

What Do Outside Experts Bring To A Committee?
Evidence From The Bank of England

Stephen Hansen and Michael McMahon

No 946

WARWICK ECONOMIC RESEARCH PAPERS

DEPARTMENT OF ECONOMICS

THE UNIVERSITY OF
WARWICK

What Do Outside Experts Bring To A Committee? Evidence From The Bank of England*

Stephen Hansen
Universitat Pompeu Fabra
stephen.hansen@upf.edu

Michael McMahon
University of Warwick
m.mcmahon@warwick.ac.uk

October 5, 2010

Abstract

We test whether outside experts have information not available to insiders by using the voting record of the Bank of England's Monetary Policy Committee. Members with more private information should vote more often against conventional wisdom, which we measure as the average belief of market economists about future interest rates. We find evidence that external members indeed have information not available to internals, but also use a quasi-natural experiment to show they may exaggerate their expertise to obtain reappointment. This implies that an optimal committee, even outside monetary policy, should potentially include outsiders, but needs to manage career concerns.

Keywords: Expert Behavior, Committees, Monetary Policy

JEL Codes: D70, E52

*We would like to thank Tim Besley, Pablo Casas, Francesco Caselli, Antonio Ciccone, Greg Crawford, Thomas Cunningham, Francesco Giavazzi, Charles Goodhart, Clare Leaver, Gilat Levy, Massimo Motta, Andrew Oswald, Morten Ravn, Karl Schlag, Kevin Sheedy, Jón Steinsson, and Thijs van Rens for insightful comments. 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, University College Dublin, Universitat Pompeu Fabra, University of Strathclyde, and University of Warwick. We are also very grateful to Grainne Gilmore and her colleagues at the Times of London Newspaper for making available the data on the Times MPC.

1 Introduction

A large fraction of countries now set monetary policy via independent committees of experts.¹ While there are substantial literatures on why central bank independence is important for macroeconomic stability and committees outperform individuals, how to best design a monetary policy committee is a more open question. One important unresolved issue is whether committees should include outside experts who are not full-time employees of the central bank. For example, the Federal Reserve Bank in the US and the Riksbank in Sweden use committees composed solely of bank employees, while the Bank of England’s Monetary Policy Committee (MPC) is made up of five Bank executives (so-called *internal* members) and four outside experts (so-called *external* members). According to the Bank of England (2010a), the purpose of external appointments is to “ensure that the MPC benefits from thinking and expertise in addition to that gained inside the Bank of England.” This paper aims to assess the extent to which the inclusion of external members achieves this goal.

In order to measure a member’s private information, we first construct a model in which the realization of an unknown state variable determines whether the correct decision is a low or high interest rate. The model predicts a member’s vote as a function of (1) the prior probability that the high state occurred; (2) the member’s monetary policy philosophy or preferences; and (3) the precision of the private information that a member receives about the state. We show that if a member has a more dovish monetary policy than another (i.e., he believes high inflation to be less likely), then he will tend to vote for lower rates no matter the value of the prior. In contrast, if a member has more private information than another, then he will tend to vote for the high rate more often when the prior favors the low state and for the low rate more often when the prior favors the high state; informally speaking, he is less likely to “follow the crowd”.

To use the model to distinguish empirically a member’s philosophy from his expertise, we construct a proxy for the prior distribution in each period with survey data that asks financial institutions in the City of London to state how likely they consider different interest rate movements. Our claim is that the average view of the market should accurately reflect all public sources of information about economic shocks. We then use this proxy variable along with members’ observed votes to structurally estimate our model’s philosophy and expertise parameters for internal and external members.²

¹Pollard (2004) reports that ninety percent of eighty-eight surveyed central banks use committees to decide interest rates.

²We treat internal and external members as homogeneous groups; nevertheless, we report Monte Carlo results that show that our estimator accurately measures the difference in group means when there is within-group heterogeneity.

Our baseline result is that external members have significantly more precise private information than internals.³ Our interpretation is that external members have different—as opposed to more—information than internal members since internals’ view on economic conditions is more likely reflected in the market’s beliefs (i.e. more likely to be public information) than externals’. This supports the idea that external members bring a unique perspective to the MPC that helps broaden the expertise on which UK monetary policy is based.

We then examine whether changes in expertise drive the voting dynamics identified in Hansen and McMahon (2008) which notes growing conflict in terms of average vote levels between internals and externals. Our finding is that there are no significant shifts in expertise levels over time. However, there is a large philosophical shift towards dovishness among experienced external members; this is a result of independent behavioral interest.

While our baseline results point towards externals having substantial private information, a well-established career concerns literature emphasizes that experts will exaggerate their private information if they care about acquiring a reputation for expertise.⁴ The paper therefore examines whether a particular type of reputational incentive might drive the baseline finding: the desire for reappointment. We make use of a quasi-natural experiment that exogenously reduced the probability that external members would be reappointed during a particular era, and find that inexperienced external members serving in this era receive private signals with an estimated standard deviation over three times larger than other inexperienced externals.⁵ Further evidence of career concerns is that non-academic externals show more precise private information at the beginning of their tenure than at the end. These results indicate that at least some of the estimated informational advantage of external members may be an artifact of career concerns. Career concerns also appear to influence the degree of philosophical divergence: external members’ philosophies shift less when reappointment is more salient and when they are non-academics.

We view these results as being of interest in contexts beyond monetary policy because the debate about whether to include outsiders on committees is currently occurring in several areas, and the reasoning of those in favor is remarkably similar to that of the Bank.

³We also find that external have a significantly more dovish philosophy than internal members. This confirms the findings of an existing literature that examines internal-external differences on the MPC (Gerlach-Kristen 2003, Bhattacharjee and Holly 2005, Spencer 2006, Besley, Meads, and Surico 2008, Harris and Spencer 2008) in terms of preferences and finds external members tend to be more dovish. None of these papers consider heterogeneity in private information

⁴Trueman (1994), Prendergast and Stole (1996), and Levy (2004) establish this result for an individual decision maker, while Levy (2007) and Visser and Swank (2007) explore the idea in the context of committees.

⁵We focus on inexperienced externals because the career concerns literature predicts that reputational incentives are strongest for new agents.

For example, the Food and Drug Administration in the US uses advisory committees formed largely of outside experts to vote on whether new drugs should be allowed to market. According to the FDA, “an academician or a practitioner representative uses his or her expert knowledge to provide state-of-the-art advice on scientific issues under deliberation.” In the field of corporate governance, some argue that the inclusion of non-executive directors (outside directors without managerial responsibility) on boards should be adopted as a basic tenet of corporate governance. Higgs (2003) writes that “a major contribution of the non-executive director is to bring wider experience and a fresh perspective to the boardroom.”⁶ In both of these cases, the informational advantages of including outsiders is clearly at the heart of the argument.

Since our empirical setting features internal and external experts repeatedly taking observable policy decisions on the same technical issue, it is also useful for addressing the general question of what outsiders bring to a committee since data is not available in the above-mentioned contexts.⁷ We believe ours is the first paper to show that externals have non-trivial private information, and this supports their inclusion in decision-making bodies. At the same time, our results on career concerns show that agency problems can also arise, meaning that committee designers should take care in arranging the institutional details of externals’ participation. Nevertheless, we view our paper as providing a qualified “yes” to the question of whether committees can benefit from a mixed composition.

2 The Monetary Policy Committee

This section describes the institutional details of the MPC. Until 1997 the Chancellor of the Exchequer (the government official in charge of the Treasury) had sole responsibility for setting interest rates in the UK. One of Gordon Brown’s first actions on becoming Chancellor in the government of Tony Blair was to set up an independent committee, the MPC, for setting interest rates. 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,

⁶Higgs (2003) goes on to write that

Ensuring that the board as a whole has an appropriate mix of skills and experience is essential for it to be an effective decision-making body. There is no standard board or standard non-executive director, nor can there be. It is the range of skills and attributes acquired through a diversity of experiences and backgrounds that combine to create a cohesive and effective board.

⁷For example, individual level votes from corporate board meetings are typically kept secret.

the committee seeks to achieve a target inflation rate of 2%,⁸ based on the Consumer Price Index. If inflation is greater than 3% or less than 1%, the Governor of the Bank of England must write an open letter to the Chancellor explaining why. The inflation target is symmetric; missing the target in either direction is treated with equal concern.

The MPC first convened on 6 June 1997, and has met every month since. Throughout the paper we analyze the MPC voting records between June 1997 and March 2009,⁹ the point at which the main focus of the decision (temporarily) shifted to asset purchase decisions related to quantitative easing. 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. We also gathered background information for each member from press releases associated with their appointment and from information provided to the Treasury Select Committee ahead of their confirmation from which we gather information on whether the member was an academic appointment to the MPC.

The MPC has nine members; five of these come from within the Bank of England: the Governor, two Deputy Governors, the Chief Economist, and the Executive Director for Markets. The Chancellor also appoints four members (subject to approval from the Treasury Select Committee) from outside the Bank. There are no restrictions on who can serve as an external member. Bar the governors, all members serve three year terms; the governors serve five year terms. When members' terms end, they can either be replaced or re-appointed. Table 1 lists the members that served on the MPC during our sample, whether they are internal or external, the period in which they served, and whether they joined the committee directly from an academic position. Our sample contains a total of 13 internal and 14 external members, with academics making up around one third of members.

Each member is independent in the sense that they do not represent any interest group or faction. The Bank encourages members to simply determine the rate of interest that they feel is most likely to achieve the inflation target,¹⁰ and majority vote determines

⁸This target changed from the RPIX to the CPI measure of inflation in January 2004, with a reduction in the inflation target from 2.5% to 2%.

⁹The data are available from the Bank of England (2010b). We use each regular MPC meeting in this period but we drop from the dataset the (unanimous) emergency meeting held after September 11th.

¹⁰According to the Bank of England (2010a)

Each member of the MPC has expertise in the field of economics and monetary policy. Members are not chosen to represent individual groups or areas. They are independent. Each member of the Committee has a vote to set interest rates at the level they believe is consistent with meeting the inflation target. The MPC's decision is made on the basis of one-person, one vote. It is not based on a consensus of opinion. It reflects the votes of each individual member of the Committee.

Table 1: MPC Members

Member	Type	Tenure	Academic
Howard Davies	INT	06/97 - 07/97	N
Edward George	INT	06/97 - 06/03	N
Mervyn King	INT	06/97 - 03/09	N
Ian Plenderleith	INT	06/97 - 05/02	N
David Clementi	INT	11/97 - 08/02	N
John Vickers	INT	06/98 - 11/00	Y
Charles Bean	INT	10/00 - 03/09	Y
Paul Tucker	INT	06/02 - 03/09	N
Andrew Large	INT	11/02 - 01/06	N
Rachel Lomax	INT	07/03 - 06/08	N
John Gieve	INT	01/06 - 03/09	N
Spencer Dale	INT	07/08 - 03/09	N
Paul Fisher	INT	03/09 - 03/09	N
Willem Buiters	EXT	06/97 - 05/00	Y
Charles Goodhart	EXT	06/97 - 05/00	Y
DeAnne Julius	EXT	11/97 - 05/01	N
Alan Budd	EXT	12/97 - 05/99	N
Sushil Wadhvani	EXT	06/99 - 05/02	N
Christopher Allsopp	EXT	06/00 - 05/03	Y
Stephen Nickell	EXT	06/00 - 05/06	Y
Kate Barker	EXT	06/01 - 03/09	N
Marian Bell	EXT	06/02 - 06/05	N
Richard Lambert	EXT	06/03 - 03/06	N
David Walton	EXT	07/05 - 06/06	N
Tim Besley	EXT	06/06 - 03/09	Y
David Blanchflower	EXT	06/06 - 03/09	Y
Andrew Sentance	EXT	10/06 - 03/09	N

the outcome. Consistent with its one-person one-vote philosophy, the MPC displays substantial dissent. 64% of the 142 meetings in the sample have at least one deviation from the committee majority. Figure 1 shows the level of interest rates that the MPC has implemented, the votes of each member around this, and highlights the periods of interest rate loosening.¹¹ Figure 2 shows how many votes were cast in each meeting in opposition to the final decision. Within the set of non-unanimous meetings, 5-4 and 6-3 decisions are not uncommon.




