The benefits of price stability: some estimates for New Zealand

Leo Bonato, Economics Department

Inflation is widely believed to undermine the performance of the economies that suffer from it. The reasons why inflation impairs economic performance are well known: it causes uncertainty, distorts the working of the price mechanism for allocating resources, and increases the costs of doing business because the effect of inflation has to be factored into business contracts and transactions.

But empirical evidence to support these propositions generally is not strong, at least not for low to moderate inflation rates. However, recent work by Martin Feldstein, of Harvard University, suggests that the costs of inflation may be significant even at low rates. He focuses on the way in which inflation distorts the tax system and how this adds to what economists call the "dead-weight cost" of taxes. This article outlines the nature of those dead-weight costs and provides some estimates of their magnitude in New Zealand.

1 Introduction

Following the high inflation experience of most OECD countries in the 1970s and 1980s, the 1990s has been a decade when the objective of price stability has appeared largely achievable to many countries. At the same time, a broader consensus on the importance of achieving price stability has emerged among economists and central bankers. These developments have had an impact on how monetary policy frameworks have evolved. This is particularly true for inflation-targeting countries, where the inflation objective is quantitatively explicit. In the cases of all inflation-targeting countries, something close to "true" price stability has been chosen[1].

The consensus about the importance of price stability is not limited to economists and central bankers. For example, a survey conducted by Shiller (1997) in the United States, Germany, and Brazil shows that the public at large is also averse to inflation. The reason may lie in the substantial costs that inflation inflicts on the economy. These costs, however, have always proved difficult to estimate with a reasonable degree of precision.

In this article, we firstly review the main costs of inflation (section 2). We then outline the methodology used in a recent work by Feldstein (1997a) (section 3) and present some estimates, based on the same methodology, of the costs of inflation in New Zealand. These are presented alongside results of similar studies for other countries (section 4). The costs of inflation in New Zealand are also compared with estimates of the temporary loss of output involved in achieving true price stability (section 5). The final section draws some conclusions.

2 Why price stability?

Inflation imposes significant costs on the economy.[2] Many of these costs are associated with unanticipated inflation. Uncertainty about inflation generates three main types of costs:

- First, it results in a **malfunctioning of the price mechanism**. Under inflation uncertainty, households and firms may easily confuse movements in the aggregate price level with changes in relative prices. Because of the wrong price signals, resources are not allocated efficiently.

1 By "true" price stability we mean net of measurement error. All consumer price indices are affected by an upward statistical bias due to a number of factors, like unmeasured changes in the quality of goods, introduction of new goods and shopping outlets, shifts in consumer spending patterns to goods and services whose relative prices are falling, asymmetries in the way in which upward and downward price changes are measured, and so on. This bias is very difficult to estimate precisely and there is no direct evidence about its size in New Zealand. The report of the Boskin Commission (1996) finds that the bias is likely to be slightly above 1 percent for the United States. Comparisons of the Statistics New Zealand CPI methodology with that used in the United States suggest that the bias is likely to be smaller in New Zealand, but 1 percent per annum is probably a good approximation.

2 There is a wide theoretical literature on the costs of inflation. A recent survey is in Briault (1995).

- Secondly, uncertainty about inflation **shortens time horizons and increases the risk premium** incorporated in long-term interest rates. Firms become more reluctant to enter into long-term commitments. They look for short-term contracts and incur higher transaction costs.
- Thirdly, surprises about inflation cause arbitrary **redistribution of income and wealth**. Many New Zealanders directly experienced these costs in the 1970s and the 1980s, when the value of their savings was eroded severely as a consequence of high inflation.

Inflation imposes costs even when it is fully anticipated. These costs arise because either a) people try to avoid the unwanted consequences of inflation, or b) many of the accounting and contracting conventions that underpin the functioning of a monetary economy assume a constant standard of value. Costs that arise from anticipated inflation include:

