

Issues related to monetary policy rules

In the recent academic literature, monetary policy rules have become a prominent feature. In this context, a policy rule is often understood as a description, for all possible contingencies, of how a policy instrument reacts to changes in the economic environment. Such a rule is typically either postulated in a simple form, linking the policy instrument to a small set of economic variables or indicators, or it is derived explicitly from an optimisation problem given a particular representation of policy objectives and the working of the economy.

This article discusses some issues related to such rules for monetary policy. It argues that simple rules linking changes in the monetary policy instrument directly to the evolution of a restricted number of indicator variables are too rigid and inefficient. At the same time, optimising rules based on some given model of the economy cannot take sufficient account of the limitations of the central bank's knowledge as regards the state of the economy and the "true" economic model.

In practice, central banks have to cope with the fact that knowledge about the economy is imperfect. For this reason, a commitment to a monetary policy strategy, which sets out the policy objective and the tools used to achieve it, is the appropriate way to take into account all relevant information, as well as model uncertainty and structural change in the economy. A monetary policy strategy provides a systematic framework for the analysis of information and the taking of policy decisions, without specific policy conclusions being predetermined in a mechanical manner.

I Rules in monetary policy

The experience of the 1970s awakened policy-makers, observers and the general public to the causes and costs of high inflation. The late 1970s thus saw a lively resurgence of interest among economists in the issue of optimal monetary policy design. In academic circles, this period witnessed a revival of the long-standing monetary policy debate on rules versus discretion. The voluminous literature on monetary policy rules that has developed since then has contributed a number of key insights into the "science of monetary policy".

The first key contribution concerned the nature, scope and limits of monetary policy-making: in short, it considered the legitimate long-term objectives of monetary policy-making. In its original form, this debate built on the observation that paper money has historically created a temptation to engineer inflation surprises on an unsuspecting public. One source of such temptation has been attempts by governments to pursue a policy aimed at maintaining output above the "natural" or "potential" level to which it will gravitate in the long run. However, monetary "surprises" aimed at boosting output in the short run become ingrained in expectations

over time. Moreover, if price setting depends on expectations of the future, higher expected inflation quickly turns into higher actual inflation. Thus, stimulative monetary "surprises" would rapidly lose their leverage – being anyhow only of a temporary nature – over the level of real economic activity. Against this background, any announcement by the central bank of its determination to keep inflation low and stable would not be believed. Inflation would be permanently higher – reflecting the so-called "inflation bias" – without any lasting gains in terms of output and employment, compared with a situation where policy-makers could commit themselves not to succumb to the temptation to surprise.

It was concluded that any appropriate policy rule should preclude attempts by the central bank to push output above its "natural" or "potential" level. One way to achieve this was to grant central banks institutional independence and assign them a clear, overriding mandate to maintain price stability on the grounds that this would serve to insulate monetary policy from pressures to pursue inappropriate objectives.

However, the problem of the credibility of monetary policy has not disappeared with the resolution of the academic dispute on objectives. Assigning a central bank a mandate based on outcomes – say, the achievement of price stability or low inflation – is no guarantee in theory that the preferred outcome will ultimately be delivered. Moreover, an understanding of the central bank's actions is important for its credibility and, therefore, its effectiveness in achieving its objectives. Thus it is now increasingly recognised that the case is strong for monetary policy to behave in a predictable and systematic way over time in order to have a stabilising effect on expectations. In this context, a clear framework, or strategy, that disciplines policy choices and keeps decisions consistently anchored to the mandated objectives can enhance the macroeconomic outcome in the medium term.

The notion of consistency in the way monetary policy responds to new occurrences has thus restated the importance of rules for monetary policy from a different angle. At the same time, the concept of monetary policy rules has been broadened considerably. In most recent literature, the scope of a rule goes beyond the long-term objectives of a central bank and embraces

the procedures and strategies that should systematically guide the conduct of monetary policy along its way.

The renewed emphasis on rule-guided monetary policy in recent academic literature is generally welcome on two grounds. First, it provides a salutary antidote to the perennial risks of a discretionary, ad hoc approach to policy-making. Second, more recent literature has begun to move in the direction of greater realism, and thus greater relevance for practical policy, by restating the role of rules in the form of strategies and procedures guiding the daily conduct of monetary policy. The challenge for monetary policy in practice is to retain the virtues of rule-based policy-making, while taking into account the complex, uncertain and constantly evolving environment facing monetary policy-makers. In this respect, an explicit framework for information processing and decision-making by central banks can provide a further disciplining element in addition to a strong commitment to the overriding policy objective. As a consequence, the public is more likely to perceive monetary policy as moving steadily in a clear direction towards the indicated end-point, facilitating the achievement of the ultimate goal of price stability.

2 Simple rules

Traditionally, central banks and academics have tended to seek simple policy rules in order to reduce discretion and foster credibility. Simple rules were seen as a safeguard against overly ambitious policies, which were likely to become a source of additional uncertainty in the presence of long, variable and uncertain effects of policy on the economy.

