# The optimal rate of inflation: an academic perspective

**By Peter Sinclair,**[1] **Houblon-Norman Fellow**[2] **and Professor of Economics at the University of Birmingham. The views expressed in this article are the author's and do not necessarily reflect those of the Bank of England or the Monetary Policy Committee.**

*In an economy free of all imperfections, inflation should be slightly negative. Prices should keep dropping, at the real rate of interest. Any higher rate of sustained inflation (or lower deflation) would reduce the benefits from holding real money. Central banks typically aim for modest positive inflation, however.*

*This article explores five types of imperfection: inertia in nominal prices, the need for distorting taxes, market power for retail banks, the value of the option to cut nominal interest rates in bad times, and menu costs. It concludes that the combined effect of these imperfections is in practice likely to justify a small positive rate of inflation.*

## Introduction

A large and growing number of central banks target inflation. Sometimes the central banks themselves set their own inflation targets. Others are given them by government. No inflation targets, anywhere, are negative. Yet several economists argue that what monetary authorities should aim for is price deflation—negative inflation—at the real rate of interest. This paper explores this recommendation (in the next section) and then considers five arguments against obeying it. It concludes that a low rate of positive inflation is probably best.

## The case for price deflation at the real rate of interest

Everyone would agree that you can easily have too much inflation. But there is often disagreement about how much is 'too much'. For some people, all inflation is simply bad, so we should aim to eliminate it. For others

'too much' might mean something in excess of perhaps 2% or 3% per year. Economic analysis gives us, however, a strong (if controversial) case against both these views, which is actually in favour of price deflation. That case is described in this section; later sections consider challenges to it.

In an otherwise perfect economy, any good should be priced at its marginal cost. Define money as currency, and its 'price', the opportunity cost of holding it, as the nominal rate of interest, for example on a treasury bill. It follows that the nominal rate of interest should be zero. At this point, the quantity of money is at its optimum; attaining this means arranging for prices to fall at the real rate of interest.

Money is held because it is useful. It is often feasible to trade and transact without money, but not always convenient. Valuable resources would be squandered. Trading partners would have to be sought out, relative

prices negotiated, and awkward inventories carried. The chief resource all this would take is surely the time that would be sacrificed—time that would otherwise have been applied to leisure or rewarding work.

If money is useful, and takes the form of token or fiat currency, which costs next to nothing to produce, surely households and firms should be encouraged to hold as much of this useful asset as they could possibly want? But the benefits from this asset depend upon its real value, not its paper quantity. Maximising the stock of real currency certainly does not imply that paper money should be expanded fast: if anything, the reverse. If the marginal cost of providing real money is negligible, as is widely assumed, currency should ideally be a free good.

If there is any good or service in the economy that is not priced at its marginal cost, it is possible, in principle, to reallocate resources so that everyone can gain. So, in that sense, prices that differ from marginal cost can only signify waste. The 'price' of holding real money is the nominal return that its holder could have earned instead on an alternative (riskless) asset, such as treasury bills. Making money costless to hold would therefore mean making the nominal rate of interest zero on such alternatives. It would encourage people to hold as much (real) money as they would like, with all the advantages this would bring in the form of greater convenience, security, production and trade.

Nominal interest is the sum of two elements: real interest, which should normally be positive, at least in the long run, and expected inflation. So bringing the nominal interest rate down to zero means creating a monetary framework where prices are expected to keep declining. The rate of that decline would be the real rate of interest.

Cutting the rate of inflation from something positive down to zero would benefit the holders of money, and hence society in general. Lucas (2000, 2003) calculates that the gain from reducing annual inflation at 10% to zero, in a modern economy such as the United States, is equivalent to about 1% of total consumption. In addition to a variety of other effects, many of them harmful, inflation wastes resources—most obviously, labour time—which could be put to better use. For the British economy, recent estimates by Bakhshi, Martin

and Yates (2002) point to a rather smaller number than Lucas's estimate. But the logic of the above argument is that getting rid of inflation would not in fact be going far enough. It would be better still to reduce inflation further, to the point where prices were expected, on average, to trend downwards at the real rate of interest. Only then would the benefits from holding money be exploited to the full.