- **Tax system non-neutralities**. Tax systems are mostly based on nominal income. Interest payments are deductible from taxable corporate profits at their nominal, rather than inflation adjusted, value. Similarly, interest income is taxed at the personal level at its nominal value. Also, firms are allowed to deduct depreciation of capital assets, and the cost of goods sold out of inventories, on the basis of historic rather than current replacement costs. As a result, inflation modifies the real tax burden, both in aggregate and across sectors, assets, and sources of finance, and exacerbates the dead-weight costs that are inherent in most taxes.[3]
- **Inefficient portfolio allocation**. Inflation distorts investment towards "real" assets, like land or property. Excess investment in these assets will move the economy away from the most efficient resource combinations. This distortion results in a lower rate of return on investment for the economy as a whole, and thus a lower level of output.
- **"Shoe-leather" costs**. Inflation erodes the value of cash balances, inducing firms and households to cut their currency holdings. This involves a higher frequency of cash withdrawals, with associated transaction costs. These costs have been traditionally named "shoe-leather" costs, because withdrawing money used to imply walking to the bank. Electronic technology is likely to have reduced their significance.
- **Menu costs**. Inflation increases the frequency with which firms must revise quoted prices and contracts. In order to do that, firms incur additional transaction costs.
- **Liquidity constraints**. A common practice for debt contracts is to fix service payments in nominal terms. Inflation, to the extent it is incorporated into nominal interest rates, "tilts" the weight of mortgage servicing towards the present, increasing the ratio of debt service to income and squeezing out those who are most cash-flow constrained.

The existence of these costs implies that inflation has a negative impact on the rate of growth, as well as on the level of GDP. However, the available evidence of a negative relationship between inflation and economic **growth** is not very robust. Sarel (1996), for example, finds no effects of inflation on growth at inflation rates below 8 percent. This suggests that low to moderate inflation does not constrain growth indefinitely. However, even a temporary negative impact on the growth rate translates into a permanently lower **level** of GDP. Andres and Hernando (1997) have estimated that an additional percentage point of inflation reduces per-capita income by 0.5 to 2 percent, even when inflation is moderate.

Volatile inflation can be as harmful as high inflation. Judson and Orphanides (1996) find there is a clear link between inflation volatility and growth, even for those countries where inflation is low. The authors estimate that a 50 percent decrease in inflation variability raises GDP growth by a quarter of a percentage point per annum. This result suggests that the negative effects of inflation uncertainty are not fully captured by focusing only on the level of the inflation rate.

Another approach followed in the literature has been to look directly at the mechanisms through which inflation impacts on the economy. A notable example is Feldstein (1997a), who has shown that the costs of inflation due to tax system non-neutralities can be quantitatively significant, even at low rates of inflation. He estimates that, at a 2 percent "true" inflation rate, they amount to as much as the equivalent of

3 The nature of this dead-weight cost is elaborated on later.

0.76 percent of GDP per annum in the United States (Feldstein, 1997b).

3 Inflation as a tax

One way to look at some of the costs of inflation is to consider inflation as a tax. By interacting with the tax system itself, inflation modifies the tax burden on savers, investors and workers, and exacerbates the distortions to relative prices that are inherent in most taxes.[4] At least some of the costs of inflation can thus be analysed by using the tools economists use to analyse the economic cost of taxes.

The standard approach to analysing the economic costs of taxes is based on market demand curves. An example of a market demand curve is depicted in figure 1, where the quantity of a certain good that is requested by consumers varies with its price. The curve is downward sloping, reflecting that a larger quantity of the good is consumed when the price is lower. This curve obviously contains information about the preferences of consumers. In particular, the area below the curve is a measure of the market's willingness to pay for the good. In turn, this can be interpreted as a measure of welfare, ie of the satisfaction that consumers derive from the good itself.[5]

The intersection of supply and demand determines the market price P_0 that is paid by consumers to acquire a quantity X_0 of the good. However, at that price, only the marginal buyer pays exactly what they are willing to. All the other consumers end up paying less; in other words, they enjoy what is called "consumer surplus". This is represented in figure 1 by the red triangle, ie the difference between consumers' willingness to pay (the area below the demand curve for a quantity X_0) and what they actually pay (the quantity P_0X_0, represented by the rectangle).

