# Committees versus individuals: an experimental analysis of monetary policy decision-making

**By Clare Lombardelli and James Talbot of the Bank's Monetary Assessment and Strategy Division and James Proudman of the Bank's Conjunctural Assessment and Projections Division.**(1)

*This article reports the results of an experimental analysis of monetary policy decision-making under uncertainty. The experiment used a large sample of economically literate undergraduate and postgraduate students from the London School of Economics to play a simple monetary policy game, both as individuals and in committees of five players. The result—that groups made better decisions than individuals—accords with a previous study in the United States with Princeton University students. The experiment also attempted to establish why group decision-making is superior: although some of the improvement was related to committees using majority voting when making decisions, there was a significant additional committee benefit associated with members being able to observe each other's voting behaviour.*

## Introduction

On 6 May 1997, the Monetary Policy Committee (MPC) of the Bank of England was established and granted operational independence in setting short-term interest rates to achieve the government's inflation target of 2.5%. This new framework replaced the previous system of a single individual—the Chancellor of the Exchequer—deciding on the appropriate level of UK base rates.

Why delegate monetary policy to a committee? The academic argument for central bank independence is well established (see, for example, Barro and Gordon (1983)). And in practice, there is strong evidence from across the world to suggest that committees are the preferred arrangement for setting monetary policy by central banks. For instance, a wide-ranging survey undertaken by Fry, Julius, Mahadeva, Roger and Sterne (2000) finds that 79 central banks out of a sample of 88 use some form of committee structure when setting monetary policy. By weight of numbers, it appears to be accepted that setting interest rates by committee is superior. And the intuitive argument that committees make better decisions than individuals—because they allow decision-makers to pool judgment—also seems plausible.

The theoretical economics literature has less to say about the consequences of delegating responsibility to a committee: the hypothesis that groups make better monetary policy decisions is difficult to test, due to a lack of comparable empirical data. This problem motivated Blinder and Morgan (2000) to adopt a different approach: carrying out a 'laboratory experiment' on a large sample of Princeton University students to test whether groups do indeed make monetary policy decisions differently.

In an experiment, the researcher can isolate the relative performance of individual and group behaviour, controlling for differences in the abilities, incentives and preferences of the decision-makers, and of the environment in which they work. The main drawback is that it is artificial—it is not possible to replicate exactly the complexities of real-world policy-making in the context of a simple experiment. But the results may still be informative when thinking about the arrangements for monetary policy making.

Although experimental techniques are relatively new to monetary economics, they are well established in other branches of economics such as asset pricing, game theory and decision-making under uncertainty.(2) In addition, psychologists have studied group behaviour for

(1) The authors would like to thank the London School of Economics for its help and assistance in this project, in particular Richard Jackman, Paul Jackson and Gill Wedlake, but also all the LSE students who took part. A longer version of this paper is published on the same day as this *Bulletin* in the Bank of England *Working Paper series, no. 165*.

(2) See Davis and Holt (1993) and Kagel and Roth (1995) for excellent surveys of this literature.

many years, and a series of experiments—for example, Hall (1971), Janis (1972) and Myers (1982)—have shown that group decisions are rarely equal to the sum of their parts. Group performance depends on the nature of the interaction between members and the task in hand, but the consensus view seems to be that for complex tasks, decisions taken by committee should be at least as good as the average of the individuals that comprise it.

This hypothesis was supported by the results of Blinder and Morgan (2000). In their experiment, groups made substantially better decisions on average than individuals. And, just as in real life, there were also disagreements between committee members over interest rate decisions. But, contrary to their expectations, groups did not make decisions more slowly than individuals.

Examining whether groups make better decisions than individuals is the main focus of this article. It describes a new experiment with students from the London School of Economics, which explored in more detail why groups are superior. One explanation is that majority voting helps to eliminate the poor decisions of a minority of members. But this experiment provided evidence that committees do more than just this, allowing members to pool information and—through communicating with each other—learn more about the game they are playing. And it explicitly tested whether the ability to exchange information through discussion improved performance.

Such a finding would not be surprising if players came to the experiment with different views about the nature of the (unknown) model of the economy. So the experiment tried to examine such differences of opinion by means of a questionnaire designed to help establish these prior beliefs. Asking participants to fill in the questionnaire again at the end indicated how much they learned about the underlying model during the experiment.

The rest of this article is organised as follows. The first section describes the economic model used and the structure of the experiment; the second section discusses the results; finally, the article concludes by trying to draw some inferences from our work for the design of monetary frameworks in the real world.