Figure 1: Votes and Decisions

The MPC meets on the first Wednesday and Thursday of each month. In the month between meetings, members receive numerous briefings from Bank staff and regular updates of economic indicators. On the Friday before MPC meetings, members gather for a half-day meeting in which they are given the latest analysis of economic and business trends. Then on Wednesday members discuss their views on several issues. The discussion continues on Thursday morning, when each member is given some time to summarize his or her views to the rest of the MPC and to suggest what vote they favor (although they can, if they wish, wait to hear the others views before committing to a vote (Lambert 2006)). This process begins with the Deputy Governor for monetary

¹¹The loosening cycle is defined as the period from the first cut in interest rates until the next increase in interest rate.




Figure 2: Deviations

policy and concludes with the Governor, but the order for the others is not fixed. To formally conclude the meeting, the Governor suggests an interest rate that he believes will command a majority. Each member then chooses whether to agree with the Governor's decision, or dissent and state an alternative interest rate. The MPC decision is announced at noon, and two weeks after each meeting, members' votes are published as part of otherwise unattributed minutes.

3 Model

We now lay out a model that captures the basic elements of MPC voting and that will serve as the basis of our empirical analysis.¹² In period t voter i must choose a vote $v_{it} \in \{0, 1\}$ where 0 represents the lower of two possible rate changes and 1 the higher.¹³ The restriction that a vote must take one of two values is not as restrictive as it might first appear since there are three distinct votes in only 7 of the 142 meetings in our sample

¹²We model voting along the lines of the jury literature. Early references are Austen-Smith and Banks (1996) and Feddersen and Pesendorfer (1998); a more recent survey is Gerling, Gruner, Kiel, and Schulte (2005).

¹³We do not attempt to model the two rates over which voting occurs; recent work on this issue of agenda setting on MPC's is found in Riboni and Ruge-Murcia (2010).

and in no meeting are there four or more distinct votes.

Member i 's payoffs over votes and states of the world is given by

$$u(v_{it}) = \begin{cases} 1 & \text{if } \omega_t - (1 - \theta) > 0 \text{ and } v_{it} = 1; \text{ or if } \omega_t - (1 - \theta) < 0 \text{ and } v_{it} = 0 \\ 0 & \text{if } \omega_t - (1 - \theta) > 0 \text{ and } v_{it} = 0; \text{ or if } \omega_t - (1 - \theta) < 0 \text{ and } v_{it} = 1. \end{cases} \quad (1)$$

Members' utility depends on $\omega_t - (1 - \theta)$ which we can think of as a measure of inflationary pressures in the economy. As the extent of inflationary pressures is greater (lesser), the member will derive utility by voting for the higher (lower) interest rate. Inflationary pressures depend on an unknown state variable $\omega_t \in \{0, 1\}$ that reflects economic conditions relevant to inflation such as demand shocks or, as in Gerlach-Kristen (2006), the output gap of the economy. We assume that members have a prior belief q_t that the high state prevails. One can think of q_t as reflecting the conventional wisdom or average view of the market about the probability the economy is in the high state. Members also have their own assessment of economic conditions that we model as the realization of a private signal $s_{it} \sim N(\omega_t, \sigma_{it}^2)$. Member i forms the belief that the high state occurred via Bayesian updating, yielding

$$\hat{\omega}(s_{it}) = \Pr[\omega = 1 \mid s_{it}] = \frac{q_t f_1(s_{it})}{q_t f_1(s_{it}) + (1 - q_t) f_0(s_{it})}. \quad (2)$$

where $f_1 \sim (1, \sigma_{it}^2)$ is the distribution of s_{it} conditional on $\omega_t = 1$ and $f_0 \sim (0, \sigma_{it}^2)$ is the distribution of s_{it} conditional on $\omega_t = 0$. We refer to σ_{it} as *expertise* because it reflects how much additional information member i has above and beyond conventional wisdom. Rather than interpreting s_{it} as additional data (as pointed out by Blinder (2007a), MPC members have access to the same economic data relevant to inflation), we interpret it as that part of a member's viewpoint on the implications of the data for inflation that is not already incorporated into standard thinking.¹⁴ A member with no expertise ($\sigma_{it}^2 = \infty$) will always hold the same beliefs as the conventional wisdom ($\hat{\omega}(s_{it}) = q_t \forall s_{it}$) while a member with more expertise will in general hold beliefs that diverge from conventional wisdom. This view of σ_{it} as expertise corresponds to the notion of expertise acquired outside central banking as referred to by the Bank of England (and as discussed above).

Inflationary pressures also depend on $\theta \in (0, 1)$, an unknown parameter that reflects

¹⁴For example, in the lead-up to the recent financial crisis, all MPC members had the same access to financial market data, but some argue that its implications were not fully understood by the Bank of England. Former external member Sushil Wadhvani claims that "little was done (by the Bank) to deal with the bubble, despite public concerns about excessive risk taking...In my time at the MPC at the bank, I was surprised by the lack of interest in issues relating to financial markets" (Bloomberg Businessweek 2010b). In our model, a member that recognized that financial market data boded poorly for future economic growth when the majority view did not would have expertise.

uncertainty about the transmission mechanism or, more generally, the model that maps interest rate decisions into inflation outcomes at time t .¹⁵ Unlike with ω_t , members do not receive information about θ ; they simply have beliefs θ_{it} at time t on its magnitude. Alternatively, one could simply view θ_{it} as a preference parameter reflecting how much evidence member i needs that the high interest rate is needed before voting for it (or, in the language of jury models, his or her “burden of proof” (Feddersen and Pesendorfer 1998)). Because the θ parameter has multiple interpretations, we will simply refer to it as a member’s *economic philosophy*. In line with the standard terminology in the monetary literature, we define a member as a *dove* if $\theta_{it} < 0.5$, a *neutral* if $\theta_{it} = 0.5$, and a *hawk* if $\theta_{it} > 0.5$.

An important assumption is that agents’ utilities depend on their votes rather than the committee decision. This is a way of imposing a sincere voting behavioral assumption. In many voting models agents behave strategically, meaning they incorporate the information conveyed by other members’ strategies in their voting decisions. Here, they vote for whichever interest rate maximizes their expected utility only conditioning on their own information. Solving and estimating a model of strategic voting would be difficult. First, members vote sequentially rather than simultaneously. Even if we could observe the order in which MPC voting occurred (which we cannot), a strategic voting model would have multiple equilibria and arguably limited predictive value. Moreover, Gerlach-Kristen (2004) finds evidence that markets react to dissenting votes by adjusting the yield curve, which itself impacts on inflation. So, unlike in standard voting models, MPC members would have to condition on all possible vote configurations, not simply the ones in which they are pivotal for the decision. Solving for the equilibria of this game would be a challenge for us, and certainly for the MPC members. The sincere voting assumption attributes to MPC members a straightforward rule of thumb for reacting to their information and allows for a clean empirical analysis.

Member i votes for the high rate if and only if $\widehat{\omega}(s_{it}) \geq 1 - \theta_{it}$, which, after some algebraic manipulations,¹⁶ implies that he votes for the high rate if and only if

¹⁵There is a large literature on model uncertainty in monetary economics starting with the seminal contribution of Brainard (1967).

¹⁶

$$\widehat{\omega}(s_{it}) \geq \theta_{it} \Leftrightarrow \frac{q_t}{q_t + (1 - q_t) \frac{f_0(s_{it})}{f_1(s_{it})}} \geq 1 - \theta_{it} \Leftrightarrow \frac{\theta_{it}}{1 - \theta_{it}} \frac{q_t}{1 - q_t} \geq \frac{f_0(s_{it})}{f_1(s_{it})}$$

Taking logs on both sides of the final inequality gives that member i votes high if and only if

$$\ln \left(\frac{\theta_{it}}{1 - \theta_{it}} \frac{q_t}{1 - q_t} \right) \geq -\frac{1}{2\sigma_{it}^2} s_{it}^2 + \frac{1}{2\sigma_{it}^2} (s_{it}^2 - 2s_{it} + 1),$$

from which expression (3) derives.

$$s_{it} \geq s_{it}^* = \frac{1}{2} - \sigma_{it}^2 \ln \left(\frac{\theta_{it}}{1 - \theta_{it}} \frac{q_t}{1 - q_t} \right). \quad (3)$$

Gathering these observations establishes the following result.

Proposition 1

$$v(s_{it}) = \begin{cases} 1 & \text{if } s_{it} \geq s_{it}^* \\ 0 & \text{if } s_{it} < s_{it}^*. \end{cases}$$

Because the normal distribution satisfies the monotone likelihood ratio property, a higher realization of s_{it} gives more evidence that a high state occurred. So, voters adopt a cutoff rule that calls on them to vote for the high rate if and only if the realization of their signal crosses the threshold s_{it}^* .¹⁷ When this threshold increases, the probability of observing $v_{it} = 1$ decreases since it becomes less likely that s_{it} will take on sufficiently high values to justify implementing the higher rate.¹⁸

3.1 Philosophy versus expertise

Our hypothesis of interest is how external and internal members compare in terms of expertise, but uncovering this difference is complicated by the fact that members also potentially differ in terms of philosophy. Thus it is crucial to understand how differing in term of expertise is empirically distinguishable from differing in terms of philosophy.

Suppose first that σ_{it}^2 increases. The associated change in s_{it}^* is

$$\frac{\partial s_{it}^*}{\partial \sigma_{it}^2} = - \ln \left(\frac{\theta_{it}}{1 - \theta_{it}} \frac{q_t}{1 - q_t} \right) \begin{matrix} \leq \\ \geq \end{matrix} 0 \iff q_t \begin{matrix} \geq \\ \leq \end{matrix} 1 - \theta_{it}. \quad (4)$$

A higher σ_{it}^2 means that the prior belief q_t will be more influential in determining member i 's vote. If q_t favors the high state, member i will vote for the high rate more often; when q_t favors the low state, he votes for the low rate more often. Intuitively speaking, if a member has less expertise, he becomes more likely to “follow the crowd”. If conventional wisdom says that rates should be high, he votes high more often; if it says rates should be low, he votes low more often. In the limit as $\sigma_{it} \rightarrow \infty$, the member becomes totally unresponsive to his signal and votes high if and only if $q_t \geq 1 - \theta_{it}$. It is important to

¹⁷We use a normal distribution for signals rather than the more common Bernoulli distribution so that the log-likelihood function that generates the voting data is continuous in the underlying parameters.

¹⁸To see this more formally, observe that the probability of voting for the higher rate is

$$q_t \left(1 - \Phi \left(\frac{s_{it}^* - 1}{\sigma_{it}} \right) \right) + (1 - q_t) \left(1 - \Phi \left(\frac{s_{it}^*}{\sigma_{it}} \right) \right)$$

where Φ is the standard normal cdf. Differentiating with respect to s_{it}^* gives the result.

stress that even if two members have the same philosophy, they will still exhibit conflict if their expertise levels differ. The member with more expertise will vote for systematically higher (lower) rates than the member with less expertise when their prior beliefs are low (high).