In particular, simple *unconditional* policy rules have a long and distinguished history in monetary economics. The gold standard and, in general, all regimes making paper money directly or indirectly convertible into a precious commodity at a fixed price are

prominent examples of this class of simple rules. Another simple unconditional rule providing a nominal anchor for the operation of a fiat currency is the constant money growth rule advocated by Milton Friedman, among others.¹ According to his proposal, the central bank should establish a constant rate of growth for the stock of money and maintain that growth rate consistently. Both a commodity currency regime, such as the gold standard, and the constant money growth rule rely entirely on a self-balancing

¹ See, for example, Milton Friedman (1956): "The quantity theory of money: A restatement" in *Studies in Quantity Theory*, Chicago University Press.

endogenous reaction in the real interest rate relevant for private sector decisions to changes in output and prices. For example, a rule keeping the growth rate of money constant would generate an endogenous rise in the real interest rate in the wake of an increase in aggregate demand above potential output. This adjustment in the real interest rate would counteract the inflationary pressures associated with the excess demand.

However, in many circumstances such self-balancing forces in the economy may not operate to a sufficient degree or at an acceptable speed. Under certain circumstances, simple unconditional rules – like an unfettered gold standard or a constant money growth rule – may lead to undesirably high volatility in prices and output. As a consequence, such rules have not, in practice, been applied in their strict form.

Somewhat more elaborate *contingency* rules linking a policy instrument to a set of indicators have therefore been proposed in the theory of economic policy. A feedback or reaction formula of this kind makes the monetary policy instrument a mechanical function of a restricted number of information variables. The idea underlying simple feedback functions is, in principle, straightforward. Borrowing from control engineering, it builds a parallel between an economy and a mechanical system, the motion of which is controllable by an instrument. The feedback function spells out the way this instrument will react, over time, to what happens to the system in order to regulate and stabilise its functioning. Applied to the theory of monetary policy, a feedback function establishes a direct link between a number of selected information variables – deemed good indicators of risks to price stability – and an instrument used for monetary policy purposes.

The analytical framework of these simple state-dependent rules varies considerably across proposals. However, most share the notion that monetary policy should be geared

towards achieving price stability or a low inflation rate in the medium term. Some rules also reflect the idea that monetary policy should minimise undue short-run fluctuations of output around its long-term potential. Differences concern the policy instrument which the central bank is assumed to adjust in reaction to the state of the economy, and the information variables taken to signal which state has emerged.

As far as the selection of the instrument is concerned, the proposed alternative is between the stock of base money and a short-term interest rate under the control of the central bank. A rule based on setting the level of base money requires that the central bank conduct open market operations to the extent necessary to enforce the quantity indicated by the rule and to accept whatever interest rate is required for base money demand to absorb that prescribed quantity. Alternatively, a rule based on the setting of interest rates implies that the quantity of base money be adjusted via appropriate open market operations to clear the market for base money at the particular interest rate implied by the rule.

Information variables – acting as the triggers of change in the instrument – are also diverse across rules. They may or may not include the variables directly representing the ultimate policy objective, i.e. first and foremost the inflation rate. In any event, information variables include macroeconomic indicators thought to be suggestive of the extent to which the goal variables risk departing from the stated targets.

One example of a simple feedback formula is the base-money rule, which was proposed by McCallum.² According to this rule, the central bank adjusts the monetary base in response to (i) deviations in the growth rate of nominal GDP from a specified target and (ii) some estimate of changes in money base velocity.

² See Bennett T. McCallum (1988): "Robustness properties of a rule for monetary policy", *Carnegie-Rochester Conference Series on Public Policy* 29.

This type of rule has attracted somewhat less attention in recent years. This may be partly due to the tendency to assign monetary policy an unambiguous role in maintaining price stability, while a nominal GDP target may tend to blur responsibilities. In addition, central banks' operational frameworks, by and large, make it more natural to think of the interest rate as the policy instrument rather than the monetary base.

The following discussion, therefore, focuses on an example of a simple rule formulated in terms of the policy interest rate, namely a type of rule widely known as the "Taylor rule".³ This rule has become rather popular both in academic literature and among professional central bank watchers in recent years.

Taylor-type rules

A conventional linear formulation of the Taylor rule is the following:

$$(1) \quad i_t = r^* + \pi^* + a(\pi_t - \pi^*) + \beta(y_t - y^*)$$

where the short-term nominal interest rate i_t decided by the central bank at time t is set to track its long-term level (which is given by the sum of the long-run equilibrium value of the short-term real interest rate r^* and a long-term inflation objective π^*), unless contemporaneous inflation π_t is out of line with its long-term objective π^* and/or output y_t deviates from its long-term potential level y^* . In this formulation, the weights a and β assigned to the inflation and output deviations measure how aggressively policy should respond to deviations in inflation from its target, and in output from its potential level.