The damage that a small positive nominal interest rate will do is negligible. But a large one will do great damage. Roughly speaking, the cost of departing from the ideal of a zero nominal interest rate is often thought of as increasing with the square of the nominal interest rate. If so, a nominal interest rate maintained in perpetuity at 16% would do about 16 times more damage than one kept at 4%. Mild inflation is a very minor irritant, therefore; but high inflation is costly.

The rate of inflation varies over time, and also across the vast range of goods and services bought and sold. The focus in this paper will be on the long-run average rate of inflation. In general, published measures of inflation often tend to overstate the true rate of inflation. There are a number of reasons for this, among them failure to correct for substitution towards goods that have fallen in relative price, and insufficient allowance for quality improvements. Measurement problems can be serious in some economies. In the United States, for example, the Boskin Report (1996)[1] found that the 'true' rate of inflation could be 1% or so below the headline figure. If it were, a 'true' inflation target of $x$% would be achieved when headline inflation exceeded $x$. But the gap between an annual inflation target of say 2.5%, and price deflation at the real interest rate (perhaps 3%), is far greater than this. That implies that inflation measurement issues are modest when compared with a gap of some 5% or more.

The main argument sketched out above was first proposed by Milton Friedman (1969).[2] His call for prices to fall at the real interest rate is known as the 'Full Liquidity Proposition'. It has provoked voluminous research, surveyed (and extended) recently by Lucas (2000).[3] Most of this research confirms Friedman's claims—at least within the confines of an idealised, simple, perfectly competitive economy, free from frictions and distortions such as monopoly,

(1) See also Boskin *et al* (1998), Deaton (1998) and Diewert (1998).
(2) In his practical policy recommendations for the United States, Friedman in fact recoiled somewhat from the logic of his argument, and limited himself to urging a policy of money growth targets expected to deliver zero inflation, not negative inflation.
(3) Woodford (1990) is the author of an extensive earlier survey.

uncertainties, price rigidities, the costs of changing nominal prices ('menu costs'), markets that fail to clear, and taxation that impairs economic efficiency.

The key question for policy-makers is how far these various complications qualify or undermine the applicability of Friedman's proposals to contemporary economies. Do the market imperfections mean that prices should decline more slowly than the real rate of interest? Might zero inflation be best? Or is at least a modicum of positive inflation the proper objective? Many central banks now try to target inflation, implicitly or explicitly; but their inflation targets are invariably positive. Why is this? Is it correct? What are the arguments for a positive rate of inflation, and how strong are they?

What all the arguments against adopting a goal of price deflation have in common is that each rests on some form of market imperfection. There are five main arguments:

(i)     markets fail to clear continuously, especially if there is excess supply, and a background of mild inflation may improve resource allocation;

(ii)    public finance considerations, stemming ultimately from some type of market failure, mean that it is quite possible that money should in effect be taxed;

(iii)   making currency less attractive to hold can undo some of the damage caused by the exercise of market power in the retail banking sector;

(iv)    occasional recessions, which are symptoms of macroeconomic market failure, need to be countered by nominal interest rate cuts, and, since nominal interest rates cannot be negative, it makes sense for nominal interest rates normally to be positive, in order to create room to cut them if necessary;  and

(v)     many firms opt to hold prices for long periods, which points to a degree of market power and systematic overpricing, which mild inflation may reduce.

Each of these arguments will be explored in turn in the five sections that follow.

## Inflation helps markets to clear more quickly

Argument (i) starts with the notion that wages and prices are sticky downwards. That means that any excess supply tends to persist. Since it is real prices that are too high, and the source of the problem, disequilibria are removed faster if other prices are climbing, rather than flat or falling.

When applied to labour markets, this view could suggest a negative relationship between the rate of inflation and the average level of unemployment, and possibly one where that negative association persists. Most observers have long concluded that any benefits from raising the rate of inflation were fleeting and trivial, in comparison with the problems created (and also, perhaps, with the cost of reversing it). Yet, no matter how clear the evidence that annual inflation does great damage above, say, 3% or 5%, it does not follow that macroeconomic stability is more imperilled by inflation in the 1% to 3% range, than by price deflation at the real rate of interest. Inflation could be a terrible curse at rapid rates, and yet actually a modest blessing at low ones.

