brackets added)⁴

Kareken and Solow thus recognize that there will be no correlation between the target variable and instrument when policy is “heroically” optimal, in the sense that all deviations of the target variable from the (constant) target are eliminated. But, since heroically optimal targeting is unrealistic, it is understandable that readers of Kareken and Solow, like they themselves, should fail to generalize this result to the realistic case, where the policy-maker has less-than-perfect foresight, so that the target variable does deviate randomly from the target. This generalization can easily be missed because, if the target variable is a constant, then, trivially, it cannot by definition be correlated with any other variable.

The generalization to imperfect foresight, and hence imperfect control of the target variable, is found in Peston (1972), who considers a model,

$$Z_t = aZ_{t-1} + U_{t-1} + ce_{t-1} + e_t$$

where Z is the target variable, U the instrument, and e a random error. He assumes that the policy-maker chooses the instrument optimally to minimize the variance of the target variable. He correctly concludes that “when the minimum variance strategy is used, correlation is zero between U_t and Z_{t+1} , and unity between U_t and Z_t ” (page 430). If we recognize (as Peston himself did not explicitly) that in his model the lagged target variable is the indicator (indeed the only indicator), then Peston’s result is the same as mine; there will be no correlation between the target variable and the instrument, and a perfect correlation between the instrument and the indicator (he does not draw the immediate conclusion that there will also be no correlation between the target variable and the indicator). Peston also says that “The issue is not ‘correct’ versus ‘incorrect’ signs [of the correlations], but that signs are no help at all” (page 430). Here, he overlooks the fact that, although the signs are indeed not helpful when interpreted in the usual way, they are

4. The first elision deletes material that makes no sense grammatically or substantively and must reflect an editorial error.

nevertheless helpful when interpreted as giving information on how the policy-maker's reaction function needs to be modified to improve targeting.

Peston considers his paper to be a supplement to an earlier paper by Worswick (1971), which sets up a simple multiplier model in which exports are exogenous and follow a sine wave over time. The policy-maker's instrument is government expenditure. The optimal reaction function requires that the policy-maker adjust government expenditure in perfect negative synchronization with exports and thus perfectly stabilize output. Worswick explores the various correlations between the target variable (output) and instrument under various assumptions about the relative magnitude and timing of fluctuations in exports and government expenditure. He concludes from these examples that

. . . even in the very simplest model, the sign of the correlation between [the instrument] and [the target variable] tells us nothing whatever about the success or otherwise of the stabilization. Ideally, the correlation ought to be zero. If we endow the authorities with perfect timing, then the issue is whether they underdo or overdo their intervention and, as we have seen, according as they do one or the other systematically, the correlation will come out either positive or negative. (page 42)

Worswick recognizes that the sign of the correlation between the target variable and instrument tells us whether the policy-maker is under- or overreacting to exogenous shocks (which represent a particular example of changes in the indicators). This insight of Worswick's seems to have been totally ignored. It is closely related to my own use of the sign of the correlation between the target variable and indicator to identify whether policy needs to be changed.

Goodhart (1975) references Peston's paper, and repeats his result that the target variable must be uncorrelated with the instrument under the optimal policy rule, but he fails to notice that the target variable must also be uncorrelated with the indicator, or indeed with any variable in the policy-maker's information set. Apart from Goodhart, Peston's paper seems to have been almost totally ignored. A computer search discovered only one citation of Peston's paper, by Henry and Desai (1975).

In the same footnote (page 192, note 1) in which Goodhart references Peston (1972), he also references Goldfeld and Blinder (1972) on the same subject. Goldfeld and Blinder construct examples where the target variable is determined contemporaneously by the instrument, the indicator, and an error term. Using our notation, their version of equation (1) is:

$$\pi_t = A + BR_t + CX_t + e_t.$$

They assume that the policy-maker has imperfect information on the variable X_t , but may have some information on the error term e_t . Thus the reaction function sets the instrument as a function of the policy-maker's expectations of X_t and e_t . The main conclusion of their paper is that if the policy-maker has some information on e_t , the instrument will be correlated with the error term, and the OLS estimator of B will therefore be biased. Simultaneous equation estimation of the structural equation and the policy-maker's reaction function may (or may not) reduce this bias.

