



Preferences or private assessments on a monetary policy committee?



Stephen Hansen^a, Michael McMahon^{b,*}, Carlos Velasco Rivera^c

^a Universitat Pompeu Fabra and Barcelona GSE, Spain

^b Department of Economics, University of Warwick, CEPR, CAGE (Warwick), CEP (LSE), CJM (LSE), and CAMA (ANU), Coventry CV4 7AL, United Kingdom

^c Princeton University, United States

ARTICLE INFO

Article history:

Received 9 November 2012

Received in revised form

5 June 2014

Accepted 26 June 2014

Available online 5 July 2014

Keywords:

Committees

Monetary policy

Private information

ABSTRACT

Using Bank of England voting data, we show empirically that members' votes are driven by heterogeneous individual assessments of the economy as well as their individual policy preferences. Estimates indicate that internal committee members form more precise assessments than externals and are also more hawkish. The estimates allow the first quantification of the gain due to information aggregation on monetary policy committees. The marginal gain from additional committee members tapers quickly after five members. There is no evidence of gains through externals' moderating internals' preferences. A relatively small committee of highly informed internal members emerges as a desirable committee structure.

© 2014 Elsevier B.V. All rights reserved.

1. Introduction

What role do private assessments of economic conditions play in explaining individual voting behavior on the Bank of England's Monetary Policy Committee (MPC)? While empirical work focuses largely on preference differences (hawkishness or dovishness) as the primary driver of individual voting differences, both committee designers and theoretical models emphasize that policymakers may differ in their assessments of economic conditions. Even if they share the same preferences, different beliefs on such conditions can lead to different votes. Pooling of private assessments also provides a natural channel for understanding why committees outperform individuals, as found in experimental work (Blinder and Morgan, 2005; Lombardelli et al., 2005). Our empirical approach allows *both* differences in preferences and in private assessments to play a role in committee decisions.

The first major contribution of this paper is to estimate a model of equilibrium voting behavior using the individual voting record of the MPC. In the model, the individual structural parameters of interest are preferences and the precision of private assessments, which we call *expertise*. The estimation follows the two-step approach of Iaryczower and Shum (2012), who estimate these quantities for US Supreme Court justices. The model is estimated both under the assumption that voting is sincere, in which case members behave as if their votes determine policy, and strategic, in which case members condition their votes on being pivotal. These assumptions do not change estimates of expertise, but do change those of preferences.

* Corresponding author. Tel.: +44 24 7652 3056.

E-mail address: m.mcmahon@warwick.ac.uk (M. McMahon).

A model evaluation exercise based on out-of-sample prediction shows that sincere voting explains the data modestly better than strategic voting.¹

The expertise estimates show that private assessments play an important role in individual voting behavior. In meetings with the greatest uncertainty about the state of the economy, an individual member is up to 40 percentage points more likely to make the correct decision relative to a model in which private views play no role. Modeling private views is therefore important for explaining voting behavior, and differences in these views can be expected to lead to different votes even when members share the same preferences.

The estimates also allow one to explore differences between externally appointed experts—members appointed solely to make the decision each month—and internal members—those who also have executive responsibilities in the central bank. Some countries (for example, the UK, Poland, and Hungary) make use of externals, presumably to take advantage of diversity, while others (for example, the USA and Sweden) do not. In line with the existing literature (e.g. Gerlach-Kristen, 2003, and references below), internals' estimated preferences are significantly more hawkish, but our novel finding is that they have higher estimated expertise. This finding implies that the justification for externals' inclusion from a purely informational perspective requires further thinking.

To try to uncover the source of these differences, we examine the behavior of members who have worked in central banking prior to their MPC appointment—insiders—in order to see whether prior experience or current position matters most for the behavior of internals. Estimates show that insiders are more hawkish and have more expertise than outsiders, which suggests that it is prior experience as central bankers rather than holding a senior position in the central bank that drives the internal–external differences.²

Optimal committee design is an important and open area in monetary policy (Reis, 2013). The second major contribution of this paper is to use the estimated parameters to assess how different committee structures affect the quality of decision making via counterfactual simulations. The first aspect of committee design considered is the effect of size. A well-known hypothesis is that committees outperform individuals because they aggregate dispersed private information (see Gerlach-Kristen, 2006 or Blinder, 2007 for recent discussions of this in the context of monetary policy). A major advantage of the estimation approach is that one can measure the extent of these gains in a real-world committee.

In periods in which the economic environment is most uncertain, moving from an individual expert to a committee structure improves decision making substantially, although the gains decline rather quickly. For example, a committee of five internal members is between 7 and 11 percentage points more likely to make the correct decision than an individual internal member. This indicates that the gain from committees' pooling imperfect knowledge is potentially significant. However, the marginal benefit of an additional member tapers quite quickly after about five members; the marginal gain from adding four additional internal members, making a committee of nine internal members, is only about 1–3 percentage points. The overall message is that the additional members on large committees, such as the 24-member Governing Council of the ECB, may improve decision making only marginally if expertise is sufficiently high. Given the potentially large costs from adding more and more members, such as information exchange problems and free-riding by members (see Sibert, 2006, for a discussion), a smaller committee is likely to be better.

The second design issue explored is whether, as argued by Blinder (2007), external members can add value through moderating internals' hawkish preferences. This would support their inclusion in committees despite their lower expertise. When the economic data broadly points to the high interest rate likely being correct, the preference moderation effect indeed arises because internals are less willing to follow private signals that low rates should be chosen. But, the moderation effect is not found to be large enough to fully overturn the effect of lower expertise, and internal and external members perform equally well. Moreover, when economic data suggests that the low interest rate is likely correct, internals' hawkish preferences work to offset externals' rigid decision making, amplifying internals' outperformance of externals. Overall, the gains and the losses from preference moderation roughly cancel out on average.

In short, a relatively small, homogenous committee of members with high expertise performs very well even if they have a hawkish bias. Of course, further work remains to be done on the exact nature of deliberation in the committee meeting, and alternative sources of value that external members might bring. Still though, the paper is the first to separately identify the role of preferences and individual assessments in monetary policymaking, and so provides potentially valuable facts for extending our knowledge of both the how and why of policymaking by committee.

The paper is related to two main strands of the existing literature. First, there is a growing literature on all aspects of the use of committees to make monetary policy decisions: this includes important summaries of the state of knowledge (Gerling et al., 2005; Blinder, 2007), issues of agenda-setting (Riboni and Ruge-Murcia, 2010), reputation-building on monetary policy committees (Sibert, 2003; Hansen and McMahon, 2013), credibility of committees (Mihov and Sibert, 2006), and the desirability of a committee over an individual when there is uncertainty about the economic situation (Gerlach-Kristen, 2006). Committee decision making has also been extensively studied by social psychologists and Sibert (2006) provides a discussion of the main findings and how it applies to monetary policy committees. Secondly, there are

¹ The results described here correspond to the sincere model, but results for both assumptions are contained either in the main text or the accompanying online appendix.

² We also examine splits based on age, education and prior career. While in no case are there significant preference (both sincere and strategic) or expertise differences, in all cases members' decisions are importantly influenced by their private assessments of the economic situation. This corroborates the view that members' individual views of the unknown state of economy are an important driver of voting heterogeneity.

a number of other papers that specifically study the MPC and internal–external differences within it including Gerlach-Kristen (2003), Bhattacharjee and Holly (2005), Spencer (2006), Besley et al. (2008), Harris et al. (2011) and Hix et al. (2010). As mentioned above, these all focus exclusively on preference differences, with the general conclusion that external members are more dovish the internals.

The rest of this paper is organized as follows. Section 2 provides background information on the MPC. Then, Section 3 presents the voting model that serves as the basis of the empirical analysis. Section 4 details the estimation strategy and Section 5 describes the data used to implement the strategy. The estimation results are presented in Section 6. Section 7 presents counterfactual results on the committee structure and Section 8 concludes.

2. The monetary policy committee

The MPC first convened on 6 June 1997, and has met every month since. Its remit, as defined in the Bank of England Act (1998), is to “maintain price stability, and subject to that, to support the economic policy of Her Majesty’s government, including its objectives for growth and employment.” In practice, the committee seeks to achieve a symmetric target inflation rate of 2%, based on the Consumer Price Index.³ At the end of each meeting, the Governor proposes an interest rate decision that he or she believes will command a majority and each member then chooses whether to agree with the Governor’s proposal, or dissent and state their preferred alternative interest rate. Each member’s vote is published as part of otherwise unattributed minutes. Throughout, the analysis considers the MPC voting records between June 1997 and March 2009, when the main focus of the decision (temporarily) shifted to asset purchase decisions related to quantitative easing.⁴ The voting records indicate both the proposed interest rate decision (such as +50 basis points), as well as the alternative preference for those who do not back the proposal (such as +25 bps).⁵

The MPC provides a useful setting for the analysis because there is a one-person, one-vote philosophy: the Bank encourages members to simply determine the rate of interest that they feel is most likely to achieve the inflation target (Bank of England, 2010) and majority vote determines the outcome. As such, the observed votes should reflect members’ genuine policy preferences. Consistent with this philosophy, the MPC displays substantial dissent: 64% of the 142 meetings in the sample have at least one deviation from the committee majority and within the set of non-unanimous meetings, 5–4 and 6–3 decisions are not uncommon.

Also useful is that the MPC is made up of five internal and four external members. Internals serve as executives of the Bank of England: the Governor, two Deputy Governors, the Chief Economist, and the Executive Director for Markets. The Chancellor of the Exchequer appoints externals (subject to approval from the Treasury Select Committee) from outside the Bank.⁶ There are no restrictions on who can serve as external members, and they have come from many different backgrounds. The sample contains a total of 13 internal and 14 external members. Every member receives all Bank of England staff briefings related to monetary policy and attends the monthly meetings in full. Bar the governors who serve five year terms, all other members serve three year terms. When members’ terms end, they can either be replaced or reappointed.