Does evidence bear this out? Akerlof, Dickens and Perry (1996, 2000) argue that, for the reasons given above, modest inflation may lower long-run unemployment, while the uncertainties that accompany rapid inflation may increase it. They find US evidence to be broadly consistent with this view: for them, the unemployment-minimising rate of annual inflation appears to be somewhere between 1.5% and 4%.[1] Wyplosz (2001) looks at data for France, Germany, the Netherlands and Switzerland. He concludes that unemployment is not completely independent of the rate of inflation, and some of the results tally with the view that a little inflation helps to cut unemployment, and not just temporarily. But his conclusion is that 'we do not know yet how high inflation should be'. Even if we could be sure that there was an inflation range that minimised unemployment in a particular economy in a particular period of time, there are many reasons[2] why we could not necessarily extrapolate that to other economies or periods.

## The public finances

The second riposte to Friedman's Full Liquidity Proposition, due originally to Phelps (1973), runs thus.

---

(1) There is also other evidence testifying to a negative long-run link between inflation and unemployment in the United States: for example, Fair (2000) and King and Watson (1994). On the broader issues of Phillips curve non-linearities and whether wage and price nominal rigidities become more serious as inflation falls, Yates (1998) provides a very comprehensive survey.

(2) Most obviously the Lucas Critique argument: behaviour depends on policies pursued.

A positive nominal interest rate is a distortionary tax on real holdings of currency. But other taxes, on income, profits, sales or value added for example, are distortionary too. Is money really so special that it should qualify for the privilege—a unique privilege, perhaps—of tax exemption? It is only lump-sum taxes that do not distort; and equity or feasibility[1] considerations will rule them out. So if public goods and transfers and public debt service have to be financed by wasteful, distortionary taxation, should money holdings really go untaxed?

Governments may have large revenue needs, entailing serious distortions throughout the economy. Taxing money can alleviate these distortions a little. It also offers a means for taxing informal-economy transactions, which are mostly conducted by cash. On the other hand, taxing money tends to hurt the poor relatively more than the rich, and taxing the money people need to pay for taxed goods out of taxed income amounts to double taxation. Unless the government's revenue requirement was very large—and even if it were—a social planner may balance accounts better by relying solely on other revenue sources, such as income tax.[2]

## Market power in retail banking

The third argument for not making prices trend down at the real rate of interest is based upon the view that the market for retail bank deposits may not be perfectly competitive. This is simply a hypothesis—there is no suggestion that this is in fact the position in the United Kingdom today.

As Edgeworth (1888) was the first to argue, the costs of operating the payments system, clearing cheques, evaluating alternative assets and holding prudential reserves tend to make banking an industry where suppliers enjoy increasing returns. That implies there will probably be room for just a few, large retail banks. Further, if banks aimed to maximise profit,[3] they would exploit their monopsony power by underremunerating deposits.

One answer to this is regulation. Banks could be forced to pay more to depositors, for example, and to charge less for loans. But such a policy has the drawback of

ultimately weakening banks' balance sheets, raising the likelihood of failure and systemic financial instability.

So another option is to levy a small tax on currency. Currency is a close but imperfect substitute for bank deposits. If bank deposits are too low from a welfare standpoint, and cannot be raised safely by direct means, then a tax on their substitutes is appealing. It would raise the volume of bank deposits. That should bring a first-order gain to social welfare (assuming no adverse repercussions elsewhere) and only second-order welfare losses due to reduced currency holdings, so long as the tax on currency is low. Furthermore, such a small tax on currency would also strengthen banks' balance sheets and reduce risks of bank failure. This counterargument to Friedman's Full Liquidity Proposition, which is explored in Mullineux and Sinclair (2003), is a straightforward application of the Theory of the Second Best to a hypothetical case where banks behave as Edgeworth argued long ago they could.

## Preserving room for manoeuvre in monetary policy

The fourth argument against Friedman's call for price deflation at the real interest rate rests on three propositions:

(a)    the official, central bank nominal interest rate is the key lever of monetary policy;

(b)    the nominal interest rate can never be negative; and

(c)    unforeseeable shocks sometimes make it right to engage in expansionary monetary policy.

Adherence to Friedman's Full Liquidity Proposition means setting nominal interest rates at zero. Together, (a), (b) and (c) imply that the central bank has closed any option to engage in temporary monetary relaxation, should circumstances appear to warrant it. Squeezing demand is still available, because this entails a temporary rise in the policy rate. Proposition (b) imposes a floor on that rate, not a ceiling. But loosening policy is impossible, if the nominal interest rate is already at the floor. It is precluded, at least, if the central bank employs its standard device—a temporary cut in the policy rate.

