Preliminary Draft - Comments Invited

The Brown Book Method:
Adaptive Learning for Inflation-Targeting Central Banks

(Chapter 3)

Nicholas Rowe
Department of Economics, Carleton University
and
Bank of Canada

Nick_Rowe@Carleton.ca

this version
21 February 2003

Abstract

This paper builds on my previous paper, Rowe (2002), and develops a method whereby an inflation targeting central bank can improve its monetary policy by learning from its past experience. I show how to estimate a recommended monetary policy reaction function, showing how to set the monetary policy instrument as a function for the indicators, which improves on the bank's previous reaction function. Estimating a simple inflation forecasting equation can reveal systematic mistakes in past monetary policy. Adapting the estimated past reaction function to reduce the size of these systematic mistakes generates an improved monetary policy reaction function designed to keep inflation closer to target. I also show how to evaluate the usefulness of various indicators of inflation, and how to evaluate the usefulness of this new source of monetary policy advice.

Acknowledgements

I began writing an earlier version of this paper while visiting the Bank of Canada in 2002, first in the Research Department, and then in Monetary and Financial Analysis. The current version concentrates solely on the method by which a "Brown Book" may be estimated; actual empirical estimates of Brown Books for the Bank of Canada have been excluded. All errors and opinions in this paper are my own; the views herein cannot be attributed to the Bank of Canada or to Carleton University. That being understood, I wish to thank the many people at the Bank of Canada who commented on earlier drafts and presentations. They include especially: Bob Amano, Don Colletti, Allan Crawford, Pierre Duguay, Chuck Freedman, Dave Longworth, Tiff Macklem, Brian O'Reilly, Jack Selody, and Carolyn Wilkins. I thank David Tulk and Tracy Chan for excellent research assistance. I apologise to those whose names I have inadvertently omitted. And I thank everyone at the Bank for their hospitality, and for making my visit so rewarding.

1. Introduction

What I call 'The Brown Book method' is a way of providing advice on monetary policy to the Bank of Canada. I have used the Bank of Canada as an example, but nearly everything I say here could apply equally well to any central bank which has a fixed inflation target. The objective of this advice is to help the Bank hit its inflation target.[1]

The Bank of Canada's announced target for monetary policy is to keep the CPI inflation rate at 2 per cent, which is the mid-point of the 1 per cent to 3 per cent target range. The Bank's announced[2] monetary policy instrument is the overnight rate, which is a very short term interest rate on loans between commercial banks.[3] Eight times a year, the Bank of Canada faces a 'Fixed Action Date' (FAD) decision. At each FAD, the Bank must decide the setting of its monetary policy instrument.[4] The Bank chooses to set the overnight rate at a level which it thinks will keep future inflation at the 2 per cent target. Since monetary policy affects inflation with a lag, the Bank's announced intention is to set the overnight rate so that the annual inflation rate 8 quarters later (which means the percentage change in the CPI from 4 quarters ahead to 8 quarters ahead) will equal the 2 per cent target.

---

[1] This paper contains almost no references to previous literature. A discussion of the historical antecedents to my approach is found in Appendix B of my previous paper Rowe (2002).

[2] Statements about what instrument the Bank is using are closer to semantic conventions than to empirical statements about the world. If two variables Y and Z are both immediately observable by the Bank and are related, then it makes little difference whether we say the Bank controls Y watching Z, or controls Z watching Y. A visiting anthropologist would probably translate the statement "Y is the Bank's instrument" as "Y rarely changes except during the short hiatus immediately following a crescendo of meetings of the Bank's elders". But if you are giving advice to someone, it is usually better to frame the advice in their own language.

[3] Strictly speaking, the Bank sets a target for the overnight rate, but the actual overnight rate usually stays very close to the announced target.

[4] In exceptional circumstances the Bank may change the overnight rate between FADs.

The Brown Book method can be used to generate a 'Brown Book', containing a recommended setting for the overnight rate, for each FAD. The Brown Book can complement the other sources of advice the Bank receives. It is unlikely that any one source of advice will dominate all other sources, and that the Bank should rely solely on one source and ignore all others. It is more likely that each source of advice will contain some useful information that is at least partly independent of other sources, and that the Bank should follow some weighted average of advice from all sources. Even if one source of advice did objectively dominate all other sources, the Bank is unlikely to know that fact, so that many sources of advice will be subjectively useful. My aim is to provide a source of advice which, because it is based on an entirely different method from the other sources, is likely to contain at least some useful independent information, so that it is very unlikely to be dominated by any other source of advice, and should therefore deserve a positive weight in the Bank's decisions.

The Brown Book method formalizes the idea of learning from experience, or learning from mistakes, or adaptive learning. It is a largely self-contained method, and imposes almost no structure from economic theory. The only strictly necessary theoretical assumption for applying the Brown Book method is that an exogenous increase in the overnight rate, other things equal, would cause a reduction in future inflation. That assumption is almost certainly non-controversial to any economist[5]. Apart from this one necessary assumption, the only role for economic theory or other prior knowledge is to help choose the set of economic indicators on which Brown Book advice is based. Since the number of possible indicator sets is much too large relative to the data sample for statistical analysis alone to suffice, an economist's judgement will be useful in practice to supplement statistical analysis in choosing the best set of indicators on which to base Brown

[5] Anyone who did question that assumption should consider the following thought experiment. Suppose a driver did not know that his car's steering wheel had been rigged the wrong way. Seeing his car approach the right side edge of the road, he turns the wheel to the left, the car turns more to the right, he turns the wheel more to the left....and the car will inevitably crash. Since the Bank has believed for a long time that an increase in the overnight rate reduces future inflation, and has acted on this belief, and has not yet crashed the economy with either hyperinflation or hyperdeflation, the Bank's belief must be correct.

Book advice.

To understand intuitively how the Brown Book method can work, imagine that you are a non-economist suddenly forced to choose monetary policy for the Bank of Canada. If you are told the instrument, the target, the indicators, and that an increase in the overnight rate will reduce future inflation, *and if you have a long enough data set*, your job is actually very easy. You need ask only two questions of the data. First, you ask "At all times in the past when the indicators were the same as they are today, what was the average setting of the overnight rate chosen by the Bank?" Suppose the answer to this first question is "4 per cent". Second, you ask "And what was the average inflation rate two years later?" If the answer to the second question is "2 per cent", so that average inflation equalled the target, you simply follow past practice and set the overnight rate equal to 4 per cent. If the answer to the second question is more than 2 per cent, you partially correct for the past errors by setting the overnight rate just a little bit higher than 4 per cent. If less than 2 per cent, you set the overnight rate just a little bit lower than 4 per cent.

Provided there have been no structural changes, and provided the data set is varied enough to contain all possible values of the indicators, and long enough so that sample averages equal population averages, a non-economist using this simple method can set monetary policy as well or better than the Bank did over the sample period. The only real danger of a worse performance is if the non-economist gets overenthusiastic and overcorrects for the Bank's past errors.

Real world application of the Brown Book method is more complicated than the above example only because the usable data set is not long enough, which means we have to use statistical methods to interrogate the data in place of the two simple questions. Because the data series is too short, it is very unlikely we will have observed in the past exactly the same combination of indicator values that we observe today, so we need statistical methods to interpolate or extrapolate between the observed data points. Also, because the data series is too short, we need statistical methods to try to distinguish between missing the inflation target because of bad luck (an unforecastable shock to inflation), and missing the inflation target because of a systematic policy mistake.

The reason the usable data set is short is that, as the Lucas Critique reminds us, the statistical relation between indicators, instrument, and the target variable cannot be assumed invariant to the monetary policy regime. Since the Brown Book Method does not try to estimate purportedly policy-invariant deep parameters, it cannot use data from before the inflation targeting period. This short data sample is a problem, but it will become less of a problem over time. Moreover, I argue in Rowe (2002) that existing methods of generating monetary policy advice may face problems when their underlying models are estimated using only data from the period of inflation targeting. If so, those existing methods may become progressively less reliable over time if inflation targeting continues, while the Brown Book Method should become progressively more reliable.

The Brown Book method can be applied to any central bank which has been targeting inflation for several years. But since it is based on learning from experience, it cannot help a central bank which is just starting to target inflation, except insofar as the experience of other inflation-targeting central banks is applicable.

2. An Outline of the Brown Book Method

The Bank of Canada has no crystal ball. It must set today's overnight rate as a function of the information it has available today. Since the Bank presumably knows today the setting of the overnight rate it is choosing today, the chosen setting of the overnight rate instrument is part of the Bank's information set. The other part of the Bank's information set we call the 'indicators'. An indicator is then any variable, either exogenous or endogenous, which the Bank can observe to help it set the overnight rate to keep future inflation on target.

Indicators are shadows. The indicator GDP is not the same as true GDP. The indicator GDP is

GDP as measured and reported by Statistics Canada. GDP the indicator is a noisy and lagged reflection of true, objective GDP. The same is true, to a greater or lesser extent, for all indicators. This means that the indicators themselves need have no direct causal effect on future inflation, and indeed will have no direct causal relation to future inflation (except insofar as published information affects the Bank's and other people's decisions). The relation between indicators and future inflation is informational, and is not required to be causal. Assuming full memory, today's indicator set may also contain any lagged indicator, including lagged values of the overnight rate instrument.

