

Dynamic Voting Patterns on the MPC*

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

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

April 4, 2010

Abstract

This paper attempts to explain the empirical finding that new external members of the Bank of England's Monetary Policy Committee vote for the same average interest rate as new internal members, but that experienced external members vote for lower average rates than experienced internal members. We derive a simple Bayesian decision rule to predict members' votes as a function of their preferences, the precision of their private information, and the prior distribution on the correct decision. This allows one to distinguish explanations for dynamic voting that rely on learning and preferences. Using market survey data to test the model, we find evidence that externals react less to market beliefs than internals in all periods and that externals' preferences are not time-invariant.

PRELIMINARY AND INCOMPLETE:

- PLEASE DO NOT CIRCULATE WITHOUT PERMISSION.
- PLEASE DO NOT CITE.
- COMMENTS AND SUGGESTIONS ARE WELCOMED.

Keywords: Voting, Committees, Monetary Policy

JEL Codes: D72, D78, E52

*The authors acknowledge, without implicating, the initial advice and guidance of Tim Besley, Francesco Caselli, Antonio Ciccone, Thomas Cunningham, Francesco Giavazzi, Charles Goodhart, Gilat Levy and Andrew Oswald. We have also benefited from the comments and suggestions of seminar participants at the Australia National University, the Bank of England, LSE, Monash University, Univeristy Pompeu Fabra, University of Strathclyde and University of Warwick.

1 Introduction

A dramatic change has occurred in how central banks around the globe determine monetary policy: responsibility for setting interest rates has shifted from individuals to committees. In fact, Pollard (2004) reports that ninety percent of eighty-eight surveyed central banks use committees to decide interest rates, underscoring their growing ubiquity. There are also growing calls for the UK to follow other countries and establish an independent committee of experts to provide input into the process of fiscal policy (Besley and Scott 2010).

Although the trend is heavily in favour of collective decision making, some fundamental issues regarding the optimal structure of committees remains unclear. One of these is whether committee members should come from heterogeneous or homogeneous backgrounds. Some central banks, like the European Central Bank and US Federal Reserve, have committees composed solely of internal members (experts employed within the bank). Others, like the Bank of England and Reserve Bank of Australia, have committees that consist of internal as well as external members (experts who are not part of central bank staff). Such appointments are often justified using a sort of Condorcet intuition. In complementary theoretical work, we explore the desirability of having a “mixed” committee (Hansen and McMahon 2010) and find that, even assuming that members behave according to the remit of the committee, the benefits of a “mixed” committee can be achieved without giving the external members a vote but rather allowing them simply to advise a chosen decision-maker.

The goal of this paper is two-fold. The first is to examine whether the members of a mixed committee, the Bank of England’s Monetary Policy Committee (MPC), do in fact behave according to the theoretical ideal given the remit of the MPC. Although numerous other papers have explored their voting behaviour, they do so using reduced form estimates which, as we shall explain, are difficult to properly interpret. Using a framework that more closely links the voting theory with our empirical methodology, we show that external members behaviour differs from those of internal members. The cause of the different behaviour appears to be related to a change in the preference parameter of voting members. This result is interesting in its own right; dynamic voting not well explored theoretically or empirically.

The second goal is to explore the reasons for such a preference shift. We focus on a reputational career concerns explanation (though we also discuss numerous potential explanations which would, in our model, appear as a change in the preference parameter such as risky shift). TO FINISH - THIS IS WORK IN PROGRESS.

In the next section we set out a brief description of the institutional details of the

MPC, as well as a providing some details on their voting behaviour. We show, as was previously established in Hansen and McMahon (2008), that the voting behaviour of external members of the committee appears to change as they gain experience on the committee. There are two main potential explanations for this change in behaviour; the first relates to the growing expertise of the external members over the early part of their tenure, while the second relates to a change in their underlying preference parameters over time. In order to try to disentangle these effects, in Section 3 we explore a simple model which we believe captures the main institutional features of the MPC. We find that simple reduced form regression estimates, as have commonly been interpreted in the literature, cannot be easily used to identify the effects of each explanation.

Instead, we are able to use a clear empirical prediction of our model to motivate an equation to disentangle the effects of learning about the environment, and changing preferences¹. Our model predicts that voting behaviour responds differently to the two explanations when the choice of interest rates is easily predicted, compared to when there is a more difficult choice to make. Using the estimated Probability Distribution Function for the change in interest rates, as approximated by market economists in the days before each meeting, we can isolate how external and internal members behaviour compares with the predictions of our simple model under each explanation. The result is clear; there is limited evidence to support the learning explanation, while the behaviour closely mimics the effects of a change in the preference parameter. This analysis is presented in Section 4.

Section 6 examines a number of possible explanations that would explain the observed preference shift for external members. We conclude with a discussion of the policy implications for committee design and the institutional arrangements for monetary, and other, committees of experts.

2 The Monetary Policy Committee

2.1 The Institutional Arrangements

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 for setting interest rates in order to make monetary policy less

¹We do not, in this version, explore the possibility that members are learning about their preferences. This would only help to explain their preferences if the starting value of the external members' preferences was consistently upward biased. We shall explore this further in a future version.

arbitrary and susceptible to election cycles. The MPC first convened on 6 June 1997, and has met every month since. Majority vote determines the rate of interest. Its remit, as defined in the Bank of England Act (1998) (<http://www.bankofengland.co.uk/about/legislation/1998act.pdf>) is to “maintain price stability, and subject to that, to support the economic policy of Her Majesty’s government, including its objectives for growth and employment.” In practice, the committee seeks to achieve a target inflation rate of 2%², 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 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 Market Operations. 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. According to the Bank of England³, 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.” 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. Through June 2008, 25 different members have served on the MPC – 11 internal members and 14 external 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.⁴

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. On the Wednesday of the meeting, members discuss their views on several is-

²This 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%.

³<http://www.bankofengland.co.uk/monetarypolicy/overview.htm>

⁴According to the Bank of England website (<http://www.bankofengland.co.uk/monetarypolicy/overview.htm>)

Each member of the MPC has expertise in the field of economics and monetary policy. Members do not 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.

sues. The discussion continues on Thursday morning; each member is given some time to summarize his or her views to the rest of the MPC, and suggest what vote they favour (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 policy, concludes with the Governor, and other members are selected in random order in between. 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 12 noon. Two weeks after each meeting, members' votes are published, along with minutes of the meeting with full, but unattributed comments ⁵.

2.2 Previous Research on MPC Voting

There is substantial research that examines the issue of monetary policy committees in general and considers such questions as to their desirability and institutional features. Blinder (2007) provides an excellent coverage of the issues covered in the literature to date.

In addition, several papers examine the empirical differences in voting behaviour among Bank of England MPC members using the data set that we use in this paper. Many of the papers focus on the distinction between internal and external members⁶; this is an obvious choice of segmentation as (i) there are differences in the members' tasks as internal and externals (though all, obviously, vote for a desired interest rate decision), (ii) it is likely that heterogeneity amongst the groups, if present on the committee, will be present between these two groups, and (iii), for the later papers, the existing literature had focused on this distinction. In this regard, Gerlach-Kristen (2003), Spencer (2006), Harris and Spencer (2008), and Gerlach-Kristen (2009) all document the tendency of external members to dissent more often and to favour lower interest rates than internal members. Bhattacharjee and Holly (2005) and Besley, Meads, and Surico (2008) consider member heterogeneity more broadly, and find that there are systematic voting differences across members.

By and large, these papers assume member preferences derive from a weighted sum

⁵In the model below, we take for granted that there is transparency of voting behaviour of MPC members and that MPC meeting minutes are published; without such a design structure, the nature of our empirical work would be impossible. As a result our paper is not contributing to general discussion of whether having a committee influences monetary policy outcomes (interested readers are pointed toward Sibert (2006), Sibert (2003) and the references therein), or on the debate about optimal degree of transparency (see, for example, Geraats (2006) and Sibert (2002)).