Typically, a is set at a value in excess of unity. This numerical constraint, known as the "Taylor principle", is thought to ensure that observed signs of inflationary pressures are met with a tightening of policy to a degree sufficient to induce an increase in the real rate of interest. This rise in the real rate of interest, in turn, is considered a guarantee

that the destabilising forces acting on contemporary consumption and production decisions via inflation expectations are countered effectively and that aggregate spending is restrained in an equilibrating fashion. Some alternative representations of the rule feature a number of additional lagged terms on the right-hand side of the above expression, including lagged terms of the nominal interest rate instrument itself.

Forecast-based variants of this rule, featuring expected inflation $E_t \pi_{t+k}$ at the horizon k coinciding with the typical lag of monetary transmission, have also been proposed. Promoters of this latter version of the Taylor rule regard the inflation forecast term as an intermediate target variable of monetary policy, and often drop the output gap expression from (1). This is done on two grounds. First, it is argued that, when framing their response pattern, monetary authorities need to be conscious of the lags between the enactment of policy and its impact on inflation. These lags are regarded as being conveniently incorporated by choosing an appropriate forecast horizon k and by regarding $E_t \pi_{t+k}$ as a leading indicator of future price pressures. Second, since the current output gap is considered useful for predicting future inflation, it is already implicitly taken into account by the inflation forecast term and is thus not needed as an additional variable in (1), unless there is an explicit output smoothing objective (see Box I for details).

Proponents ascribe a number of virtues to Taylor rules. First, they are seen as very simple, easy to execute for the central bank and easy to verify, ex post, for the private sector. They therefore seem to simplify the communication of policy orientations to the general public. Second, proponents of the forecast-based version go as far as arguing that, in using predicted inflation as a trigger for reaction, this version of the rule

³ See the seminal paper by John B. Taylor (1993): "Discretion versus policy rules in practice", *Carnegie-Rochester Conference Series on Public Policy* 39.

encompasses all the relevant information for the purpose of policy-making.

A discussion of Taylor-type rules

It is a principle of good policy management that evidence of an incipient departure of key macroeconomic indicators from the values considered compatible with the objectives should make central banks vigilant and ready to act. In this sense, any systematic rule which feeds back from signs of divergence between objectives, on the one hand, and long-run sustainable values or outturns, on the other, encapsulates features of standard practice among stability-oriented central banks. This may partly account for the apparent success of some of these simple rules in loosely tracking past policy moves by central banks.⁴

However, it would be misleading to broaden the interpretation of these tests to inferences about the actual motives behind these central banks' steps in the conduct of policy. It should be noted that a number of monetary strategies, including the pursuit of a broad money growth target, if successful in maintaining the purchasing power of the currency, may – over a long sample period – be empirically indistinguishable from a policy wedded to the Taylor rule. As a consequence, a good econometric fit of a Taylor-type rule would have little, if anything, to suggest how central banks reacted to economic data and which indicators they actually consulted in the process of framing decisions. In addition, in making such empirical assessments, it should be kept in mind that econometric results, in general, appear to be very different depending on whether real-time or, alternatively, successively revised time series for output gaps are used.⁵

Normative implications of Taylor rules are, of course, even more difficult to substantiate. The first and foremost note of caution stems from considerations of efficiency in the use of information for policy purposes and applies to all simple rules in general. It cannot be assumed that all relevant information needed

to conduct monetary policy is encapsulated in current inflation and the output gap. Other variables, such as monetary and credit aggregates, fluctuations in exchange rates, stock valuations, fiscal indicators, variations in international commodity prices and wage agreements are highly indicative of macroeconomic developments and thus help to interpret the current economic situation. Much of the daily work conducted in central banks is devoted to tackling the information problem. The collection of as large a body of statistics as necessary is undertaken to enable conclusions to be drawn about the sources and propagation patterns of business cycle shocks, their nature and duration, and their structural implications. A simple Taylor rule would be incapable of processing all the material and would ignore the great bulk of the insight that this wealth of evidence routinely provides to decision-makers.

Furthermore, different sources of shocks call for very different policy responses. The need for policy to react in the face of incoming evidence depends, *inter alia*, on whether shocks arise from the supply or demand side of the economy and whether they represent temporary disturbances to an unchanged underlying structure or a lasting alteration of economic parameters. Demand shocks are typically associated with deviations of inflation from the objective and of output from trend in the same direction. For instance, higher demand is generally associated with a hike in inflation and an upsurge in output. In these circumstances, the change in the real interest rate suggested by the rule tends to have an equilibrating impact on both prices and output. It may thus be deemed appropriate to alleviate the contemporary price pressures by facilitating a reabsorption of emerging capacity excesses, thereby removing the output conditions that could perpetuate those pressures in the future.

⁴ See, for example, Richard Clarida, Jordi Galí and Mark Gertler (1998): "Monetary policy rules in practice: some international evidence", *European Economic Review* 42, pp. 1033-68.

⁵ See, for example, Anastasios Orphanides (2000): "The quest for prosperity without inflation", *ECB Working Paper No. 15, March 2000*.

