MONETARY POLICY-MAKING IN THE PRESENCE OF KNIGHTIAN UNCERTAINTY

Adam Cagliarini and Alexandra Heath

Research Discussion Paper
2000–10

December 2000

Economic Research Department
Reserve Bank of Australia

An earlier version of this paper was presented at the Reserve Bank of New Zealand Workshop on 'Monetary Policy Decision-Making – Proactive or Reactive?', 8 June 2000. We thank participants at that conference for their comments. The authors would also like to thank Geoffrey Shuetrim, Matthew Ryan, David Gruen, Guy Debelle, Luci Ellis, Charlie Bean, Jenny Wilkinson and members of the Economic Research Department for their comments and questions. As always, we bear responsibility for the content of this paper and the views expressed are not necessarily those of the Reserve Bank of Australia.

Abstract

This paper explores the extent to which Knightian uncertainty can explain features of interest rate paths observed in practice that are not generally replicated by models of optimal monetary policy. Interest rates tend to move in a sequence of steps in a given direction, or remain constant for some time, rather than experiencing the frequent reversals that commonly arise from optimal policy simulations. We categorise the types of uncertainty that have been explored to date in terms of the decision-making behaviour they imply. From this, we suggest a more intuitively appealing formulation of Knightian uncertainty than the one that has previously been used in the analysis of monetary policy. Within a very simple optimal control problem, we show that our preferred formalisation is consistent with interest rate paths with periods of no change. This suggests that the presence of Knightian uncertainty may explain some features of monetary policy-makers' behaviour.

JEL Classification Numbers: D81, E52

Keywords: Knightian uncertainty, monetary policy

Table of Contents

MONETARY POLICY-MAKING IN THE PRESENCE OF KNIGHTIAN UNCERTAINTY

Adam Cagliarini and Alexandra Heath

1. Introduction

One of the stylised facts of monetary policy is that interest rates generally move in a succession of small steps in the same direction and can remain fixed at the same level for some time. However, models of optimal monetary policy behaviour tend to generate much more volatile paths of interest rates, in which policy reversals are frequent. Clearly, central bankers are not behaving the way the models predict they should. This indicates that there is a need to reconcile model-generated and observed behaviour.

Recently, there has been a growing body of literature that has focused on the role of uncertainty and its effects on monetary policy decisions to this end. One line of this research has shown that some of the volatility in interest rate paths generated by models can be eliminated by considering uncertainty arising from the estimation of the model's parameters. In this case, the distribution of possible outcomes is determined by the probability distribution of the parameter estimates. Consequently, the standard assumptions underlying expected utility theory apply, and the methods of solving these decision problems are essentially the same as those that are used when parameter uncertainty is not a part of the decision-making environment.

A second line of research considers the consequences for interest rates of model uncertainty, and finds that the policy-maker is likely to be more responsive to changes in the economy rather than less in this environment. Another line of research has found that policy-makers, faced with uncertainty about the data they have available, either because of measurement issues or the possibility of future revisions, will tend to act more cautiously. The feature that distinguishes data and model uncertainty from parameter uncertainty is that there is no straightforward way of characterising these types of uncertainty with a unique probability distribution. Following the literature, we label this Knightian uncertainty.

An important implication of allowing for the possibility that future outcomes for the economy cannot be characterised by a single probability distribution is that the policy-maker can no longer be characterised as an expected utility maximiser. This paper discusses two alternative ways of approaching decision-making in the presence of Knightian uncertainty and the likely consequences they would each have for the observed path of interest rates. We conclude that an appropriate formulation of the monetary policy decision-making process in the presence of Knightian uncertainty has the potential to explain the inertia in the level of the interest rate observed in actual interest rate paths.

The rest of this paper is structured as follows. In Section 2, we review the current literature that seeks to reconcile observed and model-generated interest rate paths. In particular, we discuss the different forms of uncertainty that have been investigated, and use the results of this literature to motivate our interest in Knightian uncertainty. In Section 3 we make the distinction between risk, which can be thought of as uncertainty characterised by a unique probability distribution, and uncertainty, where such a characterisation is not possible. This is followed in Section 4 by a discussion of the two existing approaches to modelling decision-making under Knightian uncertainty. In Section 5 we consider the applicability of each of these approaches to the monetary policy decision-making problem and in Section 6, we discuss how the standard tools of optimal control can be adapted to allow us to operationalise the form of Knightian uncertainty we argue is most appropriate. We also show simulations of a simple closed-economy model to formalise the intuition presented in Section 5. We conclude in Section 7 by assessing the extent to which incorporating Knightian uncertainty into optimal monetary policy models has allowed us to reconcile model-generated interest rate paths with those observed in the real world.

2. The Current Literature

In the literature, the monetary policy-maker's decision has usually been described as choosing a path of interest rates to minimise losses due to deviations of output from potential and inflation from target, subject to the dynamics of the economy. The policy-maker's objective function is usually assumed to be quadratic, and the dynamics of the economy are assumed to be linear. Using these assumptions, the monetary policy decision-maker's problem can be written more formally as:

$$\min_{\{i\}} E\sum_{t=1}^{N}\delta^{t-1}x_t{}'\Omega\, x_t \tag{1}$$

subject to:

$$x_t = A + Bx_{t-1} + Cr_{t-1} + \Lambda_t \tag{2}$$

where:

- x is a vector of variables that affect the policy-maker's losses, which generally includes deviations of inflation from the inflation target and output from potential;
- Ω summarises the preferences of the policy-maker by assigning weights to each policy objective;
- i is the path of nominal interest rates;
- r is the path of real interest rates;
- δ is the discount rate;
- B describes the dynamic structure of the model;
- C describes how the economy responds to the policy instrument; and
- Λ captures additive shocks to the economy.

One of the puzzles that arises when estimating paths of optimal policy using this framework, is that the optimal path of interest rates is much more volatile and subject to policy reversals than paths that are typically observed (Lowe and Ellis 1997). Eijffinger, Schaling and Verhagen (1999) and Guthrie and Wright (2000) introduce small menu costs into the policy-maker's environment to generate extended periods of no change in model-generated interest rate paths. While this exercise would appear to go some way towards reconciling model-generated and observed interest rate paths, it is not clear what the fixed costs of changing interest rates are that would generate this result, making it difficult to justify this approach.

Blinder (1998) suggests that actual interest rates may be less volatile than those predicted by these models because central bankers take into account the fact that they are uncertain about the models they are working with. Since this suggestion was made, a significant body of research investigating the effects of uncertainty on model-generated interest rate paths has emerged.

One important source of uncertainty arises from the fact that the models being used are estimated and consequently, policy-makers are uncertain about the true parameter values. Parameter uncertainty can be characterised by a normal distribution around the estimated parameters with the estimated variance-covariance matrix. The effects of this specification of parameter uncertainty were first investigated by Brainard (1967) who shows, in a simple static model of the macroeconomy, that adjustments to the policy instrument will be dampened if the policy-maker is uncertain about the parameters of the model. However, this result does not generalise to the dynamic case. Shuetrim and Thompson (1999) show that although uncertainty about the effect of the interest rate on output decreases the willingness of policy-makers to change interest rates, uncertainty about the dynamic structure of the model can lead to the opposite result: policy-makers may wish to change interest rates by more than they would in the absence of uncertainty. Similar theoretical results have been established by Söderström (1999).

Empirical work presented by Sack (2000) and Debelle and Cagliarini (2000) suggests that the first of these two effects dominates, and that the path of optimal policy does appear to be less volatile when parameter uncertainty is taken into account. However, other studies, notably Rudebusch (1999) and Estrella and Mishkin (1999), do not find such convincing results. Although Sack and Wieland (1999) suggest that the differences in these results may be due to the degree of richness in the dynamic structures or to the number of ‘uncertain’ parameters considered, the general conclusion appears to be that parameter uncertainty is of secondary importance for explaining the differences between model-generated and actual interest rate paths.

Policy-makers, however, may be less clear about the specific nature of parameter uncertainty than this research assumes. In fact, policy-makers may be uncertain not only about the parameters, but also about the general specification of the model being used. This concern about general model uncertainty is captured by Blinder (1998, pp 12–13) who recommended:

> Use a wide variety of models and don't ever trust one of them too much. … My usual procedure was to simulate policy on as many of these models as possible, throw out the outlier(s), and average the rest to get a point estimate of a dynamic multiplier path. This can be viewed as a rough - make that very rough - approximation to optimal information processing.

