DISCUSSION PAPER SERIES

No. 9607

FIRST IMPRESSIONS MATTER: SIGNALLING AS A SOURCE OF POLICY DYNAMICS

Stephen Hansen and Michael McMahon

INTERNATIONAL MACROECONOMICS

Centre for Economic Policy Research

www.cepr.org

Available online at: **www.cepr.org/pubs/dps/DP9607.php**

ISSN 0265-8003

FIRST IMPRESSIONS MATTER: SIGNALLING AS A SOURCE OF POLICY DYNAMICS

Stephen Hansen, Universitat Pompeu Fabra and Barcelona GSE
Michael McMahon, University of Warwick, CAGE, CEP, CfM, CAMA and CEPR

Discussion Paper No. 9607
August 2013

Centre for Economic Policy Research
77 Bastwick Street, London EC1V 3PZ, UK
Tel: (44 20) 7183 8801, Fax: (44 20) 7183 8820
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 9607

August 2013

ABSTRACT

First Impressions Matter: Signalling as a Source of Policy Dynamics*

We provide the first direct empirical support for the relevance of signalling in monetary policy. In our dynamic model, central bankers make policy under uncertain inflationary conditions and place different weights on output fluctuations. Signalling leads all bankers to be tougher on inflation initially, but to become less tough with experience. This evolution is more pronounced for members who weight output more ("doves"), which provides an additional test of our model. We structurally estimate the model using Bank of England data and confirm both predictions. Signalling increases the probability new members vote for high interest rates by up to 35%.

JEL Classification: D78, E52 and E58
Keywords: committees, monetary policy and signalling

Stephen Hansen
Universitat Pompeu Fabra
Department of Economics and Business
Carrer de Ramon Trias Fargas, 25-27
Barcelona, 08005
SPAIN

Email: stephen.hansen@upf.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=171533

Michael McMahon
Department of Economics
University of Warwick
Coventry
CV4 7AL
UK

Email: m.mcmahon@warwick.ac.uk

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=153555

*We would like to thank Francesco Caselli, Antonio Ciccone, Jan Eeckhout, Andrew Oswald, Anne Sibert, Francesco Squintani and Carl Walsh for helpful comments. We also benefited from comments and suggestions by seminar attendees at Birkbeck, Exeter University, Singapore Management University, INSEAD, CESifo Applied Microeconomics Conference, Stanford University, University of California at Santa Cruz, the Society for Economic Dynamics annual conference 2012, and the Money-Macro Finance Group annual conference 2012.

Submitted 13 August 2013

1 Introduction

Ever since the seminal work of Barro and Gordon (1983a, 1983b), economists have been aware that, in the absence of commitment devices, central bankers find it difficult to achieve low inflation due to time consistency problems. As a result, many aspects of modern monetary policy are about managing inflation expectations (King, Lu, and Pastén 2008). Policy makers' trying to build credibility as inflation fighters, the establishment of independent central banks, and the Federal Reserve's recent adoption of forward guidance all reflect the importance of managing inflation expectations. The change of senior central bank personnel, such as the Chairperson or Governor, is often a period of particular importance for the control of inflation expectations. At such times, there is uncertainty and speculation among journalists and market economists about how the new person will behave relative to the out-going policy maker and how the change might affect inflation expectations. For example, Cottle (2012) contemplates whether incoming Bank of England Governor Mark Carney will be a "hawk" or a "dove". Speculation has already begun about who, and how hawkish, the next Fed Chairperson will be even before Chairman Bernanke has announced he is stepping down.[1]

Given the importance of inflation expectations for monetary policy, new policy makers may try to signal that they will be tough on inflation. The idea of signalling as a means of establishing a reputation in monetary policy has a long history.[2] The main prediction in this literature is that policy makers signal to the public how tough they are on inflation early in their careers, but become less tough on inflation over time. Having a reputation for being tough on inflation is arguably even more important nowadays given the expansion of central bank balance sheets through unconventional monetary policy; to expand money supply and liquidity without de-anchoring inflation expectations requires a great deal of credibility. Flanders (2011), in advance of the anticipated appointment of Mario Draghi as ECB President, speculated that because Draghi was Italian, he might have to go out of his way to rebut national stereotypes by being especially tough on inflation, with less expansionary unconventional policies, immediately following his appointment.[3] Whether such signalling actually takes place in monetary policy making and whether it affects monetary policy choices are important open questions that we address in this paper.

[1] The Economist (2013) and Appelbaum (2013) are recent articles. In reference to Janet Yellen, Appelbaum (2013) note that "they state that observers have gathered from her speeches that she puts a higher weight on output than inflation".

[2] For example, see Backus and Driffill (1985a, 1985b), Barro (1986), Cukierman and Meltzer (1986), Vickers (1986), Faust and Svensson (2001), Sibert (2002, 2003, 2009) and King, Lu, and Pastén (2008). This literature built on Kreps and Wilson (1982) and Milgrom and Roberts (1982).