Table 1 splits the members who served on the MPC during the sample by internal and external (their tenure is reported in brackets). For later use in the analysis, it also reports whether or not members had worked within the Bank at some point prior to their appointment to the MPC (*insiders* versus *outsiders*).

3. Model

This section presents a reduced-form model of monetary policy decision making under uncertainty on which the estimation exercise is based. Members must choose one of two interest rates—one higher than another—to implement in response to an unobserved inflationary state. The two rates in each meeting are called the agenda. All members prefer the higher (lower) rate when there is more (less) inflationary pressure. However members differ in how averse they are to wrongly choosing low interest rates and such member-specific preferences or biases determine how much evidence they need that the economy is inflationary in order to vote high. Members also form individual assessments of the state of the economy on the basis of public information and the realization of a privately observed signal, whose precision measures expertise. The model yields a voting rule in which members vote for high rates if and only if they are sufficiently convinced that the economy is in an inflationary state. The formal structure of the model is closest to that in Duggan and Martinelli (2001), who study voting over binary outcomes with continuous signals. The reduced form is for clarity and the online

³ There was a change from RPIX to CPI as the measure of inflation in January 2004, and with this change, the inflation target was reduced from 2.5% to 2%.

⁴ These data are available from the Bank of England website. We use each regular MPC meetings in this period but we drop from the dataset the (unanimous) emergency meeting held after 9/11.

⁵ Before June 1998 there is information about whether members preferred higher or lower interest rates compared with the decision, but not about their actual preferred rate. In these cases, we treat a member’s vote as either 25 basis points higher or lower than the decision, in the direction of disagreement. Given how we use the voting data, discussed below, this assumption has no implications for our analysis.

⁶ The Bank’s explicit aim in appointing external members is their heterogeneity vis-à-vis internal members. According to Bank of England (2010), external appointments “ensure that the MPC benefits from thinking and expertise in addition to that gained inside the Bank of England.”

Table 1
MPC Members.

| | Appointment | |
|----------------------------|---|---|
| | External | Internal |
| <i>Central Bank Career</i> | | |
| Outsider | C. Allsopp (06/00-05/03) K. Barker (06/01-03/09) M. Bell (06/02-06/05) T. Besley (06/06-03/09) D. Blanchflower (06/06-03/09) A. Budd (12/97-05/99) W. Buitert (06/97-05/00) D. Julius (11/97-05/01) R. Lambert (06/03-03/06) S. Nickell (06/00-05/06) A. Sentance (10/06-03/09) S. Wadhvani (06/99-05/02) D. Walton (07/05-06/06) | C. Bean (10/00-03/09) D. Clementi (11/97-08/02) J. Gieve (01/06-03/09) A. Large (11/02-01/06) R. Lomax (07/03-06/08) J. Vickers (06/98-11/00) |
| Insider | C. Goodhart (06/97-05/00) | E. George (06/97-06/03) H. Davies (06/97-07/97) P. Fisher (03/09-03/09) S. Dale (07/08-03/09) I. Plenderleith (06/97-05/02) M. King (06/97-03/09) P. Tucker (06/02-03/09) |

Notes: This table shows committee members serving on the MPC for the period 06/97-03/09 and splits them into whether they are internal or external (horizontal axis), and whether they are insiders or outsiders (vertical axis). The former distinction is based on how they are appointed while the latter is determined by whether they had prior experience in the Bank of England. For example, Eddie George served from 06/97 to 06/03 as an internal member who was also an insider on account of his career at the Bank of England. In contrast, Kate Barker, who served from 06/01 to 03/09, was appointed as an external member and had no prior experience in the Bank before her appointment.

appendix shows that threshold voting rules also emerge within a standard New Keynesian model based on Clarida et al. (1999) and Galí (2008).

A more big-picture issue is how to interpret preferences and private signals in the monetary policy context. As Blinder (2007) points out, there are several admissible interpretations. Members might literally have different preferences over the trade off between inflation and unemployment, but they might also have different models of the macroeconomy, models with a particular bias, in their heads. A member who analyzes an issue using a model that more typically leads to a recommendation for higher interest rates would have a hawkish bias, which is observationally equivalent to hawkish preferences. In terms of private assessments, it could be that members literally have private information from their personal contacts which they add to the common data all members receive prior to voting. Alternatively, it could be that the diverse models and forecasting methods that members use give them imperfect, heterogeneous views about underlying economic conditions. Or, it could be that even with the same preferences and information set, heterogeneity in cognitive processing ability, of the sort emphasized in the limited attention literature (discussed, for example, in Sims, 2010), leads each member to come to a different view.

3.1. Member preferences

In each period t the committee implements a decision $d_t \in \{0, 1\}$, where 0 represents the lower of two possible rate changes and 1 the higher. The agenda is exogenous.⁷ The restriction to a two-decision agenda rarely binds since there are three unique votes in only 7 of the 142 meetings in our sample and no meetings with four or more unique votes (the appendix describes how these anomalies are treated). The committee has N (odd) members. Each one chooses a vote $v_{it} \in \{0, 1\}$ each period and $d_t = 1$ if and only if $\sum_i v_{it} \geq (N+1)/2$.

Member i 's preferences over d_t depend on a state ω_t , and are represented by $u_i(d_t|\omega_t)$. ω_t represents unknown economic conditions relevant to inflation, for example the magnitude of a demand shock or the output gap of the economy. $\omega_t = 1$ is the high interest rate state and $\omega_t = 0$ is the low interest rate state. Preferences are such that $u_i(0|0) = u_i(1|1) = 0$, $u_i(0|1) = -\theta_i$, and $u_i(1|0) = -(1-\theta_i)$. In this formulation, all members agree that decision $d_t = \omega_t$ is best, but have different payoffs from mismatches in different states. A member with a higher θ_i suffers more when the committee incorrectly chooses the lower rate than when it wrongly chooses the higher rate. We therefore interpret him as being more hawkish while a member with a lower θ_i is more dovish.

⁷ See Riboni and Ruge-Murcia (2010) for a model of endogenous agendas.

3.2. Member beliefs and expertise

Prior to voting, members form beliefs on ω_t by relying on two sources. First, there is public information about the current state of the economy like market data, staff forecasts, and each others' stated opinions. Let $q_t \equiv \Pr[\omega_t = 1]$ denote the common prior belief that the economy is in the high state that is consistent with this information. Second, member i privately observes the signal $s_{it} | \omega_t \sim N(\omega_t, \sigma_i^2)$. These signals are independent conditional on ω_t . σ_i measures member i 's expertise or the precisions of his private assessment.

Member i 's posterior belief on the state— $\hat{\omega}_{it} \equiv \Pr[\omega_t = 1 | s_{it}]$ —is formed via Bayes' Rule. Basic manipulations of the normal density yield the expression

$$\ln \left[\frac{\hat{\omega}_{it}}{1 - \hat{\omega}_{it}} \right] = \ln \left[\frac{q_t}{1 - q_t} \right] + \frac{2s_{it} - 1}{2\sigma_i^2}. \quad (1)$$

Whenever $\sigma_i < \infty$, then $\hat{\omega}_{it} \neq \hat{\omega}_{jt}$ generically for any two distinct members i and j . Even if members draw identically distributed signals, they will have divergent private assessments. A separate point is that if members differ in their expertise, they will put different weights on their private signals such that members with more accurate assessments will rely more on their own view and less on the public signal.

Several clarifications about the assumed information structure are worth making. One concern is that members have heterogeneous priors q_{it} . In this framework, such heterogeneity would be impossible to distinguish from heterogeneity in θ_i —both are individual characteristics that bias member i towards one of the interest rates relative to his colleagues. One can see this formally in the voting rules derived in the next subsection: q_{it} and θ_i would enter symmetrically in them. So, the main implication of heterogeneous priors is that they require a broader interpretation of the θ estimates to admit sources of bias apart from pure preferences. Importantly, our estimates of σ_i should remain unaffected.

Second, q_t represents members' prior belief at the moment at which they vote. If members reveal to colleagues their private signals prior to voting, as in Gerlach-Kristen (2006), then q_t pools all private assessments and becomes the common, shared belief. In this sense, the estimates of $1/\sigma_i$ represent the upper bound on precision.

A final concern is that there is serial correlation in individual votes driven by persistent private signals. But serial correlation in votes can be also be driven by serial correlation in the common prior. Modelling persistent private signals requires a significantly more complex information structure and decision rule, and estimating such models is not yet feasible. On the other hand, persistence in the prior is not incompatible with the model because the decision rules are already written conditional on q_t . In order to try to address this concern, we have explored this issue empirically (the full details are provided in the online appendix). While members' votes are serially correlated, once we control for the prior, the AR(1) coefficient is no longer significant. As such, the paper proceeds under the assumption that signals are serially uncorrelated.

3.3. How members vote given their preferences and beliefs

There are two assumptions the voting literature makes about voting behavior in committees (Austen-Smith and Banks, 1996). First, when members vote *sincerely* they behave as if they get utility from matching their vote to the state. Under this assumption member i 's expected utility from $v_{it} = 1$ is $-(1 - \theta_i)\Pr[\omega_t = 0 | s_{it}]$, while his expected utility from $v_{it} = 0$ is $-\theta_i\Pr[\omega_t = 0 | s_{it}]$. He thus chooses $v_{it} = 1$ whenever

$$\ln \left[\frac{\hat{\omega}_{it}}{1 - \hat{\omega}_{it}} \right] \geq \frac{1 - \theta_i}{\theta_i} \quad (2)$$

which implies choosing $v_{it} = 1$ whenever

$$s_{it} \geq \frac{1}{2} - \sigma_i^2 \left[\ln \left(\frac{\theta_i}{1 - \theta_i} \right) + \ln \left(\frac{q_t}{1 - q_t} \right) \right] \equiv s_{it}^*(\text{SIN}). \quad (3)$$

In other words, member i adopts a threshold voting rule in which he votes for high rates if and only if his signal provides sufficient evidence for the high state having arisen. This threshold is both time- and member-specific, and depends on both preferences and expertise. As mentioned above, if the model had heterogeneous priors, the q_{it} term would enter the threshold alongside θ_i .