⁶Most papers that focus on the internal and external groupings do not explore the normative implications of including internal and external members on the same committee. This is the focus of another of our papers - Hansen and McMahon (2010).

of inflation and output, with different members having different weights; however, these different weights are assumed to be constant. Hansen and McMahon (2008) uncovers the effect of experience on external members⁷; we argue in that paper that this “delayed dovishness” provides important information to try to understand what drives the differences in voting behaviour between internal and external members. The type of reduced form regressions estimated in that paper are reproduced in Table 1 below using the same data on MPC voting that is used in the literature; these data are described fully in the Section 4. The regressions use as the dependent variable a member-date specific variable that measures of whether member i voted for the higher of the possible interest rates considered in meeting t :

$$D(\text{votehigh})_{it} = \begin{cases} 0 & \text{if member } i \text{ in time } t \text{ votes for the lower interest rate} \\ 1 & \text{if member } i \text{ in time } t \text{ votes for the higher interest rate} \end{cases}$$

Column (1) reports the estimated coefficients, where ψ_1 gives the estimated probability voting for the higher interest rate, from the following regression model:

$$\text{Probit} [D(\text{votehigh})_{it}] = \alpha + \psi_1 \cdot D(\text{EXT})_i + \varepsilon_{it} \quad (1)$$

where

$$D(\text{EXT})_i = \begin{cases} 0 & \text{if member } i \text{ in time } t \text{ is an internal member} \\ 1 & \text{if member } i \text{ in time } t \text{ is an external member} \end{cases}$$

Column (2) and (3) decompose the estimated effect of being an external MPC member into the effect of being new and experienced. The estimated model is:

$$\begin{aligned} \text{Probit} [D(\text{votehigh})_{it}] = & \alpha + \psi_2 \cdot D(\text{EXT})_i \\ & + \psi_3 \cdot D(\text{NEW})_{it} + \psi_4 \cdot D(\text{EXT})_i \cdot D(\text{NEW})_{it} + \varepsilon_{it} \end{aligned} \quad (2)$$

where

$$D(\text{NEW})_{it} = \begin{cases} 1 & \text{if member } i \text{ in time } t \text{ has been on the committee 12 months or less} \\ 0 & \text{if member } i \text{ in time } t \text{ has been on the committee more than 12 months} \end{cases}$$

As Table 1 shows, external members are less likely to vote for the higher interest rates (column 1); this is the standard “dovish externals” result found in the literature described above. Column (2) confirms the “delayed dovishness” findings in Hansen and McMahon (2008); *ceteris paribus*, it is only after they develop some experience on the committee

⁷Relatedly, however, Gerlach-Kristen (2003) discusses a delay in a member’s first dissent: on average, it occurs after nine months.

that externals begin to vote for lower rates. Column (3) shows that this result is robust to the inclusion of time-fixed effects⁸, while Column (4) includes member specific intercepts via a random effects estimation. In each case, we cannot reject two null hypotheses: (i) the hypothesis that new external members have the same probability of voting for the higher rate as internal members, and (ii) the hypothesis that experienced external members have a lower probability of voting for the higher rate.

	Dependent Variable is D(High Vote)			
	(1)	(2)	(3)	(4)
D(External)	-0.328*** (0.0842)	-0.468*** (0.105)	-1.765*** (0.218)	-0.482** (0.200)
D(New)		-0.0416 (0.131)	-0.144 (0.254)	-0.0339 (0.138)
D(External) x D(New)		0.355* (0.184)	1.724*** (0.337)	0.362* (0.200)
Constant	0.136** (0.0559)	0.146** (0.0640)	-0.154 (0.490)	0.145 (0.138)
Observations	908	908	531	908
Number of members				24

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 1: Reduced Form Results in the Spirit of Hansen-McMahon (2008)

These reduced-form estimates are indicative of an interesting dynamic voting pattern for external members. In particular, two main explanations are usually suggested; the first is that, on joining the MPC, external members are inexperienced and thus lack the confidence to express their own views in the meetings. Over time, they begin to gain experience and confidence and so express themselves more. The second explanation is that there is something that affects the preferences of external members over time such as career concerns. While Spencer (2006) and Harris and Spencer (2008) find no evidence of career concerns on the MPC, Meade and Stasavage (2008) find evidence of career concerns on the Federal Open Market Committee in the US.

⁸As our dataset contains a limited number of cross-sectional observations (24 members), and a large number of time-series observations (though the sample is strongly unbalanced), we have a time-series, cross-sectional analysis (rather than a panel). As such, the results rely on asymptotics in T and therefore there reason to worry that inclusion of time-effects is subject to the incidental parameters problem ((Cameron and Trivedi 2009)).

Each of these explanations has a different impact on policy and the design of the MPC, in the absence of some theory to help guide us, it is difficult to know how to interpret these results. In particular, it is not clear how the coefficients from these reduced form models should change in the event that it is one explanation compared with the other. We therefore, in the next section, explore a very simple model of voting behaviour.

3 Model

Our goal is to write down the simplest possible model that captures the main elements of voters' decision each month and that allows one to distinguish between alternative explanations for dynamic voting behavior. The major simplification we make is to treat each member's vote as an individual decision problem, an assumption that we discuss in some detail below.

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. The restriction that a vote must take one of two values is not as restrictive as it might first appear, since one only observes three different votes in 7 of the 142 meetings in our sample and never observes four or more distinct votes. The utility that person i receives from his vote depends on an unknown state of the world $\omega_t \in \{0, 1\}$ where $\Pr[\omega_t = 1] = q_t$ is common knowledge for all members. The following table represents member i 's preferences over votes and states of the world:

		vote	
		0	1
state	0	1	-1
	1	$-\theta_i$	1

All members have the same policy goal: choose the vote that corresponds to the state of the world. However, differences arise over how to treat errors. A member with $\theta > 1$ considers voting 0 when the correct state is 1 to be a worse error than voting 1 when the correct state is 0, and vice versa for a member with $\theta < 1$.⁹

Before choosing a vote in period t , member i receives a private signal $s_{it} \sim N(\omega_t, \sigma_{it}^2)$ where the variance term σ_{it}^2 measures the accuracy of s_{it} . After observing s_{it} member i solves a simple Bayesian decision problem whose solution provides our first result.

Proposition 1

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

⁹We could also parametrize the upper right element of the above table, but this would not add anything substantive to the analysis.

where

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

Proof. First let

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

be member i 's updated belief on ω_t after observing s_{it} . Here $f_1 \sim (1, \sigma_{it}^2)$ 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 $\omega_t = 0$.

The expected utility that member i derives from voting 1 is $\widehat{\omega}(s_{it}) - (1 - \widehat{\omega}(s_{it}))$ and the expected utility he derives from voting 0 is $-\theta_i \widehat{\omega}(s_{it}) - (1 - \widehat{\omega}(s_{it}))$. So member i votes 1 whenever $\widehat{\omega}(s_{it}) \geq \frac{2}{3 + \theta_i}$. This is equivalent to voting for 1 whenever

$$\frac{1 + \theta_i}{2} \frac{q_t}{1 - q_t} \geq \frac{f_0(s_{it})}{f_1(s_{it})}. \quad (4)$$

Since the normal distribution satisfies MLRP, $\frac{f_0(s_{it})}{f_1(s_{it})}$ is strictly decreasing in s_{it} , implying a cutoff rule for voting. To derive the threshold value, one can take logs on both sides of (4) to yield

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

from which the result derives. ■

Votes are determined by a cutoff rule in which members vote for the high rate if and only if their private signal exceeds a threshold. Since the normal distribution satisfies the monotone likelihood ratio property, higher values of s_{it} provide stronger evidence that the correct state is 1. Of course, the threshold value s_{it}^* determines how high s_{it} must be in order for member i to choose 1, and this depends on several factors. Firstly it depends on preferences. An increase in θ_i decreases s_{it}^* : if member i becomes more averse to falsely voting 0, he or she requires less evidence that the state is 1 to vote 1. Secondly it depends on the prior q_t . When q_t increases, s_{it}^* decreases: if there is a higher chance that the state of the world is 1, member i becomes more likely to vote 1. Moreover, as $q_t \rightarrow 0$, $s_{it}^* \rightarrow \infty$ since member i will ignore his private information and simply vote 0 whenever the state of the world is known to be zero; likewise, as $q_t \rightarrow 1$, $s_{it}^* \rightarrow -\infty$. Finally, the threshold depends on the variance of the signal, which we discuss below.

Before proceeding it is important to discuss what relevance an individual decision problem has for understanding MPC voting. In particular, one might be concerned that our model misses out some important group effects. As we have already mentioned, the MPC was founded on the basis of a one-person one-vote philosophy, so that members

should simply vote for the rate that they believe comes closest to hitting the inflation target. Of course this does not rule out the sharing of information, and members indeed spend a great deal of time in discussion prior to voting. In our model q_t reflects all information that members have at the point at which they vote, and this includes any information revealed to the group during discussions. So, although we do not model discussion, it is not incompatible with our behavioral prediction.

Another potential problem arises if members' votes are influenced not just by others' stated views but also by others' actual votes. These effects might arise for either informational or strategic reasons. The voting order on the committee is ostensibly random, and to the extent that this is the case, one might not expect large systematic deviations from our theoretical voting rule. More to the point, we do not observe the voting order on the MPC, so we could not examine a model that relied on voting order; for this reason, we do not attempt to construct one.

Finally one might wonder why members' preferences depend on their votes matching the state of the world rather than the committee decision matching it. This assumption is equivalent to stating that each member behaves as if he or she is pivotal in each period. Strictly speaking, this is not the case; once 5 people have voted for one particular rate change, all remaining votes do not affect the committee decision. However, Gerlach-Kristen (2004) shows that markets react to dissenting votes by adjusting the yield curve. For example, suppose that in a nine-member MPC the first seven voters choose not to raise rates. The eighth voter may not be able to influence the committee decision, but by voting to cut rates, he or she has the opportunity to push down anticipated future rates, which has the same directional effect on inflation as an actual cut in rates. In this sense, each member is contributing to monetary policy independently of the decision, giving support to our assumption on preferences.

Although simple, the value of our model comes via its ability to discriminate between the effects of the variance and preference parameters on voting behavior. The next result provides the key for connecting our model to the data.