There is an alternative, and clearer, way to describe what Goldfeld and Blinder are doing. Let us distinguish between the information set available to the policy-maker when it sets the instrument, and the information set available to econometricians when they estimate the structural equation. By definition, the term "indicator" refers to the policy-maker's information set.

If the policy-maker has more information than the econometrician—as would be the case, for example, if the policy-maker has some information on the econometrician's error term, e_t —then this means effectively that the econometrician has failed to include a relevant variable (one or more of the indicators) in estimating equation (1). Naturally, the omitted variable will be correlated with the instrument, and so the estimated coefficient for the instrument will be biased. (And the estimated coefficients on the included indicators may also be biased.) But this bias should be understood to be an omitted-variable bias, and not a simultaneous-equation bias.

Next, consider the case where the econometrician has more information than the policy-maker, or rather, the econometrician *uses* more information to estimate the structure than was available to the policy-maker when it set the instrument. This occurs, for example, when the econometrician, in estimating equation (1), includes a right-side variable that is dated later (or was published later) than the date on the instrument. For example, the econometrician estimates:

$$\pi_t = A + BR_{t-k} + CX_{t-k} + DZ_t + e_t, \quad (1f)$$

but the reaction function sets R_{t-k} as a function of X_{t-k} only. The policy-maker will want to include within its set of instruments, X_{t-k} , not only those variables that have a direct causal impact on the target variable (those with a non-zero coefficient C), but also those variables that matter only because they help forecast Z_t . In general, estimating (1f) will not be an unproblematic way of determining the optimal policy. First, the instrument will be perfectly correlated with the indicators via the policy-maker's reaction function, and so multicollinearity will make it impossible to estimate the parameters B and C . Second, the policy-maker's optimal reaction to a change in an indicator will depend not only on how that indicator directly affects the target variable (C), but will depend also on how that indicator helps the policy-maker to predict Z_t . This second problem can be overcome by independently estimating a forecasting equation for X_{t-k} on

Z_t , but the first problem cannot be avoided unless there are no such indicators that directly affect the target variable, and all indicators are used by the policy-maker only for their ability to forecast Z_t .

Goldfeld and Blinder's Proposition 1 is based on exactly this implicit assumption. It states that "A stabilization policy that follows a *Theil reaction function* poses no problems for reduced-form (or, for that matter, structural) estimation" (ibid. page 596, original italics). They base their conclusion on an example where the econometrician's information set exceeds the policy-maker's set of indicators, so that the instrument is uncorrelated with the econometrician's error term, but they overlook the potential multicollinearity problem, because their example contains no indicators that have a direct effect on the target variable ($C=0$). If there exist indicators that affect the target variable directly, or if the econometrician has the same information set as the policy-maker, Goldfeld and Blinder's Proposition 1 is false, not because of simultaneous equation bias (which they are alert to), but because of multicollinearity (which they ignore).

A computer search reveals that Goldfeld and Blinder's paper gets 46 citations vs. Peston's single citation. But Goldfeld and Blinder miss the point.

In summation: Peston got it right (though he did not explore the implications fully), but was ignored; Goldfeld and Blinder got it wrong, but were widely read (or at least cited), with the result that for the past 30 years economists have believed that endogenous policy creates no problems that cannot be handled by simultaneous equation estimation.