Now suppose that θ_{it} increases. The associated change in s_{it}^* is

$$\frac{\partial s_{it}^*}{\partial \theta_{it}} = -\sigma_{it}^2 \ln \left(\frac{q_t}{1 - q_t} \right) \frac{1}{\theta_{it}(1 - \theta_{it})} < 0. \quad (5)$$

When a member becomes more hawkish, he requires less evidence that a high shock occurred to vote for the high rate. So independently of his prior belief, he becomes more likely to vote for high rates. If two members have the same level of expertise but differ in terms of philosophy, then the more hawkish member votes for systematically higher rates for all values of prior beliefs.

One can summarize the preceding discussion with the result that

Proposition 2 *When σ_{it}^2 increases, the probability that member i chooses $v_{it} = 1$ increases if and only if $q_t > 1 - \theta_{it}$; when θ_{it} increases, the probability that member i chooses $v_{it} = 1$ increases for all q_t .*




Figure 3: Distinguishing Information and Preferences

The bottom line is that differences in expertise and differences in philosophy are both independent sources of voting disagreement, but how this disagreement expresses itself depends on prior beliefs. One can visualize the difference in figure 3, which plots the probability of voting for the high rate for different values of the prior. An increase in σ induces a rotation in the predicted probability, while an increase in θ induces a shift.

4 Econometric Methodology

Our model generates the likelihood function L_{it} that member i votes for the high rates in time t as

$$L_{it} = \begin{cases} q_t \left(1 - \Phi \left(\frac{s_{it}^* - 1}{\sigma_{it}} \right) \right) + (1 - q_t) \Phi \left(1 - \left(\frac{s_{it}^*}{\sigma_{it}} \right) \right) & \text{if } D(\text{High Vote})_{it} = 1 \\ q_t \Phi \left(\frac{s_{it}^* - 1}{\sigma_{it}} \right) + (1 - q_t) \Phi \left(\frac{s_{it}^*}{\sigma_{it}} \right) & \text{if } D(\text{High Vote})_{it} = 0 \end{cases} \quad (6)$$

where $D(\text{High Vote})_{it}$ is a dummy variable equal to 1 if $v_{it} = 1$. If one could observe q_t and $D(\text{High Vote})_{it}$,¹⁹ one can take this likelihood function directly to the data to estimate expertise and philosophy. Before doing so, one must address two theoretical challenges.

The first is that one could obviously never separately identify θ and σ parameters for each member in each period. Instead, we assume that these parameters are constant within subgroups of voters and compare differences across subgroups either by running our estimator on a restricted sample or by allowing the θ and σ parameters to depend linearly on categorical variables. We show, using a Monte Carlo exercise, that assuming within-group homogeneity in the presence of within-group heterogeneity does not bias the estimates of the group mean.

The second difficulty in implementing the estimator is that the likelihood function is nearly flat for large and small values of q_t when σ is large. When σ increases, the probability of voting for the high rate as a function of q_t approaches a step function as figure 4 shows. This implies that our maximum likelihood estimator will have difficulty converging when q_t takes on extreme values. We detail below how we address this issue.

4.1 Constructing prior beliefs

In spite of these theoretical challenges, the main estimation difficulty is that we do not observe q_t and $D(\text{High Vote})_{it}$. Instead we use a Reuters survey that is carried out just before each MPC meeting to construct a proxy measure for the prior, which we will call \hat{q}_t , and for $D(\text{High Vote})_{it}$, which we will call $\hat{D}(\text{High Vote})_{it}$.²⁰ The survey asks around

¹⁹Although we can observe a member's vote, in periods with unanimous voting it is not clear whether the alternative under consideration was higher or lower. Thus we do not observe $D(\text{High Vote})_{it}$ directly.

²⁰As Reuters did not have the survey results stored in their database, they were unable (or unwilling) to provide the data for us. Instead, we have been able to collate copies of the survey results for most periods in the sample; the 8 exceptions are February 2000, March 2000, October 2003, December 2005, September 2007, October 2007, November 2007, and July 2008. In addition, we are unable to use the data for periods April 2000, August 2008, and November 2008 (details of why not are in the appendix). This leaves 131 out of the 142 sample periods in which we can construct market beliefs, although these meetings do not get dropped from the sample as the methodology outlined below allows us to fill in these




Figure 4: The Probability of Voting for the High Interest Rate

30-50 market economists from financial institutions in the City of London to predict the outcome of MPC voting by writing a probability distribution over possible interest rates choices. 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. The biggest problem with the survey is that it asks respondents about their beliefs about what will—as opposed to what should—happen in MPC meetings. To answer the survey question accurately, respondents would have to use their beliefs about economic conditions, along with their beliefs about the philosophy and information of members, to compute the probability of different configurations of MPC votes. We believe that it is unlikely that they go through this process and more plausible that they predict what the committee will do based on what they themselves would do if they were committee members.

In periods with two votes we set \hat{q}_t equal to the average probability placed on the higher rate over the total average probability placed on both rates and $\hat{D}(\text{High Vote})_{it} = 1$ if and only if member i voted for the higher of the two rates. In periods with one vote, we identify the two rates on which the market places the highest average probability, and set \hat{q}_t equal to the average probability placed on the higher rate over the total average probability placed on both rates. We set $\hat{D}(\text{High Vote})_{it} = 1$ if and only if member i voted for the higher of the two identified rates.

missing values, though dropping them does not alter our empirical results.

Table 2: Example of Survey Data

	+50bps	+25bps	0	-25bps	-50bps
UBS	15%	80%	5%		
Goldman Sachs	20%	75%	5%		
JP Morgan	45%	45%	10%		
AIB	15%	85%			
Average	23.75%	71.25%	5%		

Table 2 illustrates an example of the survey. The two outcomes with the highest average probability are a rise of 50 basis points and a rise of 25 basis points. So, if we observed a unanimous vote of +25, our proxy measure for the prior would be $\hat{q}_t = \frac{23.75}{23.75+71.25} = 0.25$ and we would set $\hat{D}(\text{High Vote})_{it} = 0$ for all members. In fact the survey data has somewhat different formats for different periods within our sample and in some cases respondents were not able to write their beliefs over a full probability distribution. In these cases we follow a slightly different methodology, but the construction is similar. Appendix A contains the full details of the construction in all periods, including how we treated anomalies in the data.

Of course, \hat{q}_t is not a perfect measure of the unobservable q_t ; we acknowledge it is simply a proxy and it is subject to measurement error. For example, the survey respondents predict the outcome of MPC meetings rather than what should be done. Also, the respondents, even answering in the week of the decision, may not have exposure to the full set of economic indicators that will be available to the MPC; the committee is regularly given advance access to data that will only be released subsequently to the wider public, and the information about the MPC's own quarterly forecast is not known by the market respondents until the middle of the month in which it is published (about 2 weeks after the decision). As such, they may not form the correct beliefs about the implications for monetary policy from economic data, or may simply not carefully consider their responses. In order to try to purge our proxy variable of some the noise related to these concerns, we run the regression

$$\ln\left(\frac{\hat{q}_t}{1-\hat{q}_t}\right) = \alpha_q + \beta_q \cdot f_t^{\text{MPC}} + \varepsilon_t$$

where f_t^{MPC} is the fraction of period t votes on the MPC in favor of the high rate. For example, if the MPC split 6-3 for the high rate, $f_t^{\text{MPC}} = \frac{6}{9} \approx 0.66$. This regression predicts the market's beliefs if it could observe the realized split of MPC votes in period t , assuming the average historical relationship. We choose the log odds transformation of \hat{q}_t as the dependent variable since Bayes' rule implies that the log odds of posterior

beliefs are linear in the log odds of the prior beliefs and log odds of new information.²¹ We use the fitted values from (4.1) to generate our cleaned proxy variable which we call \hat{q}_t^{MPC} . Specifically, we use the relationship

$$\ln \left(\frac{\hat{q}_t^{\text{MPC}}}{1 - \hat{q}_t^{\text{MPC}}} \right) = \hat{\alpha}_q + \hat{\beta}_q \cdot f_t^{\text{MPC}} \quad (7)$$

and back out \hat{q}_t^{MPC} . Figure 5 plots \hat{q}_t^{MPC} and $\hat{D}(\text{High Vote})_{it}$ for all sample periods. One can see that both proxies take high and low values for both loosening and tightening cycles.




Figure 5: Key Empirical Variables

In addition to cleaning \hat{q}_t , this transformation has two more advantages. Because we have voting data every period, it allows us to fill in values for the prior in the periods in which the Reuters data is missing. Also, as can be seen in figure 5, it reduces the range of beliefs from $[0, 1]$ (the range of \hat{q}_t) to $[0.2, 0.8]$ (the range of \hat{q}_t^{MPC}). This aids our maximum likelihood estimator in separately identifying the σ and θ parameters. Appendix A also shows that \hat{q}_t^{MPC} behaves in ways consistent with the theoretical quantity q_t : it strongly predicts the probability members will vote for the high rate, and has a concave relationship with the standard deviation of members' votes. While our baseline results below employ \hat{q}_t^{MPC} , we show in section 6 that they are robust to alternative constructions of prior beliefs. We also show using Monte Carlo simulations that our estimator is robust to measurement error in q_t .

²¹We also included a squared term in f_t^{MPC} to account for non-linearities in the relationship but these proved to be insignificant.

5 Results

This section presents estimates of expertise and philosophy parameters that arise from maximising the likelihood function (6) by taking²²

$$q_t = \hat{q}_t^{\text{MPC}} \text{ and} \\ D(\text{High Vote}_{it}) = \hat{D}(\text{High Vote})_{it}.$$

We first apply our estimator to the cross-sectional sample and measure overall differences between internal and external members. We then explore voting dynamics and show how expertise and philosophy evolve with tenure. Finally, we analyze whether the possibility of reappointment leads external members to exaggerate their expertise.

5.1 Baseline results

Table 3 displays the cross-sectional results. Column (1) shows estimates for the the entire sample (that is, assuming that $\theta_{it} = \theta$ and $\sigma_{it} = \sigma$). The average value of σ is 1.24 and of θ is 0.43. Both values are precisely estimated with fairly tight confidence intervals. The estimated θ is significantly different from 0.5, meaning the average members has a dovish philosophy.

Table 3: Estimates of θ and σ

	(1) Whole Sample	(2) Internal	(3) External	(4) Difference
θ	0.43*** [0.41 - 0.45]	0.51*** [0.49 - 0.54]	0.35*** [0.31 - 0.38]	-0.17*** [-0.21 - -0.12]
σ	1.24*** [1.12 - 1.35]	1.72*** [1.50 - 1.94]	1.03*** [0.86 - 1.20]	-0.69*** [-0.97 - -0.41]

95% confidence interval in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Next, we model $\theta_{it} = \alpha_\theta + \beta_\theta D(\text{EXT})_i$ and $\sigma_{it} = \alpha_\sigma + \beta_\sigma D(\text{EXT})_i$ where $D(\text{EXT})_i$ is a dummy variable equal one if member i is external. Columns (2) and (3) report the associated estimated values of σ and θ for internal and external members, and column (4) reports the differences between them.²³ The results are striking and indicate that the overall sample results mask large heterogeneity between the different member types. We find that external members have substantially more expertise than internal members. The

²²Our estimation is done in Stata and the programs are available on request. In order to facilitate faster estimation, we actually estimate $\ln(\sigma_i)$ and $\ln\left(\frac{\theta_i}{1-\theta_i}\right)$ and then convert our estimates and standard errors back appropriately.