## Experiment outline

### (i) The model

Participants were asked to act as monetary policy makers by attempting to 'control' a simple macroeconomic model that was subject to randomly generated shocks in each period, as well as a structural shock that occurred at some point during the game. The model used in the experiment (see Appendix 1 for further details) has two equations—a Phillips curve and an IS curve—and is of a type that is widely used for policy analysis in modern macroeconomics (see, for example, Fuhrer and Moore (1995)). Although the model is stylised, where possible it was calibrated with a view to matching UK macroeconomic data (see Bank of England (1999, 2000) for more details of the calibration of such models).

Players were asked to choose the path for the short-term interest rate after observing the response of the endogenous variables—output and inflation—in the previous period. The model has an 'optimal policy rule' (see Appendix 1) that provides a useful benchmark against which to compare individual and group decisions.

### (ii) Modelling prior beliefs

An intriguing feature of Blinder and Morgan's (2000) results was that committee members frequently disagreed about their decisions, despite having identical incentives and information. But even without observing such differences in voting—whether experimentally, or in real life—it seems entirely plausible that committee members can think differently about how to respond to the same economic news.

This should be especially true of a committee where members have diverse backgrounds and beliefs. At the beginning of the experiment, players filled in a questionnaire that attempted to reveal these prior beliefs.(1) The questionnaire was designed so that answers could be directly compared with the parameters of the model and the coefficients of the optimal rule.

During the experiment, players should learn about the structure of the economy—just like real-world policy-makers—by observing the response of inflation and output to changes in interest rates, updating their prior beliefs, and changing their perception of the (unknown) actual model accordingly. Participants

(1) See Appendix 2 for a copy of the questionnaire.

revealed some of what they learned by completing the same questionnaire again at the end of the experiment.

### (iii) Information flows and incentives for players

Players received a clear mandate at the beginning of the experiment: their objective was to maximise a 'score' function that penalised deviations of output and inflation from their 'target' values. The participants knew that at the end of the game they would be paid in pounds according to the following formula:

Payoff = 10 + Average score/10

where the score was averaged over the 16 rounds of the game. The maximum payoff was £20; and the minimum, £10. In practice, most students earned around £15–£16.

As in real life, the participants did not know with certainty the exact structure of the economy they were attempting to analyse. The only information given to participants about the model was that it was linear and broadly characterised the structure of the UK economy. They were also told that the economy was subject to random shocks in each period, and that a structural change occurred at some point during each game. The challenge for players was to extract the signal from the noise and change their behaviour accordingly in order to maximise their score.

### (iv) Outline of the experiment

The participants played the game under a number of different decision-making structures. The sequencing of the experiment is summarised in Table A below. But first it is perhaps helpful to define some terminology.

A period refers to a unit of time corresponding to one interest rate decision and an observation of output and inflation. The players were also given a score in each period. Each round consisted of ten periods. At the end of each round, individuals were given a final score (corresponding to the average score over its ten constituent periods), the game was reset to its initial state and the next round (ten-period game) would begin. There were four rounds in a stage. Stage 1 corresponded to four individual rounds (numbered 1–4). In Stages 2 and 3 (rounds 5–8 and rounds 9–12 respectively) individuals set interest rates together in committees of five players. Some committees were allowed to discuss their decisions in Stage 2 while others were not. Those arrangements were reversed in Stage 3. Stage 4 (rounds 13–16) consisted of a further four individual rounds, with participants playing separate games.

**Table A**
**The structure of the monetary policy experiment**

| | | |
|---|---|---|
| Read instructions sheet | | |
| Fill in 'Priors Questionnaire' | | |
| Practice rounds | No score recorded | |
| Stage 1: Rounds 1–4 | Played as individuals | |
| Stage 2: Rounds 5–8 | Played as a group | (i) No discussion<br>(ii) With discussion |
| Stage 3: Rounds 9–12 | Played as a group | (i) With discussion<br>(ii) No discussion |
| Stage 4: Rounds 13–16 | Played as individuals | |
| Fill in 'Priors Questionnaire' | | |
| Students are paid according to their average score across the four stages | | |

Participants were allocated into groups of five. They were given a set of instructions and asked to fill in the questionnaire. Players had about ten minutes to practise on their own with the actual version of the game used in the experiment before starting to play 'for real'. In each period participants had to decide what interest rate to set in response to developments in the 'economy'.