Each FAD, lacking a crystal ball, the Bank of Canada must set today's overnight rate (which will usually be the overnight rate for the next six weeks) as a function of the indicator set available today, which contains more or less noisy data available at various lags. We call this function the monetary policy reaction function.[6] The central question of monetary policy for an inflation-targeting central bank, and the question we want to answer, is the following: which particular reaction function will keep future inflation as close as possible to target? If we assume a symmetric loss function and symmetrically distributed forecast errors, we can rewrite this question as follows: which particular reaction function, if followed, will keep the rational expectation of future inflation, conditional on the available indicator set, equal to target? Call this (unknown) reaction function the 'optimal' (that is, optimal for inflation targeting) reaction function.

The Brown Book method seeks to estimate a recommended reaction function which is an improvement over the actual reaction function that the Bank has used in the past, in the sense that the recommended reaction function will be closer to the optimal reaction function, and will keep (the rational expectation of) future inflation closer to target, than the actual reaction function.

[6] An alternative name for the reaction function is "indicator rule", except that calling something a rule implies some degree of commitment, and there is no commitment to any particular reaction function. The reaction function is purely a means to an end. The end, to which there is a commitment, is keeping inflation on target..

The key to understanding the Brown Book method is to understand that deviations of inflation from target constitute the Bank's forecast errors, and under rational expectations forecast errors should be uncorrelated with anything in the Bank's information set when it made that forecast two years previously. Non-zero correlations between deviations of inflation from target and two year lagged indicators are thus evidence of systematic, and hence correctable, mistakes in past monetary policy.

Intuitively, the Brown Book method can be broken down into the following four steps:

First, for a given indicator set, estimate the actual reaction function, to show how the Bank actually set the overnight rate, as a function of that indicator set, over the sample period.

Second, estimate a 'forecasting equation', which forecasts deviations of future inflation from target conditional on the indicator set. If the Bank was already using the optimal reaction function, then these deviations of future inflation from target would represent forecast errors around the rational expectation, and would be uncorrelated with any information available to the Bank. If the estimated forecasting equation fails to forecast future inflation, then we infer that the actual reaction function cannot be improved. But if the forecasting equation lets us predict future deviations of inflation with any degree of confidence, we can have the same level of confidence that the actual reaction function is biased in a particular direction. The forecasting equation estimates the systematic errors made by the Bank's actual reaction function. For example, if an indicator has a positive correlation with future inflation, then an improved reaction function would have a more positive (larger positive, or smaller negative) coefficient on that indicator.

Third, calculate a recommended reaction function by taking the right hand side of the forecasting equation, multiplied by a small positive constant, and adding it to the estimated actual reaction function. The recommended reaction function is thus an example of adaptive learning, which partially corrects the systematic errors generated in the past by the actual reaction function. At one extreme, as the 'small positive constant' approaches zero, the recommended reaction function approaches the estimated actual reaction function. As the small positive constant gets larger, the recommended reaction function gives a bigger improvement by correcting more of the systematic error, but too large and it overcorrects, and eventually causes a deterioration.

Fourth and finally, the recommended reaction function can be used to generate a recommended setting for the overnight rate at any FAD by simply substituting into it the latest values of the indicators.

There is an alternative way to estimate a recommended reaction function, which gives an identical answer, by combining steps one and two above, and estimating a 'Brown Book equation' directly. On the right hand side of the Brown Book equation is the indicator set. On the left hand side is the overnight rate, plus a small positive constant times the deviation of future inflation from target. The estimated coefficients generate the recommended reaction function. The advantage of estimating the Brown Book equation directly, apart from simplicity, is that doing so provides useful statistics like standard errors for the recommended reaction function. It also helps us to choose the best indicator set, as discussed in section 5 below.

3. A More Formal Analysis

Assume everything is linear. Assume an increase in the overnight rate, ceteris paribus, will cause a reduction in future inflation. Assume we can place an upper bound on the magnitude of this effect. Assume the Bank wants to set the overnight rate so that the rational expectation of inflation at an 8 quarter horizon equals 2 per cent. Then everything else follows.

Let X_t designate the information available to the Bank at time t, with the exception of the overnight rate at time t, which we designate R_t. We call X_t the Bank's indicator set. In other words, $I_t = \{R_t, X_t\}$.

For notational simplicity, let P_t designate inflation minus 2 per cent (the deviation of inflation from target). Given the assumption that everything is linear we can write:

1. $E(P_{t+8}/I_t) = E(P_{t+8}/\{R_t , X_t \}) = -bR_t + cX_t$ $\qquad b>0$

Note that by writing equation (1) in terms of expectations, we don't need to include an error term, and so we don't need to make any assumptions about the serial correlation of that error term.

The Bank's problem is to set the overnight rate, R_t, as a function of its available information, X_t, so that its expectation of the future deviation of inflation from target, conditional on R_t and X_t, is equal to zero. Setting $E(P_{t+8} /I_t) = 0$ in equation (1) we get:

2. $0 = -bR_t* + cX_t$

where R_t* designates the overnight rate the Bank wants to set (call it the "optimal" overnight rate).

Rearranging equation (2), we solve for the "optimal" reaction function, where "optimal" is defined as that (unknown) reaction function which, if followed, would set the overnight rate as some function of the indicator set which would make the rational expectation of future inflation, conditional on the indicator set, equal the target.

2a. $R_t* = (c/b)X_t$

Substituting (2a) into equation (1) we can express the expected deviation of inflation from target as a function of the deviation of the actual overnight rate from the optimal overnight rate.

1a. $E(P_{t+8} /It) = E(P_{t+8} /\{R_t,X_t\}) = -bR_t + cX_t = -b(R_t - R_t*)$

Define the rational expectations forecast error e_t as:

3. $e_{t+8} = P_{t+8} - E(P_{t+8} /It)$

Substitute equation (3) into equation (1) to get:

4. $P_{t+8} = -bR_t + cX_t + e_{t+8}$

It might appear at first glance that we can estimate equation (4) directly, and thereby get estimates of the parameters b and c which define the optimal reaction function (2a). Unfortunately it is not possible to do this, because if the Bank is following any deterministic reaction function, which sets R_t as a function of X_t, there will be perfect multicolinearity between R_t and X_t. This is discussed further in Rowe (2002).

Given rational expectations, e_{t+8} must be uncorrelated with I_t, and is therefore uncorrelated with R_t and X_t. Therefore e_t may be an MA process, up to an MA(7), but it cannot be an AR process, if P_t is part of the information set X_t, for then e_{t+8} would be correlated with X_t, and this would violate rational expectations.

Rearrange equation (4) to get:

4a. $(R_t + (1/b)P_{t+8}) = (c/b)X_t + (1/b)e_{t+8}$

We would like to estimate equation (4a), because our estimate of the coefficient on X_t would give us an estimate of the parameter ratio (c/b) which defines the “optimal” reaction function. But we cannot estimate (4a), because we don’t know the parameter ‘b’, so we cannot construct the left hand side.

Suppose we take some small number ‘a’, chosen so that $0<a<(1/b)$. [Remember our assumption that we can put an upper bound on ‘b’, and hence a lower bound on (1/b)]. Then we estimate:

5. $(R_t + aP_{t+8}) = rX_t + av_{t+8}$

Equation (5) is the Brown Book equation.[7] Our estimate of the parameter ‘r’, call it r^, gives us our “recommended” reaction function: set $R_t = r\hat{}X_t$.

Remember that ‘a’ lies between 0 and (1/b). In one limit, as ‘a’ approaches (1/b), equation (5) approaches equation (4a), and our estimate r^ approaches an estimate of the parameter (c/b) in the “optimal” reaction function, conditioned on an indicator set X_t. In the other limit, as ‘a’ approaches 0, equation (5) approaches an estimate of the actual reaction function, conditioned on the indicator set X_t. In other words, when we estimate equation (5) we are estimating a convex combination of the optimal reaction function and the actual reaction function. The small positive constant ‘a’ is our adaptive learning adjustment parameter. It measures how much we adjust in response to learning from experience, or past mistakes.[8]

Suppose we estimate the Brown Book equation, and then set monetary policy according to the recommended reaction function, $R_t = r\hat{}X_t$. Substituting the recommended reaction function into equation (1), we get:

---

[7] Remember that X_t refers to information that is available at time t, when the overnight rate is set. But the normal practice is to date indicators according to the date to which they refer, not the date at which they are published. If there is a one month delay in publishing data, for example, we should estimate the Brown Book equation using last month’s unemployment, exchange rate, etc. The six weeks between Fixed Action Dates also complicates the timing, since on average today’s overnight rate was set 3 weeks ago, using data that was available 3 weeks ago, which in turn reflected what was happening in the economy some additional weeks before that. Adding together the 3 week decision lag and the data lag, a reasonable assumption would be to use a lag of a little more than one month on the indicator set. Remember also that it is permissible also to include indicators at longer lags, since they too are part of the Bank’s information set at time t.