As this quote suggests, general model uncertainty, unlike the approach to parameter uncertainty discussed above, cannot be conveniently summarised by a single probability distribution. Allowing for a less structured form of uncertainty requires more sophisticated techniques. Recently, a number of papers have used robust control methods to deal with general model uncertainty, which includes general parameter uncertainty as a special case. In order to obtain solutions, this literature makes the extra assumption that policy-makers make interest rate decisions that are robust to a specified range of model mis-specifications. This means that, having decided the range within which the true model can deviate from the estimated model, the policy-maker will modify policy settings to minimise the losses associated with the worst possible outcome. Policy-makers who are more averse to uncertainty and wish policy to be more robust will consider a wider range of possible deviations from the estimated model.

Analysis that considers very general forms of model uncertainty finds that policy reacts more to deviations of inflation from target and output from potential than would be the case if this uncertainty were not taken into account (Sargent 1999; Onatski and Stock 2000; Tetlow and von zur Muehlen 2000). Consequently, the path of interest rates appears to be more volatile, rather than less, when general model uncertainty is taken into account. In Section 5 we discuss the intuition of this result in more detail, in the context of decision-making under different forms of Knightian uncertainty.

The final form of uncertainty that has received attention is data uncertainty. This form of uncertainty becomes particularly important when policy responds to deviations of output from potential. One source of uncertainty about the output gap is the fact that output data are frequently revised. Policy made on the basis of first-release data may not look optimal by the time the data have been revised. It is also possible that with the help of coincident indicators, which are not subject to the same measurement problems, policy-makers could have a better view of the true state of the economy than the output data and consequently, their actions may look better-considered as the data are revised. The second, and more significant, problem facing measurement of the output gap is that potential output is not directly observed, and there is little agreement as to how it should be estimated.

Orphanides (1998) examines the effects of measurement error on the performance of efficient rules. Orphanides compares data that were available at the time monetary policy decisions were made with *ex post* data to evaluate the size and nature of measurement error in the United States. He then estimates the impact this data noise has on the behaviour of policy-makers. He finds that in the presence of data noise, a central bank following an optimal Taylor rule will act more cautiously in response to output and inflation data than they would if they were certain that the data were correct. Similar results have been found by Smets (1998) and Orphanides *et al* (1999) using more sophisticated macroeconomic models.

In summary, several types of uncertainty facing policy-makers have been explored as possible inclusions to the standard optimal policy framework to help reconcile model-generated and observed interest rate paths. Parameter uncertainty is the best understood of these, partly because it is relatively straightforward to implement. Unfortunately, parameter uncertainty characterised by a single probability distribution does not appear to be sufficient to explain differences between model-generated and observed interest rates. Both model and data uncertainty can be thought of as examples of Knightian uncertainty insofar as it is not clear how to capture them with a single probability distribution. Despite this, they appear to have different implications for the behaviour of interest rates. In Sections 3 and 4 we formalise the distinction between risk and uncertainty and the different ways decision-making under Knightian uncertainty can be modelled in order to understand these results more clearly.

3. Risk versus Uncertainty

One of the first economists to make a distinction between risk and uncertainty was Frank Knight (1933). Knight's interest in this subject was spurred by the desire to explain the role of entrepreneurship and profit in the production process. Knight held the view that profits accruing to entrepreneurs are justified and explained by the fact that they bear the consequences of the risks (uncertainties) inherent in the production process that cannot be readily quantified.

> There is a fundamental distinction between the reward for taking a known risk and that for assuming a risk whose value itself is not known. It is so fundamental, indeed, that, as we shall see, a known risk will not lead to any reward or special payment at all (p 44).

Known risks arise in situations where outcomes are governed by physical laws, e.g. a dice roll, or the factors affecting future outcomes remain more or less constant over time, e.g. mortality tables. In these cases, past observations of the distribution of outcomes will be a good guide for the distribution of outcomes that can be expected in future. At some point, however, the processes that need to be forecast become sufficiently complex and interrelated with the outcomes of the decisions of other agents, that the past does not provide such reliable information about the likelihood of future events occurring.

> The fact is that while a single situation involving a known risk may be regarded as 'uncertain', this uncertainty is easily converted into effective certainty; for in a considerable number of such cases the results become predictable in accordance with the laws of chance, and the error in such prediction approaches zero as the number of cases is increased (p 46).

LeRoy and Singell (1987) summarise Knight's distinction between risk and uncertainty by defining uncertainty as a situation where no objective, or publicly verifiable, probability distribution exists. In situations where a single objective probability distribution does not exist, LeRoy and Singell argue that Knight's exposition is consistent with the idea that the decision-maker forms some subjective probability. When a single probability distribution is available, it is straightforward to evaluate the expected value of pay-offs to different actions. However, forming a unique subjective probability distribution may not be straightforward as the Ellsberg paradox helps to illustrate.

Suppose we have a box of 300 balls, 100 of which are red and the rest are blue and green in undisclosed proportions. A ball is chosen at random from the box. Suppose we are offered the choice of betting on whether a red ball or a blue ball would be selected. Which should we choose to gamble on? Now suppose we are faced with a different gamble. We have to choose between betting on whether the ball is not red or not blue. Which gamble do we select in this case?

In most instances, individuals will pick red and not red in response to these two questions (Kreps (1990) based on Ellsberg (1961)). If red is selected in the first gamble, the participant has implicitly evaluated their subjective probability of getting a blue ball to be less than the probability of a red ball. The paradox here is that in order for the same individual to act rationally according to the axioms of expected utility theory when faced with the second choice, they should bet that a blue ball will not be chosen. The Ellsberg paradox highlights the fact that people prefer situations when the probabilities are known to those where they are unknown. This, in turn, implies that individuals do not always employ a single subjective probability distribution to resolve their uncertainty. Red may be chosen in the first gamble, not because we believe the probability of getting a red ball is greater than a blue ball, but because we know the exact probability of getting the red ball.

The question is, how should an individual make rational decisions when a range of probability distributions are possible? The Bayesian approach to this problem would be to decide on a probability distribution over the possible probability distributions, which essentially reduces the problem to one in which there is a single probability distribution. This solution assumes that the decision-maker is willing to make definite statements about the distribution of future outcomes. The problem with using a uniform distribution over alternatives or a diffuse prior, is that there is no information to justify using this distribution. Schmeidler (1989) argues that it is unreasonable to apply equal probabilities to unknown events when the information about these events is limited. He argues that only when symmetric information about these events is abundant should equal probabilities be applied. Therefore, the interest in the problem arises from the fact that decision-makers are either unwilling or unable to make such a strong assumption about probabilities at the first stage of the decision-making process.

Another approach is to recognise that some of the axioms underlying expected utility theory may not apply in an environment where a range of probability distributions is possible. The corollary of this is that it will not be possible to use standard expected utility solutions for these decision problems. Two ways of relaxing the assumptions underlying expected utility theory to allow for Knightian uncertainty are discussed in the next section and the implications for the conduct of monetary policy are discussed in Section 5.

4. Knightian Uncertainty and Expected Utility Theory

In cases where there is a single probability distribution over future events, decisions can be made by choosing the action with the maximum expected utility. Savage (1954) and Anscombe and Aumann (1963) provide an axiomatic foundation for this expected utility representation and therefore justify the use of the maximum expected utility decision rule.

To highlight different approaches to modelling decision-making under Knightian uncertainty, we consider a simple decision problem in which there are two possible states of the world and decision-makers can hold assets that yield utility U_1 in state 1 and U_2 in state 2. Assuming that the probability of state 1 occurring is π, the combinations of U_1 and U_2 that yield the same expected utility will lie on a straight indifference curve with slope $\frac{\pi}{1-\pi}$, as shown in Figure 1. With gambles along the 45° line, such as *C,* the decision-maker is certain of getting a particular level of utility regardless of the state of nature that occurs. These gambles will be referred to as constant acts.