[3] She writes "If you're sitting in Spain and Portugal, you might well wonder whether you would have been better off with a German in charge, trying to show off his inner Italian - than an Italian desperate to prove he's German underneath."

This paper makes three contributions. The first is to incorporate signalling into a model of monetary policy decision making that captures that policies are typically determined by a committee and that committee members may disagree on the (uncertain) state of the economy. The second contribution is to structurally estimate the model using voting data from the Bank of England's Monetary Policy Committee (MPC) and to examine whether the model's key predictions on voting dynamics are supported empirically. The third is to demonstrate when the signalling incentive alters interest rate choices. We now provide brief details of each of these contributions.

The main departure of our model of monetary policy decision making from the existing theoretical literature is that decisions are made by a committee and that committee members may differ in their preferences as well as in their private assessments of the state of the economy. Differences in monetary policy preferences are captured in the usual way. While all central bankers dislike deviations of inflation from a target rate, some place less weight on the output gap than others. A type that puts more (less) weight on output is more "dovish" ("hawkish").

To this standard setup we add that policy makers do not know the exact state of the economy, which can either be inflationary, requiring higher interest rates, or not, requiring lower interest rates. Before each interest rate decision, policy makers receive common information about the economy, such as economic data releases and staff forecasts, and also have their own idiosyncratic interpretations that they use to form their views on the state of the economy. We model these private assessments as each member's observing a private signal correlated with the state of the economy. A central banker's expertise is the precision of this signal.[4] We also depart from the existing theoretical literature on monetary policy decision making in that committees can have more than two members of more than two types, and that all types are strategic in the sense of internalizing the effect of their behavior on inflation expectations. The previous literature assumes that there are two committee members each of whom have two potential types, a mechanistic hawk that votes for zero inflation and a strategic dove that tries to build a hawkish reputation (Sibert 2003).

In our model, signalling generates policy dynamics whether there is a single central banker or whether monetary policy is decided by a committee. In equilibrium, more hawkish preference types vote for high rates more often in all periods, so if the public sees a policy maker vote for the high rate, it will put more weight on his being hawk-

[4] Modelling expert decision problems in this way is common in many areas of economics; for example, Ottaviani and Sorensen (2000) examine career concerns of experts; Gerling, Gruner, Kiel, and Schulte (2005) look at committee decision making; and Besley (2006) focuses on politicians' behavior. However it has not been incorporated into models of signalling in monetary policy. Static models that use it include Hansen, McMahon, and Velasco (2012), who show that private assessments are an important driver of decision making on the MPC, and Gerlach-Kristen (2006), whose model allows committee members to each hold a different view of potential output.

ish. *All* preference types, not just dovish ones as in the previous literature, have an additional incentive to vote for high rates when new in order to reduce future inflation expectations. This means that members' tendency to vote for low rates—their dovish bias, which depends both on their preference type and reputational incentives—increases over time. But we also show that the incentive to signal when new is greater for more dovish preference types. This delivers the prediction that the increase in the bias is higher for more dovish preference types.

A key feature of our theoretical model is that it fits into an empirical framework that allows for structural estimation using MPC voting data. Our estimator separately identifies policy makers' dovish biases from their expertise. This allows us to verify the two main dynamic predictions of the model. First, the average experienced member on the MPC has a significantly higher dovish bias than the average new member. Second, using four different measures of hawkish and dovish preference types, we show that this pattern of increasing bias is more pronounced for dovish preference types. To the best of our knowledge, these results represent the first empirical validation of models of signalling in monetary policy.

Finally, we examine when new members' signalling actually translates into their choosing high rates more often. We show that the effect of signalling on a member's decisions decreases as his expertise increases. This is because all members want to get the decision right (vote for high rates when the economy is inflationary) and their bias only affects the decision because of uncertainty. As the most expert members see the state of the economy nearly perfectly, they simply select the appropriate interest rate in all periods. For similar reasons, signalling has less impact when the common information about the state (the prior) gives a clearer picture of economic conditions. As previous papers did not allow for uncertainty or differences in expertise among central bankers, these potential mitigations of the signalling channel has been overlooked until now. Overall, we find an average new MPC member can be between 0% and up to 35% more likely to choose high rates in a given meeting depending on the prior and their expertise.

A natural alternative hypothesis for generating the dynamics we observe is learning about some persistent parameter of the macroeconomy. While learning could in theory generate the estimated policy dynamics, we argue that it would need to be somewhat unusual to capture the main features of our data. In particular, the prediction on the type-dependent evolution in the dovish bias arises from a signalling model but not from standard learning models, strengthening the message that signalling is potentially important. Also, the estimated biases diverge over time rather than converge as would be expected under standard learning models.