Proposition 2 $\frac{\partial \Pr[v_{it}=1]}{\partial \sigma_{it}} > 0$ if and only if $q > q^* = \frac{2}{3+\theta}$. In contrast $\frac{\partial \Pr[v_{it}=1]}{\partial \theta_i} > 0$ for all q .

Proof. For simplicity of notation we will drop the member i and time t subscript. The probability of voting for 1 is equal to

$$\begin{aligned} & \Pr[s > s^* | \omega = 0] \Pr[\omega = 0] + \Pr[s > s^* | \omega = 1] \Pr[\omega = 1] \\ &= \left[\int_{s^*}^{\infty} f_0(s) ds \right] (1 - q) + \left[\int_{s^*}^{\infty} f_1(s) ds \right] q \end{aligned} \quad (5)$$

where

$$s^* = \frac{1}{2} - \sigma^2 \ln \left[\frac{q}{1-q} \frac{1+\theta}{2} \right].$$

Differentiating (5) with respect to σ^2 gives

$$\begin{aligned} & - \frac{\partial s^*}{\partial \sigma^2} [f_0(s^*)(1-q) + f_1(s^*)q] \\ &= - \frac{\partial s^*}{\partial \sigma^2} \left[f_1(s^*) \frac{1+\theta}{2} q + f_1(s^*)q \right] \end{aligned} \quad (6)$$

since by definition $\frac{f_0(s^*)}{f_1(s^*)} = \frac{1+\theta}{2} \frac{q}{1-q}$. So (6) is positive if and only if $-\frac{\partial s^*}{\partial \sigma^2} = \ln[\frac{1+\theta}{2} \frac{q}{1-q}]$ is positive. This holds if and only if $q > q^* = \frac{2}{3+\theta}$. ■

Suppose that member i becomes more informed about the state of the world in the sense that the variance of his private information decreases. The effect that this has on the probability with which he votes for the higher rate depends crucially on the value of the prior q_t . A decrease in σ_{it}^2 makes him place more weight on his private information and less weight on the prior when choosing a vote. The result is that a better informed member i is more likely to vote against whichever state of the world q_t favors. When q_t is high, member i is more likely to vote for the lower rate and when it is low, he is more likely to vote for the higher rate. In contrast, if member i 's preferences change so that θ_i increases, he is more likely to vote for the high rate no matter what is the value of q_t .

Proposition 2 allows for a clear discussion of the effect of experience on voting. If members acquire more expertise with experience (in other words, the variance of their private information decreases), they do not become systematically more likely to vote for higher or lower rates; instead they become more likely to follow their private signal and less likely to rely on the prior, so that the effect on voting for the high rate depends on the magnitude of the prior. On the other hand, if members change their preferences with experience and become more or less concerned with falsely voting 0, they become systematically more or less likely to vote for the high rate.

This can be seen clearly in the following simulated plots of the theoretical decision rule. In these figures, we compare how the probability of voting for the higher interest rate varies with the common prior. In the baseline case we set $\sigma = 1$ and $\theta = 1$ (symmetric preferences); as the common prior increases, the probability of voting for the higher rate increases. In 1, we simulate the acquisition of greater private information by lowering the variance of the private signal from $\sigma = 1$ to $\sigma = 0.25$. The effect of greater precision of private knowledge is to rotate the decision rule around 0.5 (or a higher (lower) of preferences reflect more hawkish (dovish) views) meaning that the probability of voting for the high rate is less driven by the common prior. In 2, we simulate a change

in preferences (toward more dovish behaviour) by lowering the preference parameter from $\theta = 1$ to $\theta = 0.25$. In this case, the decision rule shape is broadly maintained but the line shifts out to the right indicating that doves are less likely to vote for the high rate for every value of the common prior.

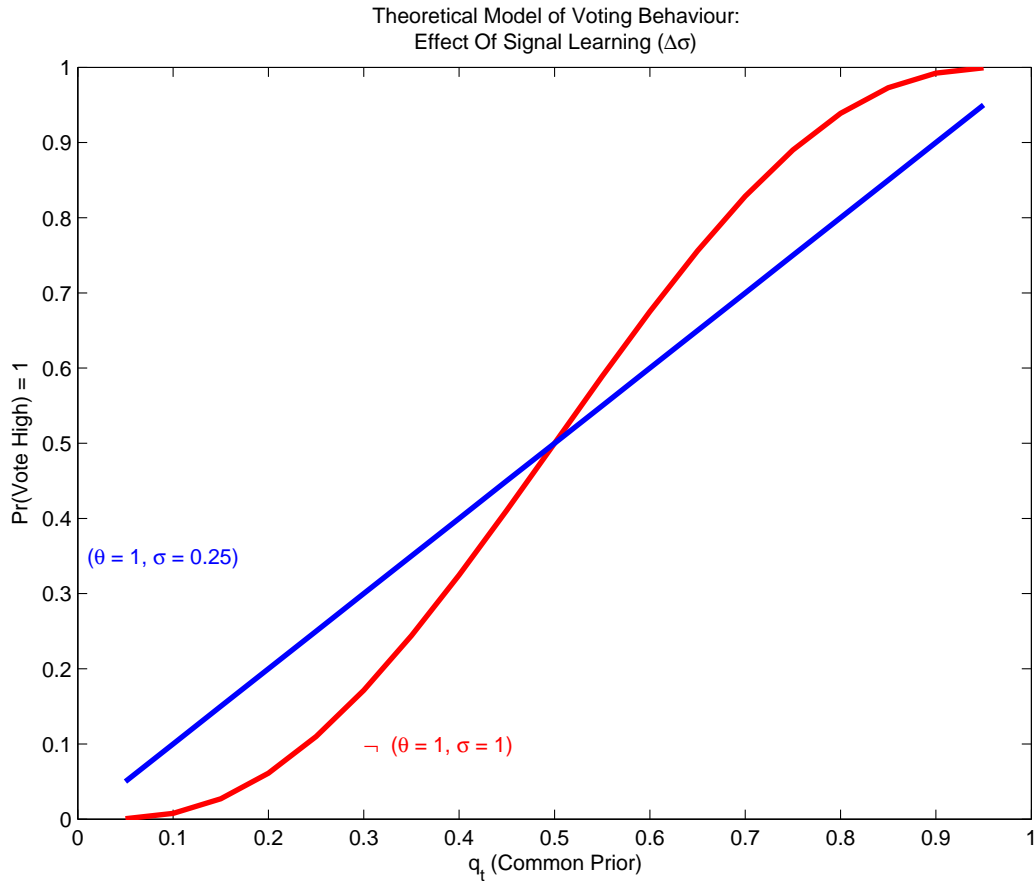


Figure 1: The Effect of Growing Private Expertise

4 Empirical Analysis

4.1 Identification and Reduced Form Regressions

As shown above, and using the fact that $s_{it} = \omega_t + \epsilon_{it}$ where $\epsilon_{it} \sim N(0, \sigma_{it}^2)$, we can rederive the decision to vote for the higher interest rate ($D(HighVote) = 1$) in terms of

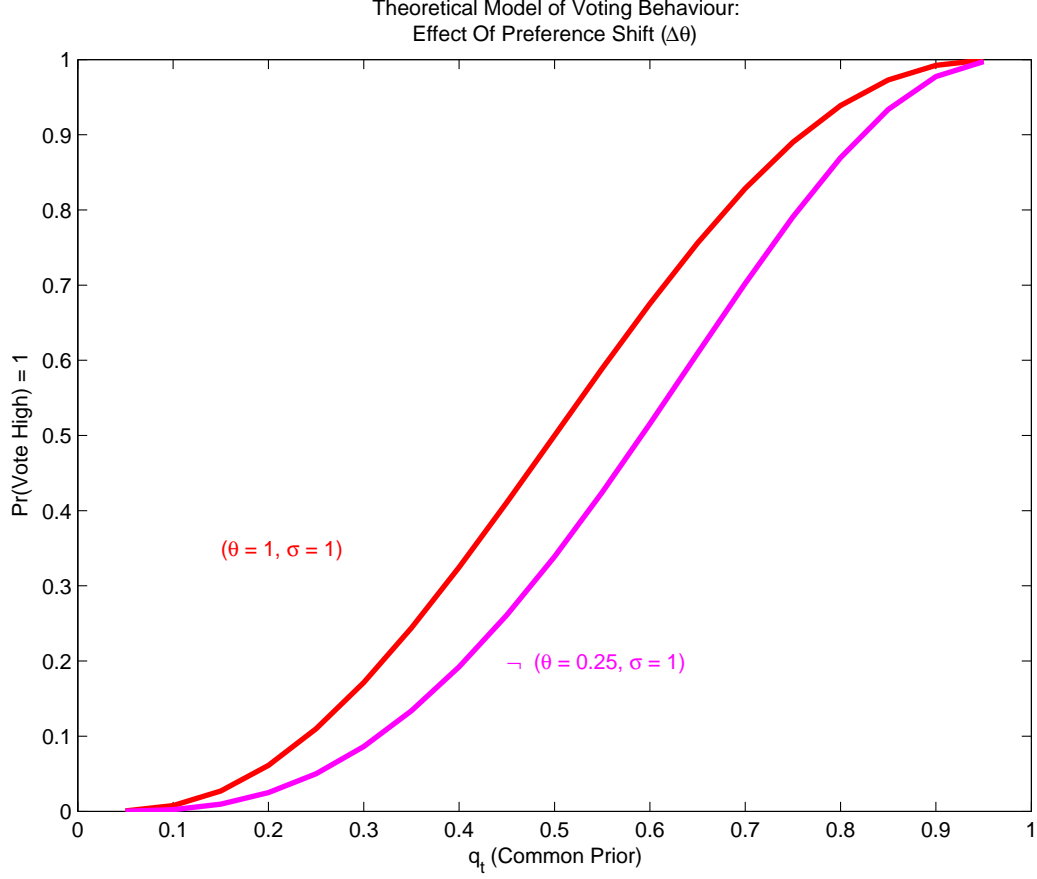


Figure 2: The Effect of Different Preferences

ϵ_{it} ¹⁰:

$$\epsilon_{it} \geq \frac{1}{2} - \sigma_{it}^2 [\ln(1 + \theta_i) - \ln(2) + \ln(q_t) - \ln(1 - q_t)] \quad (7)$$