[8] If the Bank used (say) base money as the instrument, increases in which have a positive effect on future inflation, ceteris paribus, then ‘a’ should be a small negative number instead.

6. $E(P_{t+8} / I_t) = E(P_{t+8} / \{R_t, X_t\}) = -bR_t + cX_t = (c\text{-}br\hat{}\,)X_t$

Suppose the actual reaction function, that had previously been used to set monetary policy, is $R_t = sX_t$. Substituting the actual reaction function into equation (1) we get:

7. $E(P_{t+8} / I_t) = E(P_{t+8} / \{R_t, X_t\}) = -bR_t + cX_t = (c\text{-}bs)X_t$

The Brown Book recommended reaction function (set $R_t = r\hat{}X_t$) is designed to "improve" monetary policy compared to the actual reaction function previously followed by the Bank. By "improve" we mean that the expected deviation of inflation from target, conditional on X_t, is smaller (in absolute value) if the recommended reaction function is followed than if the actual reaction function is followed. Comparing equations (6) and (7), we see that the recommended reaction function will improve monetary policy if and only if (c-br^) is smaller in absolute value than (c-bs).

Substituting the actual reaction function into the Brown Book equation (5) we get:

8. $(R_t + aP_{t+8}) = rX_t + av_{t+8} = sX_t + a(c\text{-}bs)X_t + av_{t+8} = [s+a(c\text{-}bs)]X_t + av_{t+8}$

so that the parameter 'r' in the Brown Book equation, which the econometrician seeks to estimate, will equal the parameter [s-a(c-bs)]. Given that v_{t+8} is equal to the rational expectations forecast error $e_{t+8,}$ and is therefore uncorrelated with X_t, estimating the Brown Book equation by OLS will result in an unbiased estimate r^ of the true parameter r.

Comparing equations (6) and (7), which give the rational expectation of the deviation of future inflation from target, under Brown Book recommended policy for (6), and under actual policy for (7), we see that following the recommended reaction function will improve monetary policy if and only if

9. (c-br^) =(c-br-b(r^-r)) = c-b[s+a(c-bs)]-b(r^-r)

$= (c-bs)(1-ba) -b(r^{\wedge}-r)$ is smaller in absolute value than $(c-bs)$

which will be true provided the sampling variance of the estimator is small, so that $(r^{\wedge}-r)$ is small, and provided the small positive constant 'a' is small enough so that $0<a<2(1/b)$. Notice that the recommended reaction function gives the maximum improvement in inflation targeting when we set the small positive constant $a=(1/b)$.

As an aid to the intuition, notice that with trivial algebraic manipulation we can rewrite the parameter 'r', estimated in the Brown book equation, in the following three ways:

$r = s+(a/b)[(c/b)-s] = (c/b)+(ab-1)[(c/b)-s] = (a/b)(c/b)+[1-(a/b)]s$

Remember that 's' is the coefficient in the actual reaction function, and that (c/b) is the coefficient in the optimal reaction function. The first formulation interprets the Brown Book coefficient as an estimate of the actual reaction function, with an adjustment for any deviation [(c/b)-s] between the optimal and actual reaction functions. The second formulation interprets the Brown Book coefficient as an estimate of the optimal reaction function, with an error, which error is itself the product of two other errors–our own error in letting the adjustment parameter 'a' differ from its optimal value (1/b), and the Bank's error in letting its actual reaction function 's' differ from the optimal value (c/b). The third formulation interprets the Brown book coefficient as a weighted average of the optimal and actual reaction functions, with weights (a/b) and [1-(a/b)].

Provided we have not set the adjustment parameter 'a' too large, so that $0<a<2(1/b)$, there are only three reasons why the Brown Book's recommended reaction function might fail to improve on actual policy:

1. Sampling variance of the estimator may cause the estimate $r^{\wedge}$ to differ from the true parameter r.
2. The indicator set chosen by the econometrician estimating equation (5) is smaller than the indicator set followed by the Bank's actual policy, and ignores some useful indicators.
3. The actual reaction function is already optimal, in which case the estimate of the recommended reaction function will be the same as the estimate of the actual reaction function (which can be

tested by seeing whether the estimate r^ is sensitive to the choice of ‘a’, or equivalently by estimating the forecasting equation).

Except for these three caveats, the Brown Book recommended reaction function must always lie between the actual reaction function and the optimal reaction function, and so must represent an improvement over the actual reaction function.

There is a second way to estimate the Brown Book recommended reaction function. This second method will give an identical estimate to the first method described above, but it lets us understand the intuition behind the Brown Book from a new perspective.

First we estimate the actual reaction function $R_t = sX_t$ by regressing the overnight rate on the indicator set. Second we estimate the “Forecasting Equation”, by regressing the deviation of inflation from target on the indicator set. Substituting equation (3) into equation (7) we can interpret the “Forecasting Equation” as:

10. $P_{t+8} = (c\text{-}bs)X_t + e_{t+8}$

Next we multiply equation (10) by our small positive constant ‘a’, and add it to the estimate of the actual reaction function to get an alternative route to derive an estimate of the same Brown book equation[9]:

8'. $(R_t + aP_{t+8}) = [s+a(c\text{-}bs)]X_t + ae_{t+8}$

Look again at the Forecasting Equation. If the actual reaction function were optimal, so that

[9] I cannot formally prove that the two ways of estimating the Brown Book equation must give identical estimates (though presumably a competent econometrician could), but I have found that in practice they do. It is also a good idea to use both methods, and check to see that they give the same estimate, to guard against programming errors.

s=(c/b), then the coefficient (c-bs) would be zero. A non-zero estimated coefficient in the Forecasting Equation is evidence of systematic mistakes in past monetary policy, because if monetary policy were optimal for inflation targeting then deviations of inflation from target should be unforecastable. A positive coefficient in the forecasting equation means that s<(c/b), which means that the bank was not reacting strongly enough to the indicators. A negative coefficient means that the Bank was reacting too strongly. The Brown Book recommended reaction function can be interpreted as what we get when we take the actual reaction function, and modify it to reduce its systematic errors as revealed by the Forecasting Equation.

4. The Problem of Missing Indicators

The Bank of Canada can observe hundreds of potentially useful indicators. When we include lagged indicators, it can observe thousands of potentially useful indicators. In practice, given the degrees of freedom allowed by our data sample, the Brown Book equation can only include a handful of indicators out of these thousands. What happens if, as is almost inevitable, the Brown Book equation we estimate ignores a useful indicator?

I now show that the Brown Book method is robust to the problem of missing indicators, in the sense that the recommended reaction function constitutes a constrained improvement, or second best solution, conditional on the restricted indicator set. This robustness is a major strength of the Brown Book method. How many alternative methods of targeting inflation can claim a similar immunity to the omitted variable bias, which must, in practice, contaminate all macro-econometric models?

Let us divide the indicator set X_t into two subsets: Z_t and Y_t, and rewrite equation (1) accordingly as:

1b. $E(P_{t+8} / I_t) = E(P_{t+8} / \{R_t, X_t\}) = -bR_t + cX_t = -bR_t + dZ_t + fY_t$

Similarly, we can rewrite equation (4a) as:

4b. $R_t + (1/b)P_{t+8} = (c/b)X_t + (1/b)e_{t+8} = (d/b)Z_t + (f/b)Y_t + (1/b)e_{t+8}$

Suppose that the econometrician estimating the Brown Book equation observes only the subset of indicators Z_t, and does not observe the remaining indicators Y_t (or else the econometrician does not have enough degrees of freedom to include Y_t). Y_t thus becomes an omitted variable in the Brown Book equation. Clearly, if Y_t is uncorrelated with Z_t, the presence of the omitted variable will not bias the econometrician's estimate of the parameter (d/b). But, in general, the omitted indicators Y_t will usually be correlated with the included indicators Z_t, and then the standard omitted variables bias will mean that the estimate of the parameter (d/b) will now be biased. Is this a problem with the Brown Book method?