Perhaps the most fundamental contribution of the paper is to show that reputation effects on independent monetary policy committees should be treated as of first-order importance. A large literature (summarized in Drazen (2001)) highlights the difficulties that politicians have in establishing credible monetary policy, and the establishment of independent committees was a direct response to this insight. While we agree that this no doubt eased inflationary traps, our paper shows committees cannot be viewed as simply replicating the policies of the metaphorical "hard-nosed" banker often invoked to discuss the behavior of central banks. Instead, our estimates imply that a model in which preferences are heterogenous and establishing credibility is crucial fits voting data very well in an institutional context admired for its independence. Once this fact is acknowledged, a host of interesting questions (such as optimal term lengths and committee composition) for designing monetary policy committees to exploit reputation become immediately relevant, and we hope our paper can stimulate thinking about them.

Beyond the particular application of our model to the MPC, we view our framework as a natural one for quantitatively assessing the impact of signalling on voting dynamics. It could, for instance, be directly applied to voting data on other committees, and also to the choices of a single policy maker.[5] Outside monetary policy, the mechanism we identify should be relevant in any context in which the policy maker's desired outcomes depend on the public's expectations about her actions. For example, several countries are currently establishing new macro prudential and regulatory bodies following the financial crisis. As the willingness of banks to engage in risky practices presumably depends on their beliefs that authorities will punish such behavior, regulators can signal their intention to crack down on these practices by taking tough stances at the beginning of their careers. This would discourage banks from future bad behavior, meaning regulators can achieve their policy objectives without further actions later in their tenures.

Our paper is related to a number of existing literatures, most obviously that on signalling in monetary policy. In one branch of this literature (Backus and Driffill 1985a,b, Vickers 1986, Sibert 2002), policy makers choose a level of inflation and policy choices can perfectly reveal bankers' types (there are at least as many levels of inflation as there are preference types). These "pure" signalling models generally feature multiple equilibria, for example a separating and pooling equilibrium. Another branch (Cukierman and Meltzer 1986, Faust and Svensson 2001, Sibert 2009) employs a modelling device whereby the public observes a noisy measure of the policy choice, which makes it impossible to perfectly infer the banker's type and reduces equilibrium multiplicity. Here we introduce an alternative way of limiting the public's learning, which in our setting actually generates unique equilibrium predictions. First, while there is a rich type space, there are only two available interest rates each period. Second, all policy makers want to match the interest rate to the state, and so will vote high when their private signals surpass a threshold.

[5] The latter analysis would require the use of an alternative estimator; the methodology for this alternative is presented in Hansen and McMahon (2013) but has not yet been used empirically.

But different preference types have different thresholds, meaning that observed policy is informative about preferences without ever fully revealing them.[6]

There is also a literature in which agents in the economy learn about policy behavior, but the central bank is not strategic in choosing policy to take advantage of this learning. Bianchi and Melosi (2012) consider the dynamics when the central bank switches exogenously between more and less active monetary policy and the public have to assess the monetary regime that is being used. Erceg and Levin (2003) also explore learning by households about the monetary policy regime. The central bankers in our paper differ in that they react strategically to affect inflation expectations and the dovishness of their policy stance is determined endogenously.

Besides the signalling literature, there is a large literature on the policy dynamics generated by policy makers' learning about some aspect of the macroeconomy as more information becomes available; Evans and Honkapohja (2001) is the seminal reference. This literature has examined the effects of policy makers' learning about the behavior of inflation (Sargent 1999, Cho, Williams, and Sargent 2002, Primiceri 2006), the natural rate of unemployment (Orphanides and Williams 2005), and the level (and growth) of potential output (Bullard and Eusepi 2005). We are not aware of any paper in the macro learning literature that yields all of the dynamic patterns that we identify in this paper on the evolution of the dovish bias. Furthermore, we are skeptical that straightforward extensions of the current generation of learning models could do so.

Finally, the two-step methodology we use for estimating bias and expertise was introduced by Iaryczower and Shum (2012), who apply it to the voting record of the US Supreme Court. Hansen, McMahon, and Velasco (2012) use a variation of this approach[7] to study heterogeneity in bias and heterogeneity among MPC members, but do not consider dynamic behavior. The paper proceeds as follows. Section 2 lays out a theory of signalling for single policy makers and committees. Section 3 explains the institutional setting of the MPC and presents reduced form evidence on dynamics, before we turn to a structural analysis in section 4. Section 5 then uses the estimated parameters to quantify the impact of signalling on policy choices. Section 6 asks whether learning could generate our results, and section 7 concludes.

[6]Papers in the career concerns literature such as Levy (2007) and Visser and Swank (2007) study voting on committees when members care about their reputation in addition to policy. But, whereas career concerns models assume policy makers place some exogenous weight on their reputation, in this paper (and in those cited above) reputation concerns emerge endogenously in equilibrium due to the structure of the macroeconomy.

[7]Hansen and McMahon (2013) show that using an alternative first stage specification that interacts meeting and individual characteristics allows one to use variation in the prior distribution from which the state is drawn to identify expertise differences.

2 Reputation and Policy Dynamics

This section lays out a model of how monetary policy makers' desire to control inflation expectations gives them an incentive to build a hawkish reputation early in their careers. It begins by considering the case of a single policy maker choosing interest rates before turning to the case of a committee of voters. Proofs of results are in appendix A.