In the case of Knightian uncertainty, the decision-maker applies a range of subjective probability distributions over the possible states of nature. Assuming that the range of probabilities being considered for state 1 is continuous, the lower bound for this probability can be defined as π_L and the upper bound as π_U. The difference between π_U and π_L can be thought of as some measure of the degree of Knightian uncertainty facing the decision-maker. In Figure 2, the points that are indifferent to *C* with respect to a particular probability distribution in this range are represented by the shaded area. Note that the area shaded will become larger the higher the degree of uncertainty.[1]

Figure 1

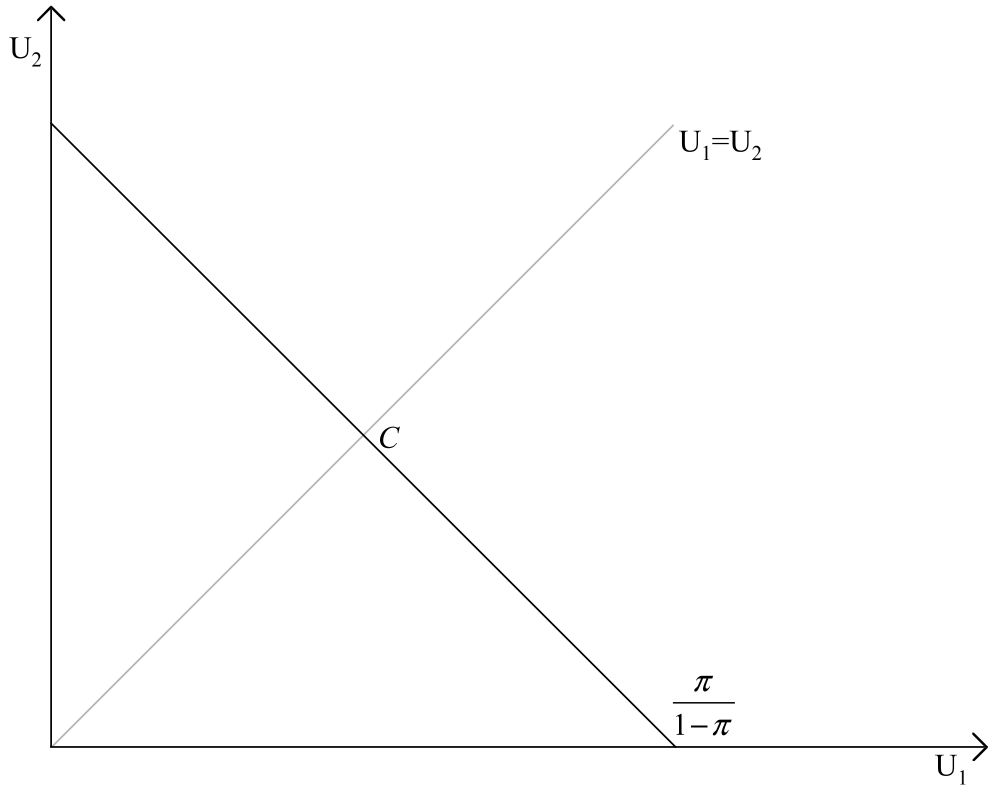

Without a unique probability distribution, decisions can no longer be made by maximising expected utility: actions that maximise utility under some probability distributions in the range will not be optimal under others. With this in mind, there have been two main approaches to modelling decision-making under Knightian uncertainty. We present each of these approaches in turn and discuss their implications for the decision rules that should apply. Their implications for monetary policy-making are discussed in Section 5.

[1] For a detailed discussion of aversion to uncertainty, see Epstein (1999).

Figure 2

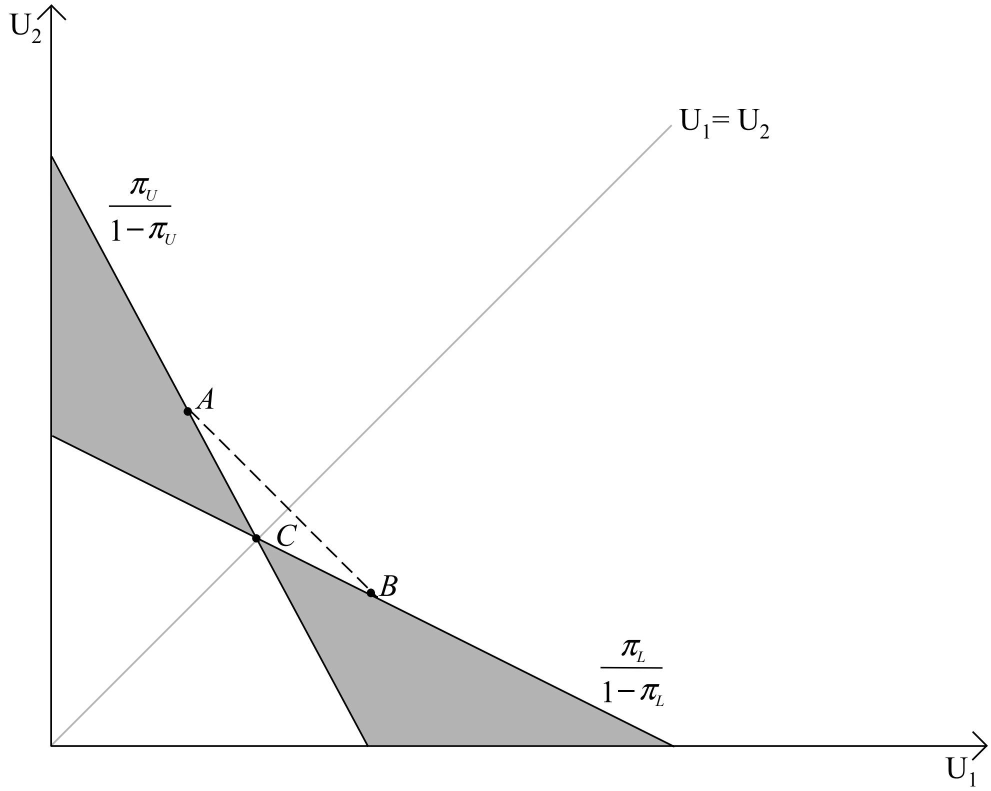

4.1 Relaxing the Independence Assumption

The independence assumption has perhaps been the most controversial axiom of expected utility theory since it was formalised by Savage (1954) as the 'Sure Thing Principle'. The independence axiom states that if *A* and *B* are two gambles, and *A* is preferred to *B*, that is $A \succ B$, a mixture of *A* with a third gamble *D* will be preferred to a mixture of *B* and *D* in the same proportions. That is $\alpha A + (1-\alpha)D \succ \alpha B + (1-\alpha)D$, where α is strictly between 0 and 1. The indifference curves in Figure 1 are linear if and only if the independence axiom holds.

Gilboa and Schmeidler (1989) argue that the independence axiom will not hold in situations where there is no unique probability distribution, but that a weaker version of independence, which they label constant-independence, will hold.

Constant-independence states that if $A \succ B$, then a mixture of *A* with a constant act, *C*, will be preferred to a mixture of *B* with the constant act *C*, that is $\alpha A + (1-\alpha)C \succ \alpha B + (1-\alpha)C$.

The other axiom Gilboa and Schmeidler introduce is uncertainty aversion. Schmeidler (1989) explains that this 'means that "smoothing" or averaging *utility* distributions makes the decision-maker better off'. This property arises from the way in which the independence assumption has been relaxed. In situations where there is a single probability distribution, and independence holds, it will be true that if *A* is indifferent to *B*, that is $A \sim B$, then $\alpha A+(1-\alpha)B \sim \alpha B+(1-\alpha)B=B$. Gilboa and Schmeidler show that if the independence assumption is relaxed, it is possible that for $A \sim B$, that $\alpha A+(1-\alpha)B \succ B$ for α strictly between 0 and 1 (see below).

Using the weaker definition of independence and introducing uncertainty aversion, Gilboa and Schmeidler (1989) prove that this description of preferences provides a theoretical foundation for Wald's (1950) minimax decision rule. This rule formalises the idea that in situations where there is uncertainty about the probability distribution to apply, decision-makers may choose to minimise the impact of the worst case scenario.[2] More formally, Gilboa and Schmeidler specify the preference relationship between two gambles, *A* and *B* as:

$$A \succ B \text{ if and only if } \min_{\pi} E_{\pi}U(A) > \min_{\pi} E_{\pi}U(B).$$

This decision rule allows points that are indifferent to a given constant act, *C*, to be represented by kinked indifference curves such as the one represented in Figure 3. The transition from Figure 2 to Figure 3 can be explained as follows. Above the 45° line, utility of state 2 is greater than the utility of state 1. The minimum expected utility will be obtained by applying the lowest probability for state 2, that is the highest probability of state 1. Hence, the appropriate slope of the indifference curve above the 45° line is $\frac{\pi_U}{1-\pi_U}$. Using the same argument, the slope of the indifference curve satisfying the minimax rule below the 45° line is $\frac{\pi_L}{1-\pi_L}$. The

[2] For monetary policy decision-making, this is similar to minimising the maximum expected loss.

consequence of the minimax decision rule, therefore, is that Gilboa–Schmeidler preferences are represented by kinked indifference curves, where the kink occurs on the 45° line.