By contrast, monetary authorities often need to react differently to a shock on the supply side, which causes output and prices to move in opposite directions. In this instance, the potential of such a shock to develop into a self-perpetuating destabilising force has to be assessed primarily on the basis of labour and goods market information that cannot be adequately processed using the Taylor formula. Price-setting habits and wage agreements have to be attentively monitored and taken into account to ensure that the shock remains temporary and does not affect inflation expectations in a permanent fashion.

In short, driving forces of different natures, possibly associated with the same inflation outturn or forecast, require offsetting actions of varying intensity and duration, as they set in motion quite different dynamics and are associated with possibly opposite tendencies in the evolution of real variables. Taylor rules, by unduly restricting the universe of information brought to bear upon policy decisions, are not a reliable guide for policy from this perspective.

Furthermore, despite their much-popularised practical orientation, conventional Taylor rules are not as straightforward to implement as is sometimes argued. The output gap and the equilibrium real interest rate – both crucial to a normative usage of the rules – are non-observable variables. Their estimation is a very delicate task, which makes their systematic use as guides for daily policy management a perilous undertaking. Moreover, depending on the estimating methods employed, the resulting Taylor interest rates may vary over a wide range and thus not provide clear policy signals.

The output gap concept has proved elusive and available estimates are widely dispersed.⁶ Conventional detrending methodologies used to estimate the excess of actual output over capacity are notoriously prone to real-time mismeasurement and suffer from a lack of theoretical foundation. Similarly, attempts to make intensive use of economic theory in

constructing measures of the “natural” level of output, at which price pressures are supposedly absent, also rely on questionable assumptions and do not constitute a reliable basis for policy decisions. In all respects, linking policy steps to output gap estimates that are highly vulnerable to ex post revisions or sensitive to specific proxying hypotheses appears to be a hazardous experiment. In such circumstances, the risk of bad policy outcomes is significant.

The equilibrium real interest rate is also hard to conjecture. In a meaningful Taylor rule, the equilibrium real interest rate should be an index of a wider array of underlying financial conditions, the realisation of which is considered compatible with stable prices. However, in the absence of a reliable and uncontroversial model gauging actual asset valuations against their “fundamentals” and spelling out the economic mechanisms connecting fundamental asset prices to the natural level of activity, any quantification of this equilibrium concept is bound to be a crude guess.

A further problem with Taylor rules of the forecast-based type arises if the output gap is included in addition to an inflation forecast. In this case, such a rule would appear to be inconsistent with the general notion that the primary objective of monetary policy shall be price stability, and could instead be seen as reflecting two distinct objectives in their own right.

Finally, the stabilising properties of Taylor-type rules can also be questioned. The criticism has two dimensions. In the first place, Taylor rules – being interest rate-centred – are particularly vulnerable to the standard problem that results from the fact that nominal interest rates cannot be forced to be negative. Consequently, Taylor rules become rapidly ineffective in keeping the system anchored to the policy objectives in

⁶ See, for example, the article entitled “Potential output growth and output gaps: concepts, uses and estimates” in the October 2000 issue of the *Monthly Bulletin*.

situations where nominal interest rates decline to very low values.

excessively sensitive to arbitrary revisions of expectations.

Another instance in which Taylor-type rules fail as equilibrating devices and may, in fact, become an independent source of instability is when they are formulated in a forecast-based fashion. As argued in greater detail in Box I, Taylor rules of this sort can exacerbate the tendency of economic systems to be

Partly as a reflection of the above observations, decision-making bodies in central banks cannot mechanically apply the Taylor-type rules assumed in the theory. The informational basis, upon which they are designed to function, is simply too narrow to be of practical assistance in conducting policy.

Box I

Are forecast-based Taylor rules always stabilising?

The Taylor rule has found extensive use as a convenient analytical device to formalise policy behaviour within the framework of a dynamic general equilibrium model of the private sector. It is in this analytical context that its capability to anchor macroeconomic magnitudes to policy objectives has been studied.

Dynamic general equilibrium models formalise the motion of the economy over time by means of a set of analytical expressions, which stem directly from optimising conditions concerning the consumption, saving, investment and production decisions of representative economic agents and firms. An extremely simplified version of this class of private sector model can be reduced to just two summary conditions, which express the current state of the private economy as functions of current shocks and expectations about the future:

$$(2) \quad y_t = \gamma_0 - \gamma_1 (\dot{i}_t - E_t \pi_{t+1}) + E_t y_{t+1} + e_t$$

$$(3) \quad \pi_t = \delta_0 E_t \pi_{t+1} + \delta_1 (y_t - y^*) + u_t$$

In (2) and (3), e_t and u_t are stochastic error terms and $\gamma_0, \gamma_1, \delta_1 > 0$ and $0 < \delta_0 < 1$ are given parameters. Equation (2) states that current output decisions react negatively to the contemporaneous real interest rate ($\dot{i}_t - E_t \pi_{t+1}$) and positively to expectations regarding future output conditions, $E_t y_{t+1}$. This is consistent with the observation that a higher real cost of borrowing impacts negatively on firms' production, whereas rosier prospects for future production encourage investment and thus expand current output. Equation (3) assumes that observed price adjustments – as captured by current inflation π_t – react to expectations of future inflation and to the current level of resource utilisation, as proxied by the output gap, $(y_t - y^*)$. This condition reflects the assumption that firms, operating in an imperfectly competitive market, face costs in revising prices, so that, at any time, only a fraction of them post new prices. This fraction thus has an incentive to look ahead to future inflation, knowing that it will be costly to modify the current pricing decision again. Equation (2) is often referred to as portraying the evolution of aggregate demand for a given policy stance. Equation (3), by contrast, captures the supply side. Both equations are usually enriched by a complex structure of lagged terms for y_t and π_t , which are ignored in the formulation given above for the sake of simplicity.