²³In the preceding tables, the stars indicate that the number is significantly different from zero.

difference between the estimated σ parameters is -0.69 , which is significantly different from zero at the 1% level. In line with rest of the MPC voting literature, we also find that external members are more dovish. The average internal member has a neutral philosophy ($\theta = 0.51$) while the average external member is dovish ($\theta = 0.35$). The difference is again significantly different from zero at the 1% level. Taken together, these results predict that external members are less likely to vote for the high rate especially as q_t grows. This can be seen in figure 6, which is the estimated counterpart to figure 3. At $q_t = 0.25$, both have a predicted 3% chance of voting for the high rate; at $q_t = 0.5$, the predicted probability that the internal votes for the high rate is 53% while for the external member it is 0.29; at $q_t = 0.75$, the internal has a 98% chance of voting for the high rate, while the external member has a 75% chance.




Figure 6: The Estimated Probability of Voting for the High Interest Rate

Our findings on the σ differences support the idea that external members bring a different perspective to the MPC that allows it to “benefit from thinking and expertise in addition to that gained inside the Bank of England”. Nevertheless, we want to emphasize that this does not imply that external members have more or better information than internal members in any absolute sense. Simply put, our estimates indicate that internal members and market economists form beliefs about inflationary pressures that are closer together than those of external members and the market. This may be because many City economists have backgrounds in central banking or because the thinking of the internal members is the established norm that drives the market’s views.

One reason that internals and market economists might have similar views is that a major source of public information produced by the Bank is the Quarterly Inflation Report, which contains forecasts of future inflation. The thinking behind these forecasts may be driven in large part by the internal members and the staff who produce them. Former external member David Blanchflower writes that (Bloomberg Businessweek 2010a)

During my time on the Bank of England’s Monetary Policy Committee, which makes quarterly economic prognoses, Governor Mervyn King controlled the hiring and firing of the forecast team, who did his bidding. They had to produce a result that was consistent with King’s views, or else they would be history.

While the Bank’s forecast is informative and highly influential, it is also not perfect. The role of external members is then to add their own views so that the MPC does not rely solely on the main forecast.²⁴ In other words, the policy conclusion is not that one should do away with internal members altogether to maximize the available information, but that one should potentially mix them with externals.

5.2 Voting dynamics

As mentioned in the introduction, Hansen and McMahon (2008) show there is evidence from reduced-form regressions that the average voting difference between internal and external members grows with time. An advantage of our structural approach is that it can extend this finding by attributing it to either expertise and philosophy. One might imagine that external members, being new to monetary policy, might take time to learn how to read the state of the world and follow the prior (and therefore internal members) more initially than later. This learning effect would be reflected by similar σ estimates for new internal and external members, but then lower σ estimates for experienced members.

In order to explore this question, we model our voting parameters as

$$\theta_{it} = \alpha_{\theta} + \beta_{1\theta} \cdot D(\text{EXT})_i + \beta_{2\theta} \cdot D(\text{EXP})_{it} + \beta_{3\theta} \cdot D(\text{EXT})_i \cdot D(\text{EXP})_{it} \quad (8)$$

$$\sigma_{it} = \alpha_{\sigma} + \beta_{1\sigma} \cdot D(\text{EXT})_i + \beta_{2\sigma} \cdot D(\text{EXP})_{it} + \beta_{3\sigma} \cdot D(\text{EXT})_i \cdot D(\text{EXP})_{it} \quad (9)$$

where $D(\text{EXP})_{it}$ is a dummy variable equal to one if member i has served for at least twelve months on the MPC up to time t . This choice of the twelve month cut-off is

²⁴Although the forecast is “the best collective judgement” of the MPC, in practice it is not always possible to incorporate all differing views in the central case. As evidence of the fact that not all views of MPC members are contained in the forecast, the Inflation Report sometimes shows the central case under some of these alternative assumptions (Table 6.B of the Inflation Report).

somewhat arbitrary (it corresponds to $\frac{1}{3}$ of the term length for an external member) but the results are unchanged for thresholds of 9 and 18 months.²⁵

Table 4: MPC Members Distinguished by Experience

	(a) σ Estimates			(b) θ Estimates			
	(1) Internal	(2) External	(3) Difference	(1) Internal	(2) External	(3) Difference	
New	1.91*** [1.29 - 2.53]	1.08*** [0.77 - 1.39]	-0.83** [-1.53 - -0.14]	New	0.48*** [0.43 - 0.54]	0.49*** [0.42 - 0.55]	0.0025 [-0.080 - 0.085]
Experienced	1.70*** [1.46 - 1.94]	1.23*** [0.98 - 1.49]	-0.47*** [-0.82 - -0.12]	Experienced	0.52*** [0.49 - 0.55]	0.29*** [0.26 - 0.33]	-0.23*** [-0.27 - -0.18]
Difference	-0.21 [-0.88 - 0.46]	0.15 [-0.25 - 0.55]	0.36 [-0.42 - 1.14]	Difference	0.037 [-0.025 - 0.098]	-0.19*** [-0.26 - -0.12]	-0.23*** [-0.32 - -0.14]

95% confidence interval in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 4 shows how θ and σ evolve for internal and external members. In terms of the effect of experience on external members, there is no evidence of learning: the estimated σ for externals actually grows with time, albeit insignificantly so. Thus it appears that changes in expertise do not play a large role in voting dynamics. In contrast there are strong dynamics in θ . Internal and external members begin their time on the MPC with neutral economic philosophies (both θ estimates are insignificantly different from 0.5) and internal members remain neutral. On the other hand, the estimated θ coefficient for experienced external members drops to 0.29. Although our paper focuses on expertise differences, we view this result as being of independent interest. As external members gain experience, they become systematically less tough on inflation, indicating that their later incarnations disagree with the decisions that their earlier selves made.²⁶ Figure 7 shows the effect of experience on the predicted probability that internals and external vote for the high rate.

It might be tempting to conclude that externals want to blend in at first by mimicking internal members, but then feel more comfortable contradicting them once they are settled on the MPC. While this may be true to some extent, we want to highlight that our results imply that new externals are not afraid of standing out. Even though we measure nearly the same θ for internals and externals, the fact that they differ in terms of σ means

²⁵Another interesting dimension of the experience analysis is whether the differences we capture are the effect of being new or of beginning a new term. Unfortunately, there are not enough 2nd or 3rd term external members for us to carry out the analysis for that group, while for internal members the results show no differences.

²⁶Note that one cannot generate the result simply by allowing beliefs about θ to evolve as members refine their view on the transmission mechanism. One would need to add an assumption that new external members had a prior on θ with an upward bias. Otherwise, by the martingale property, the average value of θ would remain constant over time.

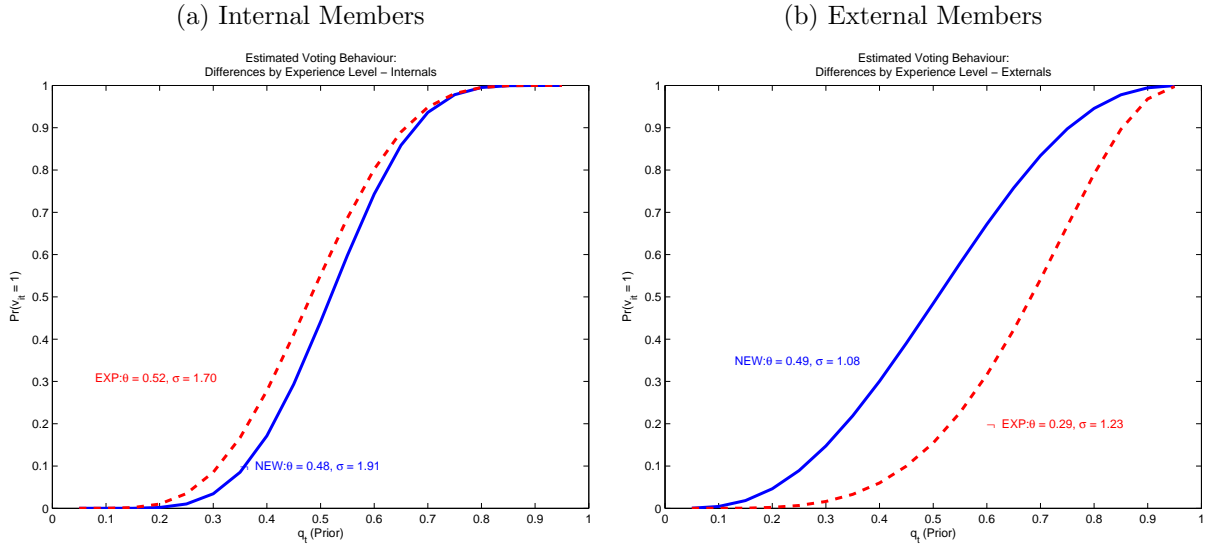


Figure 7: The Effect of Experience on Voting Behavior

there is disagreement.²⁷ As we did above, we can calculate, based on the estimated parameters, the predicted probability that new externals and internals vote for the high rate for different levels of q_t ; this gives an indication of the likelihood of disagreement. When the prior is 0.30 (a 30% chance that the high state of the world prevails), a new internal member has only a 3% chance of voting for the high rate whereas a new external has a 15% chance. If q_t were 0.7, the new internal would almost certainly vote for the high rate (94% likely) whereas the new external only votes for the high rate with 83% probability.

5.3 Voting and reappointment

We conclude our main empirical results by examining a particular type of career concern: the desire for reappointment to the committee. Appointment to the MPC is undoubtedly a prestigious position in the UK policy world that brings with it media attention and other forms of public exposure that are important for building professional stature. One might well then imagine that members would want to behave in such a way as to maximize the probability that they are reappointed, or offered reappointment. Since external members are appointed specifically to bring a fresh perspective to the MPC, and because the observable behavioral corollary of expertise is voting against the conventional wisdom, it may be the case that external members are contrarian simply to build a reputation as

²⁷We should note that conflict on the σ dimension is difficult for reduced form regressions to capture because two members with the same θ will vote on average for the same interest rate even if they differ in expertise.

experts rather than because they actually have an informative private viewpoint.

To formalize this idea, one can easily extend the model from section 3. Suppose that a member has expertise captured by the parameter $\sigma \in \{\sigma_L, \sigma_H\}$ that he knows, while the government only knows that $\Pr[\sigma = \sigma_L] = p$, and that θ and q are common knowledge.²⁸ Let $R(v) = \Pr[\sigma = \sigma_L | v]$ be the government's belief that the member is the low variance type after observing his vote v . To capture reputational concerns, suppose that the member's expected utility from voting for the high rate is now

$$-(1 - \theta)(1 - \widehat{\omega}) + \beta R(1) \quad (10)$$

and his expected utility from voting for the low rate is

$$-\theta \widehat{\omega} + \beta R(0), \quad (11)$$

where $\beta \in [0, 1)$ measures the strength of reputational concerns. Here utility is increasing in R because we assume that the government is more likely to reappoint external members whom it believes to have more expertise. The following result shows how the presence of reputational concerns affects voting behavior.

Proposition 3 *There exists a Perfect Bayesian Equilibrium in which a member with $\sigma = \sigma_L$ chooses $v = 1$ if and only if $s \geq s_L^*(\beta)$ and a member with $\sigma = \sigma_H$ chooses $v = 1$ if and only if $s \geq s_H^*(\beta)$, where $s_L^*(\beta) \leq s_H^*(\beta)$ and $\frac{\partial s_L^*(\beta)}{\partial \beta}, \frac{\partial s_H^*(\beta)}{\partial \beta} \leq 0$ as $q \geq 1 - \theta$.*

Proof. See Appendix B. ■

The intuition for the result is straightforward. Suppose that $q > 1 - \theta$ and let $s_L^*(0)$ and $s_H^*(0)$ be the voting thresholds without career concerns. We know from proposition 2 that $s_H^*(0) < s_L^*(0)$ since the probability that the member with $\sigma = \sigma_L$ votes for the high rate less often when the prior probability is high. This means that $R(0) - R(1) > 0$, so both member types accrue reputation from choosing $v = 0$. So, when one moves to a situation in which $\beta > 0$, both members have an incentive to vote for 0 more often than before and so there exists an equilibrium in which $s_L^*(\beta) > s_L^*(0)$ and $s_H^*(\beta) > s_H^*(0)$.²⁹

In terms of our empirical analysis, proposition 3 makes clear that it is impossible to distinguish whether we estimate a lower σ parameter for external members than for internal members because the former genuinely have more expertise or because they want to signal to the government that they have expertise. Instead we must find an exogenous source of reputational concerns, which we have in the form of a quasi-natural experiment. The Act that created the MPC allows for the reappointment of all members, internal and

²⁸Here we drop the *it* subscripts for notational simplicity.