An alternative assumption is that committee members behave *strategically*. This requires modeling players' voting rules as strategies in a Bayesian game. The main modification from the sincere case is that members only condition their votes on the event in which they are pivotal for the committee's decision—that is, that there are exactly $(N - 1)/2$ votes for $d_t = 0$ and $(N - 1)/2$ votes for $d_t = 1$. In the spirit of the results in Duggan and Martinelli (2001), all voters continue to adopt a cutoff voting rule such that i votes 1 whenever

$$\frac{\Pr[\text{PIV}_i | s_{-i}^*, \omega_t = 1]}{\Pr[\text{PIV}_i | s_{-i}^*, \omega_t = 0]} \ln \left[\frac{\hat{\omega}_{it}}{1 - \hat{\omega}_{it}} \right] \geq \frac{1 - \theta_i}{\theta_i}, \quad (4)$$

where $\Pr[\text{PIV}_i | s_{-i}^*, \omega_t]$ is the probability that he is pivotal given other members' cutoffs s_{-i}^* and the inflation state ω_t . Thus, member i selects $v_{it} = 1$ if and only if

$$s_{it} \geq \frac{1}{2} - \sigma_i^2 \left[\ln \left(\frac{\theta_i}{1 - \theta_i} \right) + \ln \left(\frac{q_t}{1 - q_t} \right) + \ln \left(\frac{\Pr[\text{PIV}_i | s_{-i}^*, \omega_t = 1]}{\Pr[\text{PIV}_i | s_{-i}^*, \omega_t = 0]} \right) \right] \equiv s_{it}^*(\text{STR}). \quad (5)$$

An equilibrium is a collection of these cutoffs $\{s_{it}^*(\text{STR})\}_{i=1}^N$ such that all N equations described in (5) are satisfied. Clearly in general $s_{it}^*(\text{STR}) \neq s_{it}^*(\text{SIN})$, and obtaining a closed form solution for $s_{it}^*(\text{STR})$ is not possible.⁸ Since both appear in the literature, the paper estimates both models, and compares the goodness-of-fit of both models in Section 6.2.

4. Econometric methodology

In order to derive the likelihood function for observed votes, consider a meeting at time t . Under the high inflation state, member i votes for the low rate with probability $(\Phi(s_{it}^*(\cdot) - 1)/\sigma_i)$ and for the high rate with probability $1 - \Phi(s_{it}^*(\cdot) - 1)/\sigma_i$. Under the low inflation state, the corresponding expressions are $\Phi(s_{it}^*(\cdot)/\sigma_i)$ and $1 - \Phi(s_{it}^*(\cdot)/\sigma_i)$. The probabilities that the economy is in the high and low inflation states are q_t and $(1 - q_t)$, respectively. Therefore, the likelihood function given $s_{it}^*(\cdot)$ and σ_i is

$$\prod_t \left[q_t \prod_i (\kappa_{1it})^{v_{it}} (1 - \kappa_{1it})^{1 - v_{it}} + (1 - q_t) \prod_i (\kappa_{0it})^{v_{it}} (1 - \kappa_{0it})^{1 - v_{it}} \right] \quad (6)$$

where

$$\kappa_{1it} \equiv 1 - \Phi \left(\frac{s_{it}^*(\cdot) - 1}{\sigma_i} \right) \quad \text{and} \quad \kappa_{0it} \equiv 1 - \Phi \left(\frac{s_{it}^*(\cdot)}{\sigma_i} \right). \quad (7)$$

The term in square brackets in (6) is the likelihood of observing the votes in meeting t , which is then multiplied across all meetings to form the overall likelihood.

To obtain estimates of θ_i and σ_i , we follow the two-step estimation approach of Iaryczower and Shum (2012) and the literature that estimates games more broadly (see the papers cited in Pesendorfer and Takahashi, 2013). The general idea is to estimate choice—in this case vote—probabilities in a first stage with flexible functional forms that depend on observed covariates. These first-stage coefficient estimates are not themselves linked to any underlying economic model. Then, in the second stage, the estimated probabilities are linked to an economic model that explains them, from which structural parameters are backed out. It is important that these structural parameters are identified given the estimated probabilities although we delay a discussion of the intuition for such identification until the next subsection.

The model for the prior (or, in statistical language, the mixing probability) is

$$q_t = \frac{\exp(\alpha_0 + \alpha_1 q_t^R + \alpha_2 q_t^M)}{1 + \exp(\alpha_0 + \alpha_1 q_t^R + \alpha_2 q_t^M)}, \quad (8)$$

where q_t^R and q_t^M are proxy variables, described below, correlated with the true q_t . In principle, the true q_t depends on all available public information at date t , including macroeconomic aggregates and financial variables. Following Imai and Tingley (2012), we favor parsimony in the model and therefore use two statistics that should capture all of the available information as it pertains to expectations of what the MPC will do. We include both a market survey measure and data on market expectations derived from option prices because, as we describe in more detail below, both have advantages and disadvantages.

The model for the κ terms is

$$\kappa_{0it} = \frac{\exp(\beta \cdot S_{it})}{1 + \exp(\beta \cdot S_{it})} \quad \text{and} \quad \kappa_{1it} = \frac{\kappa_{0it} + \exp(\gamma \cdot S_{it})}{1 + \exp(\gamma \cdot S_{it})} \quad (9)$$

where $S_{it} = (\mathbf{1}, D_i, q_t^R, q_t^M, D_i \cdot q_t^R, D_i \cdot q_t^M, Z_t)$ is a vector of covariates. Here D_i are dummy variables that group members into categories across which we wish to compare preferences and expertise. For example, the main estimates compare internal and external members, so the model contains a single dummy that equals one if and only if member i is an internal. Z_t is a vector of meeting-specific variables that potentially affect voters' tradeoff between errors in states 0 and 1 without influencing their beliefs on economic conditions. The interactions between D_i and the proxies for the prior control for members with different signal precisions reacting differently to changes in the prior. The dependence of κ_1 on κ_0 ensures that $\kappa_1 \geq \kappa_0$, which is implied by the model and necessary for identifying the first stage parameters β and γ . Without the restriction that $\kappa_1 \geq \kappa_0$, assigning individual votes to the cluster corresponding to the correct inflationary state is not possible.

Given this model, we estimate the α , β , and γ via maximum likelihood applied to (6), as do Iaryczower and Shum (2012). From these estimates, one obtains fitted values \hat{q}_t , $\hat{\kappa}_{0it}$, and $\hat{\kappa}_{1it}$. Using the definitions in (7), one can recover estimates for

⁸ In principle there can be multiple equilibria of the strategic voting game corresponding to different cutoffs, but specifying which one is being played is not important for the estimation strategy.

signal accuracy and equilibrium voting threshold under both the sincere and strategic voting models as follows:

$$\hat{\sigma}_{it} = \frac{1}{\Phi^{-1}(1 - \hat{\kappa}_{0it}) - \Phi^{-1}(1 - \hat{\kappa}_{1it})} \quad \text{and} \quad \hat{s}_{it}^* = \frac{\Phi^{-1}(1 - \hat{\kappa}_{0it})}{\Phi^{-1}(1 - \hat{\kappa}_{0it}) + \Phi^{-1}(1 - \hat{\kappa}_{1it})}. \quad (10)$$

Obtaining estimates for the preference parameter θ requires specifying the sincere or strategic model. Under the former, $\hat{\theta}_{it}$ can be obtained by plugging \hat{q}_{it} , $\hat{\sigma}_{it}$, and \hat{s}_{it}^* into the equations defined in (3) and solving directly. Under the latter, plugging \hat{q}_{it} , $\hat{\sigma}_{it}$, and \hat{s}_{it}^* into (5) generates a system of equations to solve for θ_{it} .

The estimated parameters in the econometric model are the first stage coefficients, which do not vary over time. The mapping of these parameter estimates into the second stage, however, delivers an estimate of σ_i and θ_i for each separate meeting because of variation in the proxies q_t^R and q_t^M . We construct estimates of σ_i and $\theta_i - \hat{\sigma}_i$ and $\hat{\theta}_i$ —by averaging $\hat{\sigma}_{it}$ and $\hat{\theta}_{it}$ over time, which should cancel out measurement error in the relationship between q_t^R and q_t^M and the true q_t . Monte Carlo exercises, presented in the online appendix, show that $\hat{\sigma}_{it}$ and $\hat{\theta}_{it}$ indeed vary over time when recovered from voting datasets generated with constant σ_i and θ_i parameters, but that averaging produces estimates centered on the true values.

4.1. Intuition for identification of θ and σ

The key link between the behavioral model and empirical exercise comes via the extraction of θ and σ from the estimated state-contingent voting probabilities $\hat{\kappa}_0$ and $\hat{\kappa}_1$. This section provides an intuitive argument for why they are identified. Fig. 1 plots κ_0 and κ_1 —the theoretical probabilities of voting for the high rate in states 0 and 1, respectively—for different member types as a function of the prior q_t in the case of sincere voting. The intuition for the strategic model is essentially the same.