Or, writing in terms of a standard normal error term:

$$\eta_{it} \geq \frac{1}{2\sigma_{it}} - \sigma_{it} \left[\ln(1 + \theta_i) - \ln(2) + \ln \frac{(q_t)}{(1 - q_t)} \right] \quad \eta_{it} \equiv \frac{\epsilon_{it}}{\sigma_{it}} \sim N(0, 1) \quad (8)$$

¹⁰We rewrite the likelihood ratio of the signal in terms of ϵ_{it} giving:

$$\begin{aligned} \frac{f_0(s_{it})}{f_1(s_{it})} &= \frac{f_0(\epsilon_{it})}{f_1(1 + \epsilon_{it})} \\ \Rightarrow \ln \left[\frac{f_0(s_{it})}{f_1(s_{it})} \right] &= \ln \left[\frac{1}{\sqrt{2\pi\sigma_{it}^2}} e^{-\frac{\epsilon_{it}^2}{2\sigma_{it}^2}} \right] - \ln \left[\frac{1}{\sqrt{2\pi\sigma_{it}^2}} e^{-\frac{(1+\epsilon_{it})^2}{2\sigma_{it}^2}} \right] = \frac{1}{2\sigma_{it}^2} + \frac{\epsilon_{it}}{2\sigma_{it}^2} \end{aligned}$$

The problem of identification is immediately obvious from 8; the expertise variable ($\frac{1}{\sigma_{it}^2}$) and the preferences variable (θ_{it}) enter multiplicatively. In particular, when we include the $D(Ext)$, $D(New)$ and $D(Ext) \times D(New)$ variables, we can capture the average change with experience in the intercept of the decision rule (Equation 8) for internals and externals. However, it is not possible to identify whether this change in the reduced form coefficient is driven by changes in expertise or preference changes.

This is the power of Proposition 2; it provides clearly identifiable predictions on the effects of learning versus preference shifts on observed outcomes. However, in order to explore the predictions of Proposition 2, we use the same data as above but we augment it with a measure of the common prior. Before conducting the analysis, we describe the voting data and then the measure of common prior that we use.

4.2 Voting Data

We use the MPC voting records between July 1997 and July 2008¹¹; these data are freely available from <http://www.bankofengland.co.uk/monetarypolicy/decisions.htm>. The data contain a record of every decision ($decision_t$) taken by the MPC, as well as each member's vote in each meeting ($vote_{it} = \Delta r_{it}$).¹² 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. The Bank website also provides information on which members were external appointments and which were internal. For every member we gathered biographical information, including previous occupation, educational background, and age from press releases associated with their appointment and from information provided to the Treasury Select Committee ahead of their confirmation.

We drop from our dataset the emergency meeting held after September 11th, at which the vote was unanimous in favour of lowering interest rates, for the programming convenience of having only one meeting per month. Howard Davies served on the MPC for the first 2 meetings and is the only member who voted exclusively on unanimous committees but we include him in our baseline regressions nonetheless. Lord George, the Governor for most of our sample, always voted with the majority regardless of his starting position; as a result we think that these voting records do not represent his own

¹¹We stop at this point only because of availability of the data on the common prior (described below). We are currently working to extend this to March 2009 at which point the voting over interest rate decisions (temporarily) ceased and shifted to asset purchase decisions related to quantitative easing.

¹²We express members' votes in terms of their preferred change in interest rates rather than their preferred level. This makes no difference to the results.

views in all cases. Even under the governorship of Mervyn King, the Governor has only deviated twice since taking office in July 2003. Nonetheless, we include the observations for the Governor in the regression results presented below, though all of the results stand if we exclude the data on the Governor at each meeting.

In Table 2 we provide summary statistics of the individual members on the MPC. Of the 24 MPC members that we consider in our sample, 14 are external and 10 are internal as indicated by the variable $D(\text{External})$ ¹³. It is also clear from Table 2 that the educational background of both groups is heterogenous as shown by the $D(\text{Educ})$ variable:

$$D(\text{Educ})_i = \begin{cases} 1 & \text{if } i\text{'s highest degree is a Bachelors degree} \\ 2 & \text{if } i\text{'s highest degree is a Masters degree or Professional qualification} \\ 3 & \text{if } i\text{'s highest degree is a PhD or they held a Professorship} \end{cases}$$

The fact that the MPC has a changing membership, and that different members have served for differing amounts of time is clear from the statistic on “Total Service”. In Figure 3 we explore the overlapping nature of different members more closely with only the current Governor, Mervyn King, having served throughout our sample.

The MPC members also differ in their propensity to deviate from the committee; two members, Richard Lambert and Sir Eddie George, never cast a dissenting vote in their time on the committee. David Blanchflower and Sir Alan Budd dissent most frequently with around 40% of their votes as dissenting votes. Finally, as it relates to our model most closely, the table explores how often each member votes for the higher interest rate on offer. Again, there is large variation between members; Andrew Sentance, Sir Andrew Large and Tim Besley all voted for the higher rate in more than 70% of their meetings, while David Blanchflower voted for the higher rate in only 6% of his meetings.

As this illustrates, disagreement is quite common; we find that 58% of the 104 meetings in our sample have at least one deviation from the committee majority. Figure 4 shows the level of interest rate chosen by the MPC, where the markers indicate the votes of individual members; deviations from the majority are those that are off the MPC decision line. These deviations occur regularly and not just around turning points in the interest rate cycle (marked with shading on the figure). Figure 5 shows how many deviations occur in each meeting.

The key variable in our analysis, defined on page ??, is $D(\text{High Vote})$. In those meetings in which there are two unique votes, the “high vote” is calculated as the higher of the two votes cast by members. In situations in which there is only one vote cast in the

¹³No member has so far served as both an external member and an internal member, though there is nothing that prohibits this from happening in the future.

Member	D(External)	D(Educ)	Total Service	Deviations		High Votes	
				Total	%	Total	%
Andrew Sentance	1	3	14	4	29	11	79
Charles Bean	0	3	78	4	5	39	50
Charles Goodhart	1	3	22	3	14	12	55
Christopher Allsopp	1	2	31	10	32	8	26
David Blanchflower	1	3	18	7	39	1	6
David Clementi	0	2	42	4	10	19	45
David Walton	1	2	10	2	20	8	80
DeAnne Julius	1	3	29	6	21	3	10
Ian Plenderleith	0	2	41	2	5	18	44
John Vickers	0	3	16	4	25	10	63
Kate Barker	1	1	73	4	5	37	51
Marian Bell	1	2	35	5	14	13	37
Mervyn King	0	3	104	13	13	65	63
Paul Tucker	0	1	63	6	10	38	60
Rachel Lomax	0	2	50	4	8	24	48
Richard Lambert	1	1	32	0	0	16	50
Sir Alan Budd	1	3	9	4	44	6	67
Sir Andrew Large	0	2	38	8	21	29	76
Sir Edward George	0	1	54	0	0	28	52
Sir John Gieve	0	1	21	2	10	11	52
Stephen Nickell	1	3	64	14	22	27	42
Sushil Wadhvani	1	3	27	8	30	5	19
Tim Besley	1	3	15	4	27	11	73
Willem Buiter	1	3	22	7	32	12	55
Average			53	6	14	27	50