²⁹A symmetric argument goes through for the case in which $q < 1 - \theta$.

external. When the first group of externals and internals served on the MPC, they would have operated under the assumption that reappointment to the committee was possible, although uncertainty still existed about how the reappointment system would function. Then, on 18 January 2000, Willem Buiter wrote an open letter to then Chancellor Gordon Brown that laid down forceful arguments for not reappointing external members. To quote:

With the end of my term approaching, I have given considerable thought to whether I should be a candidate for re-appointment. I have come to the conclusion that both the appearance and the substance of independence of the external members of the MPC are best served by restricting their membership to a single term - three years as envisaged in the Bank of England Act 1998.

Whether or not this letter swayed Brown's decision is unclear, but he did not reappoint a single external member from the original group, even though they included some of the most prominent macroeconomists in the UK. A clear precedent was set that the government would be reluctant to reappoint external members. By February 2003 this view was again modified due to the reappointment of Stephen Nickell. Since then, Kate Barker has been reappointed twice, and Andrew Sentance has been reappointed once.

In terms of the above model, one can think of the β being lower from February 2000 to February 2003 for external members than in other periods since the perceived probability of being reappointed conditional on any voting history was presumably lower. Table 5 tabulates the number of each member's votes that fall into these eras of lower and higher reappointment probability. A total of 322 out of 1246 votes lie in the era of lower reappointment probability and eight internal and external members vote in both eras.

An important insight from the career concerns literature is that reputational incentives should be expected to vary over time. The most common view is that they should be strongest at the beginning of the career when there is the most uncertainty about an agent's type and when the time horizon over which reputation pays off is relatively long (Prendergast 1999), but one might also argue that members care more about their reputations in the periods leading up to reappointment. Regardless, we continue to model θ and σ as in (8) and (9) to allow for dynamic effects.

Table 6 shows how θ and σ evolve for external members across the different eras. Column (1) reports the results for February 2000 - February 2003 and column (2) reports the results for the rest of the sample. In terms of the estimated expertise of new external members, there is an economically and statistically significant difference across the eras. When the reappointment probability is lower, the estimated σ is over three times as large as when it is higher. This is consistent with the view that external members initially

Table 5: Members Votes by Reappointment Probability

(a) Internals			(b) Externals		
Member	Votes		Member	Votes	
	Low Prob.	High Prob.		Low Prob.	High Prob.
Davies	0	2	Buiter	3	33
George	36	37	Goodhart	3	33
King	36	106	Julius	15	30
Plenderleith	27	33	Budd	0	18
Clementi	30	30	Wadhvani	27	9
Vickers	7	21	Allsopp	33	3
Bean	29	73	Nickell	33	39
Tucker	9	73	Barker	21	73
Large	5	35	Bell	8	28
Lomax	0	60	Lambert	34	0
Gieve	0	37	Walton	0	12
			Besley	0	31
			Blanchflower	0	34
			Sentance	0	30
Total	179	517	Total	143	407

exaggerate their private information in order to signal their expertise, and that some of the estimated expertise difference between internal and external members may not be driven by actual differences in information at all, but by external career concerns.³⁰

Table 6: External Behavior Across Reappointment Periods

	(a) σ Estimates			(b) θ Estimates		
	(1)	(2)	(3)	(1)	(2)	(3)
	Reappointment Low	Probability High	Difference	Reappointment Low	Probability High	Difference
New	2.70*** [0.80 - 4.61]	0.88*** [0.53 - 1.22]	-1.83* [-3.76 - 0.11]	0.49*** [0.42 - 0.56]	0.48*** [0.39 - 0.57]	-0.0093 [-0.12 - 0.10]
Experienced	1.82*** [0.93 - 2.71]	1.19*** [0.92 - 1.45]	-0.63 [-1.56 - 0.29]	0.24*** [0.20 - 0.28]	0.33*** [0.29 - 0.37]	0.089*** [0.030 - 0.15]
Difference	-0.88 [-2.99 - 1.22]	0.31 [-0.12 - 0.74]	1.19 [-0.95 - 3.34]	-0.25*** [-0.33 - -0.17]	-0.15*** [-0.25 - -0.051]	0.098 [-0.030 - 0.23]
95% confidence interval in parentheses *** p<0.01, ** p<0.05, * p<0.1			95% confidence interval in parentheses *** p<0.01, ** p<0.05, * p<0.1			

One validity check on the natural experiment is to examine whether the estimated expertise for internal members changes across eras in the same way as for external members. This would indicate that some common factor made perceiving the state of the world more difficult for new members during February 2000 - February 2003. Table 7 reveals that in fact there are no such differences for internal members. If anything, their initial variance is lower from February 2000 - February 2003 than in other periods. The astute reader will notice that the estimated value of σ for experienced external members is

³⁰The estimated variance for experienced external members is also higher in the era of lower reappointment probability, but this difference is not statistically significant. External members also show higher variance in February 2000 - February 2003 in the cross section, but again the difference is not significant.

nearly the same as for experienced internals from February 2000 - February 2003. Rather than interpret this result as evidence that career concerns explain the entire measured σ difference in the baseline, we believe it is an anomaly generated from the restricted sample. It is not robust to alternative specifications of the prior belief (although the column difference in table 6 is). Also, the votes undertaken by external members in this period are cast disproportionately by non-academics, who as we show below tend to have a higher estimated σ .

Table 7: Internal Behavior Across Reappointment Periods

	(a) σ Estimates			(b) θ Estimates			
	(1) Reappointment Low	(2) Probability High	(3) Difference	(1) Reappointment Low	(2) Probability High	(3) Difference	
New	1.59*** [0.48 - 2.70]	2.08*** [1.24 - 2.91]	0.49 [-0.90 - 1.87]	New	0.57*** [0.44 - 0.71]	0.46*** [0.40 - 0.52]	-0.11 [-0.26 - 0.036]
Experienced	1.93*** [1.34 - 2.51]	1.64*** [1.38 - 1.91]	-0.28 [-0.92 - 0.36]	Experienced	0.56*** [0.51 - 0.61]	0.50*** [0.47 - 0.54]	-0.056* [-0.11 - 0.0027]
Difference	0.34 [-0.92 - 1.59]	-0.43 [-1.31 - 0.45]	-0.77 [-2.30 - 0.76]	Difference	-0.014 [-0.16 - 0.13]	0.041 [-0.029 - 0.11]	0.055 [-0.10 - 0.21]

95% confidence interval in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Although the expert career concerns literature has focused primarily on the signaling of expertise, table 6 also reveals a correlation between the size of the fall in θ and the probability of reappointment. When the probability is higher, the magnitude of the fall is 0.17 and when it is lower the magnitude is 0.32. One might imagine that external members would not be reappointed if the government believed that their philosophy was too extreme in either a dovish or hawkish direction, so that those with career concerns remained more neutral than their colleagues that cared less about reappointment.³¹

Occupational background also provides a source of heterogeneity for career concerns.³² Since the most important activity that academic research economists carry out is the production of basic research, external academics may not care as much about reappointment as external non-academics since recommencing an active research career is substantially more difficult after six years than after three. The fact that academics do not face future

³¹Some basic support for this view comes from the fact that the percentage of votes that an external member makes that deviate upward from the committee decision positively and significantly predicts the probability he will get reappointed, and the percentage of votes made that deviate downward from the committee decision negatively (but insignificantly) predicts the probability.

³²We have also tried to examine whether the age of the member captures any of the estimated differences but we find no differences in the behavior of either internal or external members when they are older or younger (we tried a number of different cutoffs). While one might think that younger members have more reputation to build, and the payoffs from such reputation would be enjoyed over the longer period remaining in their careers, it is not clear whether any differences between, for example, a 50 year old member and a 60 year old member, would be large enough to identify.

occupational uncertainty is a further reason why their reputational incentives may be lower.

Table 8 shows how θ and σ evolve for external academics and non-academics. When new, both academic and non-academic externals have an estimated σ near one. While academics maintain this level of expertise, non-academics actually have a significant *decrease* in their expertise, which is again consistent with the idea that the desire for reappointment influences behavior. This result also suggests that academic members have more expertise than non-academic members in general. Academic externals also display a significantly larger fall in θ than non-academic externals. This gives additional evidence for the idea that reputational concerns may limit the extent to which external members are willing to exhibit philosophical differences with internal members.

Table 8: The Behavior of External Members by D(Academic)

(a) σ Estimates				(b) θ Estimates			
	(1)	(2)	(3)		(1)	(2)	(3)
	Non-academic	Academic	Difference		Non-academic	Academic	Difference
New	1.18*** [0.76 - 1.59]	1.00*** [0.52 - 1.47]	-0.18 [-0.81 - 0.46]	New	0.45*** [0.37 - 0.53]	0.53*** [0.43 - 0.63]	0.080 [-0.047 - 0.21]
Experienced	1.75*** [1.22 - 2.28]	0.92*** [0.60 - 1.23]	-0.83*** [-1.45 - -0.21]	Experienced	0.31*** [0.28 - 0.35]	0.25*** [0.19 - 0.32]	-0.061 [-0.14 - 0.015]
Difference	0.57* [-0.10 - 1.25]	-0.080 [-0.65 - 0.49]	-0.65 [-1.54 - 0.23]	Difference	-0.14*** [-0.22 - -0.052]	-0.28*** [-0.40 - -0.16]	-0.14* [-0.29 - 0.0065]
95% confidence interval in parentheses *** p<0.01, ** p<0.05, * p<0.1				95% confidence interval in parentheses *** p<0.01, ** p<0.05, * p<0.1			

6 Robustness

In this section we explore the robustness of the estimates presented in section 5. We address three main concerns. First, since correctly measuring the prior distribution is crucial for our results, we explore the effect of using alternative constructions and how our estimator performs with measurement error in beliefs. Second, we discuss why we have not allowed for individual heterogeneity and whether our estimator can accurately measure between-group differences in the presence of heterogeneity. Finally, we ask whether all Bank outsiders necessarily bring strong private information.

6.1 Alternative constructions of the prior

The main estimation difficulty is that we do not observe the q_t variable which captures the prior belief that the correct interest rate is the high interest rate. The baseline measure that we use involves combining two proxies for the information that the MPC

has available when it makes its decision; this is our \hat{q}_t^{MPC} variable. While we believe that we have pursued a reasonable approach to calculating this variable, we also acknowledge that it remains an imperfect measure. In order to address this concern, we explore two avenues.