Fig. 1a and b compare two members who have the same signal precision ($\sigma = 0.5$) but one is a hawk ($\theta^H = 0.65$) and one is a dove ($\theta^D = 0.35$). As the prior increases, both have a higher chance of voting for the high rate. However, the hawk always has higher probability of choosing the high rate than the dove because he is inherently more inflation-averse. Formally, $\kappa_0^H \geq \kappa_0^D$ and $\kappa_1^H \geq \kappa_1^D$.

Now consider two members who have the same preferences ($\theta = 0.5$) but different signal precisions. There is a more-expert ($\sigma^{ME} = 0.4$) and a less-expert ($\sigma^{LE} = 0.6$) member. The probabilities they vote high in the two states are plotted in Fig. 1c and d. The more-expert member can better identify the state ω_t , and so his vote tends to match it more often than that of the less-expert member. So, he is on average more likely to choose $v_{it} = 0$ ($v_{it} = 1$) when $\omega_t = 0$ ($\omega_t = 1$). In short, expertise is reflected in a flipping across states in terms of who is most likely to vote for the high interest rate: $\kappa_0^{ME} \leq \kappa_0^{LE}$, but $\kappa_1^{ME} \geq \kappa_1^{LE}$. Thus, the difference in estimated voting probabilities across states and across members ($\hat{\kappa}_{i0}$ and $\hat{\kappa}_{i1}$) allows one to distinguish preference versus expertise differences.

4.2. Construction of confidence intervals

To compute confidence intervals, the paper uses a Monte-Carlo approach that is similar in spirit to boot-strapping. 500 different values for first-stage coefficients are drawn from a multivariate normal distribution whose mean is the vector of estimated coefficients, and whose variance-covariance matrix is the inverse of the negative Hessian matrix. These draws are then used to generate a distribution over the structural parameters (and their differences) from which confidence intervals are constructed. We provide the full details of our simulation approach, which is suggested in King et al. (2000), in the online appendix.

5. The data

This section describes the construction of the period t voting agenda and the proxies for the prior q_t^R and q_t^M introduced in the previous section. Further information on these constructions is in the online appendix.

5.1. Construction of the agenda

In periods with two unique votes by MPC members (64% of the meetings), we consider the agenda to be the two observed votes, and set $v_{it} = 1$ if member i voted for the higher rate. A complication arises in meetings with unanimous votes since we do not directly observe which alternative was under consideration. To address this, we use a survey of around 30–50 market economists conducted by Reuters in the days leading up to the MPC meetings. The survey asks respondents to predict the outcome of MPC voting by writing a probability distribution over possible interest rate choices.⁹ Because of the fairly large cross-sectional sample size and the prominence of the participating institutions, the average beliefs in the survey data can be taken as a good measure of conventional wisdom about inflationary pressures. When votes are unanimous, we

⁹ In fact the exact question varied over time. In the online appendix, full details are provided of the surveys and the five meetings (out of 142 for which we have voting data) for which we have no Reuters survey data.

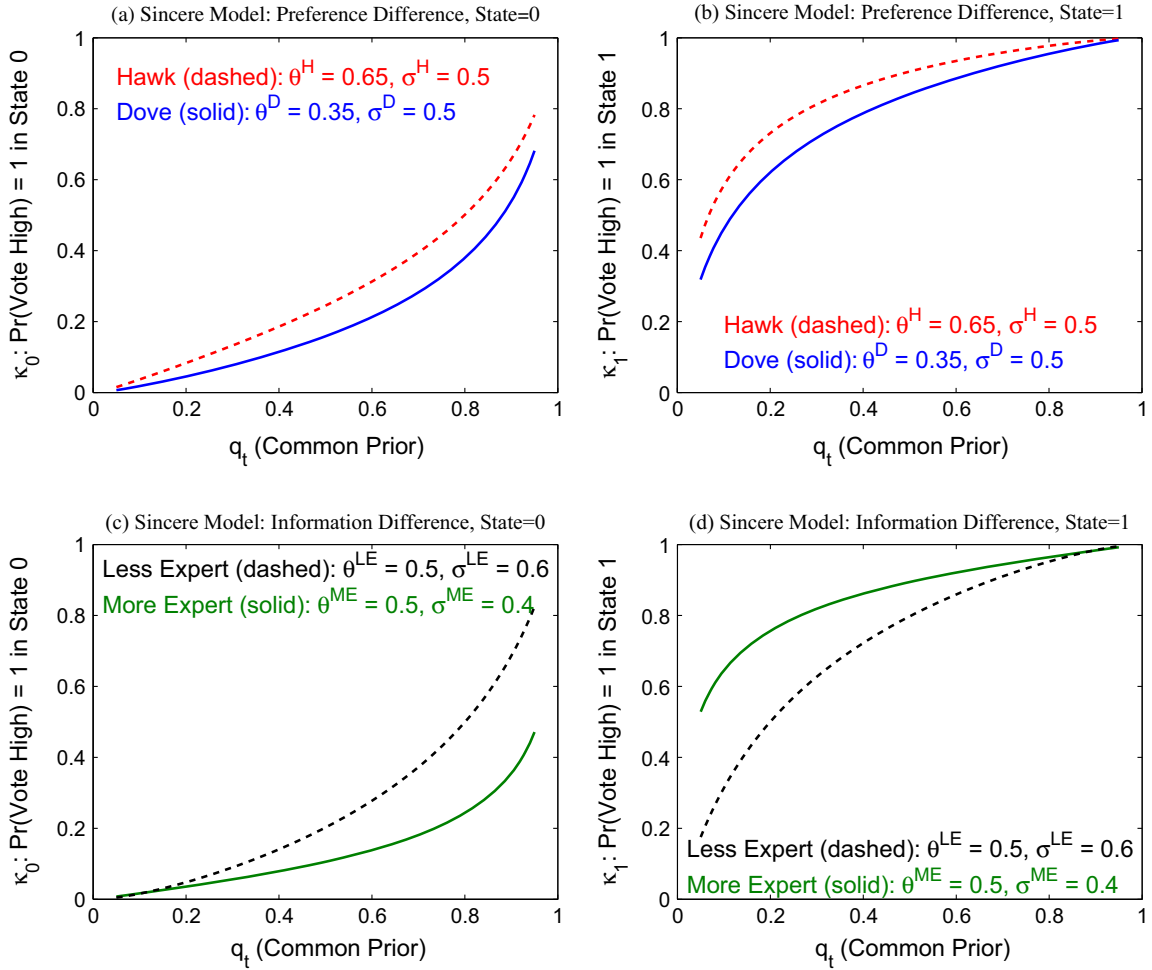


Fig. 1. Distinguishing information and preferences: Sincere voting. *Notes:* This figure shows the theoretical probability, conditional on the unknown state of the economy, that a member votes for the high interest rate ($\Pr(v_{it} = 1)$). These probabilities correspond to κ_0 and κ_1 in the model. The figures plot these probabilities as a function of the prior belief that the economy is in an inflationary state (q_t) and for different values of preferences and expertise. Figures a and b show that more hawkish individuals (higher θ), given a fixed value of expertise (σ), are more likely to vote for the high rate across both inflationary states ($\kappa_0^H \geq \kappa_0^D$ and $\kappa_1^H \geq \kappa_1^D$). Figures c and d show that expertise has a different effect on the probability that a member votes high in each state. Individuals with greater expertise, given a fixed value of bias, are more likely to vote for a rate that matches the inflationary state such that the more expert member is less likely to vote for the high rate in the low state and more likely to vote for the high rate in the high state ($\kappa_0^{ME} \leq \kappa_0^{LE}$ but $\kappa_1^{ME} \geq \kappa_1^{LE}$).

take the agenda to be the two rates on which the survey places the highest average probabilities.¹⁰ Fig. 2 shows the voting data for MPC members classified by whether they are internal or external. This figure illustrates that both internal and external members vote high ($v_{it} = 1$) and low ($v_{it} = 0$) across the whole sample.

5.2. Proxies for the common prior

The first proxy q_t^R comes directly from the Reuter's survey: q_t^R is the average probability the survey placed on the higher rate in the agenda over the total average probability placed on both rates in the agenda. The second proxy q_t^M comes from the cross-section of prices for short sterling futures options the first day (Wednesday) of the MPC meeting.¹¹ This data aggregates the opinions of a large number of agents (all traders in the sterling options market) and, in contrast to the Reuters data, these opinions are backed by real money and so potentially less subjective and manipulable. Short sterling futures contracts are effectively an option on 3 month LIBOR. The Bank of England computes the expected value of 3 month

¹⁰ We confirm that the unanimous decision reached by the MPC is one of the interest rates on which the market puts highest probability, which is itself an important test of the quality of the Reuters survey.

¹¹ Full details about these data, as well as the data itself, are provided by the Bank of England (see Bank of England, 2011; Lynch and Panigirtzoglou, 2008). There are four periods in which options price data are missing due to thin or illiquid short-sterling options markets.

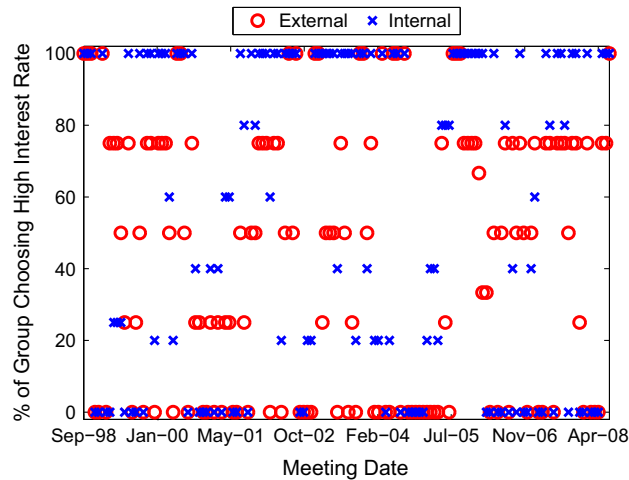


Fig. 2. Percentage of Internal and External Members who choose the higher interest rate. *Notes:* This figure shows the percentage of internal (blue X) and external (red dot) members within a given MPC meeting who vote for the higher interest rate calculated using the empirical data corresponding to the model variable v_{it} . Both types of member regularly vote for high and low interest rates. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this article.)