Figure 3: Gilboa-Schmeidler Preferences

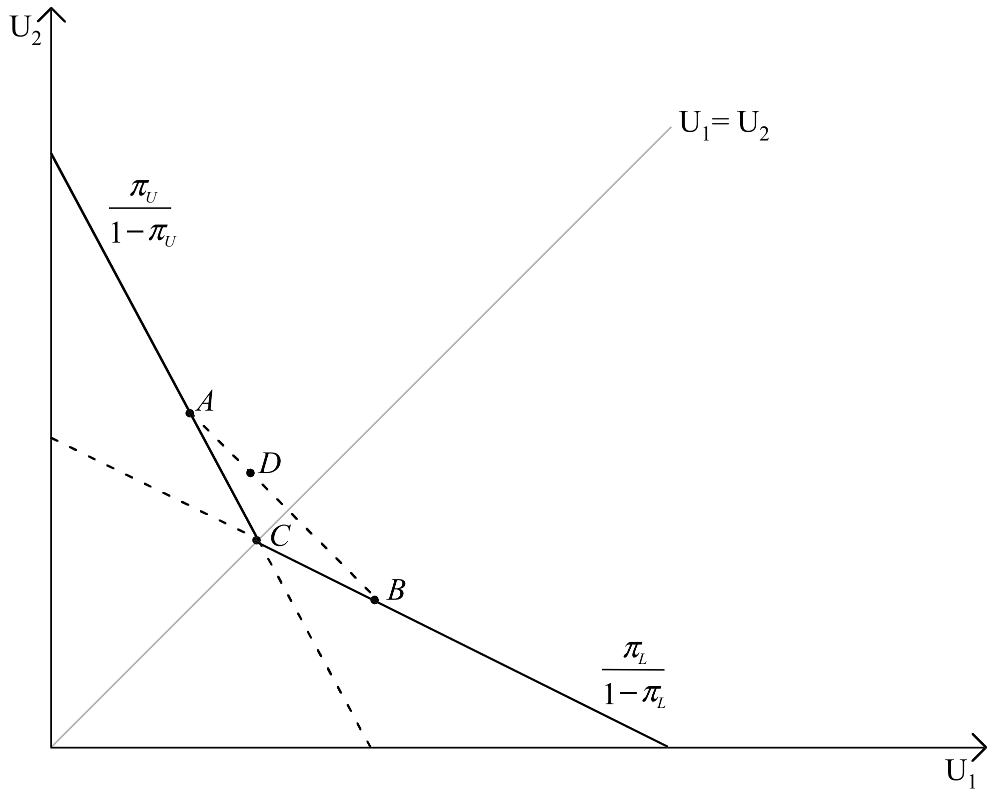

The region above the curve contains points strictly preferred to the points on the indifference curve since each point above the curve gives a higher minimum expected utility. The reverse is true for the area below the curve.

Constant-independence will hold in Figure 3: given that *A~B*, it is true that

$$\alpha A+(1-\alpha)C\sim\alpha B+(1-\alpha)C \tag{3}$$

when *C* is a constant act that lies on the 45° line. To understand how these preferences embody uncertainty aversion, we can see that although the decision-maker is indifferent between *A* and *B,* any linear combination of *A* and *B*, such as *D*, is preferred to both *A* and *B*. That is, *D* will yield a higher minimum expected utility than either A or B, despite the fact that A and B will give the decision-maker the same minimum expected utility.

4.2 Relaxation of Completeness

Another approach to describing preferences and decision-making when there is a range of possible probability distributions is to maintain the independence assumption, but relax the completeness assumption. Completeness implies that the decision-maker faced with two gambles will be able to compare them and state a preference relation between them.

Gilboa-Schmeidler preferences are complete because the decision-maker is willing to state a preference between any given alternatives. If, however, a decision-maker thinks that it is important to consider alternatives under the full range of possible probability distributions, it is not possible to reduce the points indifferent to a given constant act to a linear, albeit kinked, indifference curve. Therefore, some points may not be comparable. The possibility of capturing the effects of Knightian uncertainty by dropping the completeness axiom was introduced by Bewley (1986).[3]

One of the implications of dropping completeness is that it is possible for decision-makers to make a series of decisions that are not consistent with each other. In order to prevent such intransitivity, Bewley introduces the inertia assumption. This assumption states that the decision-maker remains with the *status quo* unless an alternative that is clearly better under each possible probability distribution is available. This can be thought of as a form of caution: if you are uncertain about whether you will be better off with the alternative, why change your current position?

Bewley preferences imply that faced with two alternatives, X and Y, X will be preferred to Y if and only if $E_{\pi}U(X) > E_{\pi}U(Y)$ for all π considered. That is to say only when X yields better expected outcomes compared to Y under every distribution will X be preferred to Y.

---

[3] In fact, Bewley asserts that, in his framework, preferences will be complete if and only if there is only one subjective probability distribution.

More generally, it can be shown that any point in the shaded region of Figure 4 will not dominate the *status quo*, *Z*. For example, a point such as *W* is better than *Z* under some probability distributions but worse than *Z* under others. In this case, the decision-maker does not have any clear preference for either *Z* or *W* but because of the inertia assumption, the decision-maker will remain with *Z*. Point *V* is clearly better than *Z* since it gives a higher expected utility under every probability distribution. Should *V* become available, the decision-maker would choose *V* over *Z*. The diagram also indicates that by mixing two gambles, both incomparable to *Z*, a gamble that is comparable to *Z* can be produced. So, although the decision-maker will prefer neither *W* nor *Y* over *Z*, the decision-maker will prefer some mixtures of *W* and *Y* to *Z*. The willingness of the decision-maker to make a decision will depend on how much uncertainty exists. The less uncertain the environment, that is, the closer are π_L and π_U, the larger is the set of gambles that are directly comparable to the *status quo*, and the more likely the decision-maker will be to give up the *status quo* for an alternative.

Figure 4: Bewley Preferences

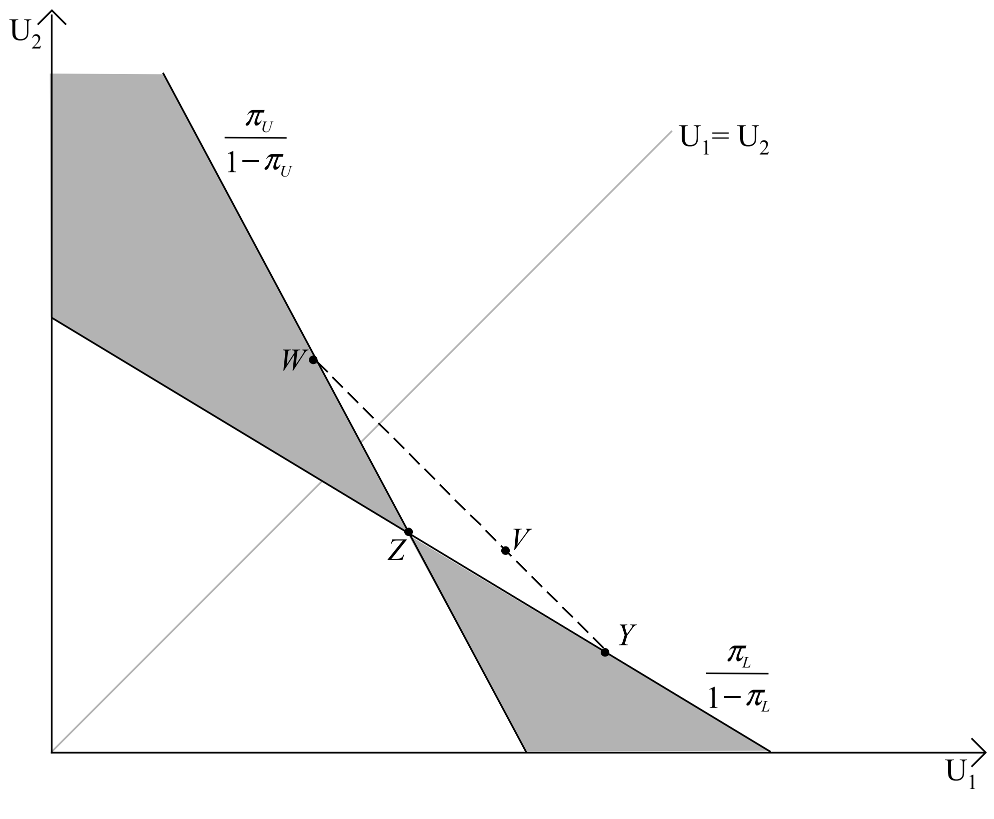

Although the inertia assumption gives the decision-maker direction when the decision-maker is faced with incomparable alternatives, it does not provide guidance in situations where there is more than one alternative that strictly dominates the current position. There are a number of possible decision rules that could be used, ranging from rules that recognise the fact that decisions today affect the timing and magnitude of future decisions, to more simple, atheoretical rules such as picking some 'midpoint'.