To see that this 'bias' is not in fact a problem, suppose that Y_t and Z_t are related according to:

11. $Y_t = gZ_t + u_t$

where u_t is a mean zero random error uncorrelated with Z_t, so that taking expectations over equation (11) we get

12. $E(Y_t /Z_t) = gZ_t$

Now let us take expectations over equation (1b) conditional upon the information set Z_t:

1c. $E[E(P_{t+8} /I_t)/Z_t] = -bR_t + cE(X_t /Z_t) = -bR_t + dZ_t + fE(Y_t /Z_t) = bR_t + dZ_t + fgZ_t$
$= -bR_t +(d+fg)Z_t$

The optimal reaction function for inflation targeting, subject to the constraint that it include only the subset of indicators Z_t, now becomes (by setting equation (1c) equal to zero):

2b. $R_t^* = ((d+fg)/b)Z_t$

We can think of equation (2b) as the second best reaction function, as compared to the first best reaction function, which can be written as:

2c $R_t^* = (d/b)Z_t + (f/b)Y_t = ((d+fg)/b)Z_t + (f/b)u_t$

Substituting equation (11) into equation (4b) we can rewrite it as:

4c. $R_t + (1/b)P_{t+8} = ((d+fg)/b)Z_t + (f/b)u_t + (1/b)e_{t+8}$

The Brown Book equation, as estimated by the econometrician on the restricted indicator set Z_t, now becomes:

5a. $(R_t + aP_{t+8}) = rZ_t + av_{t+8}$

Notice that the residual v_{t+8} in equation (5a) is a linear combination of u_t and e_{t+8}, each of which is uncorrelated with Z_t, so v_{t+8} is also uncorrelated with Z_t, thus satisfying the requirements for unbiased estimation of 'r'. Notice also that as the small positive constant 'a' in the Brown Book equation approaches (1/b), the coefficient 'r' estimated by the econometrician approaches an unbiased estimate of the parameter ((d+fg)/b). In other words, as 'a' approaches (1/b), the econometrician estimating the Brown Book equation will be generating an unbiased estimate of the second best optimal reaction function, constrained by the smaller indicator set Z_t.

What we conclude from this is that omitting useful indicators from the Brown Book equation does not in fact create a problem of biased estimates, even if the omitted indicators are correlated with the included indicators. It is correct to say that the estimated coefficient will be a biased estimate of the parameter on the included indicators in the first best reaction function, which includes all indicators, but that fact is irrelevant. What is important is that the estimated coefficient be an unbiased estimate of the parameter in the second best reaction function, subject to the constraint that the reaction function include only the indicators included by the econometrician in the Brown Book equation.

This does not mean that we can ignore the problem of missing indicators altogether. First, because the Brown book equation is 'noisier' with omitted indicators, the estimates of the second best reaction function may sometimes be less efficient, having a greater sampling variance, even if they are unbiased. Second, the second best reaction function, even if estimated perfectly, will obviously be a worse reaction function than the first best optimal reaction function, in the sense that following the second best reaction function will lead to a higher variance of inflation around the target than following the first best.

This second point above has practical importance. The Brown Book method works by looking for systematic mistakes in how the Bank responds to indicators, and improving its response to those indicators. If the econometrician estimating the Brown Book equation includes all the indicators which the Bank observes, as was assumed in section 3 above, then we know that the Brown Book method will improve inflation targeting (subject only to the caveats listed at the end of section 3). But if the econometrician includes a smaller set of indicators than the Bank observes and responds to, our conclusion is less certain. Following the Brown Book advice now has disadvantages as well as advantages compared to existing monetary policy. The advantage is improving the response to those indicators included in the Brown book equation. The disadvantage is ignoring the additional useful information contained in those indicators excluded from the Brown Book equation.

There are three things we can do to try to ensure that the advantages of following Brown Book advice outweigh the disadvantages. First, we need to ensure that the indicator set used in the Brown Book equation is a good one, so that omitted indicators contain little useful additional information. Second, even though direct estimation of the Brown Book equation is simpler than estimating the actual reaction function, estimating the forecasting equation, and adding them together, we still need to estimate a forecasting equation, because the forecasting equation gives us a measure of the magnitude of the advantages of following Brown Book advice. If the forecasting equation has little explanatory power, that would mean that the Bank is already responding close to optimally to the econometrician's indicator set, and so there would be only a small advantage to be gained from improving that response. With a small advantage, the disadvantage from omitted indicators would be more likely to outweigh the advantage from improved response. Third, we need to test Brown Book advice against actual policy by running a 'Horse Race' test, as described below in section 6, to test whether the advantages of following Brown Book advice would have outweighed the disadvantages over the sample provided by recent history, and to decide on the optimal weight to place upon Brown Book advice.

5. Choosing the Best Indicator Set

There is a large number of possible indicators of future inflation that are available to the Bank. The number of possible indicator *sets*, which are combinations of various current and lagged indicators, is even larger. Brown Book recommendations for any given FAD will vary according to the indicator set chosen. Choosing the best indicator set to use is the most important practical problem in applying the Brown Book method.

Suppose we know that (say) current core inflation is a good indicator of future CPI inflation. What does this mean? From the perspective of a private sector inflation forecaster, this would mean that current core inflation is strongly correlated with future CPI inflation.[10] But the perspective of a private sector forecaster, who is trying to *forecast* inflation, is very different from the perspective of a central bank, that is trying to *control* inflation. If the bank is reacting optimally to core inflation, then rational expectations implies that there will be zero correlation between that indicator and future deviations of inflation from target, so that core inflation will be useless for a private sector forecaster. Indeed, when monetary policy is optimal for inflation targeting, and

[10] This would be a simple correlation if the forecaster is using only one indicator, and a partial correlation if the forecaster is using multiple indicators.

responds optimally to all available indicators, *all indicators will be useless for a private sector forecaster seeking to forecast inflation at the Bank's two year targeting horizon.*[11]

From the perspective of a private sector forecaster, a good indicator is one which *does have* good explanatory power for future inflation. From the perspective of an inflation targeting central bank, a good indicator is one which *would have* good explanatory power for future inflation *if the bank were to ignore that indicator.*

Suppose that core inflation is a good indicator from the perspective of an inflation targeting central bank. If the bank ignored this indicator, then core inflation would have zero explanatory power in the actual reaction function, but have high explanatory power in the forecasting equation. But if the bank responded to this indicator optimally, then core inflation would have high explanatory power in the actual reaction function, but have zero explanatory power in the forecasting equation.

Suppose that astrological portents are a useless indicator from the perspective of an inflation targeting central bank. If the bank ignored this indicator (responded to it optimally), then astrological portents would have zero explanatory power in the actual reaction function, and zero explanatory power in the forecasting equation. But if the bank responded to this useless indicator, then astrological portents would have high explanatory power with a positive coefficient in the actual reaction function, and high explanatory power with a negative coefficient in the forecasting equation.[12]

To measure the usefulness of an indicator to an inflation targeting central bank, what we need is an appropriately weighted sum of the explanatory power of that indicator in the two estimated

---

[11] Unless of course the private sector forecaster has private indicators that are not observed by the Bank.

[12] The negative coefficient is because the bank raises the overnight rate when it sees astrological portents of inflation, which causes future inflation to fall below target. Note that astrology would then become a very useful (contrary) indicator for a private sector forecaster.

equations: the actual reaction function; and the forecasting equation. A positive coefficient implies a positive weight, and a negative coefficient implies a negative weight. With the weights chosen correctly, a useless indicator, like astrological portents, would have a zero sum. A useful indicator, like core inflation, would have a positive sum.[13] And a useful contrary indicator, like the unemployment rate, should have a negative sum.[14] The absolute value of the appropriately weighted sum of the explanatory powers would be our measure of the indicator's usefulness. But what are the appropriate weights? The appropriate weights should presumably be one on the actual reaction function, and (1/b) on the Forecasting Equation, which are the same as the optimal weights for constructing the Brown Book equation. Then if the Bank were responding to astrological portents, for example, the weighted average of the estimated coefficients in the reaction function and forecasting equations should exactly sum to zero.

A further problem is that the Bank is not restricted to using a single indicator; it can use a set of many indicators. The usefulness of an indicator considered in isolation may not necessarily reflect the marginal usefulness of adding that indicator to a set of other indicators. For example, two different measure of core inflation may each be good indicators considered in isolation, but if the two are highly correlated with each other then the set of the two indicators will be little better than either of them alone. Conceivably also, two indicators might each be useless alone, but both indicators together might be a useful indicator set. What we really need is some measure of the usefulness of an indicator *set*, so we can choose the most useful set of indicators with which to estimate our Brown Book recommended reaction function.

---

[13] Note that if the bank were to overreact to a useful indicator, for example by raising the overnight rate too much when core inflation increased, then the coefficient in the actual reaction function would be larger, but the coefficient in the forecasting equation would now become negative, while an appropriately weighted sum would stay constant, reflecting the same underlying usefulness of the indicator regardless of whether policy is optimal or not.

[14] Because either the bank responds optimally to unemployment, giving it a negative coefficient in the actual reaction function and zero coefficient in the forecasting equation, or it ignores unemployment, giving a zero coefficient in the actual reaction function and a negative coefficient in the forecasting equation.

We know that, in practice, insufficient degrees of freedom will ensure that we will estimate the Brown Book equation on less than the full set of potentially useful indicators. There will be missing indicators. We are in the world of second best. Actually, we are not even in the second best world, because sampling variation will cause the estimated second best reaction function Rt=r^Zt to differ from the true second best reaction function Rt=rZt..We are in the world of the third best. We face a trade-off. The more indicators we include in the Brown Book equation, the more information is contained in the indicator set, and the second best reaction function gets closer to the first best. But given a small data sample, adding additional indicators may also reduce the efficiency of our estimate, and our third best estimate may get further away from the second best truth.