A model such as that represented by equations (2) and (3) can be “closed” (i.e. solved for the relevant endogenous variables y_t and π_t) by appending a Taylor rule of, say, the forecast-based type described by the following equation (1a):

$$(1a) \quad \dot{i}_t = r^* + \pi^* + a (E_t \pi_{t+k} - \pi^*)$$

As explained in the text proper, advocates of (1a) stress the advantages attached to having policy decisions at time t react to the rate of inflation anticipated to prevail at a horizon k equal to the typical lag between the

taking of monetary measures and their impact on price determination. Existing macroeconomic models – it is maintained – can rationalise a transmission lag k of up to two years.

The issue of whether a policy programme based on a rule such as (1a) can provide an adequate anchor for nominal and real magnitudes in an economy described by equations (2) and (3) has attracted growing attention in recent years from both a theoretical and an applied perspective. Results are not always encouraging, however. Studies using numerical simulations have noted that rules such as (1a) sometimes turn out to be an important source of instability in that they render the economy prone to arbitrary revisions of expectations which are not justified by the structural fundamentals of the system as represented by preferences and technologies. The likelihood of encountering these instability problems proves to rise with the length of the chosen horizon, i.e. with k .

The ultimate origin of instability is twofold. First, if point inflation projections are surrounded by a wide area of uncertainty (as is likely the case in real forecasting exercises), which grows as the forecasting horizon lengthens, then the high sensitivity of policy to these forecasts (and their revisions) may induce excessive volatility in inflation and output outturns. Second, instability may stem from a more general problem which typically affects systems – like economies – the current state of which depends in crucial ways on expectations about their future state. Since there are generally multiple ways to form expectations, it is possible that shocks to expectations – even those completely divorced from changes in fundamentals – may lead to a number of different plausible current states of the economy. Rules that link policy action to forecasts make the current state of the economy especially sensitive to changes in expectations and thus make the system vulnerable to this sort of multiplicity. Instability arises from the fact that, in these circumstances, it is not fully determined how the system may respond to an exogenous shock such as e_t or u_t in equations (2) and (3).

In conclusion, the stabilising properties of Taylor-type rules such as (1a) deteriorate in response to certain events. Even if all the information and efficiency problems discussed in this article are left unconsidered, the capacity of such rules to provide the anchor that the economy needs to be firmly attached to the intended policy objectives may depend on the absence of destabilising shocks to expectations. The economy may thus need an anchor that the Taylor-type policy in (1a) appears incapable of offering.

3 Optimising rules

In order to assess the performance of simple rules such as the Taylor rule, their properties need to be examined in the context of some model (or range of models) of the economy, such as that illustrated in Box I. In addition, a measure of society's welfare or a policy objective function needs to be specified to allow a comparison and a ranking of the macroeconomic outcomes associated with the use of different policy rules.⁷ Given such a model and a specification of the objective function, it is then – at least theoretically – an obvious step to try to find a fully optimal rule that maximises the objective function, rather than to implement simple rules, which are likely to be suboptimal. If the fully optimal

rule can be derived from such an optimisation procedure, there would thus seem not to be a case for considering simple rules such as Taylor's.

Indeed, a popular approach to modelling economic policy in the past decades has been to derive the optimal path for the setting of policy instruments starting from a specification of the objective function and a model describing the working of the economy. This approach rests on the

⁷ See, for example, the analysis of Taylor-type rules in the various models contained in John Taylor (1999): "Monetary policy rules", NBER Conference Report, University of Chicago Press.

assumption that there is a model of the economy in which policy-relevant variables such as inflation and output respond to policy measures according to a known pattern of reactions. Central banks should thus implement the optimal time path of the policy instrument derived from the model, namely the time path of the short-term interest rate which maximises a given specification of a policy objective function. That path would yield a macroeconomic outcome which, by construction, would be preferable to any alternative policy scenario.

In its canonical characterisation, this approach would amount to solving a standard constrained optimisation problem, with the model equations summarising the dynamics of the system acting as the constraints. The analytical solution to this procedure would yield a quite complex reaction formula whereby the policy instrument would be expressed as a function of all the state variables figuring in the model. In this way, the “best” policy move would be made contingent on the entire history of shocks relevant to monetary policy according to the structure of the model.