---

(1) Truth-telling issues: lump-sum taxes increasing with earning ability reported to the tax authority would, for example, tempt the able to lie.

(2) If income tax rates were non-linear, as urged in the pioneering paper by Mirrlees (1971), it turns out that the conditions under which money should be taxed for fiscal reasons become even more stringent.

(3) With at most limited powers of price discrimination, and immune from the threat of entry.

What gives this observation special significance is the fact that models of the macroeconomic short run relate aggregate demand positively to the rate of expected inflation. Buyers tend to bring purchases forward if prices are expected to climb. Expectations of price declines encourage them to wait. So in the very circumstances when monetary reflation may be most needed—positions of exceptionally weak aggregate demand—Friedman's Full Liquidity Proposition appears to preclude it. To some, the last steps of the route to a zero nominal rate start to look like a one-way street. The nub of the argument is that a central bank should aim to keep nominal interest rates well inside positive territory, in order to preserve the option of cutting them if and when necessary.

This pessimistic view may be rather overdrawn. Proposition (a) is open to challenge. Even if and when the policy rate is zero, the central bank can reflate in other ways. It could conduct open market purchases of longer-dated bonds that still carry positive nominal redemption yields, or of equities, or foreign exchange.[1] Furthermore, when close to Friedman's optimum, but not quite there, the authorities might, as Ueda (2002) discusses, lengthen the duration for which nominal interest rates are cut. This would compensate for the fact that the size of the cut is smaller than it would have been without the zero bound. Or, as Eggertson and Woodford (2003) argue, policy should aim, in such circumstances, at raising the expected price level for some later date. Finally, as Yates (2003) discusses, there are several possible devices that might, one day, allow a central bank to set a temporarily negative yield on cash.

Nonetheless, to the extent that (a), (b) and (c) do hold under present conditions, they do at least constitute a case for less price deflation than Friedman's rule implies, and maybe for zero or even positive inflation. The gain, presumably a diminishing one at the margin, from moving towards Friedman's optimum has to be set against the cost of reducing, and in the limit eliminating, the option to reflate, should the need arise, with the standard medicine of an interest rate cut. The greater the chance of adverse demand shocks calling for that medicine, the greater the value of the option.

## Menu costs

Menu costs are the real costs of changing nominal prices.[2] They represent an imperfection in the sense that they impede the flexibility of prices. If one wished to pursue a monetary policy that minimised the average incidence of menu costs, the overall rate of inflation would be zero.

Menu costs do not imply, however, that zero inflation is ideal. There are three reasons for this. First, other factors need to be taken into account. The gains from price deflation at the real interest rate, which Friedman emphasises, the subject of the second section, are one such.[3] Second, if menu costs are present, and help to explain the widespread phenomenon of nominal price rigidities at low rates of inflation, it is not easy to think in terms of perfect competition. Perfect competitors are price-takers. It is more natural to treat firms that face costs of changing prices as price-setters. And, as we saw with banks, the power to set prices may well lead to overcharging and hence to underprovision. In an otherwise perfect economy, it must, assuming profit is maximised, price discrimination limited, and entry blocked. The third key issue is the fact that the real rate of interest is typically positive. That means that firms care more about real profits this month (or year) than next.

Suppose a firm with monopoly power knows the rate of overall inflation, and knows it is constant. If there were no menu costs, and its nominal marginal cost kept rising at that rate, so would its (profit-maximising) nominal price.[4] But menu costs would make a policy of continuous tiny price increases prohibitively expensive. Instead, the nominal price of its product would be raised by a discrete proportion infrequently, and remain fixed between revisions. With both demand and marginal cost steady in real terms, and all other things being equal, steady inflation would make the interval between those revisions constant. The firm would presumably choose both the size and the frequency of its nominal price changes to maximise its stream of discounted net profit.

Suppose that the interval between price changes was one year, when the rate of inflation and the real interest

---

(1) As urged for example, among others, by McCallum (2000).

(2) Early contributions to the literature on menu costs are Barro (1972) and Sheshinski and Weiss (1977). The main focus here was positive, not normative. Diamond (1993) constructs a search model where the real interest rate is positive, and finds that, with monopoly, optimum inflation is strictly positive, because it tends to force monopolists to charge less on average in real terms. What follows in this section is a simplified account of Sinclair (2003) which confirms Diamond's conclusion in a menu cost setting.