Suppose we estimate the Brown Book equation on a given indicator set Z_t, and then set monetary policy according to the recommended reaction function, $R_t = r\hat{}Z_t$. Substituting the recommended reaction function into equation (1), and recalling that the first best optimal reaction function is $R_t^* = (c/b)X_t$, we get:

13. $E(P_{t+8} /It) = -bR_t + cX_t = -b(R_t - R_t^*) = -b(r\hat{}-r)Z_t + b(f/b)u_t = -b(r\hat{}-r)Zt + fu_t$

Recall that fu_t represents the additional information contained in the missing indicators. Equation (13) lets us see the trade off. Increasing the size of the included indicator set Z_t means that we reduce the amount of information lost from the missing indicators, fu_t, thus reducing the expected deviation of inflation from target. But at the same time we may reduce the efficiency of our estimate r^, and so increase (r^-r), which increases the expected deviation of inflation from target. We seek to determine the size and composition of the included indicator set Z_t in such a way as to minimise the sum of these two errors. No doubt a skilled econometrician could solve this problem, but I am not a skilled econometrician. My guess is that choosing the indicator set Z_t to maximise the adjusted R- squared of the Brown Book equation is the way to choose the best indicator set. If my guess is right, then choosing the indicator set to maximise the adjusted R-squared of the Brown book equation would gives us the best available Brown Book recommendation, *if that Brown Book recommendation were going to be followed exactly*.

But it is very unlikely that it would be optimal for the Bank to follow exactly any single source of advice. More likely, each of several sources of advice would contain some useful information which is at least partly independent of other sources. If so, the Bank should optimally follow some weighted average of its several sources of advice. And if the Bank will not follow Brown Book advice exactly, but instead will place the optimal weight on the Brown Book compared to other sources of advice, then we need to redefine what we mean by choosing the optimal indicator set for the Brown Book. The optimal indicator set Z_t should be chosen to minimise the variance of inflation around the target, not when the Bank is following the Brown Book exactly, but when the Bank is placing the optimal weight w* on Brown Book advice. Again, I am not a skilled econometrician, but my guess is that this objective will be satisfied when the indicator set is chosen to maximise this optimal weight w*. In Rowe (2003) I show that w*=(h/b) , where 'h' is a coefficient estimated in the "Horse Race" equation discussed below in section 6. If my guess is correct, we should then choose the indicator set to maximise the coefficient 'h' estimated in the "Horse Race" equation, for this will maximise the optimal weight w* to place on Brown book advice, and thereby maximise the marginal usefulness of that advice over and above existing sources of advice.

6 Testing the Recommended Reaction Function - The Horse Race

The Brown Book method can be used to estimate a recommended reaction function, which is intended to improve monetary policy, in the sense that following the recommended reaction function, or putting at least some positive weight on Brown Book recommendations, should keep inflation closer to target.

There is a simple way to test any proposed new recommended reaction function, like the Brown Book, compared to historical monetary policy. I call this the "Horse Race". First we construct a historical series for the Brown Book recommended overnight rate, Bt, by substituting the indicator set, for each point in time, into the recommended reaction function. The Horse Race is a simple regression of deviations of inflation from target on two year lagged deviations of recommended overnight rate from actual overnight rate, with no constant term in the regression.

14. $P_{t+8} = h(Bt - Rt) + vt$

where Bt represents the Brown Book recommended overnight rate, and Rt represents the actual overnight rate. We hope that the estimated coefficient 'h' is positive and statistically significant.

The idea underlying the Horse Race test is very simple. Suppose the recommended reaction function specifies a higher overnight rate at time t than the actual overnight rate at time t. In effect, the recommended reaction function is saying that monetary policy was too loose, and that inflation will rise above target two years later. And if the recommended reaction function specifies a lower than actual overnight rate, then the recommended reaction function is saying that monetary policy was too tight, and that inflation will fall below target two years later. These implicit predictions are testable, and are tested in the Horse Race.

To understand how to interpret the sign of then estimated coefficient in the Horse Race, it helps to think of the Brown Book recommended overnight rate, or the deviation between the recommended and actual overnight rate, as just another indicator, to which monetary policy should or should not respond. A zero estimated coefficient 'h' in the Horse Race conforms to the requirement that rational expectations forecast errors be uncorrelated with the indicator set; it means that the Bank was implicitly already responding optimally to this indicator (the Brown Book). The implication is that policy could not be improved by explicitly recommending the Bank follow the Brown Book method. And a negative estimated coefficient 'h' suggests that Brown Book recommendations point monetary policy in the wrong direction; whenever the Brown Book says that monetary policy is too tight, it is probably too loose, and vice versa. Ironically, the Brown Book method could then nevertheless improve monetary policy, but only if the Bank did the exact opposite of what the

Brown Book recommended. For example, if the Bank was planning to set the overnight rate at 5%, and then heard that the Brown Book was recommending 6%, the Bank should then set the overnight rate below 5%.

If the estimated coefficient is positive, and statistically significant at some level of confidence, this means that monetary policy could be improved (with the same level of confidence) if the Bank put some positive weight on the Brown book recommendation. If the Bank was planning to set the overnight rate at 5%, and then heard that the Brown Book was recommending 6%, the Bank should then set the overnight rate above 5%.

To understand how to interpret the magnitude of the estimated coefficient in the Horse Race, consider the following: A positive coefficient of (say) 0.5 would mean that if the actual overnight rate is 5%, and the Brown Book recommends 6%, future inflation two years hence will on average be 0.5 percentage points above target. A positive coefficient of (say) 0.3 would mean that future inflation would be only 0.3 percentage points above target. Other things equal (in particular, for an equal deviation of Brown Book recommendation from actual policy), the bigger the coefficient, the bigger the potential gains from putting a positive weight on Brown Book advice, and the bigger the weight that should be placed on Brown Book advice.

Suppose that the Brown Book recommendation is made available to the Bank, and the Bank thereafter places some weight 'w' on the Brown Book recommendation. In other words, the Bank now sets an overnight rate R't, where:

15. $R't = wBt + (1-w)Rt = Rt + w(Bt-Rt)$

where Rt is the overnight rate that it would have set if the Brown Book recommendation were not available, as in the past.

If the Brown Book recommendation Bt contains useful marginal information, because it removes some systematic mistakes in inflation targeting, and if actual policy Rt also contains useful marginal information, because it reacts more or less correctly to indicators which are omitted from the Brown Book, then optimal policy will look at both Bt and Rt, and put a positive weight on both.

I have shown in Rowe (2003) that the Horse Race test can in principle be used to give some measure of the optimal weight to place on a new source of monetary policy advice. The optimal weight to place on Brown Book advice, w*, to minimise the variance of inflation around target, is shown to be:

16. $w^* = h/b$

Unfortunately, we need to know the parameter ‘b’ before we can say exactly how much weight we should place on Brown Book advice.

To understand the above result intuitively, subtract equation (2) from equation (4) to express the deviation of inflation from target as a function of the deviation between the optimal and actual overnight rate:

17. $P_{t+8} = b(R^*t - Rt) + et+8 = b[(R^*t - Bt) + (Bt - Rt)] + e_{t+8}$

The Horse Race equation (14) can be seen as an attempt to estimate equation (17) with an “errors in variables” problem, because the Brown Book recommendation Bt is a noisy measure of the true optimal overnight rate R*t. If the Brown Book recommendation were perfect, so that Bt=R*t, then the estimated coefficient ‘h’ in the Horse Race equation would be an unbiased estimate of the parameter ‘b’, and so the optimal weight to place on the Brown Book recommendation would then be w*=h/b=b/b=1. The Bank should then, naturally, follow Brown Book recommendations exactly. But if the Brown Book recommendation is a noisy estimate of R*t, the “errors in variables” problem will cause our estimate of ‘h’ to be biased downwards from ‘b’, so w*=h/b<1, and naturally we should put less weight on the Brown Book recommendation.

The Horse Race test does not test whether following the Brown Book recommendations exactly would have lead to better inflation targeting than actual monetary policy over the sample period. (Though this conclusion would presumably be valid if the optimal weight w* were greater than 0.5). Instead, it tests whether actual monetary policy could have been improved if some positive weight had been placed on the Brown Book recommendations. But that second question is the more relevant question. When we propose a new source of advice, what matters is whether that new advice can supplement existing sources of advice; not whether it can completely replace existing sources of advice. It is rarely optimal to listen to only one source of advice and ignore all others.