Table 9: Different Approaches to Calculating the Prior Distribution

Variable	Mean	Std. Dev.	Min.	Max.	N
\hat{q}_t^{MPC}	0.5	0.2	0.2	0.8	1246
$\hat{q}_{it}^{\text{MPC}}$	0.5	0.2	0.2	0.8	1246
\hat{q}_t^{MPC} (02 - 09)	0.6	0.2	0.4	0.9	685
$\hat{q}_{it}^{\text{TIMES}}$ (02 - 09)	0.6	0.2	0.3	0.9	685

The first is to explore alternative proxy constructions. As mentioned in section 2, members spend over a day in discussions prior to voting. During this time members share their views with each other, allowing their colleagues to update their beliefs on the likelihood of various economic shocks. \hat{q}_t^{MPC} does not capture this effect because it attributes to every member a common prior. In order to tackle this, we fit a member-specific prior belief $\hat{q}_{it}^{\text{MPC}}$ using the estimated coefficient values from equation (4.1):

$$\ln \left(\frac{\hat{q}_{it}^{\text{MPC}}}{1 - \hat{q}_{it}^{\text{MPC}}} \right) = \hat{\alpha}_q + \hat{\beta}_q \cdot f_{it}^{\text{MPC}}, \quad (12)$$

where f_{it}^{MPC} is the share of votes in time t for the high rate excluding the vote of member i . This is a measure of the extent to which the information that member i receives from other members favors the high rate. Our construction assumes that committee members update their beliefs in response to f_{it}^{MPC} in the same way the market's beliefs would update if it could observe f_t^{MPC} . We find that all our baseline results are robust to this measure of beliefs. Here we only report the results related to our central findings. As one can see in table 10a, we continue to find that external members have significantly more private information than internals. Tables 10b and 10c report estimates for our quasi-natural experiment, and the crucial finding, namely that new external members seeking reappointment appear to exaggerate expertise again arises. In the period of higher reappointment probability, experienced externals are also measured to have a smaller dovish shift.

As we mentioned previously, a worry about the Reuter's survey is that it measures beliefs about what the MPC will decide rather than what respondents feel should be done. Rather than use the actual MPC votes to purge \hat{q}_t of noise, a better method may be to use another source of information in which experts predict what should be done. We do this by using voting data from the Times MPC, a committee of experts—many of whom are former or future MPC members—that votes monthly on interest rates as

Table 10: Robustness of our Key Results to Individual Specific Priors

(a) Baseline Estimates - $\hat{q}_{it}^{\text{MPC}}$			
	(1)	(2)	(3)
	Internal	External	Difference
θ	0.52***	0.30***	-0.22***
	[0.49 - 0.55]	[0.25 - 0.36]	[-0.28 - -0.16]
σ	1.42***	0.73***	-0.69***
	[1.24 - 1.59]	[0.58 - 0.88]	[-0.92 - -0.46]

95% confidence interval in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

(b) σ Estimates - $\hat{q}_{it}^{\text{MPC}}$			(c) θ Estimates - $\hat{q}_{it}^{\text{MPC}}$				
	(1)	(2)	(3)		(1)	(2)	(3)
	Reappointment Probability		Difference		Reappointment Probability		Difference
	Low	High			Low	High	
New	1.93***	0.53***	-1.41**	New	0.50***	0.47***	-0.023
	[0.89 - 2.98]	[0.16 - 0.90]	[-2.51 - -0.30]		[0.42 - 0.58]	[0.29 - 0.66]	[-0.22 - 0.18]
Experienced	0.98***	0.89***	-0.090	Experienced	0.17***	0.30***	0.13**
	[0.43 - 1.54]	[0.68 - 1.11]	[-0.69 - 0.50]		[0.087 - 0.26]	[0.24 - 0.35]	[0.023 - 0.23]
Difference	-0.95	0.36*	1.31**	Difference	-0.32***	-0.17*	0.15
	[-2.13 - 0.23]	[-0.069 - 0.79]	[0.059 - 2.57]		[-0.44 - -0.20]	[-0.37 - 0.016]	[-0.076 - 0.37]

95% confidence interval in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

part of a monthly feature for the Times of London newspaper. We construct \hat{q}_t^{TIMES} by again running a regression as in (4.1), but replace f_t^{MPC} with the fraction of Times MPC members that vote for the high rate. Unfortunately, the Times MPC only began meeting November 2002, so using it means we lose many observations and cannot replicate our natural experiment. In order to compare like with like, we also regenerate our original \hat{q}_t^{MPC} variable on the November 2002 - March 2009 subsample.

Table 11: Robustness to the Times MPC measure

	\hat{q}_t^{MPC} : 2002 - 2009 Sample			\hat{q}_t^{TIMES} : 2002 - 2009 Sample		
	(1)	(2)	(3)	(4)	(5)	(6)
	Internal	External	Difference	Internal	External	Difference
θ	0.39***	0.22***	-0.17***	0.36***	0.19***	-0.18***
	[0.35 - 0.43]	[0.18 - 0.26]	[-0.22 - -0.11]	[0.33 - 0.39]	[0.11 - 0.26]	[-0.26 - -0.091]
σ	1.58***	1.02***	-0.57***	1.29***	0.70***	-0.60***
	[1.31 - 1.85]	[0.80 - 1.23]	[-0.91 - -0.22]	[1.07 - 1.52]	[0.47 - 0.92]	[-0.92 - -0.27]

95% confidence interval in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

Table 11 shows that our baseline results hold up when we clean \hat{q}_t with the Times MPC voting data rather than the MPC voting data. Externals' private information and dovish philosophy strongly distinguish them from internal members.³³

³³We could of course use \hat{q}_t itself to estimate the model. The problem is that when we reduce its

The second approach to account for measurement error in \hat{q}_t^{MPC} is to see whether our estimator can recover the true underlying parameters in Monte Carlo simulations when we give it noisy measures of q_t . Appendix C shows that it still performs well in the sense of generating biases that are small both in expected value and absolute value.

6.2 Individual heterogeneity

In our analysis, we have not included member fixed effects for two main reasons. The first is that, as we make use of a probit-type estimation, inclusion of fixed-effects is potentially problematic in such a framework (Baltagi 2005). Second, we are interested in the contribution of the average external and internal member to the MPC and not heterogeneity among individual members. While we accept that the θ and σ parameters may differ at the level of individual members, member fixed effects would soak up the variation that we are interested in. We would therefore need to calculate member-specific estimates and then examine the distribution of these member-specific parameters by internal and external grouping. The problem with this approach is that many external members have not served for long periods, and our likelihood estimator is unable to converge in many cases. The approach we pursue above means that we can draw on statistical power from the fact that we have over 500 external member votes and over 600 internal member votes.

Nevertheless, one still may be concerned that individual heterogeneity manifests itself in the form of biased estimates when we impose the restriction of a common θ and σ . In order to address these concerns, we use Monte Carlo simulations to study the properties of our estimator when there is individual heterogeneity in both parameters. Under these conditions, our estimator produces an accurate estimate of the difference in group means between internals and externals, which is the main quantity of interest in our paper. Thus unmodelled individual heterogeneity does not appear to be of great concern. Full details of the approach and results are in appendix C.

6.3 Externals versus outsiders

External members on the MPC in fact have two distinguishing characteristics. First, they are not selected from the staff of the central bank; second, they do not become Bank executives when they join the MPC. One might wonder the extent to which this second feature matters for the differences we find. In order to shed light on this issue, we make use

range to $[0.2, 0.8]$ through dropping data we lose around 40% of our sample and convergence is difficult to achieve in some cases. Nevertheless, we can estimate the cross-sectional results, and externals again have a highly significant difference from internals in terms of private information.

of the fact that internal members contain a number of people who were not previously employed by the Bank of England prior to their joining the MPC (in contrast to the external outsiders, these internals also become part of the Bank of England Executive Team and thus are Bank employees and have management responsibilities within the Bank). These members are Clementi, Vickers, Bean, Large, Lomax and Gieve. This is similar to the situation on the FOMC, which has had many outside experts who are appointed to the Federal Reserve staff (either the Board of Governors or a regional Fed).

Table 12: Estimates of θ and σ for Bank Insiders and Bank Outsiders

	(1) Outsider	(2) Insider	(3) Difference
θ	0.50*** [0.46 - 0.54]	0.53*** [0.50 - 0.57]	0.032 [-0.023 - 0.086]
σ	1.33*** [1.10 - 1.57]	1.53*** [1.26 - 1.79]	0.19 [-0.16 - 0.55]

95% confidence interval in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 12 shows that there are no significant differences in either philosophy or expertise between “inside” and “outside” internal members. This is not proof that not being Bank employees is important for the behavioral differences we observe for external members because it may be that outside experts similar to central bankers are more likely to be selected as internal members. If this is the case, table 12 shows that not just any type of outsider brings new information to the table—it is only outsiders with truly different backgrounds who do. On the other hand, it may be that some institutional conditioning affects the behavior of internal members. For example, internal members may be less contrarian to avoid offending the staff they manage who produce the pre-MPC analysis.³⁴

7 Conclusion

This paper makes two significant contributions to the literature on expert decision making. The first is methodological. We show how one can separately identify an expert’s philosophy (or preferences) from his private information using voting data and a measure of the public’s beliefs about the realization of an unknown state variable. This same approach could potentially be used in a variety of situations. The second, arguably more important one, is that this is the first paper, to the best of our knowledge, to show that external experts bring private information to a committee. This finding not only backs

³⁴This is a counter point to Blinder (2007b), who argues that the optimal monetary policy committee should be composed solely of government employees.

up the view of the Bank of England, but also supports their inclusion in other technical decision-making bodies.

At the same time, the findings are not unambiguously positive. Most notably, we find evidence that externals may exaggerate the information to gain reappointment, and to the extent that they engage in this behavior, they distort the final decision. However, reputational incentives are endogenous to the design of the committee. For example, if the possibility of reappointment leads to a behavioral distortion, the UK government could amend the Bank of England Act to include a one-term limit for external members or else seek to appoint more academic members. Our results show that these decisions should be carefully considered along with the decision about whether to include external members. Another potential strike against external members is that they have different philosophies than internal members, and depending on the preferred approach of the committee designer, this may harm welfare. On the MPC, for example, the government may consider externals too soft on inflation.³⁵ In this sense, allowing for reappointment may be a double-edged sword: it can both lead to information exaggeration (which lowers welfare) but also help enforce a neutral philosophy (which raises welfare, assuming that the government favors a neutral position).

In spite of these concerns, the paper shows that a mixed committee structure can indeed expand the information available to a group when it comes to take a decision, and this broadly favors the appointment of externals. The challenge for future research is to determine how best to draw on their information while minimizing other potential welfare losses arising from their inclusion.

³⁵The FDA has also recently been criticized for bringing new drugs onto the market too quickly, which may be a result of outside experts having a lower burden of proof for drug safety than the government.

References

- AUSTEN-SMITH, D., AND J. S. BANKS (1996): “Information Aggregation, Rationality, and the Condorcet Jury Theorem,” *American Political Science Review*, 90(1), 34–45.
- BALTAGI, B. (2005): *Econometric Analysis of Panel Data*. John Wiley and Sons, New York, New York, third edn.
- BANK OF ENGLAND (2010a): <http://www.bankofengland.co.uk/monetarypolicy/overview.htm>, last accessed 15 September 2010.
- (2010b): <http://www.bankofengland.co.uk/monetarypolicy/decisions.htm>, last accessed 15 September 2010.
- BANK OF ENGLAND ACT (1998): <http://www.bankofengland.co.uk/about/legislation/1998act.pdf>, last accessed 15 September 2010.
- BESLEY, T., N. MEADS, AND P. SURICO (2008): “Insiders versus Outsiders in Monetary Policymaking,” *American Economic Review*, 98(2), 218–23.
- BHATTACHARJEE, A., AND S. HOLLY (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. (2007a): “Monetary policy by committee: Why and how?,” *European Journal of Political Economy*, 23, 106123.
- BLINDER, A. S. (2007b): “On the Design of Monetary Policy Committees,” Working Paper 153, Center for Economic Policy Studies.
- BLOOMBERG BUSINESSWEEK (2010a): “King’s Men Fiddle With U.K. Forecasts: David G. Blanchflower,” <http://www.businessweek.com/news/2010-07-12/king-s-men-fiddle-with-u-k-forecasts-david-g-blanchflower.html>, last accessed 26 September 2010.
- (2010b): “Wadhvani Suggests ‘Unrepentant’ BOE May Not Deserve New Powers,” <http://www.businessweek.com/news/2010-07-13/wadhvani-suggests-unrepentant-boe-may-not-deserve-new-powers.html>, last accessed 26 September 2010.
- BRAINARD, W. (1967): “Uncertainty and the Effectiveness of Policy,” *American Economic Review*, 57(2), 411–425.