In the first stage, the participants acted as individual policy-makers, playing separate games on separate computers for four rounds. Beginning with round 1, the game started at period 1, where participants decided on the appropriate level for the interest rate after observing the initial values of inflation and output with a one-period lag.[1] This vote was entered into the computer and the game proceeded to period 2. The computer displayed output and inflation outturns for period 1, along with the score for that round and the interest rate decision. The same process was repeated until the game reached period 10. At this point, players were told their average score for round 1, the game was reset, and play continued, for a further three rounds.

The committee phase was played in two stages (Stages 2 and 3 in Table A above). Stage 2 began at round 5. The five players observed the same information in each period—the level of output and inflation of the previous period(s) as well as the history of interest rates and scores—and entered their votes while sitting at separate computers. But this time, in each period, the computer

(1) Inflation and output would always be close to their target values initially.

selected, and then set, the median vote of the group (as a proxy for a majority-voting rule). Participants observed this committee decision, as well as the response of output and inflation to it. They also saw the (unattributed) votes of their fellow committee members and overall score for the period and the round so far. Again each round lasted for ten periods. Stage 2 finished in round 8.

The committees were divided into two sets. For one set, discussion among members of the committee was not allowed in stage 2. For the other set, discussion was permitted. In stage 3 the organisation of stage 2 was reversed. The ordering of the discussion and no discussion games was organised in this way to control for learning.

Stage 4 (rounds 13–16) served as another control mechanism, to ensure that the comparison between individual and committee play was not affected by the fact that participants had had four (or more) individual rounds to learn before entering the committee stage. By returning to individual play at the end of the experiment, it was possible to verify that the improvement in scores during the committee stages (rounds 5–12) was not just an extension of the learning trend observed in rounds 1–4.

### (v) The data

The experiment was conducted on ten evenings between 12 November and 11 December 2001 at the London School of Economics. Participation in the experiment was voluntary with 170 students taking part in 34 independent experiments: that is to say 34 committees with 16 score observations for each.(1)

Chart 1 shows a breakdown of the participants by course studied: half of the students were postgraduate economists. And although a small minority (5%) was not currently studying an economics-related discipline, all students had taken at least one undergraduate-level economics course.

## Results

The main focus of the experiment was to provide evidence on the differences between group and individual policy-making and to offer some insight into explaining these differences. But indirectly, the results also allow some analysis of what players learned about both the model and how to play the game over time.

**Chart 1**
**A breakdown of players by course studied**

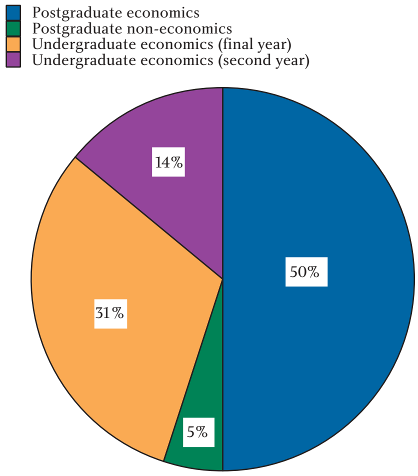


### (i) Learning about the model

Players' answers to the initial questionnaire gave some insight into their prior beliefs about the structure of the economy. All answers to the questionnaire were in numeric form, and each question was related to either the parameters of the model, or the associated optimal rule (see Appendix 1 for details).

Participants also filled in the same questionnaire again at the end of the experiment. One test of learning is therefore the extent of convergence in these views towards the actual parameter values over the course of the game.(2) To this end, a useful statistic is the mean squared error (MSE). The MSE is calculated as the average of the squared errors made by each player when responding to each question.

Over all players and questions, the total MSE statistic decreased; from 0.17 in the initial questionnaire to 0.15 at the end of the experiment. This fall is significant at the 1% level—suggesting that players' responses were closer to the actual parameters of the model at the end of the experiment. The standard deviation of responses to the questionnaire also narrowed significantly (at the 1% level) from 1.59 to 1.45, suggesting some convergence of views among players.

Can we decompose this improvement further? Chart 2 shows the change in MSE for individual questions: the

(1) A further 15 students participated in an alternative version of the experiment described below.
(2) See Lombardelli, Proudman and Talbot (2002) for a discussion of how the responses to the questionnaire can be compared with the parameters of the model and the optimal rule.