LIBOR consistent with a risk neutral trader being willing to hold the option at each observed price. This yields a distribution over risk-neutral traders' beliefs on 3 month LIBOR. The Bank then publishes the 0.05, 0.15, ..., 0.95 percentiles of this CDF. We subtract the actual value of LIBOR on the Wednesday of the MPC meeting (before the decision is made on the Thursday) to express the CDF in terms of traders' beliefs on changes in 3 month LIBOR. Since base rate changes are made in discrete 25 point movements while traders' beliefs are continuous, we consider beliefs that lie within 12.5 basis points on either side of the corresponding change to be beliefs associated with that change being more likely. As we only observe certain percentiles of the cdf, we linearly interpolate between the two percentiles in which a rate change falls.

Both q_t^R and q_t^M have weaknesses, but perhaps the biggest is that they predict the outcome of MPC voting rather than the realized inflationary shock. Consider however an observer of the committee who holds the prior belief q_t that the economy is in the inflationary state. Since the probability that each member votes high is increasing in q_t , then the observer's prediction on the MPC outcome will also be increasing in q_t . So while q_t^R and q_t^M do not directly measure q_t , at least their rank correlation should be high. Also important to emphasize is that the relationship between q_t and the proxies is estimated, not assumed.

6. Results

This section presents the main results from the estimation exercise. Ideally one would include 27 individual dummy variables in the first-stage and recover θ_i and σ_i for each voter, but there are too few observations per member to allow for this approach. Instead, the paper takes D_i to be a dummy variable that divides MPC voters into two groups A and B, and recovers and compares group-level parameters θ^A , σ^A , θ^B , and σ^B .

6.1. Differences between internal and external members

In the baseline specification, $D_i=1$ ($D_i=0$) if member i is an internal (external) member. To control for members' potentially having different disutilities from errors in states 0 and 1 depending on the agenda, we include in Z_t a dummy variable indicating if meeting t had at least one choice on the agenda to hike interest rates—the most common such meetings are those that have a choice of no change and a choice of raising by 25 basis points.

Table 2 contains point estimates and p-values for the first stage parameters in the baseline specification under the "Internal Baseline" heading. The main point of interest is the large and highly significant relationship between the prior q_t and the Reuter's proxy q_t^R —see coefficient α_1 in the table—while the relationship between q_t and q_t^M is small and insignificant. This indicates that the Reuter's survey data is a good predictor of the prior to which market price data adds little. As indicated by the positive and significant estimates for β_2 and γ_2 , higher values of q_t^R are also associated with a higher probability that members vote high in both states of the world. However, higher values of q_t^M are only associated with members voting high more often in state 0.

Table 3 presents the results of the second stage using the "Internal Baseline" specification. The σ estimates, which are invariant to sincere or strategic behavior, indicate that internal and external members form precise (though not perfect) private assessments of economic conditions. In Section 7 we provide a measure for how much the private signal improves individual decision making. For now we simply point out that heterogeneity in views is an important driver of heterogeneity

Table 2
First stage estimates.

| Regressor Coefficients | (1) Internal | | (2) Insider | | (3) Internal | | (4) Internal | |
|--------------------------------------|-----------------|---------|----------------|---------|-----------------|---------|-----------------|---------|
| | baseline | | baseline | | alternative 1 | | alternative 2 | |
| | Coeff. | p-value | Coeff. | p-value | Coeff. | p-value | Coeff. | p-value |
| α_0 (constant) | -2.35 | 0.01 | -3.41 | 0.00 | -2.50 | 0.01 | -2.43 | 0.01 |
| α_1 (Reuter's) | 5.75 | 0.00 | 5.54 | 0.00 | 5.23 | 0.00 | 4.83 | 0.00 |
| α_1 (market) | -0.55 | 0.39 | 1.85 | 0.18 | 0.31 | 0.44 | 0.47 | 0.40 |
| β_0 (constant) | -6.89 | 0.00 | -5.01 | 0.00 | -6.18 | 0.00 | -6.14 | 0.00 |
| β_1 (D_i , group dummy) | -1.02 | 0.26 | -2.01 | 0.14 | -0.92 | 0.28 | -1.29 | 0.21 |
| β_2 (Reuter's) | 2.18 | 0.02 | 2.14 | 0.05 | 3.38 | 0.01 | 3.63 | 0.00 |
| β_3 (market) | 4.77 | 0.01 | 2.6 | 0.06 | 3.61 | 0.05 | 3.21 | 0.05 |
| β_4 (Reuter's $\times D_i$) | 2.58 | 0.02 | 1.9 | 0.11 | 2.05 | 0.07 | 1.97 | 0.05 |
| β_5 (market $\times D_i$) | 0.48 | 0.43 | 2.53 | 0.2 | 0.71 | 0.40 | 1.36 | 0.30 |
| β_{61} (hike) | 1.69 | 0.00 | 0.83 | 0.11 | 1.14 | 0.07 | 1.00 | 0.10 |
| β_{62} (Status quo H) | | | | | -0.63 | 0.11 | -0.75 | 0.08 |
| β_{63} (Reuter's \times IR) | | | | | | | 0.73 | 0.06 |
| γ_0 (constant) | -1.03 | 0.09 | -0.64 | 0.17 | -2.26 | 0.01 | -2.09 | 0.01 |
| γ_1 (D_i , group dummy) | -4.1 | 0.00 | -1.74 | 0.13 | -3.22 | 0.00 | -3.32 | 0.00 |
| γ_2 (Reuter's) | 2.54 | 0.00 | 2.1 | 0.00 | 2.82 | 0.00 | 2.77 | 0.00 |
| γ_3 (market) | -0.33 | 0.40 | 0.29 | 0.39 | -0.82 | 0.26 | -1.13 | 0.20 |
| γ_4 (Reuter's $\times D_i$) | 2.74 | 0.02 | 6.13 | 0.00 | 3.91 | 0.00 | 4.39 | 0.00 |
| γ_5 (market $\times D_i$) | 8.57 | 0.00 | 1.44 | 0.33 | 5.40 | 0.02 | 5.2 | 0.01 |
| γ_{61} (hike) | 1.46 | 0.00 | 0.55 | 0.08 | 2.68 | 0.00 | 2.7 | 0.00 |
| γ_{62} (Status quo H) | | | | | 1.10 | 0.00 | 1.16 | 0.00 |
| γ_{63} (Reuter's \times IR) | | | | | | | -0.21 | 0.25 |

Notes: This table shows the results of the first stage estimation of (6). Each column represents a different specification. The first is the baseline results with members split according to whether they are internal or external. The second instead uses the split of insiders and outsider, but the baseline specification in terms of other regressors. The last two specifications again use the internal–external split, but consider additional covariates in the first stage regression. Each row in each column contains the coefficient estimates (first) and significance is reported using p -values (second).

Table 3
Baseline estimates of structural parameters.

| Structural Parameter | (1) Internal | (2) External | (3) Difference |
|-------------------------|-----------------|-----------------|-------------------|
| σ | 0.39 | 0.54 | -0.15 |
| 95% Range | (0.35 - 0.48) | (0.45 - 0.7) | (-0.3 - -0.02) |
| $\theta(SIN)$ | 0.65 | 0.34 | 0.30 |
| 95% Range | (0.52 - 0.76) | (0.26 - 0.44) | (0.18 - 0.4) |
| $\theta(STR)$ | 0.54 | 0.51 | 0.02 |
| 95% Range | (0.43 - 0.48) | (0.4 - 0.62) | (0.02 - 0.05) |

Notes: This table shows the structural estimates for internal (column 1) and external (column 2) members, as well as the difference between them (column 3). The rows report the estimates for the precision parameter (σ), as well as preferences under sincere voting ($\theta(SIN)$) and strategic voting ($\theta(STR)$). 95% confidence intervals are reported below each point estimate.

in observed votes. If two members share the same preferences and expertise, they will still in general receive different signals that may lead them to vote for different rates. It is notable that differences in private assessments persist even after the committee has met and discussed current conditions at length. This means that voting is an important mechanism to fully incorporate individual members' views in the final decision, suggesting that a committee in which all members have voting rights has advantages over a single decision maker with an advisory committee. Finally, as can be seen from the difference entry, internal members are estimated to receive more precise private signals than externals.

Extracting the θ parameters requires one to specify whether voting is sincere or strategic. In the sincere case, we estimate a preference difference that is large in magnitude and significant, with external members systematically more dovish than internals. This finding is in line with the existing literature discussed above. With strategic voting, however, external and internal members are estimated to have nearly identical preferences. Although the difference remains significant, the point estimates are very close together—the difference falls from 0.3 under sincere voting to 0.02 if voting is strategic.

These results indicate that members of the committee differ along a dimension that corresponds precisely to a variable that the government (or, more generally, the committee designer) controls directly—whether to appoint external members.

The estimates show that this institutional design feature has measurable consequences that affect decision making (we will quantify the size of the effects below). Unfortunately the results do not determine whether internal and external members behave differently because they have fundamentally different characteristics or because there is something about serving in an internal or external capacity that changes preferences and belief formation. It could be that internal members become more inflation-averse, or develop greater expertise, via their experience in central banks, and that they take this with them onto the MPC. Alternatively, it could be driven by the nature of their position. For instance, internals might have more expertise because of more direct control of the work streams (as the internal members are also senior management of the Bank), while external members, who typically serve as part-time monetary policymakers, are hindered because they have less time to spend getting on top of the large amounts of analysis that is provided.