Table 2: Sample Statistics

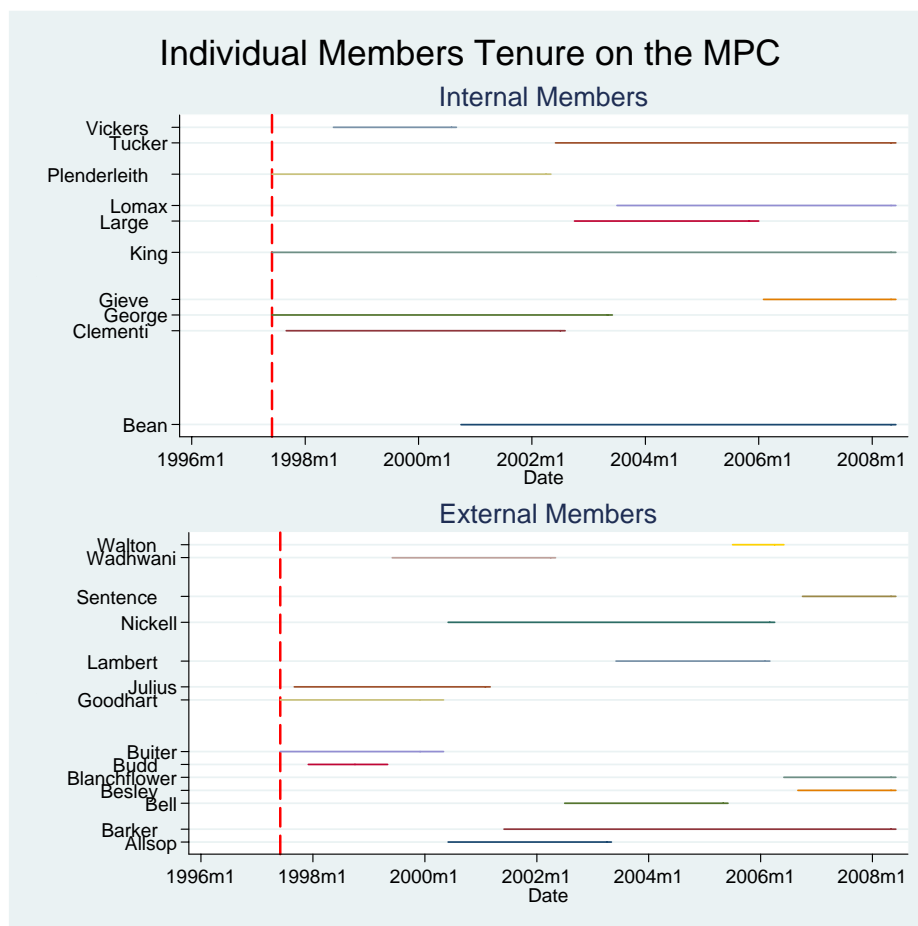


Figure 3: Tenure of MPC Members

MPC, we do not wish to choose arbitrarily that this vote is high or low; instead, we can use the Reuters data described below to indicate what, in terms of the financial markets, was the alternative possible interest rate (no matter how unlikely). Once we have our “high vote” for each meeting, it is easy to calculate the $D(\text{High Vote})$ dummy; we plot this dummy variable, which will serve as the dependent variable in the following analysis, in Figure 6. For both internal and external members, there are many observations of both high and low votes.

4.3 Data on the Common Prior, q_t

In order to test proposition 2, we need a measure of the common prior on each interest rate decision. For this, we make use of a Reuters survey of market economists that is carried out just before the MPC meeting. We use this measure as our q_t variable. The survey data runs from June 1997 through November 2008 (more recent data is currently being compiled and added). As the survey evolved, it became more sophisticated and

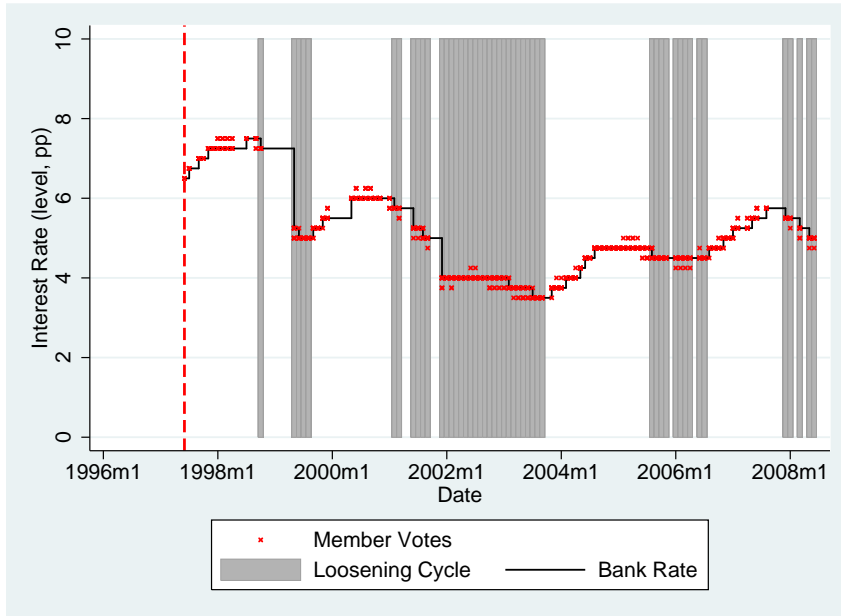


Figure 4: Votes and Decisions of the Monetary Policy Committee




Figure 5: Voting Deviations by MPC Members

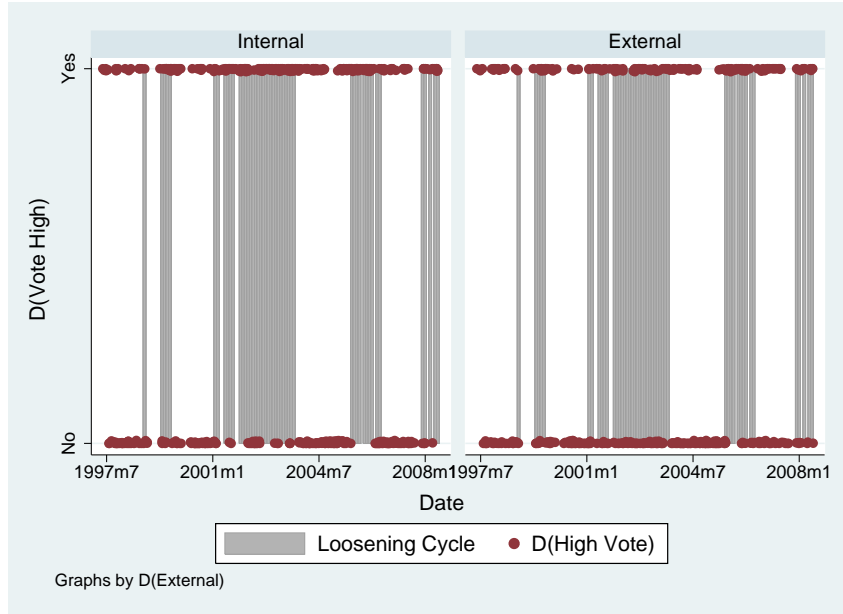


Figure 6: D(High Vote) Variable by D(External)

allowed respondents to give more nuanced opinions. There are three distinct periods, and we employed a different methodology for constructing our proxy for q_t in each one.

From June 1997 through June 1998, the survey simply asked respondents to state the change in rates that they considered most likely. All responses were for either no change or a rise of 25 basis points, except for one in October 1997 that was for a rise of 50 basis points. Moreover, during this period, all members voted for either no change or a rise of 25 basis points, except for June 1998 when members voted over a 25 point rise and a 25 point cut. We dropped this meeting from our sample. To construct q_t for the rest of the periods, we simply took the ratio of responses for a rise of 25 basis points over the total number of responses each period.¹⁴

From July 1998 through December 2001, the survey asked respondents to write down how likely they considered a cut in rates, no change, or a rise in rates, or else restricted them to write a probability distributions over no change and rise or no change and cut. The main informational limitation of this question is that it does not allow respondents to specify how likely they consider different changes. This is particularly problematic for the periods in which the disagreement on the committee was over the magnitude of the change, not whether to make a change.¹⁵ Perhaps to address this concern, in some

¹⁴We treated the response for a rise of 50 basis points as a response for a rise of 25 points since this respondent presumably believed that a rise of 25 basis points was more likely than no change.

¹⁵For example, in January 2000 every member voted to raise by 25 or 50; in April, May, and November 2001 every member voted to cut by 25 or 50; and in November and December of 1998 and February of 1999 every member voted to cut by 50 or 75.

periods the survey asked the question: if there is a change in rates, which change do you consider most likely? However, over 95% of non-missing values for this question have a magnitude of 25, so there is no clear way of constructing the relative weights that the market placed on the different outcomes for periods in which no member votes for no change. For this reason, we only keep periods in which at least one member voted for no change, and assume that during these periods, survey respondents are writing down beliefs over a cut of 25 points, no change, or a rise of 25. The only clear violation of this assumption comes in March 1999, when the two unique votes on the MPC were for no change and a cut of 40 basis points, so we drop this meeting from the sample.

The next step is to identify the two outcomes on which the market places the highest probability (of course, given our assumption on the domain over which respondents are reporting beliefs). To do this, we simply take the average probability placed on each outcome, and consider the two outcomes with the highest average. For consistency, we need to ensure that members' votes in each period lie within this set of outcomes. This is true in all but two periods. In April 2000 the combined average probability for no change and cut of 25 basis points exceeded 95%, while the committee voted over no change and a rise of 25. In December 2000 the average probabilities were nearly symmetric around no cut with a slight tilt towards a rise of 25, while the committee voted over no change and a cut of 25. We drop these two periods from the sample. The final step is to construct our proxy for q_t . For this we take the average probability placed on the higher of the two outcomes on which the market placed the highest probability over the total average probability placed on these two outcomes.