### Chart 2
### Reduction in mean squared error of responses between the initial and the final questionnaire(a)

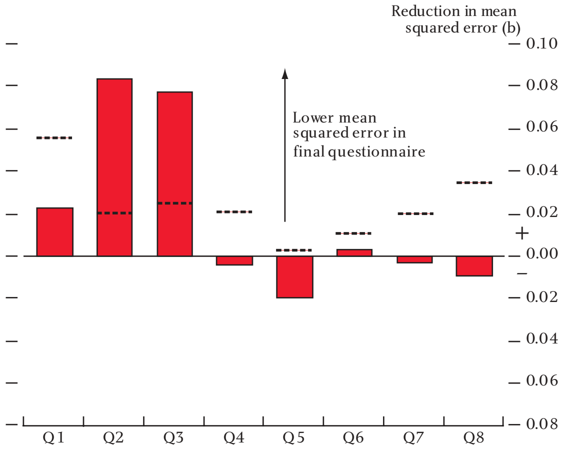


(a) See Appendix 2 for details of the questionnaire.
(b) Dashed lines indicate significance at the 5% level.

### Chart 3
### Average committee scores over time

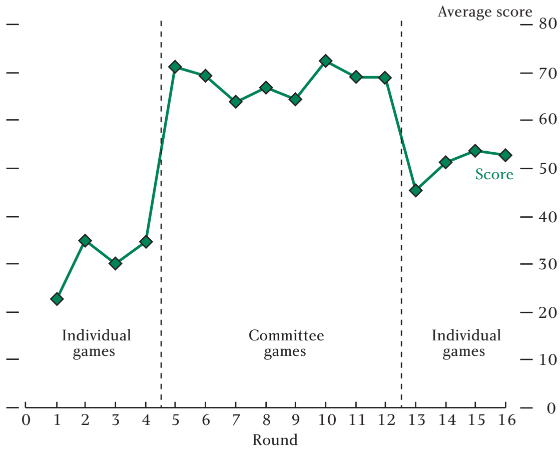


dashed lines represent the reduction in error required for a significant improvement (at the 5% level) in response to each of the questions. This implies that participants learned most about the lags in the transmission mechanism of monetary policy (Q2) and the weight they should attach to deviations of output from trend in their 'rule' (Q3). The change in response to the other questions was more mixed. Participants did less well at working out the parameters of the model (Q4–8)—particularly how much impact interest rate changes have on output (Q5) and the long-run impact of output on inflation (Q8). But each game may have been too short to learn much about these aspects—especially the long-run neutrality property of the model. There was also a fall in the MSE of responses to the question on how cautious monetary policy makers should be when setting interest rates (Q1), but this was not significant.

### (ii) Learning about the game

The results of the questionnaire provide tentative evidence of learning about certain aspects of the model and the nature of the optimal rule, but does this mean that players actually became better at playing the game over time?

Chart 3 below shows a summary of the mean scores attained by the 34 committees over time. This is broken down into the first set of individual play (rounds 1–4), committee play (rounds 5–12) and then individual play for a second time (rounds 13–16). For the individual rounds, the 'committee' score is taken to be the mean of the scores across the five individuals playing separately. For the committee rounds, this statistic is the mean score from committee decisions.

There are three striking features of the data:

(i) the significant upward trend in the results over time—indicative of the learning that occurred during the game;

(ii) the large rise in scores when players moved to committee decision-making in round 5; and

(iii) the large downward move in scores when participants returned to playing as individuals in round 13.

The dispersion of scores in any given round—measured by the standard deviation across players—more than halved during the game from 76 in round 1 to 35 in round 16. This suggests that the worst players learned most about the game: those who performed poorly in the first rounds got disproportionately better.

Chart 4 shows that it was not just the worst players who learned during the course of the experiment. We can rank the five players in each committee by their initial performance (in rounds 1–4), and calculate how much they improved by the final round. Although the worst players learned most, and only the worst two players in each committee made a significant improvement (again the dashed lines represent 5% significance levels), even the best players improved somewhat by the end of the game.

### (iii) Group versus individual performance

There was strong evidence that decisions taken by committees were superior to those of individuals.