One way to shed some light on the relevance of these channels is to use the fact that some newly appointed internals have come from careers in central banking (especially within the Bank of England), while others have come from other backgrounds, so that their tenure on the MPC coincides with their first central bank job. Analogously, one external member is a former Bank of England central banker. We redefine our internal group to instead be a group of central bank insiders by setting $D_i=1$ if member i had previously worked in the Bank of England or another central bank prior to taking their position on the MPC. The insider group pools insider-externals (former central bankers serving as externals) and insider-internals (career central bankers) while the outsider group consists entirely of people who are coming to the MPC from something other than central banking. As can be seen in Table 1, this converts one external members (Charles Goodhart) into an insider and about half the internal members (such as Charlie Bean) to outsiders.

The second column of Table 2 presents the first stage estimates for this specification and Fig. 3 shows the estimate, and 95% confidence interval, of the difference in parameters between insider and outsider members recovered from the second stage (internal and external member differences are shown for ease of comparison). Given that the groups are closely overlapping, it is no surprise that comparisons between insiders and outsiders in terms of structural parameters are similar to those between internal and external members. Of more interest is the difference-in-differences. Insiders seem to be even more hawkish relative to their outsider colleagues than internals are relative to externals. This can only be explained by the fact that insider-externals are more hawkish than outsider-internals, which suggests that it is the experience of central banking rather than the institutional responsibilities of being an internal that lead to hawkishness.¹² On the expertise side, the two alternative ways of splitting the data lead to almost exactly the same difference meaning that outsider-internals have as much expertise as insider-externals. If there is increased expertise from prior experience of central banking, this is offset by being an external member. This might be because it is a part-time position, or because these members have less control over the development of work and less information about other areas of the Bank's business such as financial stability.

6.2. Sincere versus strategic model

Given the difference in structural parameters between the sincere and strategic models, one question that arises is: which model better explains the data? To answer it we conduct a test relying on out-of-sample prediction. That is, we implement our estimation procedure on a subsample of the data, and then use the point estimates under the sincere and strategic specifications to generate predictions about the votes associated with the meetings not used in the estimation stage. We use four splits of the data, each one with the out-of-sample data corresponding to roughly one third of the meetings, to conduct the analysis.

The baseline test is to compute the mean absolute value of the difference between observed individual votes and the predicted probability of individual votes' being high given the observed meeting and individual characteristics. In each split, the sincere model performs marginally better; the mean absolute deviation under the sincere model is between 1.4 and 3.3 percentage points lower than under the strategic model.¹³ Thus, in the rest of the paper we focus on the estimates generated under the sincere model and include the ones for the strategic specification in the online appendix.

6.3. Differences between other groups

We also examine various other splits of the MPC besides internal and external. These include splits based on the members' career background prior to joining the committee (whether members worked in the private sector, or were academics), based on whether members hold a PhD, and based on whether members were older than the median age of new MPC members, 49, when joining the MPC.¹⁴ To save on space, we simply discuss the results and present the details in the online appendix. In all cases members' signals are estimated to be drawn from precise distributions similar to the results

¹² It might be that long periods of work in central banks make economists more hawkish, or that those who are inflation-averse self-select into careers in central banking (or thrive in them to the point of becoming expert enough to be selected to the MPC).

¹³ An alternative test discussed in the online appendix relies on individual vote classification. A model predicts a high (low) vote if the predicted probability of voting high is above (below) some cutoff value between 0 and 1. The main finding is that for some values of the cutoff, the prediction error of the two models is nearly identical, while for others the sincere model performs better (by around 3–5 percentage points). In contrast, for no split of the data do we find cutoffs for which the strategic model meaningfully outperforms the sincere model.

¹⁴ A table in the online appendix contains the classification of each member according to these criteria.

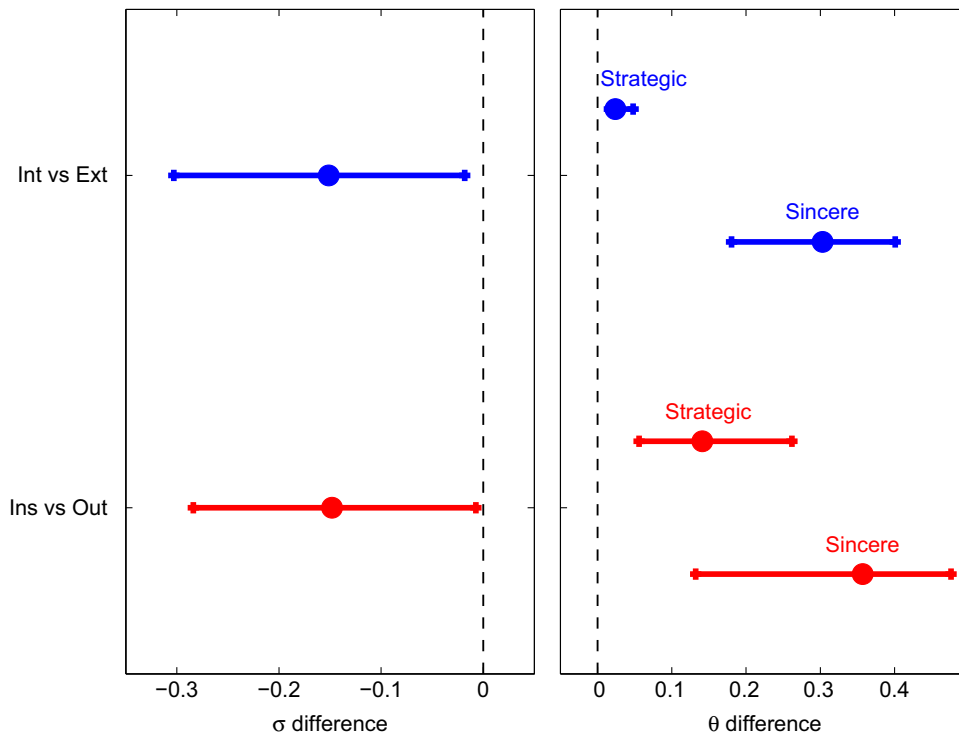


Fig. 3. Internal vs External or Insider vs Outsider? *Notes:* This figure shows, for expertise estimates (left panel) and preference estimates (right panel), the estimated difference between internal and external members (top lines of the figure labeled “Int vs Ext”), and between insiders and outsiders (bottom lines of the figure labeled “Ins vs Out”). The dot represents the mean difference for each parameter, while the 95% confidence interval is presented as the line around the dot. The estimated expertise gaps are almost identical for internal and insider members. Insider members are somewhat more hawkish relative to outsiders than internal are to external members.

for internal and external members. However, there are no statistically significant differences between the different splits in terms of preference or expertise parameters. This suggests that changing the mix of committee members in terms of career background, education, or age is unlikely to result in substantively different decision making.

6.4. Robustness of the baseline results

In order to test the robustness of the baseline results, we also examine alternative meeting controls in the first stage regression. Here we discuss two alternative specifications though we have tried many others with no change in the results. In the first we introduce a second agenda indicator variable to capture whether the high interest rate also corresponds to the status quo decision of no change in rates. Along with the first agenda indicator, whether a hike is on the agenda, this second agenda variable is indeed a significant predictor of individual votes (see estimates of β_{61} , β_{62} , γ_{61} , and γ_{62} in the third column of Table 2), but the values of the structural parameters and relevant differences are unchanged.

In the second alternative we attempt to address a concern about our q_t^R proxy, the one with the most predictive power in the first stage regressions. In every meeting, members have access to information that is also available to the entire market, as well as proprietary information from within the Bank that market participants do not have. The former information is reflected in our proxies, while the latter is not. If the Bank’s information is the main driver of voters’ beliefs, our estimation approach might be problematic. In meetings that coincide with the preparation of a quarterly Inflation Report (IR), voters have a particularly large amount of information available before voting (updated staff forecasts of macro variables) that the market does not see until after the meeting. So, if Bank information is the main source of information from which voters derive their prior beliefs, one would expect our proxies to be less correlated with voting high in IR months. To test this idea, we introduce an interaction between q_t^R and an IR month dummy into the κ_{10} and κ_{11} terms. In fact the coefficient on the interaction term is estimated to be *positive* and significant in the κ_{10} equation—meaning that economists’ predictions on the MPC decision are even more correlated with votes in the low state in IR months—while the equivalent coefficient in the κ_{11} equation is estimated negative but insignificant. Regardless, the values of the structural differences are unchanged.

To conclude our robustness exercises, we take medians of the structural parameter distributions (rather than means as described in Section 4.2), and also take means across members (rather than across time). These are done using the first stage estimates from the “Internal Baseline” specification. The qualitative features of the second stage results are identical.

Overall, our results highlight three interesting messages. First, the main driver of voting differences between most splits of the MPC is differences in beliefs on economic conditions, not differences in preferences or biases. This is notable given that preference differences (typically) receive more attention in the monetary literature. Second, for some splits considered, members are ex-ante identical in terms of both preferences and expertise. This means, for example, that replacing older members with younger members or vice versa might have little effect on decision making. Third, as already discussed, the one split on which members *do* differ in terms of expertise and preferences corresponds to a committee design feature—the appointment, or not, of external members. We now quantify the extent to which the committee aggregates private information, as well as the effect of adding external members to the MPC on decision-making quality.

7. The effect of different committee sizes

Our finding that members draw imperfect private signals on economic conditions suggests that the MPC adds value compared to individual decision making by aggregating dispersed knowledge. The main goal of this section is to quantify this gain with counterfactual simulations. A secondary goal is to assess whether externals' having different preferences from internals can improve decision making in spite of their having lower signal precisions.