Figure 1

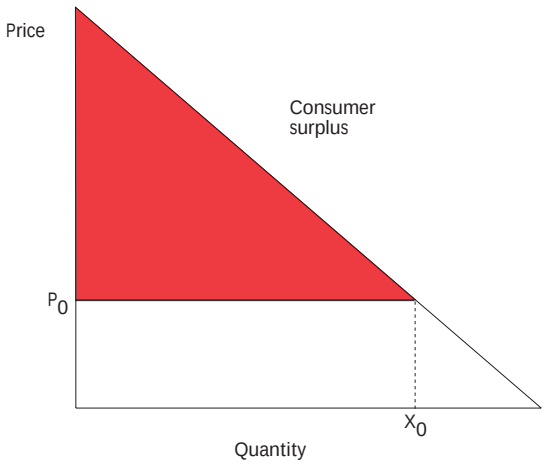

What happens when a tax is introduced?[6] This is shown in figure 2. Price increases from P_0 to P_1. As a result, consumers demand less of the good (X_1, which is smaller than X_0). The consumer surplus, represented by the red triangle, shrinks as consumers have to pay more for the same good. The reduction in consumer surplus clearly represents a welfare loss. This loss is partially offset by the tax revenue collected by the government, represented by the blue rectangle in figure 2. This revenue can be spent on welfare-enhancing activities. However, the gain does not fully match the loss. The difference amounts to a net welfare cost, represented by the area of the dark grey triangle in figure 2. This is the dead-weight cost of taxation.

So far, we have assumed there is no inflation in the economy. What happens when inflation enters the picture? Figure 3 illustrates the new situation. If the interaction of inflation with the tax system causes the effective tax rate to rise, then the market (tax inclusive) price increases further, from P_1 to

4 A simple example is the effect that inflation has on the real after-tax return to savers. Assuming a personal tax rate of 33 percent, the net-of-tax return on an investment yielding 10 percent is equal to 10(1-0.33)=6.7 percent. However, if inflation goes from zero to 2 percent, and the nominal return increases correspondingly to 12 percent, the new after-tax return in real terms becomes 12(1-0.33)-2= 6.0 percent. Thus, tax distorts the relative price of spending now rather than later (the interest rate) by 3.3 percentage points. The interaction of inflation and tax distorts it by a further 0.7 of a percentage point.

5 Market demand curves are drawn on the assumption that nominal income and the price of other goods are held constant. However, consumers' satisfaction is more correctly measured by using the so-called "compensated" demand curve. When the price of a good increases, the purchasing power (real income) of consumers is reduced as they can now buy only a smaller quantity of the same good. This is called the "income" effect, as opposed to the "substitution" effect, whereby consumers substitute away from that good towards cheaper alternatives. The effect of a price change on consumers' utility is correctly measured holding real income constant, ie compensating for the "income" effect. Compensated demand curves are built on this basis.

6 For now, the exposition is based on a tax that is levied on the good in question, rather than on income. Later we show how the same analytical framework can be applied to income taxes to estimate the effect of inflation-generated distortions.

Figure 2

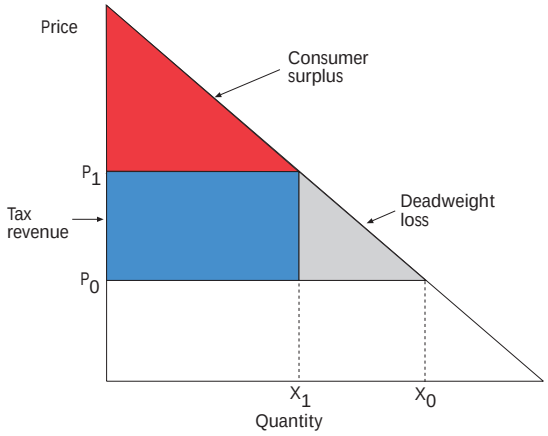

P_2. As a result, consumers reduce the quantity demanded from X_1 to X_2. The red triangle, which represents consumer surplus, shrinks further. Once again, tax revenues (the blue rectangle) partially offset this loss. The deadweight loss, however, is now larger than before and is made up of two components:

- a tax-induced component (the dark grey triangle); and
- an inflation-induced component (the light grey trapezium).

Two points can be noted. First, as the figure shows, the cost of the additional distortion caused by inflation can be quite large: the light grey trapezium is large relative to the dark grey triangle. The size of the cost depends on three factors:

- the extent to which the tax system is distortionary at zero inflation (measured by the vertical distance between P_1 and P_0);
- the extent to which the tax burden increases with inflation (measured by the vertical distance between P_2 and P_1); and
- the price elasticity of demand, represented by the slope of the demand curve in figure 3, which is a measure of consumers' reaction to a price change.