## Chart 4
## Improvements in scores for players ranked by initial performance

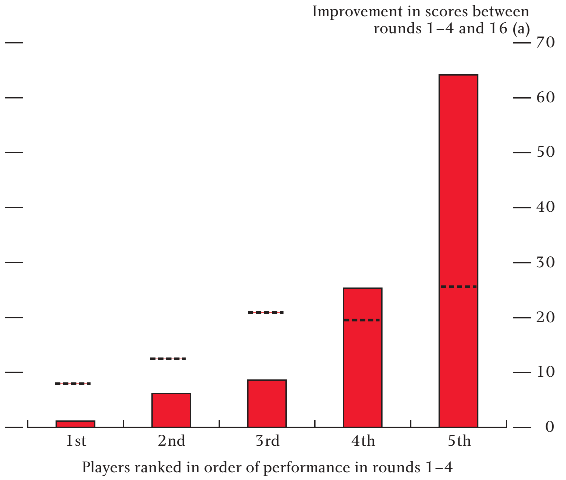


(a) Dashed lines indicate significance at the 5% level.

The average committee score of 68 (over rounds 5–12 in Chart 3 above) was nearly two-thirds better than the average score of 41 for the individual rounds (rounds 1–4 and 13–16 in Chart 3). And this increase in score is significant at the 1% level.

To give some idea of the scale of the improvement, the average score of someone following the optimal rule (see Appendix 1) would be 85, much higher than the best individual player's score (71), but only slightly better than the best committee (83). On average, moving from individual decision-making to a committee structure closed nearly two-thirds of the 'policy gap'.

What explains this improvement in committee performance? There are (at least) two distinct, competing hypotheses that can be used to explain why committee decisions are superior to those of the individuals that comprise it:

*Hypothesis 1:* A committee with 'majority' voting can neutralise the impact of some members playing badly in any given game.

*Hypothesis 2:* Committees allow members to improve performance by sharing information and learning from each other.

Chart 5 shows a visual representation of the contribution of these two hypotheses to the improvement of committees over individuals. The blue line represents the average—over the 34

## Chart 5
## Mean and median scores for committee members

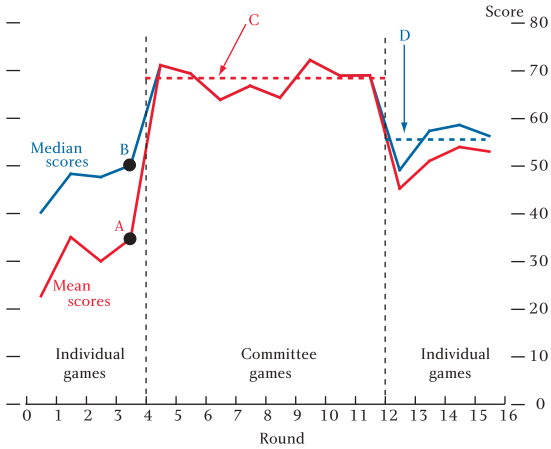


independent groups of five players—of the median player's score. The red line is simply the mean score across all players in each committee.[1] Line C is the mean score over all the committee rounds and line D is the mean score over rounds 13–16 for the median players in each of those rounds. The overall improvement in performance—generated by setting interest rates by committee—is therefore measured as the distance between C and A: the difference between the average score in the final individual round and the committee rounds.

The chart decomposes this improvement into two distinct components. The difference between the score of the mean and median player in the individual rounds (the distance B–A in Chart 5) represents the adverse effect of a minority of poor performers on the mean individual score. This is therefore the extent of improvement under Hypothesis 1 described above. And this portion of the difference in means is significant at the 1% level. So Hypothesis 1 cannot be rejected.

The remainder should represent the contribution of Hypothesis 2: C–B (the portion of the committee improvement not explained by the move to majority voting). This difference is also significant at the 1% level, so Hypothesis 2 cannot be rejected either.

Another striking feature of both these results and those of Blinder and Morgan (2000) was the significant decline in scores as participants move back to individual play, in this case at the end of round 12. By definition, this component of the committee improvement (represented by distance C–D in Chart 5) cannot be

(1) Note that the mean score in the committee rounds is the score of the committee's interest rate decision.

associated with learning about the game over time, because players know at least as much about the model in round 13 as they did before. So it therefore seems likely that this residual effect stems from the ability of committees to pool judgment, expertise and skill. This fall in scores is significant at the 1% level too: in other words, there is 'something special' about committees in addition to their ability to aid learning and to strip out the effects of 'bad' play.

Further evidence that a committee is more than just the sum of its parts is shown by asking whether the performance in the committee stages was better than the mean score of the best individual in each committee when playing alone. The mean committee score (68) was somewhat higher than the mean score of the best individual (65) (this difference is significant at the 10% level), providing evidence that committees did more than just replicate the behaviour of their best individual.