More recently, there have been new attempts in literature to apply this optimising approach to the problem of monetary policy design within the context of a rule-like institutional environment. In this more recent version, the central bank is assigned the objective of minimising the deviation of the model-projected inflation rate at some fixed horizon from a pre-specified target. This literature uses the notion of a “loss function” to represent the objectives mandated to the central bank. Typically, this loss function is assumed to be quadratic, meaning that perceived losses around target levels would be symmetric and increasing in the target misses.

Within this context, the rule-like element is represented by the central bank’s commitment to an inflation target and to an optimising procedure – rather than a direct feedback rule – which has to be employed in

the actual pursuit of the target. This procedure would be optimising in the sense that it would simulate a given model for a set of alternative interest rate paths and select that which is “best” according to the assumed loss function.⁸

The prescriptions for virtuous central banking embodied in optimising rules address some of the criticisms of simple rules discussed in Section 2 in an abstract way. Optimising rules of this type can embody a resolutely forward-looking orientation of policy, while avoiding some of the drawbacks associated with a policy reacting mechanically to a specific inflation forecast (as under a forecast-based Taylor rule). In principle, such a procedure could use state-of-the-art economic modelling to process information and inform policy decisions. In addition – at least theoretically – it may ensure that the information set, upon which decisions are based, is much broader than would be the case under simple rules.

However, the kind of optimising rules described above – if taken literally and applied mechanically – remain too restrictive for policy purposes in several respects. First, the optimising procedure, as proposed, seems to underrate the need for judgement in the use and interpretation of any economic model. Second, optimising rules mandating central banks to select a policy path which ensures that projections of goal variables are in line with their targets at *predetermined horizons* unduly restrict the relevant time frame for policy. Given that the transmission lag is variable, it is difficult to determine the relevant horizon for the projection path. Moreover, different types of models capture different elements of the transmission mechanism and are relevant at different time horizons. Restricting attention to a specific projection horizon may, in such circumstances, induce short-sighted reactions, the effects of which may have to be

8 See, for example, Lars E. O. Svensson (1999): “Inflation targeting as a monetary policy rule”, *Journal of Monetary Economics* 43.

counteracted at a later date, with associated costs in terms of instability. Third, if optimising rules or procedures are implemented period by period, there appears to be no mechanism to ensure that the resulting policy recommendations and actions are consistent over time.

More generally, any such optimising rules or procedures are likely to remain too restrictive if the design of policy is conditional on the structure of any single model used. A sequence of policy moves which may be considered optimal on the basis of one model of the economy may often turn out to be associated with bad policy outcomes if simulated on the basis of a different model, representing alternative views about the workings of the transmission mechanism. Therefore, a variety of models need to be used for different purposes within central banks and various – more or less formal – ways of interpreting data and economic developments need to coexist. As a consequence, if a suite of models and other indicators are used and if they are complemented and combined with judgement, it becomes less clear how the proposed

optimising rule could still be implemented as a strictly codified procedure. Moreover, any resulting projection path, taken by itself, will contain only very limited information. In such circumstances, it is more important to understand the factors underlying such a projection and to convey the judgement that motivates one particular view of the world rather than another.

For all these reasons, optimising rules as proposed in academic literature, while feasible in principle, remain insufficient for practical use in the strict sense of a monetary policy rule. In particular, such rules in their current form do not take sufficient account of the problem of model uncertainty and its far-reaching consequences for central banks. This has recently led another strand of literature to consider different decision criteria based on the notion of robustness of policy measures rather than the traditional principle of optimality, as discussed in Box 2 in more detail. However, at the present stage, it would be rather premature to draw any general conclusions from this emerging, alternative literature.

Box 2

Model diversity and robustness

Critics of optimising rules have emphasised their lack of robustness to the uncertainty surrounding the functioning of the economy. Notably, they stress the fact that policy recommendations stemming from the use of optimising rules are typically highly model-dependent. Two main methodological approaches have received particular attention in recent economic literature in an attempt to address issues related to uncertainty about the “true” model of the economy.

One approach to model uncertainty retains a Bayesian, probabilistic representation of uncertainty and the optimising approach commonly used in economics. If applied in the context of an optimising rule, the Bayesian approach to tackling model uncertainty would require the central bank to adopt the following procedure. First, the central bank would have to identify a class of models considered plausible representations of the functioning of the economy. Second, it would assign some probability to the lack of these various possible specifications. These probabilities should reflect the central bank’s opinion on the likelihood of different models representing the “true” model of the economy. Lastly, the central bank would proceed by finding the path of its policy instrument that minimises its loss function conditional on the various models considered plausible, weighted by their respective likelihood. The resulting optimal path for the policy instrument would be a combination, i.e. a weighted average, of the optimal paths that would be found under each of the possible models considered.

Although this procedure could, in theory, be integrated into an optimising rule as described in this article, it is doubtful that this could be achieved in practice. In addition to the difficulties already mentioned, it would be very hard for the decision-making bodies of the central banks to specify numerically and reach an agreement on the probabilities to be assigned to the various models used in such a procedure.