Finally, from January 2002 through November 2008 survey respondents were asked to write down the probabilities they attached to a rise of 50 points, a rise of 25 points, no change, a cut of 25 points, and a cut of 50 points. The one issue with this domain restriction occurs in the last meeting for which we have survey data: in November 2008 all members voted for a cut of 150 basis points, so we drop this from our sample. We then proceed exactly as above. First we compute the two outcomes with the highest average probability and check that members' vote lie in this set. There are three exceptions: in March 2007, the two outcomes that respondents perceived as most likely were no change and cut 25, while members voted on no change and raise 25; in September 2008, respondents placed highest probability on no change and cut 25 while members vote on no change and cut 50; and in October 2008, respondents placed highest probability on no change and cut 25 while members vote unanimously for a cut of 50. We drop these meetings from our sample, and for the rest of the periods we construct q_t by again taking the average probability placed on the higher of the two outcomes on which respondents placed the highest probability over the total average probability placed on these two

outcomes.

Figure 7 shows the time-series behaviour of our constructed q_t variable.

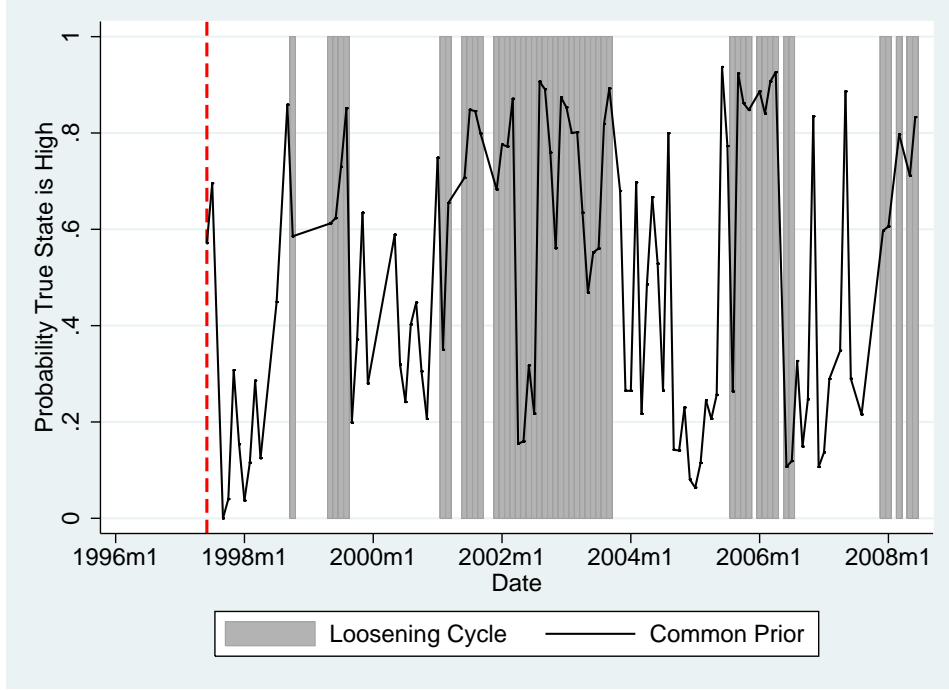


Figure 7: The Evolution of the Market Prior

We also explore a number of robustness checks on our common prior data. First, we estimate the correlation between q_t and the $D(\text{votehigh})_{it}$ variable using equation 9; a reasonable requirement of the constructed q_t data is that ψ_1 is positive. As reported in Column 1 of Table 3, this check is successfully passed.

$$D(\text{Vote High})_{it} = \alpha + \psi_1 \cdot q_t + \sum_t \tau_t \cdot \text{Time}_t + \varepsilon_{it} \quad (9)$$

Our second check on the q_t variable concerns the relationship between the common prior 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 MPC votes within a period should decline; this predicted inverted-U shape relationship can be explored using equation 10. As reported in Column 2 of Table 3, we find the predicted relationship.

$$\text{Std Deviation of Vote}_{it} = \alpha + \psi_1 \cdot q_t + \psi_2 \cdot q_t^2 + \varepsilon_{it} \quad (10)$$

	(1)	(2)
	D(Vote High) _{it}	Std Deviation of Vote _{it}
q_t	1.204*	0.190***
	(0.720)	(0.0316)
q_t^2		-0.177***
		(0.0307)
Constant	-0.363	0.0258***
	(0.470)	(0.00645)
Observations	531	908
R^2	.	0.041

*** p<0.01, ** p<0.05, * p<0.1

Standard errors in parentheses

Table 3: Behavior of Market Prior

4.4 Econometric Results

In order to test the effect of experience on external members, we can estimate equation 2 augmented with our measure of q_t , as well as interaction terms between q_t and our other variables of interest. The equation, presented below (equation 11) allows experience to have a different effect depending on the value of q_t and, in particular, allows us to see how the estimated probabilities of voting for the high interest rate is affected by the common prior for internal and external members with different levels of experience. The estimated equation is:

$$\begin{aligned}
Probit [D(votehigh)_{it}] = & \alpha + \varphi_1 \cdot D(EXT)_i + \varphi_2 \cdot D(NEW)_{it} \\
& + \varphi_3 \cdot D(EXT)_i \cdot D(NEW)_{it} + \varphi_4 \cdot q_t + \varphi_5 \cdot D(EXT)_i \cdot q_t \\
& + \varphi_6 \cdot D(NEW)_i \cdot q_t + \varphi_7 \cdot D(EXT)_i \cdot D(NEW)_{it} \cdot q_t + \varepsilon_{it}
\end{aligned} \tag{11}$$

The estimates are presented in Column (2) of Table 4; Column (1) replicates Column (2) of Table 1 above for comparison. While the coefficient estimates are indicative of differential behaviour of experience depending on the common prior, interpreting these results is made easier if we replicate the theoretical Figures (1 and 2) using our estimated decision rules. In particular, for internal and external members, both new and experienced, we can plot their estimated probability of voting high conditional on the common prior. We do this in Figure 8¹⁶.

¹⁶In Figure 8 we omit confidence intervals as it makes the graphs easier to read and compare across

	(1)	(2)
	D(High Vote)	D(High Vote)
D(External)	-0.468*** (0.105)	-0.567* (0.342)
D(New)	-0.0416 (0.131)	0.248 (0.340)
D(External) x D(New)	0.355* (0.184)	1.324*** (0.492)
q_t		4.205*** (0.330)
q_t x D(External)		-1.894*** (0.685)
q_t x D(New)		-0.313 (0.659)
q_t x D(External) x D(Experienced)		1.306 (0.872)
Constant	0.146** (0.0640)	-1.926*** (0.176)
Observations	908	908

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 4: Main Results

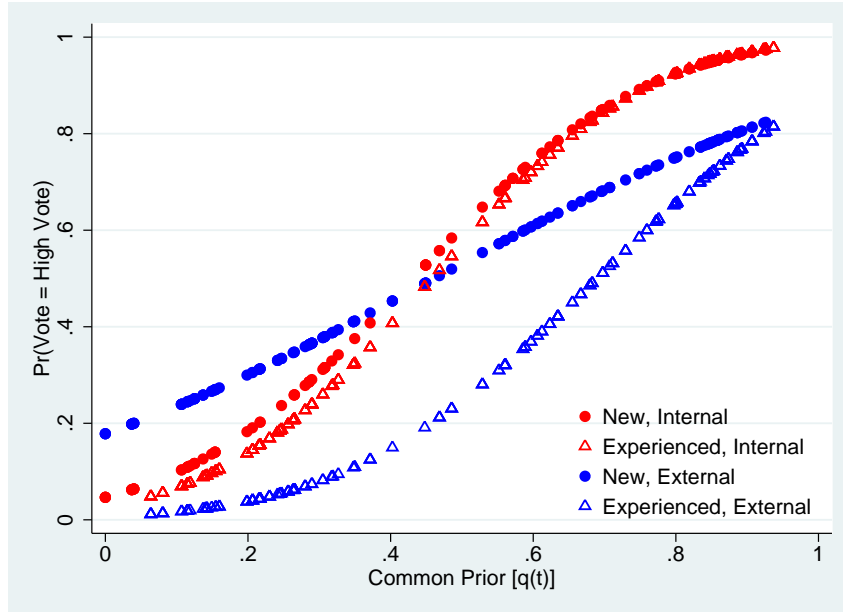


Figure 8: Main Results - Estimated Probability of Voting High Conditional on the Common Prior

A number of results stand out. The first is that it appears that internal and external, when new, have the same preferences (θ_i) but the externals appear to bring more private information to the discussion ($\sigma_{INT} > \sigma_{EXT}$); comparison of the the line of solid blue diamonds (new, externals) with the solid red diamonds (new, internals) mimics our theoretical thought experiment of greater precision (see Figure 1). There are many possible explanations for this; for example, externals are appointed to bring “thinking and expertise in addition to that gained inside the Bank of England” and they may therefore be better informed on some aspects of the information, or the common prior may not, in the early days of a new member, reflect that members information.