### (iv) What makes a good committee?

If committees improve decision-making by exploiting their members' ability to pool information and knowledge and to learn from each other, communication must be key. As discussed earlier, the experiment included two different ways of organising committee decision-making: one where participants were allowed to discuss their views and another where no verbal communication was allowed. Perhaps the most surprising result was that the ability to discuss did not significantly improve committee performance.

This result was in contrast to earlier trials of the game on Bank staff. So, in addition to the main experiment described above, a further small sample of students was asked to play a different version of the game as a robustness check. This variant was designed so as to raise the implicit benefit of discussion: committee members were told—with a lag of up to two periods—that a shock had occurred, and the length of this information lag was allowed to vary across players. The ability to discuss was therefore more valuable because committee members with more timely information could share this with others more quickly by verbal communication. The average score of discussion committees was higher than for non-discussion in this version of the game, although the small sample size—three committees—meant that the significance of this improvement could not be tested.

Another hypothesis is that people can communicate in different ways. And the benefits of different forms of communication are likely to depend on the nature of the game, as well as the individuals taking part. There are many games—for example snooker or chess—that may be easier to learn by watching, rather than through discussion. But for the main version of the game, and for this set of students, discussion did not provide more information than could be acquired by observing others' votes.

There is also some evidence from the psychology literature that discussion may not always enhance group performance. The idea of 'group polarisation'—as proposed by Myers (1982)—suggests that discussion tends to polarise any initial tendency within the group. This is because people have an innate desire to compare themselves favourably with each other, and so take increasingly extreme positions in favour of the initial group proposition. One way around this problem is to ensure that a frank and open exchange of views takes place at the beginning of the discussion—as outlined in an earlier study by Hall (1971) who showed that groups who established a common consensus quickly were often less effective.

So if discussion did not help committees to improve their scores in our experiment, what sort of behaviour does? Lombardelli, Proudman and Talbot (2002) explored this question in more detail, using an econometric analysis to model scores over time and across committees. After controlling for committee-specific features—such as the innate ability of participants to play the game—the model captured the upward trend in scores over time, and the rise in scores during the committee stages. Committee scores were positively related to the period in which the structural shock occurred in each round. Intuitively, the earlier in the game the structural shock took place, the more difficult the economy was to control over the remainder of the game—particularly if it took some time for the player to recognise that such a shock had occurred. Higher interest rate activism—as measured by the standard deviation of the interest rate in each ten-period game—was associated with lower scores for both individuals and committees.

But on the whole, the econometric analysis reinforced the results presented above: that committees performed significantly better than individuals, and that there was some evidence of participants learning about the game over time.

## Conclusions

This article discusses an experimental analysis of monetary policy decision-making. Although such a stylised experiment can never hope to capture fully the complexity of the decision problem faced by real-world policy-makers, the results provide evidence that the decisions made by committees were superior to those of a single individual. And there is also evidence to suggest that committee performance was, on average, better than the performance of the best individual. This suggests that the real-world preference for setting interest rates by committee is justified.

The experiment also tried to examine why committee decisions were superior to those of individuals. A significant portion of the improvement could be attributed to the process of majority voting. But there was also evidence that there is something 'special' about committees beyond their ability to strip out the effect of bad play. The ability of committees to allow the pooling of judgment and information (in whatever form) means that a group can be more than just the sum of its parts.

Perhaps surprisingly, committees who were able to discuss their decisions did not perform better than those where discussion was not allowed. It seems that, in the experiment, it was possible to glean the same amount of information by observing the votes of other committee members. But, as noted above, real-world policy-making is undoubtedly a more complex affair. The Monetary Policy Committee takes into account a much wider range of data and information than just lagged inflation and output when making its monthly interest rate decision.

What this simple experiment has shown is that it is not enough simply to take a majority decision among fixed views that have been reached in isolation. The pooling of knowledge among committee members—in whatever form—is one important reason why group decision-making is superior. And that reflects one feature of the practical operation of the Monetary Policy Committee—the exchange of views among the group helps to determine the votes of each individual.