The Horse Race test also allows us to reexamine the role of 'a', the small positive constant by which we multiplied the Forecasting Equation before adding it to the estimated actual reaction function to generate the Brown Book recommended reaction function. We can rewrite the Horse Race equation as:

14' $P_{t+8} = h[(Bt\text{-}At) + (At\text{-}Rt)] + v_{t+8}$

where At represents the fitted values from the estimated actual reaction function. Moreover, since Bt=At + aFt, where 'a' is the small positive constant, and Ft is the fitted values from the Forecasting Equation, we can rewrite the Horse Race equation again as:

14" $P_{t+8} = h[aFt + (At\text{-}Rt)] + v_{t+8}$

Previously we had worried about setting the appropriate magnitude for 'a'. If we set 'a' too small, the Brown Book would not be different enough from actual policy to eliminate the systematic errors revealed by the Forecasting Equation. And if we set 'a' too big, the Brown Book would over-adjust for the systematic errors, and following Brown Book recommendations exactly might make systematic errors of the opposite sign. We will now see that it makes no difference how large or small we set 'a', provided we also use the Horse Race test to recommend the optimal weight w*=h/b to place on the Brown Book. By using (14) and (14") we can rewrite (15) as:

15' $R't = Rt + w(Bt-Rt) = Rt + w^* [aFt + (At-Rt)] = Rt + (w^*/h)[E(Pt)] = Rt + (1/b)E(P_{t+8})$

And we see that R't is now independent of 'a'. The reason is that as we increase 'a', and increase the variance of (Bt - Rt), the estimated value of 'h' decreases correspondingly, and so does the optimum weight w*=h/b we place on (Bt-Rt).

To understand the Horse Race better, we can think of actual monetary policy (the historical series for the actual overnight rate) as being the sum of the fitted values of the estimated actual reaction function plus apparently "random errors" around those fitted values. These "random errors" cannot, of course, be truly random errors, because the Bank of Canada does not play dice with monetary policy. It is better to interpret these "random errors" as responses by the Bank to indicators which the Bank observes but which are not included in the econometrician's indicator set. As discussed above, there will almost always be some useful indicators which are omitted from the Brown Book's indicator set. The advantage of the Brown Book method is that it can improve the Bank's response to the variables included in the econometrician's indicator set (unless actual policy is already optimal in that regard). The disadvantage of the Brown Book method is that it ignores some potentially useful omitted indicators. The Brown Book method can make a net improvement to monetary policy if and only if this advantage outweighs this disadvantage.

A Horse Race between the fitted values of the recommended reaction function and the fitted values of the estimated actual reaction function, that is a regression of P_{t+8} on (Bt-At), or equivalently of P_{t+8} on aFt, would have a predictable winner. By construction, the recommended reaction function could not lose this race, because it adjusts for revealed systematic mistakes in the estimated reaction function; there could only be a tie if the two functions were identical. The Brown Book recommended reaction function must improve inflation targeting compared to the estimated actual reaction function (unless the Forecasting Equation reveals no systematic mistakes), and a Horse Race between those two alternatives would never be lost by the Brown Book. But the Brown Book recommended reaction function may or may not improve inflation targeting compared to actual policy, because actual policy might improve on the fitted values of the estimated reaction function, provided the Bank is responding roughly correctly to the omitted indicators. And the sign of the coefficient ‘h’ in a Horse Race between the Brown Book and actual policy also cannot be determined a priori. To see this, look at equation (14") again.

The term aFt in equation (14") is a linear function of the indicator set. The term (At-Rt) in (14") represents the residuals from a regression of Rt on the indicator set, and is therefore orthogonal to the indicator set. So (At-Rt) is orthogonal to aFt. If the Bank is responding optimally to the omitted indicators, then P_{t+8} will also be uncorrelated with (At-Rt). The estimate of ‘h’ in the Horse Race will then necessarily be positive. But if the Bank is under-reacting to the omitted indicators, the estimated coefficient ‘h’ in the Horse Race could be negative. When the Bank raises Rt, and so lowers (At-Rt), in response to inflationary signals from the omitted indicators, if it does not respond strongly enough tot he omitted indicators, future inflation will rise above target, so there will be a negative correlation between (At-Rt) and P_{t+8} . If this negative correlation is stronger than the necessarily positive correlation between P_{t+8} and (Bt-At), then the Horse Race coefficient ‘h’ would be negative. And if ‘h’ is indeed negative, then placing a positive weight on the Brown Book recommendation would indeed worsen monetary policy. The reason is that placing a positive weight on the Brown book would reduce the weight on the omitted indicators, weakening an already too weak reaction to those indicators. If the advantage from improving the reaction to the included indicators is outweighed by the disadvantage from worsening the reaction to the omitted indicators, then placing a positive weight on Brown Book advice would on balance worsen monetary policy. To guard against this possibility, it is important to estimate the Horse Race equation, and to reject Brown Book advice if the estimated coefficient h is not positive.

The Horse Race is designed to test whether the Brown Book recommended reaction function could have been used to improve past monetary policy. But since the Brown Book’s recommended reaction function was estimated over a data set, not all of which was available to past policymakers, the test is perhaps unfair. The Brown Book has the benefit of hindsight. For example, if we use data from 1990 to 2000 to estimate the Brown Book recommended reaction function, and use it to give advice on monetary policy in 1996, we have an advantage over the Bank, which was setting monetary policy in 1996 lacking data from 1997 to 2000. To circumvent this problem, an alternative procedure is to run the Horse Race using a real time data sample. This means we would compare actual policy in 1996 with a Brown Book estimated on a data sample which stops in 1996, and actual policy in 1997 with a Brown Book estimated on a data sample which stops in 1997, etc. The real time sample Horse Race is a better test of the Brown Book; but unfortunately inflation targeting is too recent a policy for this test to have much power currently. With 10 years total data, for example, if we need a minimum of (say) 6 years to estimate a Brown Book, we have only 4 years of data on which to run the Horse Race.

The discussion above assumes that the Horse Race is being used to test the usefulness of a *new* source of policy advice. A very different interpretation should be placed on the Horse Race if it is being used to evaluate the usefulness of a *previously existing* source of policy advice.

Any source of monetary policy advice, which specifies a recommended setting for the overnight rate, can be thought of as analogous to an indicator of inflation. Like any other indicator, the recommended overnight rate is just a variable which, it is argued, the Bank ought to look at when setting the overnight rate. Suppose Statistics Canada started measuring and publishing a totally new indicator. Evaluating the usefulness of a new indicator is very different to evaluating the usefulness of a previously existing indicator. For a previously existing (old) indicator, as I have argued above, we need to look both at the forecasting equation and at the estimated actual reaction function. For a new indicator, we need only look at the forecasting equation. If we wanted to evaluate the usefulness of that new indicator, we would estimate the forecasting equation on that new indicator. If it had explanatory power in the forecasting equation we would consider it a useful indicator. If not, we would consider it a useless indicator. We do not need to estimate whether that new indicator has explanatory power in the actual reaction function; by virtue of its being a new indicator, previously unavailable, the Bank could not have been responding to that indicator in the past.

This procedure is very different to the procedure for evaluating a previously existing indicator. For an indicator which was available to the Bank over the data period, if the Bank had been responding optimally to that indicator, the indicator should have zero explanatory power in the forecasting equation. But if the old indicator is a good indicator, it should nevertheless have good explanatory power in the actual reaction function.

New sources of advice are like new indicators; old sources of advice are like old indicators. For a new source of advice, we need look only at the Horse Race equation, which is simply a variant on the inflation Forecasting Equation. For old sources of advice, we need to look both at the Horse Race equation and at the explanatory power of the old advice in the estimate of the actual reaction function.

Suppose there is an old source of advice on monetary policy, which is a good source of advice, and to which the Bank had been responding optimally. By virtue of the Bank's having been responding optimally to that advice, the coefficient in the Horse Race equation should be zero. A positive coefficient would mean that the Bank should have put more weight on that advice; a negative coefficient would mean the Bank should have put less weight on that advice. And if the Bank had been responding optimally, then by virtue of it being good advice, an estimate of the actual reaction function, regressing the actual overnight rate on the recommended overnight rate, should have high explanatory power and a coefficient of one. In other words, how do we know if advice is good advice? If you do what it says (high explanatory power and unit coefficient in the actual reaction function), and if doing so does not lead to systematic errors (zero coefficient in the Horse Race or Forecasting Equation), then the advice must be good advice.[15] This topic is

[15] There is a subtle difference between the Horse Race equation and the Forecasting Equation. The Horse Race regresses deviations of inflation from target on two year lagged deviations of recommended from actual overnight rate (with no constant). A pure Forecasting

discussed more completely in Rowe (2003).

Since the Brown Book is a new source of advice, it must have had zero weight in determining previous monetary policy. A zero coefficient in the Horse Race equation would mean that the existing weight (zero) was optimal. A positive coefficient in the Horse Race equation says that more weight (and hence a positive weight) should be placed on Brown book advice. But in future, if the Brown Book has been published and available to the Bank over the sample period, and has been given some weight in policy determination, then a zero coefficient would not mean that the Brown Book is useless; instead it would mean that the Bank should continue to place whatever weight it had previously placed on Brown Book advice.

7. Leads and Lags

We have assumed that the Bank of Canada targets inflation at a two year horizon. It sets the overnight rate, as a function of the indicators, so that inflation two years from now is expected to equal the target. It is theoretically possible that today's inflation depends only on the values of the instrument and indicators precisely two years ago. But this is extremely unlikely. More plausibly, today's inflation depends on a whole distributed lag of past instrument and indicators. Recognising this fact does not invalidate the Brown Book method; but it may influence our choice of the indicator set.