The second result that stands out is that while internals display no change in behaviour with experience, the external members, on average, appear to display a clear change in preference. Experienced external members are less likely to vote for the higher interest rate for all values of q_t . This result, equivalent to the theoretical result in Figure 2, confirms our earlier “delayed dovishness” result and suggests that this change in behaviour is driven by a preference change and is not the result of growing experience. In fact, if anything, the extra private information brought by the new external members disappears; this is consistent with the idea that new members views take time to be reflected in the market survey from which we derive the common prior.

This second result has clear implications for the commonly held view in the theoretical

member-types. In the appendix on page 30, figure 11 adds 90% confidence intervals to Figure 8 although we draw a separate graph for internals and for externals in order to make it easier to read.

literature that preferences are a fixed characteristic. Moreover, given the possibility that preferences may display a dynamic pattern has important implications for committee design. We discuss these issues in the concluding section; in the next section we discuss what might cause preferences to shift.

5 Reputation and Voting - WORK IN PROGRESS

The previous model derived the voting rule for individuals as a function of their preferences, information, and prior without career concerns. If we reject this model, which one do we use in its place? There are many candidate explanations for the behaviour we observe:

- Career concerns - external members' behaviour may be influenced by reputation effects.
- Composition effects from the other members who serve together?
- Psychological factors affecting the behaviour of group members - for example, there is a concept called “group polarization” in which, relative to making decisions as individuals, groups making decisions can become more conservative, or more extreme, through their interactions as a group.

In order to explain the behaviour, we explore a career concerns explanation. There are many possible stories for how external members behaviour might be affected by career concerns. For example perhaps the members want to establish credibility in order to get reappointed by the government. Alternatively, members may wish to signal particular preferences to the private sector in order to “line up” more opportunities for themselves at the end of their time on the committee. In the language of career concerns, the external members may, in addition to their desire to get the correct interest rate, wish to signal their expertise or their preferences.

We are working on a simple career concerns model to see (1) how the presence of career concerns affects the optimal voting rule and (2) whether a career concerns model can give any predictions on the time path of voting behavior. The main insight is that effect fades through time, and is proportional to weight you attach to getting reappointed relative to signalling to the market (β_i).

In order to test this, we would ideally need a measure of β_i to distinguish between which career concerns explanations drive the behaviour we observe; however, β_i is unobservable. Therefore, instead we need to use an exogenous variation in the extent of career

concerns to help us to determine whether there is any evidence of career concerns driving the behaviour.

Fortunately, we have one such natural experiment. The Act that created the MPC allows for the reappointment of all members, internal and external. When the first group of externals and internals served on the MPC, they thus 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 (Buiter 2000). To quote from this letter:

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 some were still among the most prominent monetary policy experts in the UK. A clear precedent was set: external members would find reappointment difficult, most likely extremely so. All external members served for only one term until February 2003 (almost 6 years since the first MPC meeting), when Brown unexpectedly reappointed Stephen Nickell to the MPC (HM Treasury 2003). Since then, Kate Barker has also been reappointed twice.

If career concerns existed, one would expect different voting patterns between external members serving from February 2000 to February 2003 and those serving at other times, since the rewards to reputation presumably changed when reappointment was and was not possible. There are 322 votes cast during this period in which reappointment of external members was not possible. Of course, there is no reason to expect that this lack of reappointment opportunities for external member to affect internal member voting; during this period many external members were reappointed to the committee. To this end, we define a dummy variable $D(reappoint_{it})$ which equals 1 before February 2000 and after February 2003. We then estimate the following career concerns regression:

$$\begin{aligned}
Probit [D(\text{votehigh})_{it}] &= \alpha + \varphi_1 \cdot D(\text{EXT})_i + \varphi_2 \cdot D(\text{NEW})_{it} + \varphi_3 \cdot D(\text{EXT})_i \cdot D(\text{NEW})_{it} \\
&+ \varphi_4 \cdot q_t + \varphi_5 \cdot D(\text{EXT})_i \cdot q_t + \varphi_6 \cdot D(\text{NEW})_i \cdot q_t + \varphi_7 \cdot D(\text{EXT})_i \cdot D(\text{NEW})_{it} \cdot q_t + \varphi_8 \cdot D(\text{reappoint})_{it} \\
&+ D(\text{reappoint})_{it} \cdot \left(\varphi_9 \cdot D(\text{EXT})_i + \varphi_{10} \cdot D(\text{NEW})_{it} + \varphi_{11} \cdot D(\text{EXT})_i \cdot D(\text{NEW})_{it} + \varphi_{12} \cdot q_t \right. \\
&\left. + \varphi_{13} \cdot D(\text{EXT})_i \cdot q_t + \varphi_{14} \cdot D(\text{NEW})_i \cdot q_t + \varphi_{15} \cdot D(\text{EXT})_i \cdot D(\text{NEW})_{it} \cdot q_t \right) + \varepsilon_{it} \quad (12)
\end{aligned}$$

$$\begin{aligned}
\frac{dProbit [D(\text{votehigh})_{it}]}{dq_t} &= \varphi_4 + \varphi_5 \cdot D(\text{EXT})_i + \varphi_6 \cdot D(\text{NEW})_i + \varphi_7 \cdot D(\text{EXT})_i \cdot D(\text{NEW})_{it} \\
&+ D(\text{reappoint})_{it} \cdot \left(\varphi_{12} + \varphi_{13} \cdot D(\text{EXT})_i + \varphi_{14} \cdot D(\text{NEW})_i + \varphi_{15} \cdot D(\text{EXT})_i \cdot D(\text{NEW})_{it} \right) \quad (13)
\end{aligned}$$

Given the complex interactions and the difficulty in interpreting individual coefficients, we do not present the estimation results for this equation¹⁷. Instead, as before, we report the more intuitive figures of the estimated probability of voting for higher rates; we do this in Figures 9 and 10¹⁸.

We would not expect the possibility of reappointment for external members to affect the behaviour of internal members. Figure 9 confirms there is no difference in the internal members behaviour between the two periods and is a first cross-check of the validity of our natural experiment.

Figure 10 shows the effect on external members of lower perceived reappointment possibility (a lower β_i). The results suggest that it is only when reappointment is possible that new external members come on to the committee and appear to display greater expertise ($\sigma_{INT} > \sigma_{EXT}$); this is consistent with the idea that signalling expertise is a means to get reappointed and to be considered a good appointment which benefits your reputation.

The second finding, despite the smaller sample, and so wider standard errors, during the “no reappointment period”, is that the identified preference shift takes place primarily when reappointment is not possible. This is consistent with signalling preferences if there is no payoff to signalling expertise.

¹⁷These estimate results are available from the authors upon request.

¹⁸And as before, we use the appendix (pages 31 and 32) to present figure 9 and 10 with 90% confidence intervals.

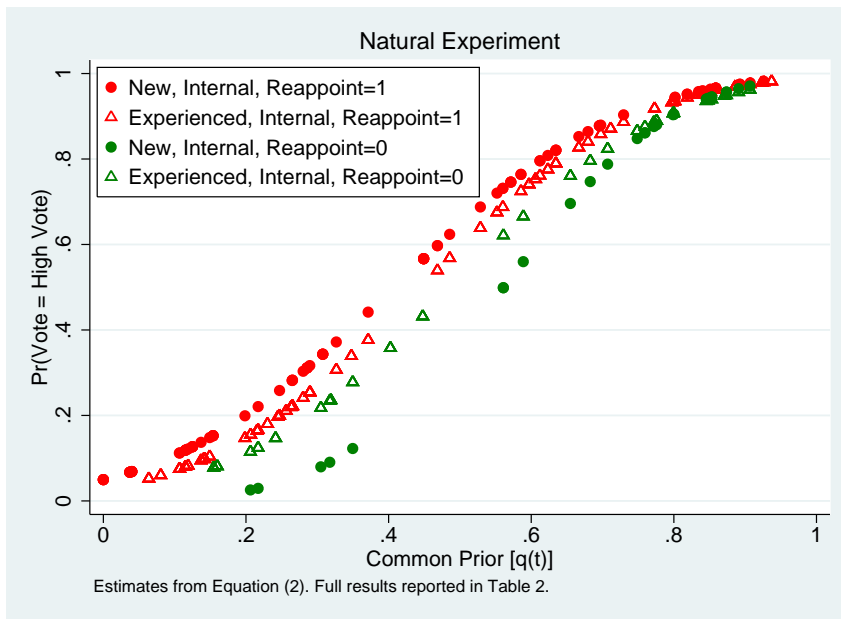


Figure 9: Natural Experiment of Reappointment of External:
Internal Estimated Probability of Voting High Conditional on the Common Prior