Eliminating inflation has a positive welfare effect because it removes the inflation-induced component (the light grey trapezium). We will refer to this as the direct effect of disinflation. However, eliminating inflation also reduces tax revenues (by the difference between the area of the blue rectangle in figure 2 and the same in figure 3). This represents the amount of tax that would have to be raised from other sources in the absence of inflation. These other taxes create their own distortions. This is what we call the indirect effect of disinflation.

Figure 3

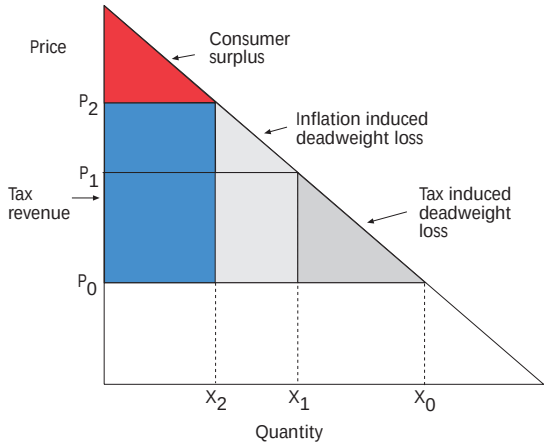

4 Some costs of fully anticipated inflation in New Zealand[7]

How can this methodology be applied in practice? We follow Feldstein (1997b) and focus on the way in which the interaction of inflation with the tax system distorts:[8]

- consumption/saving decisions; and
- housing investment decisions.

Distortions to consumption/saving decisions

Inflation distorts the consumption/saving behaviour of individuals by reducing the after-tax real return to savers. The incentive to save for retirement, already affected by taxa-

7 This section is based on Bonato (1998), to which we refer for details.

8 Inflation also causes distortions to money demand and reduces the real cost of servicing government debt. Thus eliminating inflation also generates costs and benefits through these channels. However, as Table 1 shows, these effects turn out to be very small, almost negligible. Therefore, we just ignore them in the following description of the methodology.

tion, is reduced further and young people may decide to save less than they would otherwise. Even if there is no change in saving, there will be a dead-weight cost for the economy. The fall in the after-tax return to savings implies there will be fewer resources available for future consumption.

In terms of the analytical framework stated above, we can think of future consumption as a good whose price is inversely related to the interest rate. When inflation increases, the after-tax real return to savers falls, meaning that the price of future consumption increases. This relationship can be represented as a demand curve like the one described in figure 3. By moving the price of future consumption from P_1 to P_2, inflation exacerbates pre-existing tax distortions and gives rise to an additional dead-weight cost, measured by the area of the light grey trapezium.

Disinflation removes this cost, resulting in a welfare benefit (the direct effect). However, this benefit is partially offset by the loss of tax revenue. This loss of revenue is given by the difference between the area of the blue rectangle in figure 2 and that in figure 3, and needs to be compensated by other taxes, which have dead-weight costs of their own (the indirect effect).

How do we actually measure these costs and benefits? If we imagine figure 3 as depicting a demand curve for future consumption, we have to calculate the area of the light grey trapezium (direct effect) and multiply the difference between the area of the blue rectangle in figure 2 and that in figure 3 by the dead-weight cost of other taxes (indirect effect). As in any analysis, we must make assumptions that are consistent with economic theory and complement them with estimates derived from empirical investigation. In this particular case, we need numerical estimates of:

- the real after-tax rate of return to savers, with and without income taxes and at different inflation rates (we have used 0 and at 2 percent, both net of measurement bias). This information enables us to estimate values for P_0, P_1, and P_2;[9]
- the elasticity of saving with respect to the price of retirement consumption (the slope of the compensated demand curve);
- the saving of the young[10]; and
- the dead-weight cost of alternative taxes.

On the basis of available estimates, we calculate the benefit of reducing inflation from 2 percent to zero as being equivalent to 0.39 percent of GDP per annum. This figure encompasses both the direct (0.46 percent) and the indirect (-0.07 percent) effect.

Distortion to housing investment decisions

Owner-occupied housing is generally favoured by the tax system over other capital investments because no tax is imposed on the implicit rental return to the capital invested in the property.[11] As a result, investment in housing is thought to be above its optimal level. The associated welfare loss is the opportunity cost of the resources that are not put to their best use. Inflation adds to the distortion by reducing the net-of-tax real return on alternative assets.