2.1 Model: Single Policymaker

Suppose there is a single central banker C with discount factor $\delta \in (0,1]$ who chooses interest rate $r_t \in \{0,1\}$[8] in periods $t = 1, 2$.[9] In each period, a state variable $\omega_t \in \{0,1\}$ drawn from a Bernoulli distribution with $\Pr[\,\omega_t = 1\,] = q_t$ is realized; let $\overline{q} = \mathbb{E}[\,q_2 \mid q_1\,]$. We model inflation as decreasing in the interest rate and increasing in the inflationary state of the economy; precisely, $\pi_t = \pi(r_t, \omega_t)$ where $\pi(1,\omega_t) < \pi(0,\omega_t)\ \forall\omega_t$ and $\pi(r_t,1) > \pi(r_t,0)\ \forall r_t$. To ensure the model is well behaved, we bound the degree of asymmetry in the responsiveness of inflation to monetary policy across states.

A1 $\frac{\pi(0,1)-\pi(1,1)}{\pi(0,0)-\pi(1,0)} \in \left(\frac{1-\overline{q}}{\frac{1}{\delta}-\overline{q}}, \frac{\frac{1}{\delta}-(1-\overline{q})}{\overline{q}}\right)$.

This bound becomes weaker as δ decreases, and as $\delta \to 0$ arbitrary asymmetry is possible.

Before presenting our formulation of preferences, we motivate it with an example that uses a variation of standard preferences in the monetary literature over inflation and output.

Example 1 *Suppose C has period t utility $u_t = u(\pi_t, y_t) = -(\pi_t - \pi^*)^2 + \phi(y_t - y^*)$ where π^* is an inflation target and ϕ is the weight he puts on the output gap $y_t - y^*$. By plugging in the expectations augmented Phillips curve $y_t - y^* = \beta(\pi_t - \pi_t^E)$ (where π_t^E represents the public's inflation expectations), one arrives at a utility representation over inflation and inflation expectations $u\left(\pi_t, \pi_t^E\right) = -(\pi_t - \pi^*)^2 + \chi\left(\pi_t - \pi_t^E\right)$ where $\chi = \phi\beta$.*

Our more general formulation of preferences is

$$u_t = u\left[r_t, \omega_t, \theta\right] = M\left[\pi(r_t, \omega_t)\right] + \chi(\theta)\left[\pi(r_t, \omega_t) - \pi_t^E\right]. \quad (1)$$

[8]The interest rates that constitute the binary agenda can change from meeting to meeting and the higher of the two under consideration constitutes the $r_t = 1$ policy choice. For example, "no change" is the low interest rate option ($r_t = 0$) in a meeting deciding between "no change" and "+25", but the higher option ($r_t = 1$) if the policy under consideration is "−25".

[9]The choice of a two period model is simply to make the analysis cleaner and to link more easily to our empirical analysis, in which we measure dynamics conditional on a dummy variable that measures whether members are new or experienced. It would, with added complexity but little gain in terms of insight, be possible to extend the analysis to allow T periods.

where $\theta \in \Theta$ is the central banker's type. Θ is a finite set with cardinality K, and we denote $\underline{\theta} = \min \Theta$ and $\overline{\theta} = \max \Theta$. M captures losses from deviations of inflation from its target level. Here the weight put on the output gap depends on the type—we assume $\chi(\theta) > 0$ $\forall \theta$ and that χ is strictly increasing. We will refer to a banker with a lower θ as more "hawkish" and one with a higher θ as more "dovish." These preferences capture a setting in which C is given an explicit inflation target along with a more subjective mandate to consider how policy impacts growth that more growth-oriented types can use to justify softer policies. In assuming linearity of preferences in the output gap, we follow much of the previous literature (Backus and Driffill 1985a,b, Cukierman and Meltzer 1986, Vickers 1986, Sibert 2002, 2003, 2009). In section 2.3 we discuss alternative specifications. We assume that θ is private information for C over which the public has the prior $p_0 \in \Delta(\Theta)$, where p_0^θ denotes the probability it assigns to type θ.

We believe a natural assumption is that the central banker wants to select the high (low) interest rate when the economy is in the more (less) inflationary state. That is, we think that the central banker strictly prefers $r_t = \omega_t$ over $r_t \neq \omega_t$ in all states. Let the utility gain of matching the decision to the state in terms of meeting the inflation target be $\mu(\omega_t) \equiv M\left[\pi(r_t = \omega_t, \omega_t)\right] - M\left[\pi(r_t \neq \omega_t, \omega_t)\right]$. In order to ensure that the model features just interim disagreement (i.e. all types want to match the policy to the state and would agree on r_t if ω_t were revealed to them), we make the following assumptions on preferences:

A2 $\mu(\omega_t) > 0$.

A3 $\mu(1) - \chi(\overline{\theta})\left[\pi(0,1) - \pi(1,1)\right] > 0$.