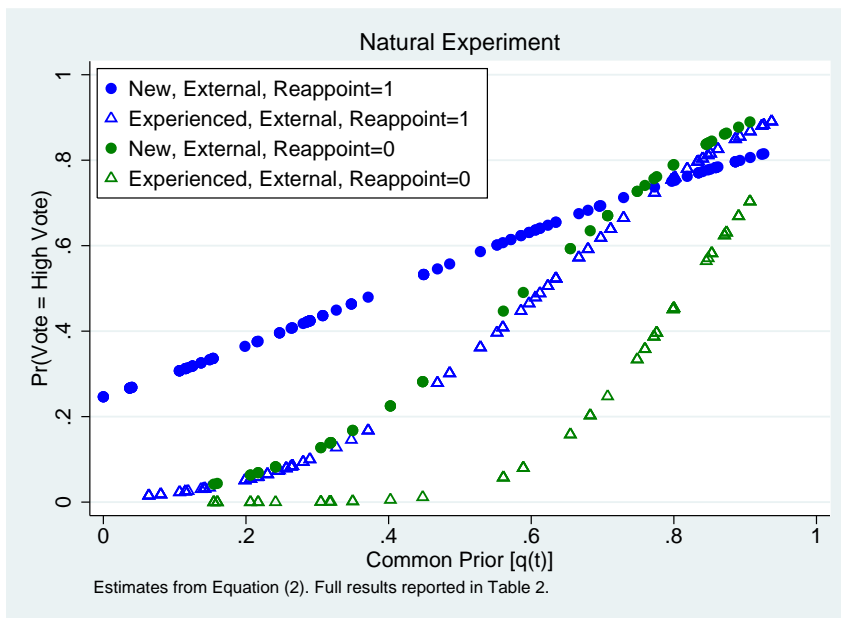


Figure 10: Natural Experiment of Reappointment of External:
External Estimated Probability of Voting High Conditional on the Common Prior

6 Conclusions and Policy Implications

TO BE ADDED.

References

- BESLEY, T., N. MEADS, AND P. SURICO (2008): “Insiders versus Outsiders in Monetary Policymaking,” *American Economic Review*, 98(2), 218–23.
- BESLEY, T., AND A. SCOTT (2010): “A new watchdog would guard us from debt,” <http://www.voxeu.org/index.php?q=node/4680>.
- 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. (2007): “Monetary policy by committee: Why and how?,” *European Journal of Political Economy*, 23(1), 106–123.
- BUIITER, 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.
- CAMERON, A. C., AND P. K. TRIVEDI (2009): *Microeconometrics Using Stata*. Stata Press, College Station, Texas, USA.
- GERAATS, P. M. (2006): “Transparency of Monetary Policy: Theory and Practice,” *CEifo Economic Studies*, 52(1), 111–152.
- GERLACH-KRISTEN, P. (2003): “Insiders and Outsiders at the Bank of England,” *Central Banking*, XIV(1), 96–102.
- (2009): “Outsiders at the Bank of England’s MPC,” *Journal of Money, Credit and Banking*, Forthcoming.
- 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.
- (2010): “The Mixed Case for Mixed Committees,” mimeo, UPF and University of Warwick.
- 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.
- HM TREASURY (2003): “Bank of England Appointments,” http://www.hm-treasury.gov.uk/newsroom_and_speeches/press/2003/press_28_03.cfm.

- LAMBERT, R. (2006): “Who sets our interest rates?,” http://www.bankofengland.co.uk/publications/other/monetary/inside_mpc.pdf.
- MEADE, E., AND D. STASAVAGE (2008): “Publicity of Debate and the Incentive to Dissent: Evidence from the US Federal Reserve,” *Economic Journal*, 118(528), 695–717.
- POLLARD, P. (2004): “Monetary Policy-Making around the World,” Presentation, Federal Reserve Bank of St. Louis.
- SIBERT, A. (2002): “Monetary policy with uncertain central bank preferences,” *European Economic Review*, 46(6), 1093–1109.
- (2003): “Monetary policy committees: individual and collective reputations,” *Review of Economic Studies*, 70(3), 649–665.
- (2006): “Central Banking by Committee,” *International Finance*, 9(2), 145–168.
- 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.

A Supplementary Figures

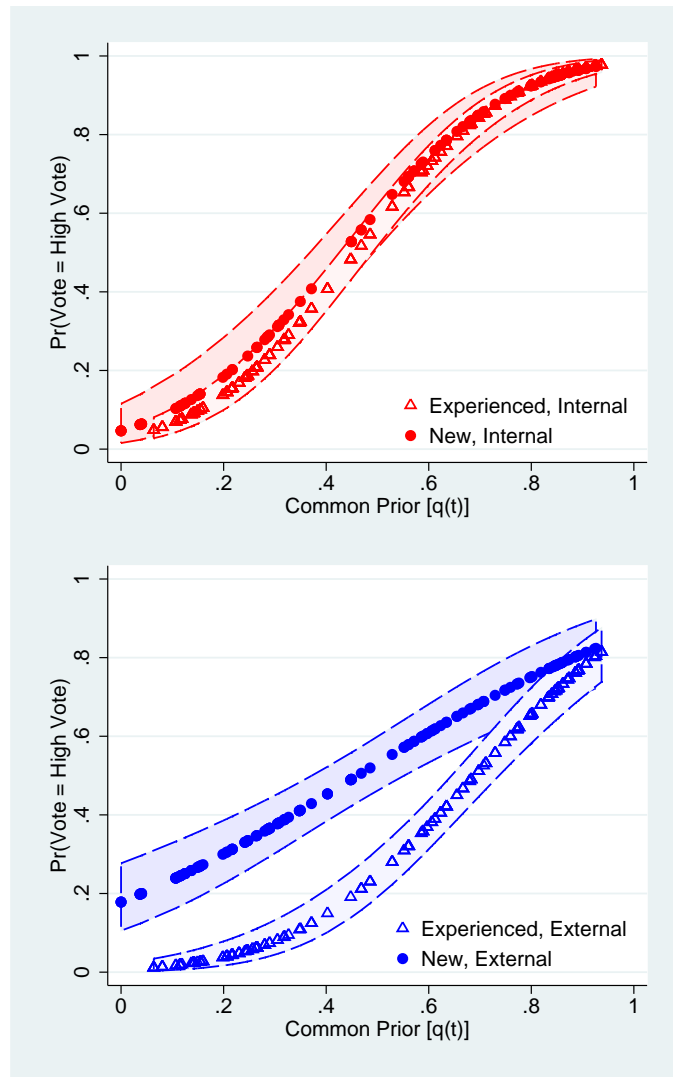


Figure 11: Figure 8 with 90% Confidence Intervals




Figure 12: Figure 9 with 90% Confidence Intervals: Internal Members

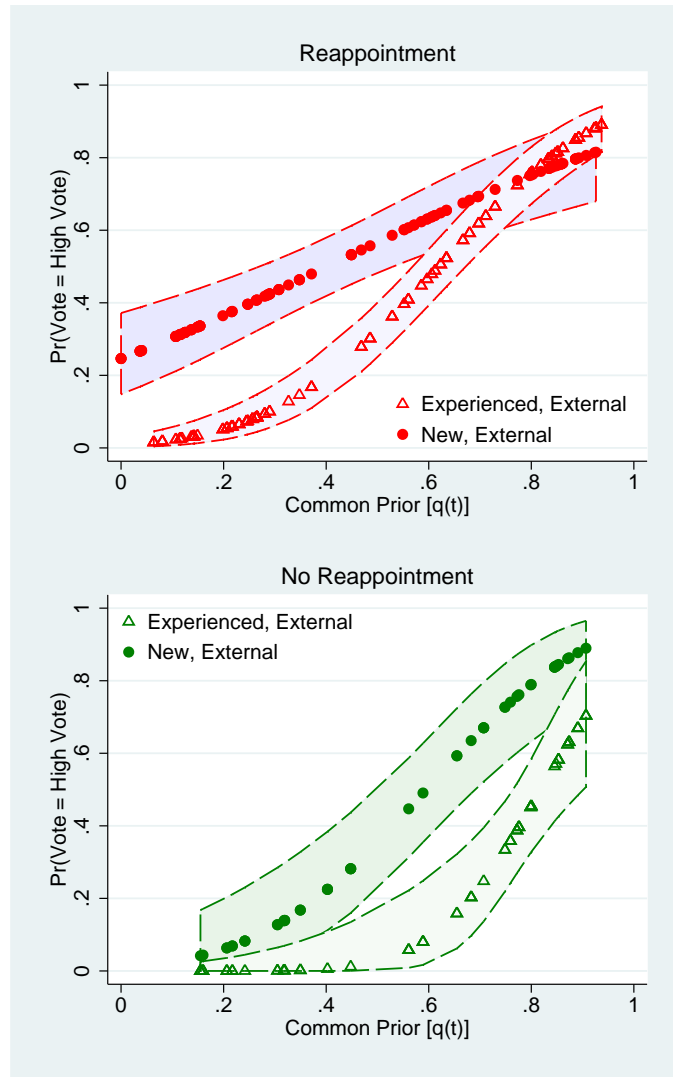


Figure 13: Figure 10 with 90% Confidence Intervals: External Members