Alternative approaches to model uncertainty do not require a well-specified probability distribution to be defined for the set of all the possible models describing the working of the economy, and thus may provide a way to deal with more pervasive, unstructured forms of uncertainty (e.g. of a “Knightian” nature). Such approaches need to consider alternative choice criteria, since traditional optimising techniques can no longer be applied. One extreme example of an alternative choice criterion is the so-called “minimax principle”. In this case, the central bank would, as before, have to identify a number of plausible models and consider a number of alternative paths for its instrument, but without attributing a probability distribution to the models considered. The policy path carrying the lowest maximum downside risk across all conceivable models – i.e. that generating the best of all worst-case scenarios – would be chosen as the “safest” and thus the most robust one. This policy action would be given the highest ranking according to the minimax procedure.

A procedure of this kind would recognise that the nature of uncertainty facing central banks often cannot be easily captured in probabilistic terms and in a way that renders the standard axiomatic choice theory and optimising approaches used in economics always applicable. However, there is no agreement on which type of alternative choice criteria should be considered in such circumstances. In general, it has been found that the implications for monetary policy often differ significantly between the standard optimising approaches and alternative ways to model “robust” decision-making.

4 Monetary policy rules and central bank practice

The above considerations have shown that, while there is a broad consensus on the importance of a systematic or rule-based approach to monetary policy for credibility, it is not possible – or at any rate not wise – for central banks, in practice, to formally commit to a specific rule prescribing the setting of policy instruments in precise terms. The basic reasons for this were illustrated in Sections 2 and 3. *Simple rules* are unable to take into account all relevant information to be considered by central banks and to offer appropriate guidance for stabilising the economy under all conceivable circumstances. Conversely, more ambitious *optimising rules*, which tend to be more complex, are difficult to implement, communicate and monitor in practice. In addition, any optimising rule is only as good as the model on which it is based. Even small changes to the model used can often lead to very different results. Such a lack of robustness may thus cause serious policy errors in an economic environment that is uncertain and subject to continuous change.

More generally, it needs to be recognised that all economic models, including those used to discuss and evaluate monetary policy rules, are by nature an abstract and incomplete representation of the economy and the behaviour of economic agents. They can only capture some particular aspects of reality, and different models are useful for different purposes. However, monetary policy has to deal with and be robust to manifold forms of uncertainty, which are only captured very imperfectly in economic modelling.⁹ In particular, model uncertainty implies that monetary policy cannot rely on any model-specific optimising rule. Instead, robust monetary policy-making needs to be compatible with different views of the structure of the economy and the monetary transmission process. In addition, the presence of model, parameter and data uncertainty – all else being

⁹ See the article entitled “Monetary policy-making under uncertainty” in the January 2001 issue of the *Monthly Bulletin*.

equal – cautions against an over-reliance on rules based on concepts or indicators (such as the equilibrium real interest rate or the output gap) which may be subject to large measurement errors and methodological dispute.

As a consequence of the degree of complexity and the nature of uncertainty involved in monetary policy, it is impossible, in practice, for central banks to write down the monetary policy decision problem in complete detail and to convey – with any degree of precision – its true “policy rule” or “reaction function” in the strict sense of the terms. That would require the central bank to specify a complete contingency plan describing the setting of policy as a function of an exhaustive list of possible events and circumstances to which central banks may react in the future.

Therefore, central banks not wishing to compromise on the efficiency and robustness of monetary policy for the sake of committing to a simple suboptimal rule or a model-specific optimising rule have no choice but to consider a broader notion of rule-governed or rule-based (rather than rule-bound) behaviour, as embodied, for example, in the commitment to an explicit monetary policy strategy. This is increasingly being taken into account in literature.

A monetary policy strategy can be defined in general terms as the central bank’s framework for the taking of monetary policy decisions and their explanation to the public. As such, a strategy comprises a set of procedures structuring the analysis of information and the decision-making process by the central bank. It provides the framework within which economic information is analysed, interpreted and explained for the purpose of monetary policy-making. The notion of a strategy as a framework or set of procedures differs from the traditional concept of a monetary policy rule.

While a strategy serving as a procedural framework will usually involve, as a rule does,

a definition of the central bank’s monetary policy objective, it will not strictly predetermine the specific policy actions required to reach that objective *ex ante*. Instead, a strategy committed to a procedure may entail a commitment to examine regularly a predefined set of economic indicators and analytical frameworks. Such a procedure may broadly set out which steps are to be followed to synthesize and cross-check information coming from various indicators and models. Thus, a procedural “rule for analysis” should, in general, be better equipped to take into account uncertainty about the nature and length of the transmission mechanism – as reflected in complementary, competing models of the economy – than a simple “rule for action” or an optimising rule based on a single model.

As a consequence, a strategy serving as a procedural framework allows greater emphasis on the interpretation of economic developments, the nature and origins of economic shocks affecting the central bank’s objectives and, ultimately, the economic “story” underlying its monetary policy decisions. In this regard, a strategy provides a framework for the systematic and consistent explanation of the considerations underpinning policy decisions, instead of representing policy as a reaction to individual indicators in isolation or the mechanical use of a specific model.