In words, A2 means that in the low (high) state the low (high) rate is more consistent with the target; and A3 that even the most dovish type gets higher period t utility from choosing $r_t = 1$ if he knows that $\omega_t = 1$. We now continue with our example to show how it is compatible with these assumptions.

Example 1 (continued) *Returning to our example, the net utility of matching the decision to the state if* $\omega_t = 1$ *is*

$$u\left[\pi(1,1), \pi_t^E\right] - u\left[\pi(0,1), \pi_t^E\right] = \left[\pi(0,1) - \pi(1,1)\right]\left[\pi(1,1) + \pi(0,1) - 2\pi^* - \chi\right]$$

and, if $\omega_t = 0$,

$$u\left[\pi(0,0), \pi_t^E\right] - u\left[\pi(1,0), \pi_t^E\right] = \left[\pi(0,0) - \pi(1,0)\right]\left[2\pi^* - \pi(0,0) - \pi(1,0) + \chi\right].$$

C strictly prefers $r_t = \omega_t$ in all states whenever

$$\frac{\pi(0,0)+\pi(1,0)+\chi}{2} < \pi^* < \frac{\pi(1,1)+\pi(0,1)-\chi}{2}.$$

As long as the weight that C puts on output is not too large, then there exists an inflation target that ensures that C will seek to match the interest rate to the state.

Before choosing r_t, the banker privately observes a continuous signal $s_t \in (\underline{s}, \overline{s})$ with distribution $G_t(s_t \mid \omega_t)$ and density $g_t(s_t \mid \omega_t)$. We denote the log-likelihood ratio as $\mathfrak{L}_t(s_t) = \ln\left[\frac{g_t(s_t|\omega_t=1)}{g_t(s_t|\omega_t=0)}\right]$ and assume it is continuous and strictly increasing, with $\lim_{s_t \to \underline{s}^+} \mathfrak{L}_t(s_t) = 0$ and $\lim_{s_t \to \overline{s}^-} \mathfrak{L}_t(s_t) = \infty$. Finally, let $\widehat{\omega}_t \equiv \Pr[\omega_t = 1 \mid s_t]$ be the banker's posterior on the state after observing signal s_t.

The timing of each period t sub-game is the following.

1. q_t is observed and ω_t is drawn by nature according to $\Pr[\omega_t = 1] = q_t$.
2. The public forms π_t^E.
3. s_t is realized and observed privately by C.
4. C chooses r_t, which the public observes.
5. π_t is publicly observed (which implicitly reveals ω_t).

The model admits a prior with structure $q_t(\omega_{t-1})$, which can capture state persistence.

A strategy for C in period t is a mapping $\rho_t^* : (\theta, s_t) \to [0,1]$ into the probability of choosing $r_t = 1$. Let $\widehat{\rho}_t$ be the public's belief on this strategy and $p_1(r_1) \in \Delta(\Theta)$ its posterior on θ after observing r_1. Strategies constitute a *Perfect Bayesian Equilibrium* when the following conditions hold:

1. p_1 is computed using Bayes' Rule given $\widehat{\rho}_1$.
2. The public forms π_t^E given $\widehat{\rho}_t$ and p_{t-1}.
3. $\rho_2^*(\theta, s_2)$ solves $\max \mathbb{E}[u_2(\theta) \mid s_2]$ given π_2^E.
4. $\rho_1^*(\theta, s_1)$ solves $\max \mathbb{E}[u_1(\theta) + \delta u_2(\theta) \mid s_1]$ given $\rho_2^*(\theta, s_2)$, π_1^E and π_2^E.
5. The public's beliefs are consistent with the banker's strategy: $\rho_t^* = \widehat{\rho}_t \; \forall t$.

Of particular interest are *threshold voting rules* in which C employs the strategy

$$\rho_t^*(\theta, s_t) = \begin{cases} 1 & \text{if } \frac{\widehat{\omega}_t}{1-\widehat{\omega}_t} > B_t \\ [0,1] & \text{if } \frac{\widehat{\omega}_t}{1-\widehat{\omega}_t} = B_t \\ 0 & \text{if } \frac{\widehat{\omega}_t}{1-\widehat{\omega}_t} < B_t \end{cases} \tag{2}$$

where B_t is the central banker's period t *dovish bias* (which we sometimes call simply bias), or how much evidence he needs that the inflationary shock has hit to justify the high rate; in simpler terms, it captures the inclination to choose the lower interest rate in a given meeting. Below we endogenize the bias as a function of θ and reputation. With threshold voting rules, different types can have different probabilities of choosing both rates but observing r_t never allows the public to perfectly infer θ.

Bayes' Rule gives

$$\ln\left[\frac{\widehat{\omega}_t}{1-\widehat{\omega}_t}\right] = \ln\left(\frac{q_t}{1-q_t}\right) + \mathfrak{L}(s_t) \tag{3}$$

so that threshold voting rules call on C to choose $r_t = 1$ whenever

$$s_t \geq \mathfrak{L}_t^{-1}\left[\ln\left(\frac{1-q_t}{q_t}B_t\right)\right]. \tag{4}$$

Since the signal satisfies MLRP, (4) implies that C chooses $r_t = 1$ if and only if his signal reaches a critical threshold. Here one can see the empirical consequences of a change in the bias: when B_t increases, the probability of voting high decreases since the threshold the signal must reach increases.