5. Implications for Monetary Policy

The two alternative approaches to dealing with Knightian uncertainty presented above have different underlying assumptions about the preferences of policy-makers and the decision rules that they use. In this section we discuss the applicability of these two specifications to the monetary policy decision-making environment.

5.1 Gilboa-Schmeidler Preferences

A decision-maker with Gilboa-Schmeidler preferences is assumed only to care about the probability distribution that delivers the worst case scenario, and chooses the action that minimises the downside of that worst case.

The main application of Gilboa-Schmeidler preferences has been in finance where Knightian uncertainty has been used to explain the behaviour of prices in equilibrium. For example, Epstein and Wang (1994) adapt the static Gilboa-Schmeidler framework to an inter-temporal setting to show that when uncertainty exists in a pure exchange economy, asset prices are indeterminate within some range defined by the degree of uncertainty. Therefore, for a given set of fundamentals, a continuum of prices may exist in equilibrium. The results Epstein and Wang (1994) generate lead them to concur with the view of Keynes (1936) that 'existing market valuation…cannot be uniquely correct, since our existing knowledge does not provide a sufficient basis for a calculated mathematical expectation'.

Recently, there has been a significant number of applications of robust control methods to the monetary policy decision-making process. This methodology assumes that the decision-maker adjusts policy to minimise the losses associated with the worst-case scenario possible within a prespecified range of outcomes. Therefore, these methods allow for Knightian uncertainty about the model being used and utilise the minimax decision rules implied by Gilboa-Schmeidler preferences. Sargent (1999) and Onatski and Stock (2000) find that when the true model is allowed to deviate from the estimated backward-looking model in very general ways, the result is a policy rule that reacts more strongly to deviations of inflation from target and output from potential than when this uncertainty is not considered. Tetlow and von zur Muehlen (2000) find a similar result using a forward-looking model.

Intuitively, this can arise because the worst-case scenario is the model in which policy has the least ability to affect output and inflation within the range of possible models. In this case, the policy-maker will have to react more to generate the same response, and will therefore deliver a more volatile path of interest rates. Consequently, it can be argued that the assumption that the policy-maker seeks to avoid the worst-case scenario underlies the result that policy is more reactive rather than less under general model uncertainty.

Several authors have also applied robust control methods to cases of model uncertainty that are more specific than the model uncertainty discussed above, but require less assumptions than parameter uncertainty, which assumes a single probability distribution, to obtain similar results. For example, Giannoni (2000) uses a forward-looking model to show that allowing for the effects of such general parameter uncertainty also increases the responsiveness of policy. Onatski and Stock (2000) and Tetlow and von zur Muehlen (2000) show that as the degree of structure placed on uncertainty increases, the responsiveness of the policy-maker decreases.

Robust control methods implement the Gilboa-Schmeidler formulation of the monetary policy decision-maker's problem by using the minimax rule. However, there are several reasons why other methods of incorporating Knightian uncertainty should be explored. First and foremost, the intuition for using a minimax rule in the context of monetary policy does not accord with the way in which monetary policy-makers talk about their decisions. Monetary policy-makers usually talk in terms of the balance of risks rather than in terms of avoiding worst-case scenarios. Second, it is difficult to generalise the results of these models, and to the extent that this is possible, they deliver the result that interest rates will be more volatile in the presence of uncertainty than without it. This runs counter to the original motivation for looking at the effects of uncertainty on monetary policy decision-making.

5.2 Bewley Preferences

Unlike Gilboa-Schmeidler preferences, Bewley preferences assume that the policy-maker cares about the outcomes under all the probability distributions they believe are possible, and do not focus solely on the worst-case scenario. The idea that policy-makers facing Knightian uncertainty cannot make sharp comparisons between some alternatives is intuitively appealing, and accords more closely with the balance-of-risks approach to making monetary policy decisions.

Given the expectations about the ranges of probability distributions that are possible over future periods, it is theoretically possible to establish a range of interest rates that are consistent with these expectations. The inertia assumption implies that the policy-maker should leave interest rates unchanged if the current interest rate lies within this range. One justification for the inertia assumption is that it prevents intransitive, apparently irrational, decisions being made. Given the importance of reputation and credibility for the effectiveness of monetary policy, this would strengthen the appeal of the inertia assumption in the context of monetary policy. Lowe and Ellis (1997) summarise several other arguments for why policy-makers may wish to avoid interest rate moves that have to be reversed, and consequently would prefer to leave interest rates unchanged unless there are sufficient reasons to change them. Some of these alternative arguments for avoiding policy reversals are that they increase financial market volatility and make long-term interest rates less predictable, and therefore weaken the transmission of monetary policy.

The consequence of the inertia assumption is that it is possible for the prevailing level of the nominal interest rate to persist for several periods if the policy-maker believes that the level of the output gap and the real interest rate will remain in similar ranges for several periods. Therefore, the Bewley formulation of Knightian uncertainty has the potential to explain one of the stylised facts of observed interest rates.

If the policy-maker with Bewley preferences has beliefs about the economy that suggest the current stance of monetary policy is no longer within the appropriate range, there will be a discrete shift in the level of the interest rate. Bewley (1986, 1987) does not discuss what the decision rule in this sort of situation should be. Ideally, the policy-maker's decision should be modelled recognising that the decision will affect future decisions, and consequently their expected losses. For example, the policy-maker may wish to choose an interest rate that decreases the amount of uncertainty they are likely to face in the subsequent period. The disadvantage of such a rule in a dynamic setting is that it increases the complexity of solving a problem that is already significantly more complex than the standard decision-making problem under uncertainty. Another alternative is to choose a simple atheoretical rule such as picking the midpoint or boundary of the range.

Figure 5 demonstrates the way in which interest rates would be expected to move if the policy-maker is faced with a given range of interest rates at each point of time. The shaded region is an example of the way in which the range of appropriate interest rates might evolve. In Section 6 we discuss some of the issues that arise for deriving such ranges, but take them as given in Figure 5. At each point in time the policy-maker must decide which interest rate to choose for the next period given their current interest rate and the range of appropriate interest rates for the next period. In Figure 5, the policy-maker does not change the interest rate if the prevailing rate is in the optimal range in the following period and, if the prevailing rate is not in the optimal range, changes the interest rate to the midpoint of the range.

Figure 5 exhibits step movements in interest rates, and the way in which the bands of appropriate interest rates are assumed to evolve delivers sequences of interest rate movements in the same direction. Another feature of this framework is that interest rates are less likely to change as the range of possible optimal interest rates increases, which will occur if uncertainty increases.

Figure 5

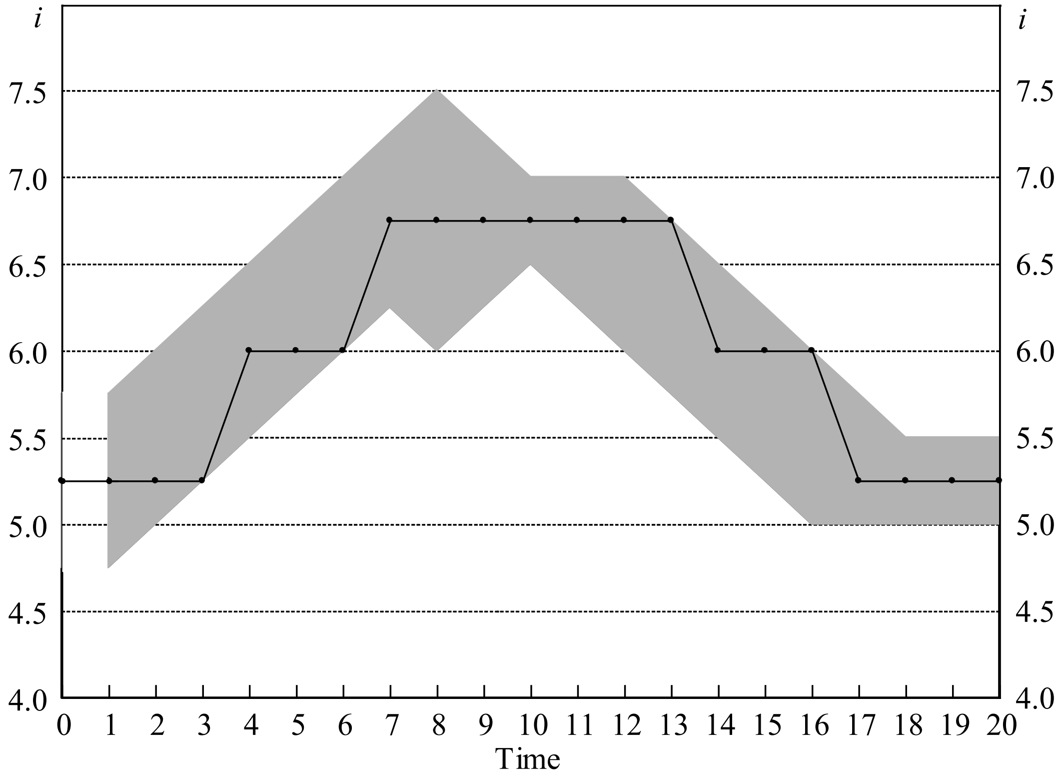

6. The Optimal Control Problem with Bewley Preferences

The difficulty in moving forward is to find a method of incorporating Bewley preferences into an optimal control framework. At a very general level, we want to minimise the policy-maker's loss function subject to the linear constraints that define the dynamics of the economy, the range of expectations the policy-maker has about the distributions of future shocks and the decision rules that are being used. By taking into account all these features, we should obtain internally consistent future paths of interest rates; that is the observed path of interest rates should not be different to the predicted path if no unexpected events occur.