Figure 4 illustrates these effects. On the vertical axis is the rental equivalent rate, or the implied rental cost of owner-occupied housing per dollar of housing capital. The rental equivalent rate is given by the sum of financial costs (the opportunity cost of equity, or the return on alternative assets, and the mortgage interest rate), maintenance and depreciation costs, and local government property tax, less

9 As in Feldstein (1997a), we assume the function relating the after-tax return to saving s and the price of retirement (future) consumption P is of the form:

$$P = (1+s)^{-T}$$

where T, the number of years savings are accumulated before retirement, is assumed to be 30.

10 The concept of saving employed here does not correspond with the standard macro-economic measure of saving. What we are interested in is the extent to which people defer consumption (save) when they are young and dissave when they are old. Personal saving, on the other hand, reflects the net saving of the whole population. For example, young people may have a high level of saving, but if old people are dis-saving at the same rate, then the net level of saving will be zero. We thus need to make assumptions about the function relating the saving of the young with personal saving. Further elaboration on this point is provided in Bonato (1998).

11 This distortion is exacerbated in some countries because home owners are allowed to deduct mortgage interest and property taxes from their tax assessable income. This is not permitted in New Zealand, meaning that the tax advantage is limited to the portion of the investment financed by equity.

Figure 4

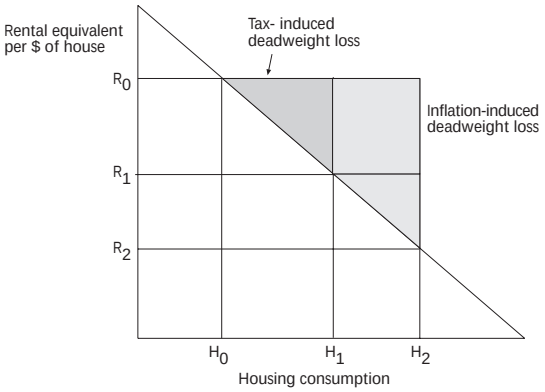

the capital gain on the property. On the horizontal axis is the demand for the services produced by owner-occupied housing capital.

The tax distortion reduces the opportunity cost of the equity invested in owner-occupied houses: the implicit income from this investment remains tax-free, while returns on alternative assets are subject to tax. Measured in rental equivalence terms, the cost of home ownership falls from R_0 to R_1, inducing excess investment of the amount (H_1-H_0). However, the real pre-tax cost of providing housing capital is R_0. Thus, the cost of supplying housing at the level H_1 is given by the quantity H_1R_0, or the area of the rectangle. This is clearly greater than the area under the demand curve, which measures consumers' satisfaction for a level of housing consumption at H_1. The dark grey triangle - the difference between the economic cost of supply, and the welfare or utility derived - represents the dead-weight cost of the tax distortion.

Inflation adds to the distortion because it reduces net-of-tax real returns on alternative financial and business assets. This is reflected in a further fall of the rental equivalent cost of owner-occupied housing from R_1 to R_2, which boosts demand from H_1 to H_2. The inflation-induced dead-weight loss is represented by the area of the light grey trapezium. This is the extent of the welfare benefit that can be captured by eliminating the excess demand for housing caused by inflation (direct effect).

Disinflation increases the after-tax return on alternative assets and pushes capital out of housing into the business sector. This flow of capital generates two opposite changes in tax revenues: a) income tax revenues increase because the effective tax rate in the business sector is higher; b) revenues from local government property tax (council rates) are reduced as the housing stock falls. (a) generally exceeds (b) and thus the indirect effect turns out to be positive.

In order to calculate the size of these effects, we need numerical estimates of:

- the real after-tax rate of return to savers (opportunity cost of equity) and the mortgage rate (cost of debt) at different inflation rates (0 and at 2 percent), with and without income taxes; maintenance and depreciation rates; and the property tax rate (council rates). From this information we can estimate values for R_0, R_1, and R_2;
- the elasticity of housing demand;
- the value of the housing stock; and
- the dead-weight cost of alternative taxes.

On the basis of the available evidence, we estimate the benefit of reducing inflation from 2 percent to zero in terms of reduced distortion to housing at just 0.04 percent of GDP per annum. This figure is given by the sum of the direct effect (0.03 percent) and the indirect effect (0.01 percent).