The measure of decision-making quality we adopt is the unconditional probability the decision equals the state. This corresponds to the utility function of a committee designer with a neutral bias of $\theta=0.5$, which appears to correspond to the Bank's explicitly stated preferences. In terms of the notation introduced in Section 4, for a committee of size N (odd), this probability is

$$q_t = \sum_{m=(N+1)/2}^N \binom{N}{m} \kappa_{1t}^m (1-\kappa_{1t})^{N-m} + (1-q_t) \sum_{m=(N+1)/2}^N \binom{N}{m} \kappa_{0t}^{N-m} (1-\kappa_{0t})^m \quad (11)$$

where κ_{ω_t} is the probability a committee member votes high in state ω_t given q_t .¹⁵ We compute the κ terms for different values of the structural parameters θ and σ under the sincere voting assumption. The results for strategic voting are very similar.

An important caveat should be heeded when interpreting the counterfactuals. We cannot measure the effect that a member's presence had on other members' voting behavior through deliberation or any other interactions. It could be that such peer effects are important, for example the presence of a particular member might affect other members' beliefs. If this were so, the following simulations should be interpreted as removing the voting rights from various subsets of members, but allowing them to participate in the rest of the committee activities as is the case on advisory committees.

To examine the effect of committee size on performance, we consider a committee whose members share the preferences and expertise of the average internal member as reported in Table 3 ($\theta=0.6$ and $\sigma=0.39$) and which faces an economy with maximum uncertainty ($q_t=0.5$). Fig. 4a presents the results, with the shaded area corresponding to the 95% confidence intervals. Despite high uncertainty, a committee of internals has a high probability of getting the decision correct. Even a single internal gets the decision right about 90% of the time (though the range of estimates is 84%–93%). As a comparison, an internal member who drew no private signal would always vote high, and so on average get the decision right in 50% of meetings.¹⁶ Private information improves the individual internal decision maker by about 40 percentage points in the most uncertain economic times. Of course, as uncertainty about the state of the economy declines, and the prior moves toward either zero or one, the value of private signals is reduced. At the same time, there is a significant improvement in decision quality as more members are added. For example, expanding the committee size from one to five members increases the probability of a correct decision by about nine percentage points.

Taken at face value, the results indicate a larger committee is always better. As a practical matter, though, larger committees also entail additional coordination and infrastructure costs. Some authors have also argued that increasing the number of members makes deliberation more difficult (Furnham, 1997), and increases the costs of free-riding in terms of information acquisition (Sibert, 2006). So, the marginal benefit of an additional member is also important to know. While we measure it to be positive and significant up to nine members, it is very small beyond five members. Intuitively speaking, internal members receive precise enough signals that most of the gains from information aggregation are realized with a small number of experts. While we do not wish to propose a theory of optimal committee size, our results are nevertheless striking. A group of five internal members over the meetings in our sample is predicted to make just a few errors. Interestingly, Napier and Gershensfeld (1999) argue that this size is optimal (or close to optimal) in most situations. However, as we see below, our finding depends crucially on the level of expertise of the appointed members.

In Fig. 4b, we simulate the performance of a committee of externals. Since they are estimated to have lower expertise than internals, this committee performs worse for all sizes. Still, a committee of nine members gets the decision right in 95% of meetings. More substantively, the marginal benefit of group size does not taper off as quickly and adding more experts can meaningfully improve decision making.

¹⁵ The committees we consider are homogenous with respect to the structural parameters, so we do not introduce an i index. We calculate the expression for odd-sized committees and use piecewise cubic interpolation between the resulting probabilities for our figures.

¹⁶ In a model with no private signals, a banker chooses the high rate if and only if $\theta_t \geq 1 - q_t$. So when $q_t < 1 - \theta_t$ the probability of a correct decision is $1 - q_t$ and otherwise is q_t .

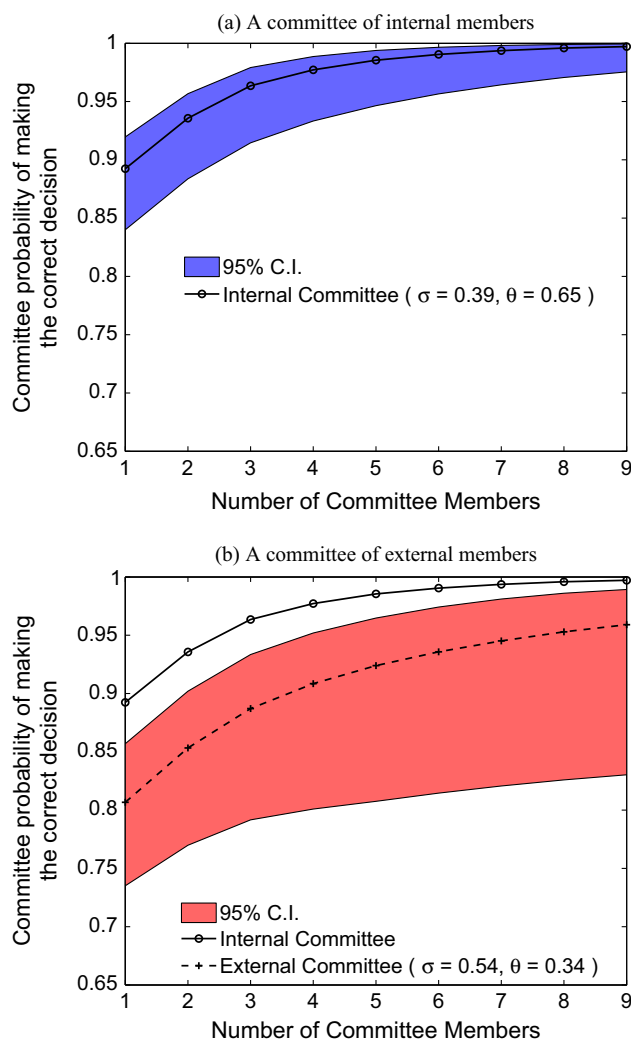


Fig. 4. The implied likelihood that different MPC Compositions get the right decision. *Notes:* This figure shows the probability that a committee makes the correct interest rate decision across different committee sizes for a fixed value of the prior ($q_t=0.5$). The probability is calculated under the assumption of sincere voting using the estimated structural parameters presented in Table 3. Figures a and b display the result for committees of internal and external members, respectively. The figure shows that internal members outperform external members in terms of the probability of making the correct decision across all committee sizes. Furthermore, the figure shows that there are gains to larger committees, though these taper off, particularly for internal committees, for levels above five members.

To push the point on lower-expertise committees further, Fig. 5 presents results for a group of neutral non-experts with $\theta=0.5$ and $\sigma=1$, whom one could imagine being members of the government or public with little background in monetary policy. This group would need around 30 members to match the performance of the five-member internal committee.

While internal committees outperform external ones due to better information, our last set of results analyzes whether they can improve decision making due to their different preferences. To examine this, we see how internals' performance relative to externals varies with the prior q_t across different committee sizes. Fig. 6a plots by how much internals outperform externals when $q_t=0.75$, and shows that, for high committee sizes, externals perform as well as internals. Intuitively, internal members are biased towards high rates, and when the prior favors high rates, externals are more likely to follow private signals that point towards low rates. But for exactly the reverse logic, internals outperform externals when $q_t=0.25$ by more than when $q_t=0.5$, as shown in Fig. 6b. In this sense, preference diversity can both improve and worsen performance from meeting to meeting. If q_t is roughly uniform, though, gains from preference diversity cancel the losses, so that externals still do worse overall.

In terms of real-world implications, our results show that combining four external members with five internal ones, as the Bank of England does, might not improve decision making much while also generating potential costs. Rather than assert that external members literally add no value, we believe our estimates show that the justification for their inclusion needs more careful thinking. For example, perhaps the presence of external members increases diversity in ways that our model does not capture.




Fig. 5. Probability a committee of non-experts gets the right decision. *Notes:* This figure shows the probability that a committee with non-expert members (i.e. $\sigma = 1$ and $\theta = 0.5$) makes the correct interest rate decision across different committee sizes for a fixed value of the prior ($q_t = 0.5$). The probability is calculated under the assumption of sincere voting. The figures shows that a non-expert committee would require around 30 members to exhibit a level of performance similar to a committee of five internal members (dashed line).

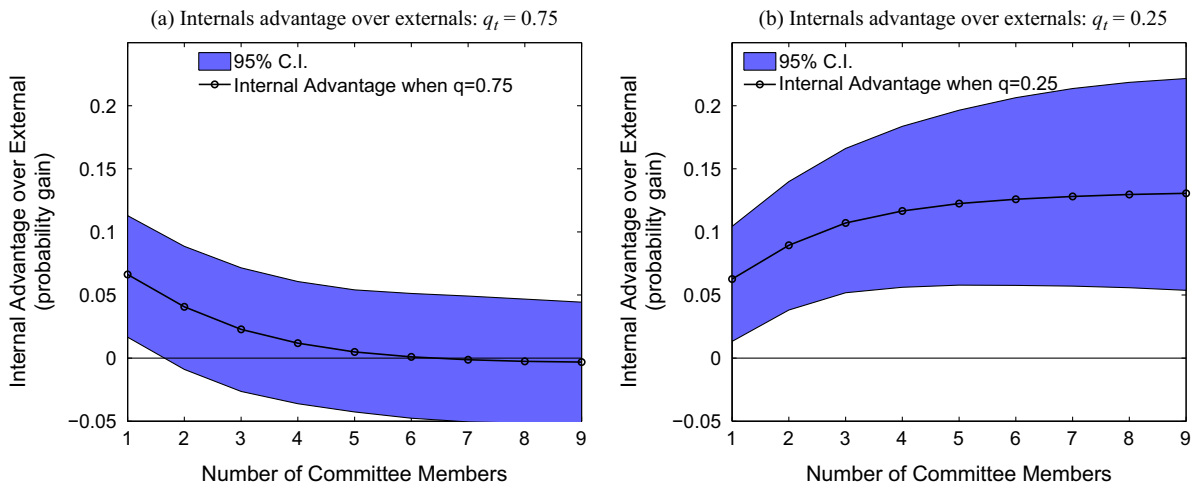


Fig. 6. The Gains from internal members. *Notes:* This figure shows the difference in the probability of making the correct interest rate decision between a committee of internal and external members across different sizes. The differences are calculated under the assumption of sincere voting and use the estimated structural parameters presented in Table 3. Figures a and b display the result when the prior is 0.75 and 0.25, respectively. The figure shows that given the estimated values of expertise, the preferences of external members are not sufficient to offset the preference bias of internal members. For instance, we observe that even for a relatively high value of the prior ($q_t = 0.75$ in Fig. a), when internal preferences may tilt this type of committees towards adopting higher rates, a committee of external members does not exhibit a better performance than a committee of internal members.