- BUITER, W. (2000): “Letter to the Right Honourable Gordon Brown MP, Chancellor for the Exchequer,” http://www.hm-treasury.gov.uk/press_16_00.htm#letter.
- FEDDERSEN, T., AND W. PESENDORFER (1998): “Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts under Strategic Voting,” *American Political Science Review*, 92(1).
- GERLACH-KRISTEN, P. (2003): “Insiders and Outsiders at the Bank of England,” *Central Banking*, XIV(1), 96–102.
- (2004): “Is the MPC’s voting record informative about future UK monetary policy?,” *Scandinavian Journal of Economics*, 106(3), 299–313.
- (2006): “Monetary policy committees and interest rate setting,” *European Economic Review*, 50(2), 487–507.
- GERLING, K., H. P. GRUNER, A. KIEL, AND E. SCHULTE (2005): “Information acquisition and decision making in committees: A survey,” *European Journal of Political Economy*, 21(3), 563–597.
- HANSEN, S., AND M. F. MCMAHON (2008): “Delayed Doves: MPC Voting Behaviour of Externals,” Discussion Paper 862, Centre for Economic Performance, London School of Economics.
- HARRIS, M., AND C. SPENCER (2008): “Decade of dissent: explaining the dissent voting behavior of Bank of England MPC members,” MPRA Paper 9100, University Library of Munich, Germany.
- HIGGS, D. (2003): “Review of the role and effectiveness of non-executive directors,” Discussion paper, UK Department of Trade and Industry.
- LAMBERT, R. (2006): “Who sets our interest rates?,” http://www.bankofengland.co.uk/publications/other/monetary/inside_mpc.pdf.
- LEVY, G. (2004): “Anti-herding and strategic consultation,” *European Economic Review*, 48(3), 503–525.
- (2007): “Decision Making in Committees: Transparency, Reputation, and Voting Rules,” *American Economic Review*, 97(1), 150–168.
- POLLARD, P. (2004): “Monetary Policy-Making around the World,” Presentation, Federal Reserve Bank of St. Louis.

- PRENDERGAST, C. (1999): “The Provision of Incentives in Firms,” *Journal of Economic Literature*, 37(1), 7–63.
- PRENDERGAST, C., AND L. STOLE (1996): “Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning,” *Journal of Political Economy*, 104(6), 1105–34.
- RIBONI, A., AND F. J. RUGE-MURCIA (2010): “Monetary Policy by Committee: Consensus, Chairman Dominance, or Simple Majority?,” *The Quarterly Journal of Economics*, 125(1), 363–416.
- SPENCER, C. (2006): “The Dissent Voting Behaviour of Bank of England MPC Members,” Department of Economics Discussion Papers 0306, Department of Economics, University of Surrey.
- TRUEMAN, B. (1994): “Analyst Forecasts and Herding Behavior,” *Review of Financial Studies*, 7(1), 97–124.
- VISSER, B., AND O. H. SWANK (2007): “On Committees of Experts,” *The Quarterly Journal of Economics*, 122(1), 337–372.

A Reuter's Survey Data

As described above, the Reuters Data that we have been able to gather is divided into three segments corresponding to data availability, and each period calls for a different construction methodology. Implicit in the construction is how we define $\widehat{D}(\text{High Vote})_{it}$.

June 1997 - June 1998: Mark I Construction

We have modal data, all of which takes a value of 0, and +25, except for October 1997 in which one bank reported +50 which we treat as +25.

- \widehat{q}_t is computed as the total number of reports for +25 over the total number of reports.
- Anomalies
 - In the May 1998 MPC meeting there were 6 votes for no change, 1 vote for -25 (Julius) and 1 vote for +25 (Buiter). We treat the -25 vote as a low vote and proceed as above.
 - In the June 1998 MPC meeting there were 8 votes for +25 and 1 vote for -25 (Julius). Again, we treat -25 as a low vote.

July 1998 - December 2001: Mark II Construction

We have partial probability distribution data over rise, no change, and cut. In some periods only two of these three options are available.

- First we compute the average probability placed on rise, cut, and no change.
- For all periods in which the observe votes that lie in the set -25,0,25 we treat the data as if the probability of rise is the probability of +25 and the probability of cut is the probability of -25.
- For periods in which the votes are +25 and 0 we construct \widehat{q}_t as the total average probability put on +25 over the total average probability placed on +25 and 0.
- For periods in which the votes are 0 and -25 we construct \widehat{q}_t as the total average probability put on 0 over the total average probability placed on 0 and -25.
- For periods in which all votes are for +25 we construct \widehat{q}_t over the total average probability put on +25 over the total average probability placed on +25 and 0.
- For periods in which all votes are for -25 we construct \widehat{q}_t over the total average probability put on 0 over the total average probability placed on 0 and -25.

- For periods in which all votes are for nochange we construct \hat{q}_t as the total average probability put on +25 over the total average probability placed on +25 and 0 if the total average probability placed on +25 is larger than the total average probability placed on -25, and construct \hat{q}_t as the total average probability put on 0 over the total average probability placed on 0 and -25 if the total average probability placed on -25 is larger than the total average probability placed on +25.
- Anomalies
 - In August 1998 there were 7 votes for no change, 1 vote for -25 (Julius) and 1 vote for 25 (Buiter). The total average probability placed on +25 is 0.39 and the total average probability placed on -25 is 0.001. So we treat +25 as high vote and compute as the total average probability place on +25.
 - In January 1999 there was one vote for no change (Plenderleith), 7 votes for -25, and one vote for -50 (Julius). We treat the vote for -25 and -50 as low votes and compute \hat{q}_t as the total average probability placed on nochange over the total average probability placed on nochange and cut.
 - In March 1999 8 people voted for nochange and 1 person (Buiter) voted for -40. We treat the -40 vote as a vote for -25 and proceed as above.
 - In April 2000 3 people voted for +25 and six for 0, but Reuters survey does not ask for probability of rise. We treat these data as missing in this period.
 - In January 2000 the votes were over +25 and +50 and we set $\hat{q}_t = 0.5$.
 - In April 2001, May 2001, October 2001, and November 2001, the votes were over -25 and -50 and we set $\hat{q}_t = 0.5$.

January 2002 - March 2009: Mark III Construction

We have full distribution data over +50, +25, 0, -25, -50.

- First we compute the average probability placed on +50, +25, no change, -25 and -50.
- For periods in which there are two unique votes we take \hat{q}_t as the total average probability placed on the higher of the two votes over the total average probability placed on both votes.
- For periods in which all votes are for +50 we construct \hat{q}_t over the total average probability put on +50 over the total average probability placed on +50 and +25.
- For periods in which all votes are for -50 we construct \hat{q}_t over the total average probability put on -25 over the total average probability placed on -25 and -50.

- For periods in which all votes are for nochange we construct \hat{q}_t as the total average probability put on +25 over the total average probability placed on +25 and 0 if the total average probability placed on +25 is larger than the total average probability placed on -25, and construct \hat{q}_t as the total average probability put on 0 over the total average probability placed on 0 and -25 if the total average probability placed on -25 is larger than the total average probability placed on +25.
- We follow a similar procedure as the above for periods in which all votes are for +25 or -25.
- Anomalies
 - In May 2006 there were six votes for no change, one vote for +25 (Walton) and one vote for -25 (Nickell). The market put probability 0.08 on -25 and probability 0.03 on +25. So we take 0 to be high vote and compute \hat{q}_t as $(\text{meannochange} + \text{meanrise25})/(\text{meancut25} + \text{meannochange} + \text{meanrise25})$.
 - In April 2008 there were votes for 0, -25, and -50. We take 0 as a high-vote and -25 and -50 as low votes, and compute the high vote as $(\text{meannochange})/(\text{meancut25} + \text{meannochange} + \text{meancut50})$.
 - In August 2008 there were votes over -25, 0, +25 but the market placed roughly equal probability on -25 and +25 so we set \hat{q}_t as missing.
 - In November 2008 there was unanimity on -150 and we set \hat{q}_t as missing.
 - From December 2008 - March 2009 we again have modal data.
 - * In December 2008 everyone vote to cut -100 and -100 was the lower bound on the modes. We set -100 as low vote and compute \hat{q}_t as all modes not equal to -100 over all modes.
 - * In January 2009 everyone votes to cut by -50 or -100 and these make up most modes. We set \hat{q}_t as the number of -50 modes over all modes equal to -50 or -100.
 - * In February 2009 everyone votes to cut -50 or -100. We take \hat{q}_t as the total number of modes not equal to -100 over the total number of modes.
 - * In March 2009 everyone votes to cut -50. We set -50 as the high vote and set \hat{q}_t as all modes -50 or greater over all modes.

A.1 Consistency of constructed beliefs

As discussed in section 4.1, we use the fitted values from (4.1) to generate our main proxy variable for the common prior. Here, we verify that our empirical measure of beliefs behaves in a way consistent with its theoretical counterpart—the unobserved q_t . To this end, we first estimate A.1 via a probit model in order to uncover the correlation between $\widehat{q}_t^{\text{MPC}}$ and $\widehat{D}(\text{Vote High})_{it}$. A reasonable requirement for our data is that ψ_1 be positive. As reported in table A.1a, this check is successfully passed.

$$\widehat{D}(\text{Vote High})_{it} = \alpha + \psi_1 \cdot \widehat{q}_t^{\text{MPC}} \quad (\text{A.1})$$

Our second consistency check concerns the relationship between $\widehat{q}_t^{\text{MPC}}$ and the within-period variability of votes. As the decision becomes more certain ($q_t \rightarrow 0$ or $q_t \rightarrow 1$), the standard deviation of votes within a period should decline; this predicted inverted-U shape relationship can be explored using equation A.2. As reported in table A.1b, we find the predicted relationship.

$$\text{Std Deviation of Votes}_t = \alpha + \psi_1 \cdot \widehat{q}_t^{\text{MPC}} + \psi_2 \cdot (\widehat{q}_t^{\text{MPC}})^2 + \varepsilon_{it} \quad (\text{A.2})$$

Table A.1: Checks on the Behavior of the Prior:
Estimates of Equation (A.1) and (A.2)

(a) Probability of voting for the high interest rate		(b) Voting dispersion	
	(1) D(High Vote)		(1) Std. Dev. Votes _t
\widehat{q}_t	5.89*** [0.24]	\bar{q}_t	1.57*** [0.086]
		\bar{q}_t^2	-1.48*** [0.079]
Constant	-3.24*** [0.15]	Constant	-0.26*** [0.021]
Observations	1246	Observations	142
		R-squared	0.643
	Standard errors in brackets		Robust standard errors in brackets
	*** p<0.01, ** p<0.05, * p<0.1		*** p<0.01, ** p<0.05, * p<0.1