Summary of the results

Table 1 brings together our above estimates of the welfare gains from reducing inflation-induced distortions to saving/consumption and housing investment. Also included are estimates of the costs of the distortions caused by inflation to money demand and government debt servicing. The estimates for New Zealand are compared with those for other countries, as derived from comparable studies. In New Zealand the benefit from eliminating inflation is estimated to be 0.39 percent of GDP, about half that calculated for the United States by Feldstein, but slightly higher than the estimate for the United Kingdom. Both Germany and Spain present much larger values, about 3.5 times larger than for New Zealand.

Clearly most of the benefits from achieving true price stability, in New Zealand and elsewhere, are generated by removing the inflation-induced distortion to consumption/saving patterns. The size of the distortion to housing demand varies widely: it is very large in Spain, and very small in New Zea-

Table 1
The net welfare benefit of reducing inflation from 2 percent to zero
(as a percentage of GDP)

	NZ	USA	Germany	UK	Spain
Consumption/saving	0.39	0.67	1.48	0.29	0.72
Housing demand	0.04	0.22	0.09	0.11	1.09
Money demand	-0.01	-0.03	-0.04	-0.02	-0.07
Government debt	-0.03	-0.10	-0.12	-0.09	-0.10
Total	0.39	0.76	1.41	0.29	1.64

Sources: NZ: Bonato (1998); US: Feldstein (1997b); Germany: Tödter and Ziebarth (1997); UK: Bakhshi et al (1998); Spain: Dolado et al (1997)

land. The inflation tax on money holdings is generally negligible and the effect from the increased real cost of servicing government debt is also smaller in New Zealand than in other countries.

Differences in tax systems explain a large part of the cross-country variation of the estimates. As we noticed in section 3, both the original tax distortion and that induced by inflation are important in determining the welfare costs of inflation. In New Zealand the tax system itself appears to be less distortionary than in the United States, Germany and Spain. One important difference relates to the integration of corporate and personal income taxation. The full dividend imputation system adopted in New Zealand eliminates the double taxation of corporate income that characterises the classical system. In the latter, currently in use in the United States, profits are taxed once at the corporate level and then again, at the personal level, as dividends in the hands of the shareholders. This increases the size of the original distortion.

A second difference is that New Zealand does not tax capital gains at the personal level and this reduces the tax burden on savings relative to other countries. At the same time, the absence of this tax prevents the inflation non-neutrality inherent in the taxation of nominal capital gains.[12]

The last main difference from other tax systems is that mortgage interest payments and local authority rates on owner-occupied housing are not tax deductible in New Zealand. This reduces the tax advantage generated by the absence of tax on the implicit rental income.

5 Price stability: permanent benefits vs temporary costs

Thus far we have discussed the benefits that can be captured by eliminating inflation. But there is also a cost: the cost of disinflation is the output loss involved in getting to price stability. This cost is measured by the "sacrifice ratio", or the cumulative output loss associated with each percentage point of disinflation.

The estimates of this ratio are typically very uncertain and show substantial variation over time. Moreover, the estimates generally fail to disentangle the effect of monetary policy from other factors. For New Zealand, we have values ranging from 0.5 (Ball, 1994) to 8.5 (Mayes and Chapple, 1994). However, it is very doubtful whether estimates obtained under a high-inflation regime can be applied to the current situation. The Forecasting and Policy System (FPS), the calibrated model of the New Zealand economy in use at the Reserve Bank, has embedded in it a sacrifice ratio of about 2, ie, each percentage point reduction in inflation is associated with a 2 percent, cumulative, loss of output. (Black et al., 1997).

A key point of the analysis is how to compare costs and benefits. Economic theory suggests that the costs of disinflation are temporary, but the benefits are permanent. To

12 For practical reasons most countries tax nominal capital gains, rather than real. A notable exception is the UK.

achieve disinflation, monetary policy must constrain economic activity, but only for the time necessary to achieve the objective of price stability. There are no permanent effects. In the long run, the level of output depends on factors that monetary policy cannot affect, that is the capital stock available to the economy, the size of the labour force and productivity.

This proposition is important for making comparisons between the costs and benefits of achieving price stability. The costs are one-off, but the benefits will keep accruing over time. The implication is that the one-off output costs must be compared with the discounted stream of future benefits.