2.1.1 Equilibrium with observable type

In order to fix ideas, we begin by discussing the equilibrium when the public observes θ. This is analogous to the original Barro and Gordon model in which the public knows the policy maker's preferences.

Proposition 1 *Suppose the public observes θ. There is a unique equilibrium in which C adopts a threshold voting rule in both periods with dovish bias*

$$B(\theta) \;=\; \frac{\mu(0)+\chi(\theta)\left[\pi(0,0)-\pi(1,0)\right]}{\mu(1)-\chi(\theta)\left[\pi(0,1)-\pi(1,1)\right]}.$$

This rule emerges from the banker's utility maximization problem treating inflation expectations as fixed. The numerator is the cost of a wrong decision in state 0 (the utility that is lost from choosing $r_t = 1$ if the realized state is 0), while the denominator is the benefit of a correct decision in state 1. A banker who puts more weight on output has a higher cost in state 0 since $r_t = 1$ implies lower output, and similarly a lower benefit in state 1. This means that he adopts a more dovish bias, and votes high less often.

Of course, in equilibrium inflation expectations are not fixed but consistent with the banker's strategy. While the bias in proposition 1 reflects the banker's behaving as if the choice of r_t affected output, in equilibrium it does not. One can easily show that the voting rule to which C would like to commit is a threshold rule with dovish bias $\frac{\mu(0)}{\mu(1)}$; this rule only considers the impact of r_t on realized inflation. In other words, the bias in proposition 1 is too soft on inflation, and every preference type except one that puts no weight on output over-inflates the economy in equilibrium.[10]

2.1.2 Equilibrium with unobservable type

We now return to the more realistic and interesting case in which θ is private information for C. We begin by stating the properties of all equilibria.

Proposition 2 *All equilibrium strategies are threshold voting rules that satisfy:*

1. $B_t(\theta)$ *is strictly increasing in* θ *for* $t = 1, 2$.
2. $B_2(\theta) - B_1(\theta) > 0 \; \forall \theta$.
3. $B_2(\theta) - B_1(\theta)$ *is strictly increasing* $\forall \theta$.

While equilibria may not be unique, all share the same qualitative features. First, types that put more weight on output have a higher dovish bias in both periods. Second, all types have a higher bias in the second period than the first. Finally, the degree to which the bias increases is increasing in dovishness. Figure 1 illustrates the predicted dynamics in the dovish bias for three different preference types ordered by dovishness (θ_H is most dovish, θ_L is least dovish). The three features together imply that, although the underlying heterogeneity in preference types stays constant over time, the heterogeneity in dovish biases increases.

The intuition for these dynamics is the following. When C chooses r_2, π_2^E is already set and the game ends afterwards. So he solves the same problem as in proposition 1 and adopts the bias

$$B_2(\theta) = \frac{\mu(0) + \chi(\theta)\left[\pi(0,0) - \pi(1,0)\right]}{\mu(1) - \chi(\theta)\left[\pi(0,1) - \pi(1,1)\right]}. \tag{5}$$

Turning to period 1, C must now consider the impact of r_1 on π_2^E in addition to its impact on contemporaneous inflation. Let $\overline{\pi}_2^E(r_1, \omega_1)$ denote the expected value of period 2 inflation expectations given (r_1, ω_1).[11] C votes high if and only if

$$\frac{\widehat{\omega}_1}{1 - \widehat{\omega}_1} \geq B_1(\theta) = \frac{\mu(0) + \chi(\theta)\left\{\left[\pi(0,0) - \pi(1,0)\right] - \delta\left[\overline{\pi}_2^E(0,0) - \overline{\pi}_2^E(1,0)\right]\right\}}{\mu(1) - \chi(\theta)\left\{\left[\pi(0,1) - \pi(1,1)\right] - \delta\left[\overline{\pi}_2^E(0,1) - \overline{\pi}_2^E(1,1)\right]\right\}}. \tag{6}$$

The key for our results is signing the terms $\overline{\pi}_2^E(0, \omega_t) - \overline{\pi}_2^E(1, \omega_t)$ and showing they are both positive. The logic follows three steps.[12] First, since $B_2(\theta)$ is strictly increasing,

[10] Here a vocabulary clarification may be helpful. Barro and Gordon refer to the equilibrium strategy as featuring an "inflation bias." The dovish bias in our model is different and captures the burden of proof applied by the central banker to his views about the state of the economy; it is a part of the decision making process rather than an equilibrium outcome for inflation.

[11] The uncertainty given (r_1, ω_1) is due to the unknown q_2.