First, suppose that today's inflation depends on the unemployment rate two years ago and three years ago. This case presents no problem. When we estimate the Brown Book equation, we need to include in the indicator set not only current unemployment but also last year's unemployment.

Equation would regress deviations of inflation from target on two year lagged recommended overnight rate (with a constant). A positive estimated coefficient in the pure Forecasting Equation says that the Bank should reacted more strongly to the recommendations from that source of advice. A positive coefficient in the Horse Race equation says the Bank should put more weight on that source of advice (and less weight on other source of advice).

The indicator set represents information available to the Bank when it sets the overnight rate. Unless there is some peculiar loss of records or memory, lagged indicators can also be included in the indicator set, and should be included if they prove useful.

Second, suppose that today's inflation depends on the unemployment rate two years ago and one year ago. This case also presents no problem, but for an entirely different reason. Unless the Bank has a crystal ball, it does not know what the unemployment rate will be one year in advance, so it is not legitimate to include it in the Brown Book equation. If we were to do so, we could not implement the recommended reaction function. Now, we could include a forecast of the future unemployment rate in the Brown Book equation, but this forecast would be a function of some current indicators. An indirect forecast of future inflation, conditional on forecasted unemployment, conditional on current indicators, would be the same as a direct forecast of future inflation conditional on current indicators. So it is simpler to estimate the Brown Book equation directly on current indicators. Conceivably, however, there may be cases where the Bank has knowledge of the future, which knowledge may not be reflected in normal indicators. An example might be an announced change in future tax policy, which could affect future inflation. In principle we could use the government's announcement itself as a current indicator, but it might be simpler just to use the future tax policy change itself.

Third, suppose that today's inflation depends on the overnight rate two years ago and three years ago, so that monetary policy affects inflation with a distributed lag. This case again presents no problem. The Bank's own past policy actions are part of its current information set. We include the lagged overnight rate in the indicator set. Tighter past monetary policy, other things equal, would presumably mean that current monetary policy should be looser.[16]

Fourth, suppose that today's inflation depends on the overnight rate two years ago and one year ago. This means that the Bank's targeting horizon is longer than the shortest lag of the monetary

[16] But see the discussion of the Wicksell effect below.

transmission mechanism. This case is more complicated, because the expected future overnight rate, and hence the recommended current overnight rate, depends on whether future monetary policy will follow the actual reaction function or the recommended reaction function. The Brown Book method works by looking at the indicators over past history and forecasting when future (two year ahead) inflation will be above (below) target and recommending a higher (lower) overnight rate. Since the recommended reaction function was not in fact followed during that historical period, the Brown Book recommended reaction function implicitly assumes that future monetary policy will follow the actual reaction function, not the recommended reaction function.

Suppose we assume that

$E(P_{t+2} /It) = -b1E(R_{t+1} /It) -b2Rt +cX_t$

Given a reaction function Rt = rXt, which we substitute into the above equation to get:

E(Pt+2/It) = -b1rE(Xt+1/It) -b2rXt +cXt

Suppose that E(Xt+1/It) = dXt + fRt

E(Pt+2/It) = -b1r(dXt + fRt) -b2rXt +cXt =-b1r(dXt + frXt) -b2rXt +cXt

Setting E(Pt+2/It) =0 and differentiating wrt Xt to solve for the optimal reaction function

0 = -b1rd -b1rfr- b2r + c

r= ?

Damn! What’s the quadratic formula?

8. A Wicksellian Perspective

The Brown Book method requires almost no theoretical knowledge. That is a major strength. But given the very large number of potential indicator sets, and the short data series over which to compare different indicator sets, it can also be a major weakness. The dangers from data mining are very real. If a given indicator set seems to work well for the existing data series, in the sense of creating a Brown Book equation with high explanatory power, can we be confident that it will continue to work well in future? As time passes, and the data series becomes longer, this weakness will become less important, but in the early years of inflation targeting it is potentially a major problem. Any theoretical guidance on what variables ought be included in the indicator set can be helpful. This section argues, from a Wicksellian perspective, that we ought to include the inflation rate and the lagged overnight rate in the indicator set.

Consider a macroeconomic model in which money is, in the long run equilibrium, neutral and superneutral. That means that an exogenous change in the growth rate of the money supply, and the concomitant change in the long run equilibrium rate of inflation, will leave the long run equilibrium values of all real variables unchanged. Within the context of such a model, it is possible to define a "natural rate" for any real variable, like the unemployment rate for example, as simply the long run equilibrium value for that variable. For a central bank which uses an interest rate as its monetary policy instrument, Wicksell's "natural rate of interest" is a very useful theoretical concept. If the bank sets the overnight rate such that the real overnight rate is permanently less than the natural rate, inflation will increase without limit. If it sets it permanently above the natural rate, inflation will decrease without limit. And if it sets the real overnight rate equal to the natural rate, the economy will be in neutral equilibrium, which is compatible with any inflation rate.

This perspective illuminates the debate between those who think of monetary policy in terms of the appropriate *level* of the overnight rate, and those who think of monetary policy in terms of the appropriate *change* in the overnight rate. To put it another way, should we estimate the actual and recommended reaction functions in levels, or in first differences? To understand the substantive issue involved, as opposed to the merely semantic, note that if inflation is on target, and the economy appears to be neither in boom or recession but in equilibrium, those who think in levels would argue that the Bank should immediately set the overnight rate at its normal or average level, say 5 per cent. But those who think in first differences will argue that the overnight rate should stay constant at whatever level it currently is, even if that is well above or below the 5 per cent average. Who is right?

Suppose that the natural rate were a known constant, equal to (say) 3 per cent, or were white noise around that known constant. With inflation equal to the 2 per cent target, setting the nominal overnight rate at 5 per cent would be, on average, compatible with inflation staying on target. But even though a 5 per cent overnight rate would be compatible with inflation staying on target, it would not ensure that inflation stayed on target. If inflation, for whatever reason, were 3 per cent, as soon as expected inflation increased to 3 per cent also, the real overnight rate would fall to 2 per cent, below the natural rate, and inflation would then rise even further.

To keep inflation from drifting away from target, the bank would then need to adjust the overnight rate more than in proportion to the deviation of inflation from target. A simple rule that would ensure this would be to set the overnight rate according to $R_t = 5 + sP_t$, where the coefficient 's' is greater than one, to ensure that the real overnight rate increases when inflation rises above target.[17]

[17] It seems plausible that 's', the coefficient on inflation, should exceed one. If it were less than one, then a permanent increase in inflation, and hence in expected inflation, would mean a fall in the real overnight rate, causing a further increase in inflation. But suppose that inflation is measured with some random error, or that there are temporary fluctuations in inflation which do not feed through one for one into expected inflation. In either case, measured inflation will be a noisy signal, and the appropriate monetary policy rule should discount fluctuations in inflation accordingly, giving a value of 's' which may be less than one. But with 's' less than one, any true increase in inflation, and expected inflation, would lead to accelerating inflation. Perhaps the resolution to this paradox would be to make the monetary policy rule non-linear in inflation, so that small (or temporary) changes in inflation would be partially discounted, but larger (or more permanent) changes in inflation would trigger a more than one for one response. Core inflation might also be a less noisy signal and a better indicator.

Now suppose instead that the natural rate of interest follows a random walk, and is not observed by the bank. The simple rule above, $R_t = 5 + sP_t$, will no longer work to bring inflation back to target. For example, if the natural rate drifted up from 3 per cent to 4 per cent, and stayed there, it would require a permanent deviation of inflation from target $P_t = 1/(1-s)$ in order to bring the real overnight rate up to the new natural rate. If inflation were equal to target, optimal policy would be to set the overnight rate at 6 per cent, to keep the real overnight rate equal to the natural rate, but the simple rule would be recommending 5 per cent instead.

An alternative simple rule that would prevent inflation from drifting away from target, when the natural rate follows a random walk, would be to set the overnight rate according to $R_t = R_{t-1} + sP_t$. If inflation is above target, the bank realizes that in the past it had set the overnight rate below the natural rate, and must increase it by enough to get the new real rate above the natural rate. So long as inflation remains above target, the bank continues to increase the overnight rate.

Notice the contrast between the case with a constant natural rate and the case with a random walk natural rate. In the first case, the appropriate response to inflation being above target is to set a *high* overnight rate; in the second case it is to set an *increased* overnight rate. The first rule is specified in levels; the second rule is specified in first differences.

Reality almost certainly lies somewhere between these two extremes. There are shocks to the natural rate of interest, and these shocks do persist, but they do not persist forever. (If they did persist forever, the natural rate would have infinite long run variance, which it does not seem to have). If we assume more reasonably that the natural rate follows an autoregressive process, which is a weighted average of white noise and a random walk, then the appropriate monetary policy rule should likewise be a weighted average of the two simple rules above, namely set the overnight rate according to $R_t = 5 + d(R_{t-1} - 5) + sP_t$, where $0<d<1$ reflects the degree of autocorrelation in the natural rate. The truth almost certainly lies somewhere between the two extremes of the level theorists and the first difference theorists. We recognise this by including the lagged overnight rate in the Brown book equation; we do not constrain its coefficient to be either zero or one, and expect its coefficient to lie strictly between zero and one.