(3) For even if Friedman's argument for price deflation goes too far, as the logic of the present paper suggests, his insights about the welfare costs of inflation remain valuable. As King (2002) stresses, above some rate, inflation is indubitably damaging.

(4) Assuming that the elasticity of demand for the firm's product is constant.

rate were both 3%. For the first six months or so, the firm would be charging more, in real terms, than the price it would have set to maximise profit in the absence of menu costs. Later on, it would be undercharging from the standpoint of maximum profit, in this sense.

This is where discounting (the real interest rate) comes in. What it loses from initial overpricing would receive more emphasis than losses from the opposite later. That means inflation should reduce its average real price somewhat. If so, on average, it would therefore produce more. On average, the price would come closer to its socially ideal value of marginal cost. This effect is strongest when the inflation rate is positive and very small, since the interval between price revisions would then be very long.

If inflation were a very small negative number, however, we would also see very infrequent price revisions. But the timing order of losses due to overpricing and underpricing would now be reversed. Negative inflation means that the real price drifts up over time between price revisions, not down. So early on the firm would see it was charging too little in real terms. So it would tend to react by raising its average real price, and therefore, on average, producing less. In an otherwise undistorted economy, first-order welfare losses would follow, amplifying the deadweight cost of monopoly.

These gains and losses from altered average real prices set by monopolists now need to be combined with the menu costs themselves (which are minimised at zero inflation) and Friedman's benefits from real money holdings (maximised when inflation is minus the real rate of interest). If the last of these were sufficiently modest, the optimum rate of inflation would be unambiguously positive. When the opposite is true, inflation should be negative, at (or quite close to) the real interest rate. In either case, a very slightly negative inflation rate could only be harmful, quite as bad for welfare as a rapid, positive inflation rate.

This fifth counterattack on Friedman's Full Liquidity Proposition is in fact the only one that establishes, under specified assumptions,[1] that small positive inflation is superior to small negative inflation. The figure above presents two possible relationships between

**Figure 1**
**Costs and benefits of inflation**
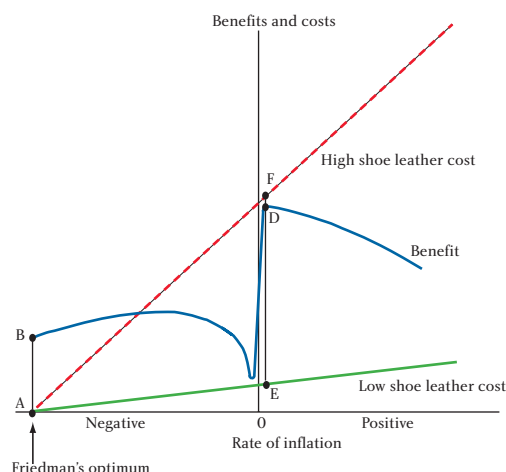


welfare and the percentage rate of annual inflation. One, the solid curve, coloured green, depicts the case when the 'shoe leather' cost[2] of departing from Friedman's ideal, maximum level of real money is quite modest. In this case welfare peaks at a positive inflation rate. The broken curve, coloured red, holds when these shoe leather costs are so massive that Friedman's Full Liquidity Proposition remains best.

The blue curve represents 'benefits'. It captures the firm's average profits, net of menu costs, and, in addition, the consumer's average surplus (the monetary difference between the utility gained from consuming the good, and what the consumer actually pays). The curve peaks at a slightly positive inflation rate, and reaches its minimum when inflation is slightly negative. The blue curve is continuous. But at relatively high and low rates of inflation, when price adjustments become more frequent, benefits tend to be less than they would be if inflation were zero, and the blue curve, as drawn, reflects this.

Society's *net* benefit at the Friedman optimum is AB. This is less than the distance DE (which represents net benefits with the green shoe leather curve), but greater than the (negative) distance DF (which gives net *cost* with the red curve). Looking at the negative inflation region, there are several factors that make up the shape of the benefit curve. A very low rate of price deflation implies a very long interval between nominal price cuts,

(1) These assumptions include horizontal marginal cost, linear (or constant-elasticity) demand, constant and known rates of inflation and real interest, and a continuum of similar profit-maximising monopolists, whose total price distribution moves evenly over time.
(2) The term 'shoe leather costs' covers the direct welfare costs of raising the rate of inflation above Friedman's optimum, in recognition of the fact that much faster inflation will encourage people to shop more often, thus wearing out their shoes at a faster rate. The green and red curves are drawn as straight lines. They could be curves, but they will always slope upwards.

and therefore a powerful stimulus to the firm to raise its average real price (to society's detriment). As the rate of *deflation* increases, this effect weakens quite quickly, because price adjustments become more frequent. This is at first beneficial (because real monopoly prices slip back). But as the rate of price deflation increases, average menu costs become more serious, and the marginal reduction in monopolists' average prices fades away.