# Appendix 1
# The model and a derivation of the 'optimal rule'

Although participants were not provided with the underlying equations the model can be described by the following equations:

$$y_t - y^* = 0.8(y_{t-1} - y^*) - 0.5(R_t - \pi_t - r^*) + \bar{g} + \eta_t \quad (1)$$
$$\pi_t = 0.7\pi_{t-1} + 0.3\pi_{t-2} + 0.2(y_t - y^*) + v_t \quad (2)$$

where: $y_t$ is log output, $y^*$ is the log of the natural rate of output (calibrated arbitrarily to 5), $\pi_t$ is inflation, $R_t$ is the nominal interest rate and $r^*$ is the neutral real interest rate (calibrated to 3% per year). $\bar{g}$ is a permanent shock, $\eta_t$ and $v_t$ are shocks corresponding to a random draw from a normal distribution ~ N(0, 0.01) in each period.

Equation **(1)** is an 'IS curve'. The current deviation of output from its natural rate ($y_t - y^*$) is a function of its one-period lag, and the deviation of the real interest rate from its neutral level in the current period ($R_t - \pi_t - r^*$). Equation **(2)** is a 'Phillips curve'. Inflation is a function of lagged values of itself and the current deviation of output from its natural rate. The coefficients on lagged inflation sum to one, reflecting the fact that although a short-run trade-off between output and inflation may exist, the Phillips curve is vertical in the long run.

Assuming that players attempt to maximise their score ($S_t$) in each period of the game, the decision problem of each player can be written as:

$$\underset{r_t}{Max} E_{t-1}\{S_t\} \quad \text{s.t. (1) and (2)} \quad \text{where: } S_t = 100 - 40|y_t - y^*| - 40|\pi_t - \pi^*| \quad (3)$$

where $\pi^*$ is the inflation target, calibrated to 2.5%.

Approximating **(3)** as a linear quadratic, we derive the optimal rule by substituting in the constraints **(1)** and **(2)** and differentiating with respect to $r_t$ to give:

$$r_t = 1.6y_{t-1} + 0.27\pi_{t-1} + 0.115\pi_{t-2} + 2\bar{g} \quad (4)$$

Obviously, the distribution of $\bar{g}$ is also unknown to participants in the experiment, so **(4)** is the 'certainty equivalent optimal rule'. Svensson and Woodford (2000) note that—under the assumption that the loss function is quadratic—the optimal policy rule under partial information is the same as its full-information counterpart. We use this optimal rule to calibrate the 'ideal' responses to the questionnaire and to conduct simulations of the model—see Lombardelli, Proudman and Talbot (2002) for further details.

# Appendix 2
# Prior beliefs questionnaire

| Date: | Group: |
|---|---|

**Please spend a few minutes filling in this questionnaire, concentrating in particular on the questions in italics. It doesn't matter if you are not familiar with the jargon in brackets: this is merely to help us calibrate your response.**

**BE AS HONEST AS YOU CAN—THERE ARE NO RIGHT OR WRONG ANSWERS!**

What is your player number?

| |
|---|
| |

1) To *what extent should monetary policy makers respond cautiously to shocks* (ie if their interest rate reaction function includes the following expression $i_t = \alpha i_{t-1}$ + ....., what weight should they place on $\alpha$)?

Not at all cautiously (ie $\alpha = 0$) Very cautiously (ie $\alpha = 1$)

| 0 | 0.1 | 0.2 | 0.3 | 0.4 | 0.5 | 0.6 | 0.7 | 0.8 | 0.9 | 1 |
|---|---|---|---|---|---|---|---|---|---|---|
| | | | | | | | | | | |

2) After *how many quarters is the maximum impact of monetary policy on inflation felt*?

| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
|---|---|---|---|---|---|---|---|---|---|---|
| | | | | | | | | | | |

3) What *relative weight should monetary policy makers place on smoothing output compared with controlling inflation* (ie if their reaction function includes the following expression $i_t = \alpha(y_t - Y) + (1 - \alpha)(\pi_t - \pi^*)$ + ....., what weight should they place on $\alpha$)?

None (ie $\alpha = 0$) All (ie $\alpha = 1$)

| 0 | 0.1 | 0.2 | 0.3 | 0.4 | 0.5 | 0.6 | 0.7 | 0.8 | 0.9 | 1 |
|---|---|---|---|---|---|---|---|---|---|---|
| | | | | | | | | | | |

4) To *what extent are shocks to output persistent* (ie if the expression for output included the following term $y_t = \alpha y_{t-1}$ + ....., what weight do you think $\alpha$ would take)?