B Proof of Proposition 3

Proof. Suppose that $q < 1 - \theta$. The member votes for the high rate if and only if $\widehat{\omega} \geq (1 - \theta) + \beta D$ where $D = R(1) - R(0)$. If the evaluator believes that the type $\sigma = \sigma_i$ expert uses the voting threshold s_i^* , then

$$D(s_L^*, s_H^*) = \frac{p \Pr[s_L \geq s_L^*]}{p \Pr[s_L \geq s_L^*] + (1-p) \Pr[s_H \geq s_H^*]} - \frac{p \Pr[s_L < s_L^*]}{p \Pr[s_L < s_L^*] + (1-p) \Pr[s_H < s_H^*]}. \quad (\text{B.1})$$

Letting s'_i denote the optimal threshold for type σ_i , one obtains

$$s'_i = \frac{1}{2} - \sigma_i^2 \ln \left(\frac{\theta - D(s_L^*, s_H^*)}{1 - \theta + D(s_L^*, s_H^*)} \frac{1}{1 - q} \right) \quad (\text{B.2})$$

which implies that

$$s'_L = \frac{\sigma_L}{\sigma_H} \left(s'_H - \frac{1}{2} \right) + \frac{1}{2}. \quad (\text{B.3})$$

Note that $s'_H > s'_L$ whenever $s'_H > \frac{1}{2}$ and $s'_H = s'_L$ when $s'_H = \frac{1}{2}$. An equilibrium exists if one can find a s_H^* for which

$$s_H^* = \frac{1}{2} - \sigma_H^2 \ln \left(\frac{\theta - F(s_H^*)}{1 - \theta + F(s_H^*)} \frac{1}{1 - q} \right) \quad (\text{B.4})$$

where

$$F = \beta \left(\frac{\frac{p \Pr[s_L \geq \frac{\sigma_L}{\sigma_H}(s_H^* - \frac{1}{2}) + \frac{1}{2}]}{p \Pr[s_L \geq \frac{\sigma_L}{\sigma_H}(s_H^* - \frac{1}{2}) + \frac{1}{2}] + (1-p) \Pr[s_H \geq s_H^*]} - \frac{p \Pr[s_L < \frac{\sigma_L}{\sigma_H}(s_H^* - \frac{1}{2}) + \frac{1}{2}]}{p \Pr[s_L < \frac{\sigma_L}{\sigma_H}(s_H^* - \frac{1}{2}) + \frac{1}{2}] + (1-p) \Pr[s_H \geq s_H^*]} \right). \quad (\text{B.5})$$

If $s_H^* = \frac{1}{2} - \sigma_H^2 \ln \left(\frac{\theta}{1 - \theta} \frac{1}{1 - q} \right) > \frac{1}{2}$, the LHS of (B.4) is bigger than the RHS since (B.5) > 0 ; if $s_H^* = \frac{1}{2}$, the LHS is smaller than the RHS since (B.5) = 0. Since (B.4) is continuous in s_H^* , by the intermediate value theorem there exists some $s_H^* \in (\frac{1}{2}, \frac{1}{2} - \sigma_H^2 \ln \left(\frac{\theta}{1 - \theta} \frac{1}{1 - q} \right))$ where (B.4) is satisfied and at which $\frac{1}{2} < s_L^* < s_H^*$.

Now consider an expert with career concerns given by $\beta' > \beta$. If $s_H^*(\beta') = s_H^*(\beta)$, then the LHS of (B.4) is bigger than the RHS since F is now bigger than before. Just as before, if $s_H^* = \frac{1}{2}$, the LHS of (B.4) is smaller than the RHS since (B.5) = 0. So there exists some equilibrium threshold $s_H^*(\beta') \in (\frac{1}{2}, s_H^*(\beta))$ implying the existence of some equilibrium threshold $s_L^*(\beta') \in (\frac{1}{2}, s_L^*(\beta))$. A symmetric argument holds for $q > 1 - \theta$. ■

C Monte Carlo Exercise

In this appendix we describe the Monte Carlo exercise that we carry out to explore the robustness of our estimation approach to both measurement error in our proxy for the unobserved prior and individual heterogeneity.

C.1 Design and Implementation

The first test of our estimator (Case 1 below) is to determine its performance when we perfectly observe the prior distribution and internal and external members are homogeneous. In order to do this, we use our model to generate a dataset with 1350 observations with 11 internal members (serving 80 meetings each) and 12 external members (serving 50 meetings each), a division that roughly mimics our sample. We attribute to internal members $\theta_{\text{Int}} = 0.5$ and $\sigma_{\text{Int}} = 1.7$ and to external members $\theta_{\text{Ext}} = 0.35$ and $\sigma_{\text{Ext}} = 1$. To generate the j th vote of individual i , we first draw q_{ij} from a $U[0.2, 0.8]$ distribution; then $\omega_{ij} \in \{0, 1\}$ from a Bernoulli distribution with parameter q_{ij} ³⁶; then s_{ij} from a $N(\omega_{ij}, \sigma_i^2)$ distribution; finally, v_{ij} is determined by the voting rule derived in section 3. We then run our estimator, with the correct values of q_{ij} and v_{ij} , and generate estimates for σ_{Int} , θ_{Int} , σ_{Ext} , and θ_{Ext} . We repeat this procedure 1,000 times.

In Case 2 we generate the dataset exactly as above, but rather than use the true value of q_{ij} in the estimation, we use the variable $q'_{ij} = q_{ij} + \eta_{ij}$ where $q_{ij} \sim U[0.25, 0.75]$ and $\eta_{ij} \sim N(0, (0.15)^2)$. This implies that the value of q'_{ij} has a 96% probability of taking a value in the range $[0.2, 0.8]$; whenever it takes a value outside that range we drop it from the estimation due to the difficulties described in section 4.

In Case 3 we draw a θ and σ parameter for each individual in our sample. For internal (external) members, we draw θ from a $U[0.45, 0.55]$ ($U[0.2, 0.5]$) distribution and σ from a $U[1.4, 2.0]$ ($U[0.6, 1.4]$) distribution and generate a dataset exactly as in Case 1. We then run our estimator under the restriction that θ and σ are constant within the group of internals and externals.

For each case we are interested in how far away the estimated values of σ_{Int} , θ_{Int} , σ_{Ext} , and θ_{Ext} are from their true values. In Case 3 we define $\sigma_{\text{Int}} = \frac{1}{11} \sum_1^{11} \sigma_i$ where i corresponds to the i th internal members. We apply similar definitions to the other three parameters of interest. For each of the 1,000 replications of each case, we calculate the

³⁶Notice that we draw a new state variable for each vote and thus do not construct a time series dimension; this should not effect the estimates in any systematic way and simplifies the programming.

Table C.1: Parameter Values for the Monte Carlo Exercises

	Case 1		Case 2		Case 3	
	Baseline		Noisy Prior		Member Variation	
	Internal	External	Internal	External	Internal	External
θ	0.5	0.35	0.5	0.35	0.45-0.55	0.2-0.5
σ	1.7	1.0	1.7	1.0	1.4-2.0	0.6-1.4
q_t^M	q_t		$q_t + \eta_{it}$		q_t	

where: $\eta_{it} \sim N(0, (0.15)^2)$

following biases:

$$\begin{aligned}
 \text{Bias}[\widehat{\sigma}_{\text{Int}}] &= \sigma_{\text{Int}} - \widehat{\sigma}_{\text{Int}} \\
 \text{Bias}[\widehat{\sigma}_{\text{Ext}}] &= \sigma_{\text{Ext}} - \widehat{\sigma}_{\text{Ext}} \\
 \text{Bias}[\widehat{\sigma}_{\text{Int}} - \widehat{\sigma}_{\text{Ext}}] &= (\sigma_{\text{Int}} - \widehat{\sigma}_{\text{Int}}) - (\sigma_{\text{Ext}} - \widehat{\sigma}_{\text{Ext}}) \\
 \text{Bias}[\widehat{\theta}_{\text{Int}}] &= \theta_{\text{Int}} - \widehat{\theta}_{\text{Int}} \\
 \text{Bias}[\widehat{\theta}_{\text{Ext}}] &= \theta_{\text{Ext}} - \widehat{\theta}_{\text{Ext}} \\
 \text{Bias}[\widehat{\theta}_{\text{Int}} - \widehat{\theta}_{\text{Ext}}] &= (\theta_{\text{Int}} - \widehat{\theta}_{\text{Int}}) - (\theta_{\text{Ext}} - \widehat{\theta}_{\text{Ext}})
 \end{aligned}$$

C.2 Results

For each case, we report in Table C.2 the mean bias as well as select points in the distribution of the biases over the 1,000 replications (we use the 10th, 25th, 50th, 75th and 90th percentiles). Introducing noise in the measurement of the prior, or individual heterogeneity for which our estimator does not account, has a negligible effect on the point estimates of parameters or the differences between internal and external members. It certainly appears that the estimates of the σ parameters are somewhat more sensitive to the precision of the underlying data than the θ estimates. Nonetheless, the effect of both additions is small on the σ parameter estimates. To give some context to the numbers, note that our estimated difference between internal and external expertise in section 5 was -0.7 and the tenth percentile in the bias is only -0.13 for case 2 and -0.14 for case 3.

Table C.2: Results of the Monte Carlo Exercises

	Variable	Mean	p10	p25	p50	p75	p90
Case 1	Bias $[\hat{\sigma}_{\text{Int}}]$	0.01	-0.13	-0.07	0.00	0.08	0.15
	Bias $[\hat{\sigma}_{\text{Ext}}]$	0.01	-0.12	-0.06	-0.00	0.08	0.14
	Bias $[\hat{\sigma}_{\text{Int}} - \hat{\sigma}_{\text{Ext}}]$	0.00	-0.19	-0.11	-0.01	0.11	0.21
	Bias $[\hat{\theta}_{\text{Int}}]$	0.00	-0.01	-0.01	0.00	0.01	0.01
	Bias $[\hat{\theta}_{\text{Ext}}]$	-0.00	-0.03	-0.01	0.00	0.01	0.03
	Bias $[\hat{\theta}_{\text{Int}} - \hat{\theta}_{\text{Ext}}]$	0.00	-0.03	-0.01	-0.00	0.02	0.03
Case 2	Bias $[\hat{\sigma}_{\text{Int}}]$	-0.06	-0.20	-0.14	-0.06	0.01	0.08
	Bias $[\hat{\sigma}_{\text{Ext}}]$	-0.12	-0.26	-0.19	-0.12	-0.05	0.02
	Bias $[\hat{\sigma}_{\text{Int}} - \hat{\sigma}_{\text{Ext}}]$	0.06	-0.13	-0.05	0.05	0.16	0.24
	Bias $[\hat{\theta}_{\text{Int}}]$	-0.00	-0.01	-0.01	-0.00	0.01	0.01
	Bias $[\hat{\theta}_{\text{Ext}}]$	-0.03	-0.06	-0.04	-0.03	-0.01	0.01
	Bias $[\hat{\theta}_{\text{Int}} - \hat{\theta}_{\text{Ext}}]$	0.03	-0.01	0.01	0.02	0.04	0.07
Case 3	Bias $[\hat{\sigma}_{\text{Int}}]$	-0.04	-0.17	-0.11	-0.04	0.03	0.11
	Bias $[\hat{\sigma}_{\text{Ext}}]$	-0.09	-0.24	-0.17	-0.09	-0.02	0.05
	Bias $[\hat{\sigma}_{\text{Int}} - \hat{\sigma}_{\text{Ext}}]$	0.06	-0.14	-0.05	0.05	0.17	0.25
	Bias $[\hat{\theta}_{\text{Int}}]$	0.00	-0.01	-0.01	0.00	0.01	0.01
	Bias $[\hat{\theta}_{\text{Ext}}]$	-0.01	-0.05	-0.03	-0.01	0.01	0.02
	Bias $[\hat{\theta}_{\text{Int}} - \hat{\theta}_{\text{Ext}}]$	0.01	-0.02	-0.01	0.01	0.03	0.05