[12] This discussion assumes that $B_1(\theta)$ is increasing in θ, which one can show must be true in any equilibrium. But, whenever $\pi(0,0) - \pi(1,0) \neq \pi(0,1) - \pi(1,1)$, it is possible that in some state ω_1 the

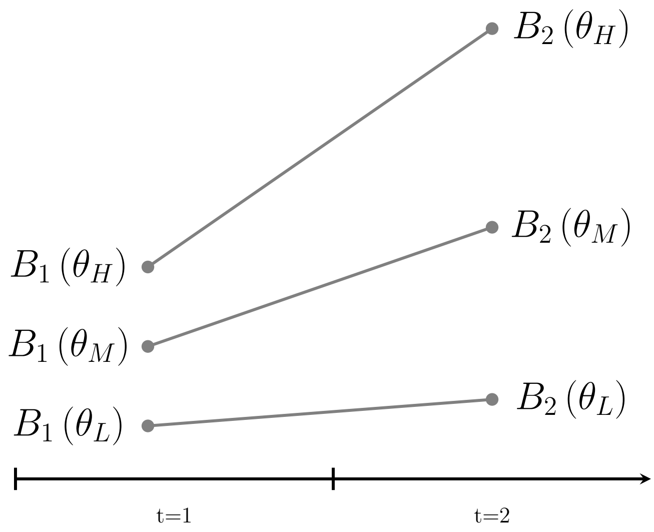

Figure 1: Predicted Dynamics in Inflationary Bias for Types $\theta_L < \theta_M < \theta_H$

Notes: This figure shows how the dovish bias evolves between periods 1 and 2 for three representative preference types θ_L, θ_M, and θ_H. All types' biases increase over time, but the bias adopted by θ_H (the most dovish central banker) increases most, followed by that of θ_M and θ_L.

the public increases π_2^E when it believes more dovish types set policy. Second, given that $B_1(\theta)$ is increasing in θ (for any value of $\overline{\pi}_2^E(0, \omega_1) - \overline{\pi}_2^E(1, \omega_1)$, in or out of equilibrium), in every equilibrium it must be the case that more dovish types are more likely to choose $r_1 = 0$. Finally, combining these two observations means that when the public observes $r_1 = 0$, it associates the banker with a more dovish type, which leads it to increase π_2^E. In short, *all* preference types have an additional incentive in the first period to vote for high rates that is absent in the second: doing so allows them to build a hawkish reputation that anchors second period inflation expectations.

The equilibrium also features a more subtle result. The signalling incentive is directly linked to the weight the banker places on future output. More dovish types by definition put more weight on output, and so, intuitively, care more about convincing the market they are hawks than hawks do. To see this point another way, consider a type that puts *no* weight on output. This extreme hawk has no incentive to signal to the market at all, because he is indifferent to the resulting change in inflation expectations. Hence the second dynamic prediction on the increase in bias being greater the more dovish the type.

change in inflation expectations is bigger than the change in period 1 output. This means that one of the terms in curly brackets in (6) might be negative. The role of assumption A1 is to rule these situations out, making the proof of equilibrium existence straightforward.

2.2 Model: Committee of Policymakers

We now consider a setting in which monetary policy is decided by a committee of N members in each period $t \in \mathbb{Z}^+$. The main economic insights are already present in the single banker case, but it is nevertheless important to derive the voting rules in committees since it is in a committee environment that we later estimate our model. We analyze an overlapping generations model in which members serve for 2 periods. A member in his first period is *young* and in his second *old*. In odd (even) periods there are N_1 (N_2) young members.[13] We denote by $Y(t)$ the set of young voters in period t, and by convention $Y(0)$ denotes the set of old members in $t = 1$. Call $\tau(i)$ the period in which member i is appointed. The committee operates under majority rule: in periods $t = \tau(i)$ and $t = \tau(i) + 1$, member i chooses a vote $v_{it} \in \{0, 1\}$ and $r_t = 1$ whenever $\sum_{i \in Y(t) \cup Y(t-1)} v_i > \frac{N-1}{2}$. We denote by $\mathbf{v}_t^Y$ the set of votes of young members in period t.

Member i has preference type θ_i over which other members and the public have a prior $p_0 \in \Delta(\Theta)$. Types are independent across members. The timing and information assumptions for a member's career are the following:

1. Member i joins committee and observes θ_i
2. $\mathbf{v}_{\tau(i)}$ is chosen by committee members and observed by the public
3. θ_i is observed by all $j \in (Y(\tau(i)) \cup Y(\tau(i) + 1)) \setminus \{i\}$
4. $\mathbf{v}_{\tau(i)+1}$ is chosen by committee members and observed by the public
5. θ_i is observed by the public

The delayed observability of members' types isolates policy makers' incentive to signal to the public from other signalling channels. If member i's colleagues did not observe his type before period $\tau(i) + 1$, then $v_{i\tau(i)}$ would serve to signal θ_i to colleagues as well as the public. If the public did not observe θ_i before $\tau(i) + 2$ and believed that young members' strategies in period $\tau(i)+1$ (when member i is old) depended on θ_i, then $v_{i,\tau(i)+1}$ would affect its update of their types. As well as isolating the signalling channel, these assumptions also capture the empirical reality that policy makers engage in numerous activities besides voting (e.g. forecasting, internal policy debates, public speeches) that serve to communicate their preferences. Since "periods" in our model capture voting over about an 18 month time horizon in our empirical context, a natural assumption is that such information eventually reveals the true type. But, a member's colleagues on the committee learn his type faster than the public because they observe more of this information.