So far there is a chance that inflation should be strictly positive; but if it is, it will in fact be tiny. (For example, it would be barely one ten thousandth of 1% per year, in fact, if the demand elasticity were constant at a value of two, and prices changed annually when the rates of inflation and real interest were both 3%.) The optimum inflation rate would become appreciable, however, if we supposed that the monopolistic firms displayed a positive productivity trend. Suppose this were 4% per year, while the rest of the economy's output accounted for half of total spending, and consisted of goods produced by perfectly competitive firms who faced no menu costs and whose productivity was unchanging.[1]

If the gains from raising real money to the Friedman optimum were small enough, it would follow that the optimum rate of overall annual inflation was close to 2%. This is because the monopolists would never have to change their nominal prices if all the other goods rose in price by 4% per year. So a rate of inflation very slightly above this (which would make price revisions very infrequent) would combine with the positive real interest rate to induce them to cut their average real prices—which were too high anyway from a welfare standpoint. If overall inflation were a little lower, however, our firms' average real prices would be higher, leading to a really sharp dip in welfare. If it were a good deal faster than 2%, benefits would fall away too, and, of course, shoe leather costs would be appreciably larger, too.

As stressed earlier, inflation is hard to measure exactly, particularly when quality is changing, new goods are being introduced, and relative price changes are inducing substitution. Productivity trends and real interest rates are also inclined to swing around, and no less open to measurement difficulty. What implication do these phenomena have? Because welfare falls away much more slowly to the right of the optimum inflation rate (assuming it is positive) than to the left, uncertainties and measurement ambiguities imply that a little overshooting is less serious than a little undershooting.

## Conclusion

Many distinguished economists have argued that prices should keep falling, at the real rate of interest. No central bank, however, operates on that principle. So who is right? Negative inflation at that rate would be ideal, normally, in a perfect world. But imperfections of many kinds tend to tilt optimum monetary policy towards less deflation, or even mild inflation. Five types of imperfection have been explored. Perhaps the most telling were the last two, discussed in the previous two sections. The first of these put the case for leaving room for cuts in the nominal rate of interest when circumstances required this. In the previous section, the Friedman arguments for price deflation were pitted against a simple account of menu costs and monopoly that could not just explain the kind of price rigidities observed, but also demonstrate how and why optimum inflation could indeed be strictly positive after all.

The combined force of all five arguments is to provide a reasonable intellectual justification for the kind of monetary policies—aiming for low, positive inflation—now conducted in many countries.

The five objections to following Friedman's call for deflation at the real rate of interest that we have considered are not alternatives. Nor are they additive: it would be wrong to say 'Let us have $x$% inflation for this reason, and $y$% for that, so $(x + y)$% is best'. Rather, they are complementary. Taken together, these objections to price deflation, and the concerns about inflation mismeasurement, suggest that, for advanced countries at least, the case for a modestly positive rate of inflation—within the range of inflation targets currently imposed—looks decidedly a wise one.

---

(1) To make this consistent with a small general equilibrium model, Sinclair (2003) finds it simplest to assume that there is a representative shareholder-consumer, who receives all profits, and whose utility is linear in the competitive sector's goods but symmetric and concave in each of the monopolist's products (with isoelastic marginal utility), and log-linear in leisure.

## References

**Akerlof, G A, Dickens, W and Perry, G (1996)**, 'The macroeconomics of low inflation', *Brookings Papers on Economic Activity*, Vol. 1, pages 1–59.

**Akerlof, G A, Dickens, W and Perry, G (2000)**, 'Near-rational wage and price setting and the long run Phillips curve', *Brookings Papers on Economic Activity*, Vol. 1, pages 1–60.

**Bakhshi, H, Martin, B and Yates, A (2002)**, 'How uncertain are the welfare costs of inflation?', *Bank of England Working Paper no. 152*.