The choice of discount rate is extremely important in determining the final result. A very small discount rate produces a strong magnifying effect on the estimated benefits. What is the correct discount rate in our case? If the market adequately reflects social preferences, the market interest rate is a good approximation. We thus use the after-tax rate of return to savers when inflation is 2 percent, which is equal to 7.4 percent.[13] Because the benefits are calculated as a percentage of GDP and therefore grow over time, the discount rate must be reduced by g, the expected growth rate of GDP. We assume that g equals 2.5 percent per annum, which corresponds to the steady-state value of growth in FPS.

The choice of the discount rate determines the cost-benefit break-even point. If the total cost (C) of going from 2 percent inflation to zero inflation is 4 percent of GDP, as implied by a sacrifice ratio of 2, the break-even point turns out to be percent in New Zealand. This compares with our estimated cost of 2 percent inflation for New Zealand of 0.39 percent (of GDP per annum), and suggests that the benefits from completely eliminating inflation would exceed the cost.

However, this result depends heavily on the value of two parameters, the discount rate and the sacrifice ratio, on which there is significant uncertainty. There may be plausible values of these parameters for which the conclusions reached would not hold. A discount rate of 14.5 percent or a sacrifice ratio of 4 would be enough to completely offset the estimated benefits. While we can safely discard the possibility of such a high discount rate, we cannot do the same for the sacrifice ratio due to the considerable uncertainty surrounding its estimates.

13 See Bonato (1998)

6 Conclusions

The objective of price stability reflects a strong aversion to inflation that is shared by different sectors of society. This aversion suggests that inflation imposes substantial costs on the economy. While the reasons for these costs are intuitively clear, their actual size is not as obvious because quantitative estimates appear particularly difficult to obtain.

In this article, after reviewing the main costs of inflation, we have focused on a particular source of costs, those generated by the interaction of inflation with the tax system. We have provided a quantitative estimate of these costs, that shows that they are non-negligible for New Zealand, even if they are in the low range when compared with other countries. The benefits arising from their elimination appear to be large enough to offset the output costs involved in reaching true price stability, at least for most of the plausible values of key parameters.

References

Andres, J and I Hernando (1997), 'Does inflation harm economic growth? Evidence for the OECD', National Bureau of Economic Research Working Paper No 6200.

Bakhshi, H, A G Haldane, and N Hatch (1998), 'Some costs and benefits of price stability in the United Kingdom', Bank of England Working Paper N .78.

Ball, L (1994), 'What determines the sacrifice ratio?', in N G Mankiw (ed.), Monetary Policy, The University of Chicago Press.

Black, R, V Cassino, A Drew, E Hansen, B Hunt, D Rose, and A Scott (1997), 'The Forecasting and Policy System: the core model', Reserve Bank of New Zealand Research Paper No 43.

Bonato, L (1998), 'Price stability: some costs and benefits in New Zealand', Reserve Bank of New Zealand Discussion Paper, forthcoming.

Boskin Commission (1996), 'Toward a more accurate measure of the cost of living', Final Report to the Senate Finance Committee.

Briault, C B (1995), 'The costs of inflation', Bank of England Quarterly Bulletin, Vol 35.

Dolado, J J, J M Gonzalez-Paramo, and J Viñals (1997), 'A cost-benefit analysis of going from low inflation to price stability', paper presented at the NBER Conference on The costs and benefits of achieving price stability, Federal Reserve Bank of New York, February.

Judson, R and A Orphanides (1996), 'Inflation volatility and growth', Board of Governors of the Federal Reserve System, Finance and Economic Discussion Paper No 96.

Feldstein, M (1997a), 'The costs and benefits of going from low inflation to price stability', in Romer, C H and D H Romer (eds) Reducing inflation, The University of Chicago Press.

Feldstein, M (1997b), 'Capital income taxes and the benefit of price stability', National Bureau of Economic Research Working Paper N .6200.

Mayes, D and B Chapple (1994), 'The costs and benefits of disinflation: a critique of the sacrifice ratio', Reserve Bank of New Zealand Bulletin, Vol 57(4).

Sarel, M (1996), 'Nonlinear effects of inflation on economic growth', International Monetary Fund Staff Papers, Vol 43(1).

Shiller, R (1997), 'Why do people dislike inflation?', in Romer, C H and D H Romer (eds) Reducing inflation, The University of Chicago Press.

Tödter, K H and G Ziebarth (1997), 'Price stability vs. low inflation in Germany: An analysis of costs and benefits', National Bureau of Economic Research Working Paper No 6170.