When there is a single probability distribution over future shocks, it is relatively straightforward to obtain the unique, internally consistent path. However, when there is more than one probability distribution, it is much more difficult to obtain this internal consistency, and the uniqueness of the path is not guaranteed. The difficulty can be illustrated through the use of a two-period problem where the

Bewley policy-maker must select nominal interest rates for the next two periods, i_1 and i_2, given their range of subjective expectations.

The policy-maker could begin by choosing an arbitrary value for $i_2 = \tilde{i}_2$. Given the current state of the economy, the range of subjective expectations and the choice of $\tilde{i}_2$, optimal control techniques can be applied with each subjective probability distribution in turn, to generate a range for i_1. By applying the decision rule at time t=1, we can determine the choice of i_1, which we denote $\hat{i}_1$. Taking this choice of $\hat{i}_1$ and the current state of the economy and expectations, we can calculate the range of optimal values for i_2. By applying the decision rule at time t=2, we can determine the choice of i_2 given $\hat{i}_1$, denoted $\hat{i}_2$. If $\tilde{i}_2 = \hat{i}_2$, then a solution has been found. If not, we can iterate on the procedure until a solution is found.

This sketch of the Bewley policy-maker's problem in a two-period setting would be made more concrete by resolving existence and uniqueness issues and proving that if a solution exists, it is internally consistent. It is also clear that it is not straightforward to extend this solution method to the multi-period case as the sequential nature of decision-making is important. The conditions for existence, uniqueness and internal consistency are straightforward in the multi-period case when there is a single subjective probability distribution.

As a first pass at this problem, we assume that the set of probability distributions being considered is convex. In order to establish the range of interest rates that could be optimal, in the next period, we solve the standard optimal policy problem for the probability distributions that form the boundaries of this range. Given convexity, we know that the two period-one interest rates generated from these problems form the boundaries of the feasible range of period-one interest rates, and subsequently the decision rules can be applied. This method ignores the effect of today's decision on the available actions in future periods and implicitly assumes that interest rates after period one will be determined optimally from the probability distribution associated with the chosen period-one interest rate. As such, this algorithm will not be optimal, but might approximate the optimal solution.

In the general model presented in Section 2, the monetary policy decision-maker chooses interest rates to minimise the loss function in Equation (1). In a finite N period problem, the future path of interest rates can be written as the solution to *M=N-(maximum interest rate lag in the output equation)* first order conditions:

$$\begin{array}{ll} \dfrac{\partial Loss}{\partial i_{t+1}} = 0 & \qquad (t+1) \\ \vdots & \qquad \vdots \\ \dfrac{\partial Loss}{\partial i_{t+M}} = 0 & \qquad (t+M) \end{array}$$

In the case where there is a single probability distribution, this will yield a unique path of interest rates. In the case where there is a set of probability distributions, it will not. However, it is possible to derive a value for i_{t+1} from a given probability distribution from this range assuming, counterfactually, that future interest rates will be determined as optimal outcomes from this probability distribution.

If the set of probability distributions is convex, the solution to this continuum of problems will yield a range of possible interest rates for i_{t+1}. In fact, given convexity, it is possible to find the range by solving for the upper and lower bounds. Once a range for i_{t+1} has been determined, the relevant decision rules can be applied to choose the specific interest rate from the range.

6.1 A Simulation of Monetary Policy in a Simple Model

The model used for the simulations is based on the small closed-economy model used by Ball (1997) and Svensson (1997). The model consists of an equation relating the output gap to its own past behaviour and past real interest rates, and a second equation relating inflation to its history and the past behaviour of the output gap, both estimated over the sample 1985:Q1–2000:Q2.

The output gap used in the model is estimated by fitting a linear trend through real non-farm output from 1980:Q1–2000:Q2. By construction, this output gap will be zero, on average, over the sample period. However, since inflation fell significantly over the sample period, we know that the output gap has been negative on average. Using the methodology outlined in Beechey *et al* (2000), we reduce the level of the output gap by 3 per cent over the entire sample period. The real interest rate in the output gap equation is calculated as the nominal interest rate, less year-ended inflation. A neutral real cash rate of 3.5 per cent in the model is accepted by the data and imposed. The estimated equation used for the simulations suggests that an increase in the *ex post* real interest rate of 1 per cent reduces the level of the output gap by 0.12 per cent one year later. The estimation results are presented below (the output gap is measured in per cent), with standard errors in parentheses.

$$y_t = \underset{(0.025)}{0.924}\, y_{t-1} - \underset{(0.031)}{0.124}\left(i_{t-4} - \pi_{t-4} - 3.5\right) + \varepsilon_t \qquad (4)$$

S.E Regression = 0.65	Adjusted R-squared = 0.92
DW = 1.88	Jarque-Bera = 0.40 (P = 0.82)

The inflation equation is estimated using four-quarter-ended inflation as measured by the Treasury underlying CPI.[4] The most significant difference between the equation estimated for these simulations and the Ball-Svensson model is that it includes changes in, rather than the level of, the output gap. These 'speed' terms generate a better fit to the data. The hypothesis that the sum of the coefficients on lagged inflation terms sum to one is accepted and imposed on the equation. The lag structure of the 'speed' terms suggests that it takes between a year and eighteen months for changes in the output gap to affect inflation.

$$\begin{aligned}\pi_t = {} & \underset{(0.116)}{1.603}\pi_{t-1} - \underset{(0.226)}{0.611}\pi_{t-2} - \underset{(0.223)}{0.246}\pi_{t-3} + \underset{(0.113)}{0.254}\pi_{t-4} \\ & + \underset{(0.054)}{0.128}\Delta y_{t-3} + \underset{(0.056)}{0.010}\Delta y_{t-4} - \underset{(0.054)}{0.023}\Delta y_{t-5} + \underset{(0.054)}{0.202}\Delta y_{t-6} + \eta_t\end{aligned} \qquad (5)$$

S.E Regression = 0.3	Adjusted R-squared = 0.99
DW = 1.79	Jarque-Bera = 3.60 (P = 0.17)

4 From 1993 to the end of 1998, the inflation target was expressed in terms of the Treasury underlying CPI series, which we have constructed after 1999.

A number of simulations are presented in this paper. In each simulation, a weight of one has been placed on the inflation objective, measured as squared deviations of year-ended inflation from target, and one-half on the square of the output gap. We have assumed that the target for inflation is 2.5 per cent per annum. As there was no inflation target in place in Australia until 1993, this is when we begin our simulations. A discount factor of one was used in the objective function and in each period, the policy-maker calculates an optimal interest rate path 160 quarters into the future.

The empirical residuals obtained from estimating each of the equations are used as stochastic shocks to the model in chronological order so that the simulated interest rate paths can be compared to the path of the cash rate we actually observe in history. In the simulations, we assume that the policy-maker chooses the nominal interest rate at the beginning of the period, and observes the output gap and inflation at the end of the period. At this time, the policy-maker 're-optimises' to select the nominal cash rate for the following period, assuming the estimated equations are structural. The simulated interest rate paths we obtain are not conditional on the history of actual cash rates observed. In this sense, each of our simulations is dynamic.

As discussed in Section 2, Orphanides (1998) has demonstrated the importance uncertainty about output gap levels and inflation, and consequently the neutral real cash rate, has had on monetary policy decisions in the US. This type of Knightian uncertainty can be translated from uncertainty about variables in the model, into uncertainty about the means of the error distributions and we assume that the policy-maker considers a bounded range of means.

The first simulation is a standard optimal policy exercise that assumes the neutral real interest rate and the state of the economy are known precisely. The second simulation assumes that the policy-maker recognises that their knowledge regarding the neutral real cash rate and the state of the economy is imprecise.