The above arguments lay out the theoretical grounds for including both the current inflation rate and the lagged overnight rate (or some closely related indicators) in the Brown Book indicator set. What other variables might one argue should also be included?

One important candidate for inclusion would be some measure of excess demand, like the unemployment rate or the output gap. Including the output gap would convert the above simple rules into generalized versions of the Taylor Rule.

Continuing with our Wicksellian perspective, the object of the monetary policy rule is to keep the overnight rate as close as possible to the natural rate, provided inflation stays on target. Ideally, if the bank could observe the natural rate directly, then the natural rate itself should be the indicator par excellence. Possibly there are some variables, like the structural fiscal deficit for example, which can help the bank infer changes in the natural rate as soon as they occur. If so, they could be useful indicators for the Brown Book. But otherwise, the bank can only know that the natural rate has fallen below the overnight rate after the fact, when the economy starts to go into recession. If the timing of recessions is such that output and employment fall first, and inflation falls later, then that is an argument for including some measure of excess demand in the indicator set. Doing so allows the Bank to anticipate future deviations of inflation from target, rather than simply waiting for them to occur. Provided there exists a good real time measure of excess demand, it could help the bank detect changes in the natural rate more quickly than if it waited for inflation to change, and should belong in the Brown Book indicator set. But given the difficulties in real time estimation of the natural rates of output and unemployment, it is not obvious a priori if this proviso is met.

A second important candidate for inclusion in the indicator set would be the exchange rate. Clearly, the exchange rate is an important price in an open economy. But it is not obvious whether, for a given overnight rate, exchange rate appreciation indicates a rising or a falling natural rate.[18]. Exchange rate appreciation originating in the capital account of the balance of payments would cause Canadian goods to become more expensive, reducing the demand for net exports, would indicate a fall in the natural rate of interest, to which the Bank should respond by cutting the overnight rate. But an exchange rate appreciation might also be a symptom of increased demand for Canadian net exports. If so, provided the appreciation is not enough to choke off the increased demand for net exports which brought it about, the appreciation indicates a rise in the natural rate, to which the Bank should respond by raising the overnight rate. Thus whether the exchange rate should be included in the indicator set, and what the expected sign on the coefficient should be, cannot be decided a priori. It is an empirical question whether the exchange rate, possibly in combination with other indicators, can serve as a useful indicator.

A third possible candidate for inclusion in the indicator set is the quantity of money. It is almost axiomatic in monetary theory that an exogenous increase in the supply of money, ceteris paribus, will be inflationary. But that does not mean that the Bank, on observing an increase in the quantity of money, should increase the overnight rate to prevent future inflation. The textbook thought experiment, with its ceteris paribus clause, and in which the changed quantity of money is caused by an exogenous change in monetary policy by the central bank, is very different from the thought experiment we are conducting here. We can fully agree with the textbook that an exogenous cut in the overnight rate, which is how our central bank would implement an increased supply of money, would, ceteris paribus, be inflationary. But that is irrelevant to the role of money as an indicator for an inflation targeting central bank with an overnight rate instrument. We are asking a different question.

Suppose we start in equilibrium, with the overnight rate constant, inflation on target, and the quantity of money growing at a constant rate. Then suddenly the bank sees the quantity of money

---

[18] It was the recognition of this point, and the difficulty it created for explaining the stance of monetary policy, which lead to the abandonment of the Monetary Conditions Index (MCI) as an announced measure of the stance of monetary policy.

start to grow more quickly. Should the bank raise or lower the overnight rate in order to keep future inflation on target? Unlike the textbook thought experiment, we are clearly not talking about an exogenous change in monetary policy. Nor are we talking about a ceteris paribus experiment; something must have caused the quantity of money to change. Nor are we necessarily talking about a change in money *supply*; it might equally have been a change in money *demand* which caused the observed quantity of money to increase.

There are different stories we could tell to explain the increased quantity of money, and they will have different implications for whether the bank should raise, lower, or leave unchanged the overnight rate, to keep inflation on target. Possibly there was an increased demand for loans from commercial banks to finance increased expenditure. If so, the increased quantity of money is a reflection of an increased natural rate, and the bank should respond by increasing the overnight rate. Another possibility is that there was a decrease in consumption demand matched by increased savings in the form of money; or a shift in portfolio demand away from consumer durables, equity, or commercial bonds, and into money. If so, the increased quantity of money is a reflection of a decreased natural rate, and the bank should respond by decreasing the overnight rate. Finally, a third possibility is that there was a shift in portfolio demand away from government bonds and into money, which indicates no change in the natural rate (unless government spending is interest sensitive) and the bank should do nothing.

Thus whether the quantity of money should be included in the indicator set, and what the expected sign on the coefficient should be, cannot be decided a priori. It is an empirical question whether money, possibly in combination with other indicators, can serve as a useful indicator.

For indicators like money, or the exchange rate, an increase in the indicator could require either an increase or a decrease in the overnight rate. It would be tempting to say that we must first interpret the underlying reasons for the change in the indicator before we decide how monetary policy should respond. That statement would be true, but unhelpful. The exchange rate does not tell us why it is changing. The only way we can identify the underlying reasons for the change in one indicator is by looking at what is happening to other indicators. But that simply means that we must look at combinations of indicators, or what is happening to a whole indicator set, to decide monetary policy, and we have been assuming that all along.

9 The Lucas Critique

The Brown Book method does not try to derive policy recommendations from an economic model with "deep" parameters which are allegedly policy-invariant. Does that mean that Brown Book policy advice is vulnerable to the Lucas Critique?

One defence would be to argue that the Brown Book recommended reaction function is estimated only on data from the period when the Bank of Canada was targeting 2 per cent inflation, and that the policy advice is not intended to change monetary policy from targeting 2 per cent inflation, but only to better implement that same policy. Since there is no true change in monetary policy regime envisaged, the Lucas Critique does not apply.

There is probably some truth to the above argument, in that the change in policy envisaged is not very big, but ultimately the above argument is a semantic sleight of hand. Whether we call it a true policy regime change or not, if Brown Book recommendations are implemented, there will be a change in the correlation between the overnight rate and the indicators, and this fact may change agents' pricing behaviour for a given overnight rate and indicators, so that the parameters 'b' and 'c' in equation (1) may change as a result of Brown book advice being acted on. How big a problem is this?

One basic Brown Book assumption is immune to this problem. Provided an optimal reaction function does exist, so that there exists a setting for the overnight rate, conditional on the indicators, such that the expectation of future inflation equals the target, then deviations of inflation from target must be uncorrelated with the indicators. That follows directly from rational expectations.

Think of the optimal monetary policy reaction function as the summit of a hill. The Brown Book method can tell us whether we are at the summit, regardless of the Lucas Critique. If we are not at the summit, the Brown Book method can tell us that also, and if the Lucas Critique can be ignored, it can also tell us in which direction the summit lies, though not how far away we have to go to get there. The Brown Book method works by gradient-climbing. Continuing the analogy, the Lucas Critique then warns us that our very act of climbing the hill may shift the location of the summit. But provided the summit shifts more slowly than we are climbing, we will still eventually reach the summit by always climbing uphill. Only if the Lucas Problem is very powerful and works in an unfortunate direction, so that the summit shifts away from us more quickly than our climbing towards it, would we fail to reach the summit by climbing uphill; for then we would do better by climbing downhill, letting the summit come to us from behind.[19]

I do not see any reason to believe this is a real danger. Surely, if the Lucas Critique is to be more than merely nihilistic, the onus should be on the skeptic to argue, in the particular economic application, why we have good reason to believe that a change in policy will cause the parameters to change in a particular way. It is not enough simply to say that they might change. Anything *might* change. No method of giving monetary policy advice is infallible. And if the Brown Book's forecasting equation reveals the existence of systematic mistakes, that also shows that existing methods of monetary policy advice are less than perfect. It would not be reasonable to rely solely on existing methods, even if they are shown to be flawed, and ignore the Brown Book advice, just because it might hypothetically be flawed.

10 Conclusion

---

[19] In a natural rate model, if we tried to use monetary policy to target any real variable, like the unemployment rate, it would be exactly as if the summit moves away from us at the very same speed as we climb towards it. But when Lucas put forward his critique, he did not merely say that this might conceivably happen; he wrote down a more or less plausible model to argue that something like this was very likely to happen. Lucas himself did not use his Critique in a nihilistic way.

Well, that just about wraps things up.

11 References

Rowe N. (2002) "How to Improve Inflation Targeting at the Bank of Canada" Bank of Canada Working Paper.

Rowe N. (2003) "How to Evaluate Monetary Policy Advice for an Inflation Targeting Central Bank, and How to Evaluate the Bank's Use of that Advice" unpublished.