Central banks have in practice largely eschewed commitment to specific policy rules and they differ in the degree to which they have announced an explicit strategy. They also vary in the degree of precision with which objectives are defined and in the emphasis given to particular benchmarks or indicators. The remainder of this section briefly reviews the salient features of the ECB’s monetary policy strategy. This provides one illustration of a commitment to a procedural framework, which may overcome some of the limitations and risks associated with an over-reliance on more narrowly defined monetary policy rules, as discussed in Sections 2 and 3.

The monetary policy strategy of the ECB sets out a systematic framework for monetary policy focused on maintaining price stability over the medium term. This framework can be interpreted as being rule-based, in particular with regard to the following elements. First, the strategy includes a clear commitment to the goal variable, i.e. the primary objective of price stability. Second, the strategy sets out a “framework for analysis” in the form of a procedural rule. This entails a prior commitment to conduct analysis and to explain policy in a systematic and structured way. Third, the use of benchmarks and “prompters” for further analysis is to some extent present in the ECB’s strategy. This pertains, in particular, to the function of the monetary reference value within the broader analysis conducted under the first pillar. Lastly, the strategy may also be interpreted as broadly setting out the main features of the policy resulting from the procedural framework of analysis. In this context, the strategy, together with continuous explanation of monetary policy decisions under the strategy, should enable the public over time to trace and broadly anticipate how monetary policy reacts to observable data and indicators in a systematic manner.

The ECB’s monetary policy strategy as a whole can be seen as addressing some of the pitfalls of the simple and model-specific optimising rules identified in the previous sections, namely the need for information efficiency, a nominal anchor in all circumstances and robustness.

First, the strategy provides a framework ensuring an efficient use and effective structuring of all information needed to take monetary policy decisions aimed at maintaining price stability over the medium term. The relevant set of information clearly extends beyond those variables typically included in simple monetary targeting or Taylor-type rules. For example, financial variables, such as bond yields, asset prices (including exchange rates), credit developments and balance sheet positions,

provide additional information that is useful for monetary policy. Similarly, a host of survey and confidence indicators, as well as fiscal and labour market developments, are regularly scrutinised. Such a detailed and extensive range of information cannot be reduced to, or fully captured by, a few simple summary indicators of the kind typically featured in simple feedback rules.

In this context, the medium-term orientation of the ECB’s monetary policy strategy also implies that policy does not feed back from a forecast at a particular fixed horizon (as would be the case with a simple forecast-based rule). Instead, the entire transmission process over a number of years and the nature of shocks influencing price developments need to be taken into account in deciding on the appropriate monetary policy response. More generally, the medium-term focus recognises the presence of different transmission channels affecting price developments with long, variable and uncertain lags. Money growth, in particular, is associated with inflation at a medium to long-term horizon.

Second, the clear commitment to the maintenance of price stability over the medium term supplies a stable nominal anchor to the economy in all circumstances. The prominent role for money in the ECB’s strategy provides an additional safeguard in this regard, which is not present in standard Taylor rules.

Third, the two-pillar structure of the ECB’s monetary policy strategy takes explicit account of the need for robustness in monetary policy-making.¹⁰ Recognising different existing models of the structure of the economy and the nature of the monetary transmission mechanism, the ECB has chosen to organise its analysis into two pillars. The first pillar represents a group of models and analytical frameworks which embody a view

¹⁰ See the article entitled “The two pillars of the ECB’s monetary policy strategy” in the November 2000 issue of the *Monthly Bulletin*.

of price level determination that accords an important role to money. The second pillar encompasses a range of alternative models of the inflation process, predominantly those which emphasise the interplay between supply and demand in the goods and labour markets.

The two-pillar structure reduces the scope for discretion, as it makes it more difficult for policy-makers to disregard or gloss over contradictory evidence (as may happen with a single summary device, such as a single

inflation forecast). The two pillars of the strategy represent a commitment to always consider and base monetary policy decisions on a careful analysis of a wide range of information variables under both pillars of the strategy. In addition, if several plausible models (or, more broadly, modelling approaches) of the economy exist, taking this fact into account is likely to be superior to picking any particular “optimising” policy suggested by a specific modelling approach in isolation.

5 Concluding remarks

Commitment to a monetary policy strategy as described above places much higher demands in terms of transparency and effective communication on the central bank when explaining monetary policy decisions than would be the case with simple instrument rules as guides for policy decisions. However, given the shortcomings of the simple and optimising rules considered in this article, the use of any of these rules, even if only as benchmarks, would in many circumstances be misleading and not contribute to a better understanding of monetary policy. Thus, there is no convincing alternative to explaining monetary policy

decisions in a way that corresponds closely to the internal framework of analysis underlying the central bank’s decision-making process, rather than presenting them in terms of policy rules of the type discussed in this article and commonly used in academic literature. Moreover, a genuine understanding of the ECB’s monetary policy approach cannot be gained from these policy rules. Such understanding is more likely to be promoted over time if the ECB’s monetary policy is assessed on the basis of the systematic framework that the ECB has itself provided through the announcement of its monetary policy strategy.