Although we impose the restriction that the neutral real cash rate is equal to 3.5 per cent, the uncertainty that we allow for in the output equation is equivalent to the policy-maker believing the neutral real cash rate could be as low as 3 per cent or as high as 4 per cent. For the inflation equation, we have allowed the means of the shocks to be within one standard deviation either side of zero to incorporate the uncertainties regarding the measurement of inflation and the output gap. The ranges we have selected are somewhat arbitrary and, of course, the ranges could be made wider, smaller, or asymmetric. We also assume for this simulation that the policy-maker chooses the midpoint of the optimal range of interest rates when the inertia assumption cannot be applied.

The results of these two simulations are presented in Figure 6. Regardless of whether Knightian uncertainty is included or not, the cash rate paths we obtain from our simulations are quite volatile relative to the path of actual cash rates. Nevertheless, when Knightian uncertainty is allowed for, we obtain extended periods of time where the cash rate remains constant although there are some immediate policy reversals.[5]

Figure 6: Paths of Interest Rates

Optimal policy with and without Knightian uncertainty

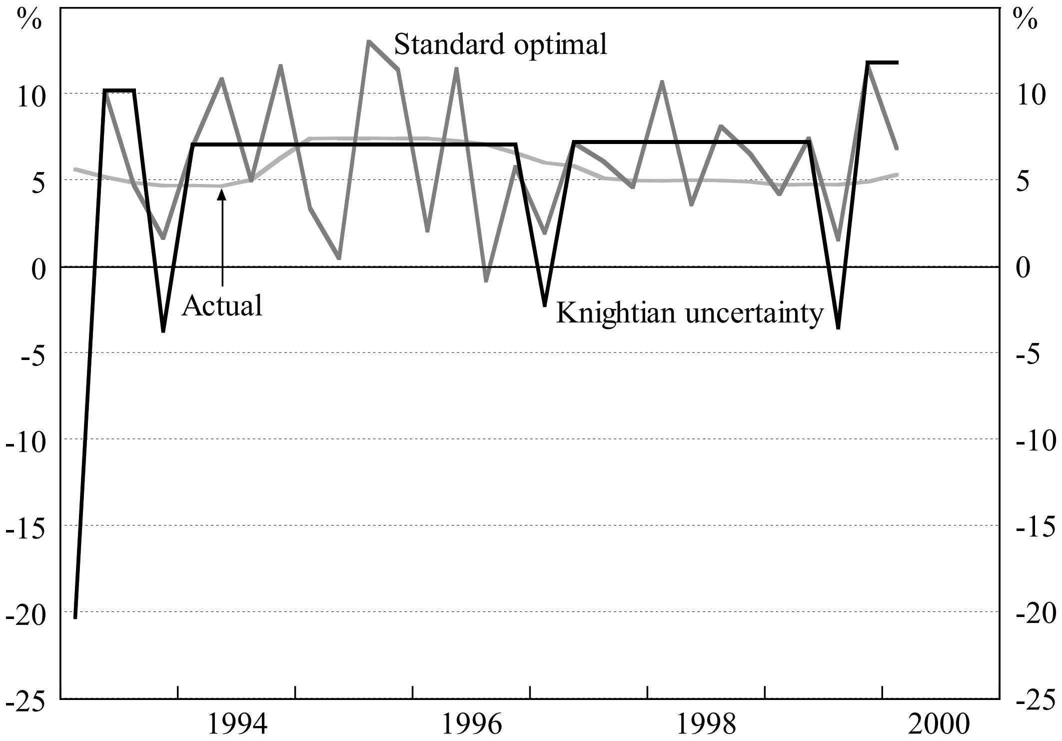

As a point of comparison, we also show, in Figure 7, what the simulated interest rate path would be if we allowed for Knightian uncertainty, but changed interest

5 Note also that there are no constraints preventing nominal interest rates from being negative.

rates to the nearest boundary of the optimal range, rather than the midpoint, when the inertia assumption does not apply. This rule, that minimises movements in the interest rates, can be argued to be more consistent with the intuitive justification for the inertia assumption in the monetary policy context. This rule also generates an interest rate path with extended periods of no change and no immediate reversals.

Figure 7: Paths of Interest Rates

Comparing actual and the closest-boundary rule

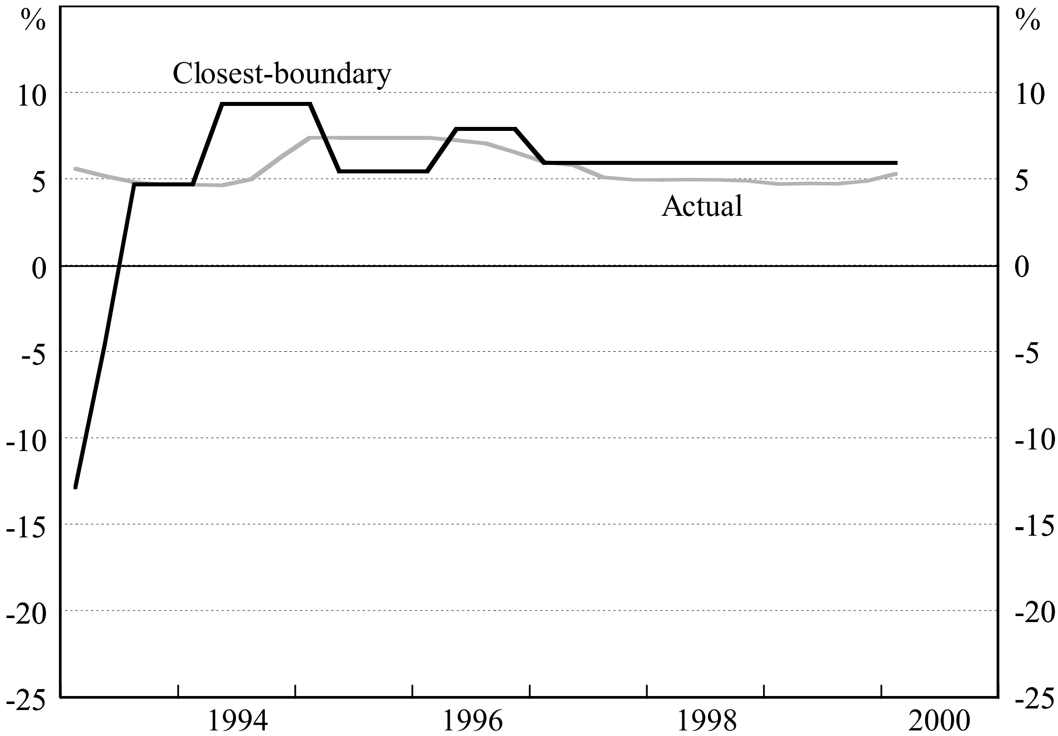

Figure 8 demonstrates the effect of 'halving' the degree of uncertainty – that is, halving the range of possible means for the error distribution. In general, less uncertainty implies more changes in interest rates. In the case of the midpoint rule, this means shorter periods of inaction with more frequent, but smaller, policy reversals. In the case of the closest-boundary rule, there are also fewer periods of no change when uncertainty decreases and the length of time between reversals shortens. In practical situations, the degree of uncertainty is likely to change over time. To the extent that points of inflection in the business cycle are periods of relatively high uncertainty, these results can help to explain why interest rates are more stable around turning points.