Not at all persistent (ie $\alpha = 0$) Completely persistent (ie $\alpha = 1$)

| 0 | 0.1 | 0.2 | 0.3 | 0.4 | 0.5 | 0.6 | 0.7 | 0.8 | 0.9 | 1 |
|---|---|---|---|---|---|---|---|---|---|---|
| | | | | | | | | | | |

5) How *sensitive is output to changes in interest rates* (ie if the expression for output included the following term $y_t = \alpha i_t + ....$, what weight do you think $\alpha$ would take)?

Not at all sensitive (ie $\alpha = 0$) Very sensitive (ie $\alpha = 1$)

| 0 | 0.1 | 0.2 | 0.3 | 0.4 | 0.5 | 0.6 | 0.7 | 0.8 | 0.9 | 1 |
|---|---|---|---|---|---|---|---|---|---|---|
| | | | | | | | | | | |

6) To *what extent are shocks to inflation persistent* (ie if the expression for inflation included the following term $\pi_t = \alpha\pi_{t-1} + ....$, what weight do you think $\alpha$ would take)?

Not at all persistent (ie $\alpha = 0$) Completely persistent (ie $\alpha = 1$)

| 0 | 0.1 | 0.2 | 0.3 | 0.4 | 0.5 | 0.6 | 0.7 | 0.8 | 0.9 | 1 |
|---|---|---|---|---|---|---|---|---|---|---|
| | | | | | | | | | | |

7) To *what extent is inflation sensitive to deviations of output from trend in the short run* (ie if the expression for inflation included the following term $\pi_t = \alpha(y_{t-1} - Y) + ....$, what weight do you think $\alpha$ would take)?

Not at all sensitive (ie $\alpha = 0$) Highly sensitive (ie $\alpha = 1$)

| 0 | 0.1 | 0.2 | 0.3 | 0.4 | 0.5 | 0.6 | 0.7 | 0.8 | 0.9 | 1 |
|---|---|---|---|---|---|---|---|---|---|---|
| | | | | | | | | | | |

8) To *what extent is inflation sensitive to deviations of output from trend in the long run*?

Not at all sensitive (ie $\alpha = 0$) Highly sensitive (ie $\alpha = 1$)

| 0 | 0.1 | 0.2 | 0.3 | 0.4 | 0.5 | 0.6 | 0.7 | 0.8 | 0.9 | 1 |
|---|---|---|---|---|---|---|---|---|---|---|
| | | | | | | | | | | |

9) *What course are you studying*?

……………………………………………………………………………………………………………………….

10) *Are you….*

| Undergraduate: 2nd year | Undergraduate: 3rd year | Graduate student |
|---|---|---|
| | | |

## References

**Bank of England (1999)**, *Economic models at the Bank of England*.

**Bank of England (2000)**, *Economic models at the Bank of England*, September 2000 update.

**Barro, R J and Gordon, D B (1983)**, 'Rules, discretion and reputation in a model of monetary policy', *Journal of Monetary Economics*, Vol. 12, No. 1, pages 101–21.

**Blinder, A S and Morgan, J (2000)**, 'Are two heads better than one: an experimental analysis of group vs individual decision making', *NBER Working Paper*, No. 7909, September.

**Davis, D D and Holt, C A (1993)**, *Experimental economics*, Princeton University Press.

**Fry, M, Julius, D, Mahadeva, L, Roger, S and Sterne, G (2000)**, 'Key issues in the choice of monetary policy framework', in Mahadeva, L and Sterne, G (eds), *Monetary frameworks in a global context*, Routledge.

**Fuhrer, J C and Moore, G R (1995)**, 'Inflation persistence', *Quarterly Journal of Economics*, Vol. 110, No. 1, pages 127–59.

**Hall, J (1971)**, 'Decisions, decisions, decisions', *Psychology Today*, November.

**Janis, I L (1972)**, *Victims of groupthink*, Boston: Houghton Mifflin.

**Kagel, J H and Roth, A E (1995)**, *The handbook of experimental economics*, Princeton University Press.

**Lombardelli, C, Proudman, A J and Talbot, J I (2002)**, 'Committees versus individuals: an experimental analysis of monetary policy decision-making', *Bank of England Working Paper no. 165*.

**Myers, D G (1982)**, 'Polarizing effects of social comparison', in Bransdatter, H, Davis, J H and Stocker-Kreichgauer, G (eds), *Group decision-making*, New York: Academic Press.

**Svensson, L E O and Woodford, M (2000)**, 'Indicator variables for optimal policy', *NBER Working Paper*, No. 7953, October.