**Barro, R J (1972)**, 'A theory of monopolistic price adjustment', *Review of Economic Studies*, Vol. 39, pages 287–303.

**Boskin, M J, Dulberger, E R, Gordon, R J, Griliches, Z and Jorgenson, D W (1996)**, 'Towards a more accurate measure of the cost of living', Final Report to the (US) Senate Finance Committee, 4 December.

**Boskin, M J, Dulberger, E R, Gordon, R J, Griliches, Z and Jorgenson, D W (1998)**, 'Consumer prices, the consumer price index and the cost of living', *Journal of Economic Perspectives*, Vol. 12, pages 3–26.

**Deaton, A (1998)**, 'Getting prices right:  what should be done?', *Journal of Economic Perspectives*, Vol. 12, pages 37–46.

**Diamond, P A (1993)**, 'Search, sticky prices and inflation', *Review of Economic Studies*, Vol. 60, pages 53–68.

**Diewert, W E (1998)**, 'Index number issues in the consumer price index', *Journal of Economic Perspectives*, Vol. 12, pages 47–58.

**Edgeworth, F Y (1888)**, 'The mathematical theory of banking', *Journal of the Royal Statistical Society*, Vol. 51, pages 113–27.

**Eggertson, G B and Woodford, M (2003)**, 'The zero bound on interest rates and optimal monetary policy', *Brookings Papers on Economic Activity*, Issue 1, pages 139–211.

**European Central Bank (2001)**, *Why price stability?*, Frankfurt am Main.

**Fair, R (2000)**, 'Testing the NAIRU model for the United States', *Review of Economics and Statistics*, Vol. 82, pages 64–71.

**Friedman, M (1969)**, 'The optimum quantity of money', in *The optimum quantity of money and other essays*, Macmillan-Aldine, Chicago.

**King, M A (2002)**, 'The inflation target ten years on', *Bank of England Quarterly Bulletin*, Winter, pages 459–74.

**King, R G and Watson, M W (1994)**, 'The post-war US Phillips curve:  a revisionist econometric history', *Carnegie Rochester Conference Series on Public Policy*, Vol. 41, pages 157–219.

**Lucas, R E (2000)**, 'Inflation and welfare', *Econometrica*, Vol. 67, pages 247–74.

**Lucas, R E (2003)**, 'Macroeconomic priorities', *American Economic Review*, Vol. 93, pages 1–14.

**McCallum, B T (2000)**, 'Theoretical analysis regarding a zero lower bound to nominal interest rates', *Journal of Money, Credit and Banking*, Vol. 32, pages 870–904.

**Mahadeva, L and Sinclair, P (2002)** (eds), *Monetary transmission in diverse economies*, Cambridge University Press, Cambridge.

**Mirrlees, J A (1971)**, 'An exploration in the theory of optimum income taxation', *Review of Economic Studies*, Vol. 38, pages 175–208.

**Mullineux, A and Sinclair, P (2003)**, 'Oligopolistic banks: theory and policy implications', University of Birmingham, *mimeo.*

**Phelps, E (1973)**, 'Inflation in the theory of public finance', *Swedish Journal of Economics*, Vol. 75, pages 67–82.

**Sheshinski, E and Weiss, Y (1977)**, 'Inflation and costs of price adjustment', *Review of Economic Studies*, Vol. 44, pages 287–303.

**Sinclair, P (2003)**, 'Can menu costs justify inflation?', University of Birmingham, *mimeo.*

**Ueda, K (2002)**, 'The transmission mechanism of monetary policy near zero interest rates: the Japanese experience, 1998–2000', in Mahadeva, L and Sinclair, P (*op cit*), pages 127–36.

**Woodford, M (1990)**, 'The optimum quantity of money', in , Friedman, B and Hahn, F H (eds), *The handbook of monetary economics*, North Holland, Amsterdam.

**Wyplosz, C (2001)**, 'Do we know how low should inflation be?,' in European Central Bank (*op cit*), pages 15–33.

**Yates, A (1998)**, 'Are prices and wages sticky downwards?', *Bank of England Quarterly Bulletin*, Vol. 38, No. 3, pages 233–47.

**Yates, A (2003)**, 'Monetary policy and the zero bound to nominal interest rates', *Bank of England Quarterly Bulletin*, Spring, pages 27–37.