Figure 8: Paths of Interest Rates

'Half' uncertainty

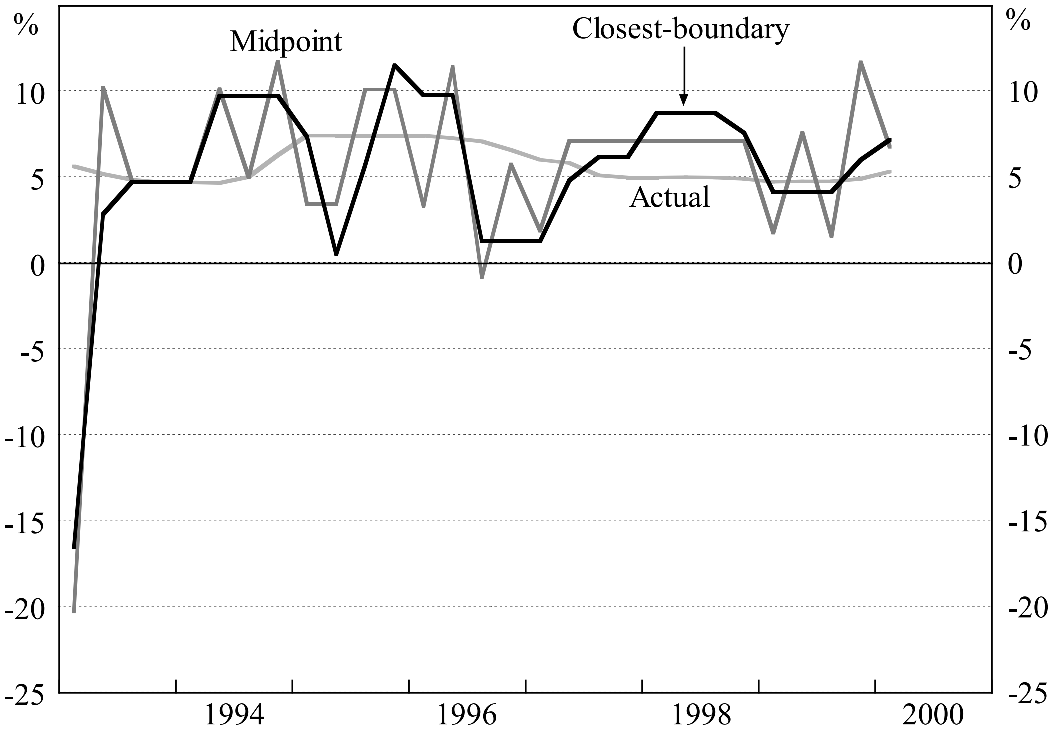

7. Conclusions and Further Research

This paper has explored the extent to which Knightian uncertainty can help reconcile differences between interest rate paths generated by models of optimal policy and those observed in practice. We began by providing a framework for discussing different types of uncertainty and the decision rules that they imply.

Recent analyses of model uncertainty have explicitly used minimax decision rules, which directly result from one particular formalisation of Knightian uncertainty. We advocate an alternative formalisation of Knightian uncertainty due to Bewley (1986) that assumes the policy-maker wishes to compare alternatives under all possible probability distributions in order to make their decisions. This formalisation provides monetary policy decision-makers with a range of possibly optimal interest rates which, when combined with intuitively appealing additional decision rules that are consistent with this decision-making environment, can generate paths of interest rates with extended periods of no change. Simulations that include examples of these simple decision rules are used to demonstrate their effects.

Although the results of our simulations are promising, there are several issues that warrant further research. Perhaps most importantly, the algorithm we have proposed is not fully optimal. In particular, we have only investigated two simple sets of decision rules. Another direction of further research is to explore more sophisticated models that contain more realistic lag structures and, perhaps, incorporate a richer set of explanatory variables. We believe extensions along these lines will be fruitful.

References

Anscombe F and R Aumann (1963), 'A Definition of Subjective Probability', *Annals of Mathematical Statistics*, 34, pp 199–205.

Ball L (1997), 'Efficient Rules for Monetary Policy', NBER Working Paper 5952.

Beechey M, N Bharucha, A Cagliarini, D Gruen and C Thompson (2000), 'A Small Model of the Australian Macroeconomy', RBA Research Discussion Paper No 2000–05.

Bewley TF (1986), 'Knightian Decision Theory: Part I', Cowles Foundation Discussion Paper No 807.

Bewley TF (1987), 'Knightian Decision Theory: Part II: Intertemporal Problems', Cowles Foundation Discussion Paper No 835.

Blinder AS (1998), *Central Banking in Theory and Practice*, MIT Press, Cambridge.

Brainard W (1967), 'Uncertainty and the Effectiveness of Policy', *American Economic Review*, 57, pp 411–425.

Debelle G and A Cagliarini (2000), 'The Effect of Uncertainty on Monetary Policy: How Good are the Brakes?', Reserve Bank of Australia Research Discussion Paper 2000–07.

Eijffinger S, E Schaling and W Verhagen (1999), 'A Theory of Interest Rate Stepping: Inflation-Targeting in a Dynamic Menu Cost Model', Centre for Economic Policy Research Discussion Paper No 2168.

Ellsberg D (1961), 'Risk, Ambiguity, and the Savage Axioms', *Quarterly Journal of Economics*, 75, pp 643–669.

Epstein LG (1999), 'A Definition of Uncertainty Aversion', *The Review of Economic Studies*, 66(3), pp 579–608.

Epstein LG and T Wang (1994), 'Intertemporal Asset Pricing Under Knightian Uncertainty', *Econometrica*, 62(2), pp 283–322.

Estrella A and FS Mishkin (1999), 'Rethinking the Role of NAIRU in Monetary Policy: Implications of Model Formulation and Uncertainty' in JB Taylor (ed), *Monetary Policy Rules*, University of Chicago Press, Chicago, pp 405–436.

Giannoni MP (2000), 'Does Model Uncertainty Justify Caution? Robust Optimal Monetary Policy in a Forward-Looking Model', Princeton University, November, mimeo.

Gilboa I and D Schmeidler (1989), 'Maxmin Expected Utility with Non-Unique Prior', *Journal of Mathematical Economics*, 18, pp 141–153.

Guthrie G and J Wright (2000), 'The Optimal Design of Interest Rate Target Changes', paper presented at Workshop held at the Reserve Bank of New Zealand, 'Monetary Policy Decision-Making – Proactive or Reactive?', Wellington, 8 June 2000.

Keynes JM (1936), *The General Theory of Employment, Interest and Money*, Macmillan, London.

Knight FH (1933), *Risk, Uncertainty and Profit*, Houghton Mifflin Co, Boston.

Kreps DM (1990), *A Course in Microeconomic Theory*, Harvester Wheatsheaf, London.

LeRoy SF and LD Singell (1987), 'Knight on Risk and Uncertainty', *Journal of Political Economy*, 95(2), pp 394–406.

Lowe P and L Ellis (1997), 'The Smoothing of Official Interest Rates' in P Lowe (ed), *Monetary Policy and Inflation Targeting*, Proceedings of a Conference, Reserve Bank of Australia, Sydney, pp 286–312.

Onatski A and JH Stock (2000), 'Robust Monetary Policy Under Model Uncertainty in a Small Model of the U.S. Economy', NBER Discussion Paper No 7490.

Orphanides A (1998), 'Monetary Policy Evaluation With Noisy Information', Federal Reserve Board Finance and Economics Discussion Series 1998/50.

Orphanides A, RD Porter, D Reifschneider, R Tetlow and F Finan (1999), 'Errors in the Measurement of the Output Gap and the Design of Monetary Policy', Federal Reserve Board Finance and Economics Discussion Series 1999/45.

Rudebusch GD (1999), 'Is the Fed Too Timid? Monetary Policy in an Uncertain World', Federal Reserve Bank of San Francisco Working Paper in Applied Economics Theory 99/05.

Sack B (2000), 'Does the Fed Act Gradually? A VAR Analysis', *Journal of Monetary Economics*, 46(1), pp 229–256.

Sack B and V Wieland (1999), 'Interest-Rate Smoothing and Optimal Monetary Policy: A Review of Recent Empirical Evidence', Federal Reserve Board Finance and Economics Discussion Series 1999/39.

Sargent TJ (1999), 'Discussion of "Policy Rules for Open Economies" by Laurence Ball' in JB Taylor (ed), *Monetary Policy Rules*, University of Chicago Press, Chicago.

Savage L (1954), *The Foundations of Statistics*, John Wiley and Sons, New York.

Schmeidler D (1989), 'Subjective Probability and Expected Utility Without Additivity', *Econometrica*, 57(3), pp 571–587.

Shuetrim G and C Thompson (1999), 'The Implications of Uncertainty for Monetary Policy' in Hunt B and A Orr (eds), *Monetary Policy under Uncertainty: Workshop held at the Reserve Bank of New Zealand 29-30 June 1998*, Reserve Bank of New Zealand, Wellington, pp 259–292.

Smets F (1998), 'Output Gap Uncertainty: Does it Matter for the Taylor Rule?' in Hunt B and A Orr (eds), *Monetary Policy under Uncertainty: Workshop held at the Reserve Bank of New Zealand 29-30 June 1998*, Reserve Bank of New Zealand, Wellington, pp 10–29.

Söderström U (1999), 'Monetary Policy with Uncertain Parameters', Sveriges Riksbank Working Paper No 83.

Svensson LEO (1997), 'Inflation Targeting: Some Extensions', NBER Working Paper No 5962.

Tetlow RJ and P von zur Muehlen (2000), 'Robust Monetary Policy with Misspecified Models: Does Model Uncertainty Always Call for Attenuated Policy?', Federal Reserve Board Finance and Economics Discussion Series No 2000/28.

Wald A (1950), *Statistical Decision Functions*, John Wiley and Sons, New York.