The above-noted papers were written before rational expectations became a familiar working assumption for macroeconomists. It may not be surprising, then, that these economists failed to recognize that deviations of target variables from the target could be interpreted as forecast errors, and that under rational expectations these forecast errors should be uncorrelated with any variable in the policy-maker's information set. But it is surprising that no economist seems to have noticed or exploited this orthogonality condition in the past 30 years. Perhaps the reason is that, 30 years ago, it was taken for granted that monetary and fiscal policy-makers were targeting real output, just as today it is taken for granted that many monetary policy-makers are targeting inflation. But in the intervening period, the 1970s and 1980s, there was little agreement about what monetary and fiscal policy-makers should, or indeed could, be targeting, and so nobody thought about correlations between targets, instruments, and indicators.

Yet the idea has not disappeared altogether. Milton Friedman, interviewed by John Taylor in May 2000, provides the following analogy:

The temperature in a room without a thermostat but with a heating system will be positively correlated with the amount of fuel fed into the heating system and may vary widely. With a thermostat set at a fixed temperature, there will be zero correlation between the intake of fuel and the room temperature, a negative correlation between the intake of fuel and the external temperature. Also, the room temperature will vary little. (Taylor 2001, page 129)

The thermostat represents a good monetary policy, the fuel represents the quantity of money, and the room temperature represents GDP. Poole (1994) explains that the weaker correlation between money and output is the result of monetary policy being chosen optimally to smooth output. Razzak (2001) uses Friedman's thermostat analogy in his study of the relationship between money, output, and inflation in New Zealand under different monetary policy regimes.

Bank of Canada Working Papers

Documents de travail de la Banque du Canada

Working papers are generally published in the language of the author, with an abstract in both official languages. *Les documents de travail sont publiés généralement dans la langue utilisée par les auteurs; ils sont cependant précédés d'un résumé bilingue.*

2002

2002-22	The Usefulness of Consumer Confidence Indexes in the United States	B. Desroches and M-A. Gosselin
2002-21	Entrepreneurial Risk, Credit Constraints, and the Corporate Income Tax: A Quantitative Exploration	C.A. Meh
2002-20	Evaluating the Quarterly Projection Model: A Preliminary Investigation	R. Amano, K. McPhail, H. Pioro, and A. Rennison
2002-19	Estimates of the Sticky-Information Phillips Curve for the United States, Canada, and the United Kingdom	H. Khan and Z. Zhu
2002-18	Estimated DGE Models and Forecasting Accuracy: A Preliminary Investigation with Canadian Data	K. Moran and V. Dolar
2002-17	Does Exchange Rate Policy Matter for Growth?	J. Bailliu, R. Lafrance, and J.-F. Perrault
2002-16	A Market Microstructure Analysis of Foreign Exchange Intervention in Canada	C. D'Souza
2002-15	Corporate Bond Spreads and the Business Cycle	Z. Zhang
2002-14	Entrepreneurship, Inequality, and Taxation	C.A. Meh
2002-13	Towards a More Complete Debt Strategy Simulation Framework	D.J. Bolder
2002-12	Modelling Financial Instability: A Survey of the Literature	A. Lai
2002-11	Risk, Entropy, and the Transformation of Distributions	R.M. Reesor and D.L. McLeish
2002-10	La fiabilité des estimations de l'écart de production au Canada	J.-P. Cayen et S. van Norden
2002-9	The Microstructure of Multiple-Dealer Equity and Government Securities Markets: How They Differ	T. Gravelle
2002-8	Restructuring in the Canadian Economy: A Survey of Firms	C. Kwan
2002-7	Contribution of ICT Use to Output and Labour-Productivity Growth in Canada	H. Khan and M. Santos

Copies and a complete list of working papers are available from:

Pour obtenir des exemplaires et une liste complète des documents de travail, prière de s'adresser à :

Publications Distribution, Bank of Canada
234 Wellington Street, Ottawa, Ontario K1A 0G9
E-mail: publications@bankofcanada.ca
Web site: <http://www.bankofcanada.ca>

Diffusion des publications, Banque du Canada
234, rue Wellington, Ottawa (Ontario) K1A 0G9
Adresse électronique : publications@banqueducanada.ca
Site Web : <http://www.banqueducanada.ca>