8. Conclusion

Taken together, our results give an empirically novel view of monetary policymaking by committee and address important issues in committee design. Reis (2013) argues that committees are preferred to individual decision makers for any of a four main virtues: (i) pooling of private information, (ii) providing a diversity of views which generates discussion of the evidence for and against different views, (iii) guarding against autocratic power, and (iv) making less volatile decisions. Our paper directly addresses the first of these, and shows that groups significantly outperform individuals due to information aggregation. Another insight is that adding dovish members (externals) to a group of hawkish ones (internals) does not necessarily improve average decision making quality through preference moderation.

The results and analysis can also shed light on the structure of other important committees like the Federal Open Market Committee (FOMC) and European Central Bank Governing Council (ECB). At the Bank of England, a small committee appears desirable because the marginal benefit of additional members due to information pooling declines rather quickly. From this perspective, the 19-member FOMC and 24-member ECB are potentially too big when one considers the costs of more

members. Of course, in economies with more regional heterogeneity like the US or euro area, drawing on a greater number of views may be more important for determining the right interest rate.

We find one type of appointee (internals) forms, on average, more precise assessments than another (externals). This raises the question of whether such differences exist on other committees with different member types, and whether the composition of these committees is optimal. For example, Reis (2013) argues that one justification for the inclusion of regional Fed Presidents on the FOMC is that they contribute new ideas. An approach similar to ours using policy preference data from the FOMC could help determine whether Presidents indeed have more (or less) expertise than the Fed Governors based in Washington DC. One argument in favor of Presidents having more expertise is that each has a reasonably large staff and budget which might encourage competition in the market for economic analysis (Goodfriend, 1999). The Fed Governors, by comparison, have to rely mostly on the analysis of the Board of Governors' staff, which is shared amongst FOMC members. At the Bank of England, it is the externals who only have a small staff and these are the group found to have less expertise.

One note of caution is that the correct decision might depend on multiple dimensions, some of which externals have more expertise on than internals. A model that explicitly built in such diversity would be a natural extension of our work. Moreover, the committee structure may affect deliberation and therefore the evolution each member's views. We leave these issues to future research. Nevertheless, the paper is, to our knowledge, the first to decompose voting heterogeneity on a monetary policy committee into distinct preference and signal components, and as such provides potentially valuable facts for moving the debate on monetary policymaking by committee forward.

Acknowledgments

We are very grateful to the editor Ricardo Reis and an anonymous referee for their insightful suggestions. We would also like to thank Tim Besley, Alan Blinder, Pablo Casas, Fabio Canova, Francesco Caselli, Antonio Ciccone, Greg Crawford, Thomas Cunningham, Wouter den Haan, Nathan Foley-Fisher, Francesco Giavazzi, Charles Goodhart, Matias Iaryczower, Kosuke Imai, Clare Leaver, Gilat Levy, Massimo Motta, Andrew Oswald, Morten Ravn, Karl Schlag, Kevin Sheedy, Jón Steinsson, Thijs van Rens, and members of Kosuke Imai's reading group for insightful comments on this and earlier versions of this work. We have also benefited from the suggestions of seminar participants at the Australia National University, the Bank of England, London School of Economics, Monash University, Queen Mary University, University College Dublin, Universitat Pompeu Fabra, University of Strathclyde, and University of Warwick.

Appendix A. Supplementary material

Supplementary data associated with this article can be found in the online version at <http://dx.doi.org/10.1016/j.jmoneco.2014.06.004>.

References

- Austen-Smith, D., Banks, J.S., 1996. Information aggregation, rationality, and the Condorcet Jury Theorem. *Am. Polit. Sci. Rev.* 90 (1), 34–45.
- Bank of England, 2010. (<http://www.bankofengland.co.uk/monetarypolicy/overview.htm>), last accessed 29 July 2012.
- Bank of England, 2011. Notes on the Bank of England option-implied probability density functions. (<http://www.bankofengland.co.uk/statistics/Pages/implicitpdfs/default.aspx>) (last accessed 12 March 2011).
- Bank of England Act, 1998. (<http://www.bankofengland.co.uk/about/legislation/1998act.pdf>) (last accessed 15 September 2010).
- Besley, T., Meads, N., Surico, P., 2008. Insiders versus outsiders in monetary policymaking. *Am. Econ. Rev.* 98 (May (2)), 218–223.
- Bhattacharjee, A., Holly, S., 2005. Inflation targeting, committee decision making and uncertainty: the case of the Bank of England's MPC. Discussion paper, Department of Applied Economics, University of Cambridge.
- Blinder, A.S., 2007. Monetary policy by committee: why and how?. *Eur. J. Polit. Econ.* 23.
- Blinder, A.S., Morgan, J., 2005. Are two heads better than one? Monetary policy by committee. *J. Money Credit Bank.* 37 (October (5)), 789–811.
- Clarida, R., Gali, J., Gertler, M., 1999. The science of monetary policy: a new Keynesian perspective. *J. Econ. Lit.* 37 (December (4)), 1661–1707.
- Duggan, J., Martinelli, C., 2001. A Bayesian model of voting in juries. *Games Econ. Behav.* 37 (November (2)), 259–294.
- Furnham, A., 1997. *Psychology of Behaviour at Work*. Psychology Press, Hove, East Sussex.
- Gali, J., 2008. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*. Princeton University Press, Princeton, USA and Oxford, UK.
- Gerlach-Kristen, P., 2003. Insiders and outsiders at the Bank of England. *Cent. Bank. XIV* (1), 96–102.
- Gerlach-Kristen, P., 2006. Monetary policy committees and interest rate setting. *Eur. Econ. Rev.* 50 (February (2)), 487–507.
- Gerling, K., Gruner, H.P., Kiel, A., Schulte, E., 2005. Information acquisition and decision making in committees: a survey. *Eur. J. Polit. Econ.* 21 (September (3)), 563–597.
- Goodfriend, M., 1999. The role of a regional bank in a system of central banks. *Carnegie-Rochester Conference Series on Public Policy* 51 (December (1)), 51–71.
- Hansen, S., McMahon, M., August 2013. First impressions matter: signalling as a source of policy dynamics. CEPR Discussion Papers 9607, C.E.P.R. Discussion Papers.
- Harris, M., Levine, P., Spencer, C., 2011. A decade of dissent: explaining the dissent voting behavior of Bank of England MPC members. *Publ. Choice* 146 (March (3)), 413–442.
- Hix, S., Hoyland, B., Vivyan, N., 2010. From doves to hawks: a spatial analysis of voting in the Monetary Policy Committee of the Bank of England. *Eur. J. Polit. Res.* 49 (6), 731–758.
- Iaryczower, M., Shum, M., 2012. The value of information in the court: get it right, keep it tight. *Am. Econ. Rev.* 102 (February (1)), 202–237.

- Imai, K., Tingley, D., 2012. A statistical method for empirical testing of competing theories. *Am. J. Polit. Sci.* 56 (January (1)), 218–236.
- King, G., Tomz, M., Wittenberg, J., 2000. Making the most of statistical analysis: improving interpretation and presentation. *Am. J. Polit. Sci.* 44, 341–355.
- Lombardelli, C., Proudman, J., Talbot, J., 2005. Committees versus individuals: an experimental analysis of monetary policy decision-making. *Int. J. Cent. Bank.* 1 (May (1)).
- Lynch, D., Panigirtzoglou, N., March 2008. Summary statistics of option-implied probability density functions and their properties. Bank of England Working Paper 345, Bank of England.
- Mihov, I., Sibert, A., 2006. Credibility and flexibility with independent monetary policy committees. *J. Money Credit Bank.* 38 (February (1)), 23–46.
- Napier, R.W., Gershenfeld, M.K., 1999. *Groups: Theory and Experience*. Houghton Mifflin Co, New York.
- Pesendorfer, M., Takahashi, Y., 2013. Testing for equilibrium multiplicity in dynamic Markov games. Unpublished Manuscript.
- Reis, R., 2013. Central bank design. *J. Econ. Perspect.* 27 (Fall (4)), 17–44.
- Riboni, A., Ruge-Murcia, F.J., 2010. Monetary policy by committee: consensus, chairman dominance, or simple majority? *Q. J. Econ.* 125 (February (1)), 363–416.
- Sibert, A., 2003. Monetary policy committees: individual and collective reputations. *Rev. Econ. Stud.* 70 (July(3)), 649–665.
- Sibert, A., 2006. Central banking by committee. *Int. Financ.* 9 (2), 145–168.
- Sims, C.A., 2010. Chapter 4—rational inattention and monetary economics. In: Friedman, B.M., Woodford, M. (Eds.), *Handbook of Monetary Economics*, vol. 3. Elsevier, Amsterdam, pp. 155–181.
- Spencer, C., March 2006. The dissent voting behaviour of Bank of England MPC members. Department of Economics Discussion Papers 0306, Department of Economics, University of Surrey.