[13] This implies that in period 1 there are $N - N_1$ old members who do not serve in period 2.

Let $\mathbf{p}_t$ denote the vector of beliefs the public has on period-t committee members' types; this will be p_0 for all young members, and depend on $\mathbf{v}_{t-1}^Y$ for old members. Let $\boldsymbol{\theta}_t^O$ denote the vector of old members' types. Prior to choosing v_{it}, member i privately observes signal s_{it} with distribution $G_Y(\cdot \mid \omega_t)$ if i is young and $G_O(\cdot \mid \omega_t)$ is he is old. To solve the game, we focus on symmetric Markov strategies whose state variable is the set of beliefs on types. Let $x(t) = t \mod 2$ denote odd and even periods. In every period t, all members $i \in Y(t)$ adopt the strategy $\rho_Y^* : \theta_i, s_{it}, \boldsymbol{\theta}_t^O, x(t) \to [0,1]$, and all old members $i \in Y(t-1)$ adopt the strategy $\rho_O^* : \theta_i, s_{it}, \boldsymbol{\theta}_{-it}^O, x(t) \to [0,1]$. Let $\widehat{\rho}_Y$ and $\widehat{\rho}_O$ be the public's beliefs on these strategies. We analyze *Markov Perfect Bayesian Equilibria*, whose definition is omitted but standard. It is important to keep in mind that the public forms π_t^E given $\mathbf{p}_t$, $\widehat{\rho}_Y$ and $\widehat{\rho}_O$, and members strategically react to this. Threshold voting rules will again play an important role; we denote the bias used by young members as $B_Y[\theta_i, \boldsymbol{\theta}_t^O, x(t)]$ and by old members $B_O[\theta_i, \boldsymbol{\theta}_{-it}^O, x(t)]$.

2.2.1 Equilibrium dynamics with sincere voting

The theoretical voting literature assumes that committee members can adopt two types of behaviors, *sincere* and *strategic* (Austen-Smith and Banks 1996). We first assume the former, which means that members vote for the rate that is consistent with their private signal. More concretely, member i has preferences in periods $t \in \{\tau(i), \tau(i)+1\}$ given by

$$u_{it} = u\left[v_{it}, \omega_t, \theta_i\right] = M\left[\pi(v_{it}, \omega_t)\right] + \chi(\theta_i)\left[\pi(v_{it}, \omega_t) - \pi_t^E\right] \tag{7}$$

On the other hand, the public forms π_t^E by taking expectations over the policy r_t and policymakers strategically react to its doing so.

Just as with the single banker, the model with sincere voting features unambiguous policy dynamics.

Proposition 3 *Under sincere voting, all equilibrium strategies are threshold voting rules that satisfy:*

1. $B_x[\theta_i, \cdot]$ *is strictly increasing in* θ_i *for* $x = Y, O$.
2. $B_O[\theta_i, \cdot] - B_Y[\theta_i, \cdot] > 0$ *for all* $\theta_i, \boldsymbol{\theta}_{\tau(i)}^O, \boldsymbol{\theta}_{-i,\tau(i)+1}^O$
3. $B_O[\theta_i, \cdot] - B_Y[\theta_i, \cdot]$ *is strictly increasing in* θ_i *for all* $\boldsymbol{\theta}_{\tau(i)}^O, \boldsymbol{\theta}_{-i,\tau(i)+1}^O$

Under sincere voting, old members are in exactly the same position as the single central banker in period 2 since future inflation expectations are independent of their votes and they do not react strategically to colleagues' votes. So, their bias only depends on θ_i and is simply

$$B_O[\theta_i] = \frac{\mu(0) + \chi(\theta_i)\left[\pi(0,0) - \pi(1,0)\right]}{\mu(1) - \chi(\theta_i)\left[\pi(0,1) - \pi(1,1)\right]}. \tag{8}$$

The difference arises for young members, who know that while their votes affect period $t+1$ inflation expectations, so too do those of their young colleagues. Let

$$\overline{D}\left(\omega_t, x(t)\right) = \mathbb{E}\left\{ \pi^E_{t+1}\left[p_i(0), \mathbf{p}_{-i,t+1}, \omega_t\right] - \pi^E_{t+1}\left[p_i(1), \mathbf{p}_{-i,t+1}, \omega_t\right] \right\} \tag{9}$$

be the expected change in period $t+1$ inflation expectations under the posterior member $i \in Y(t)$ induces over his type by voting $v_{i\tau(i)} = 0$ instead of $v_{i\tau(i)} = 1$ when the realized state is ω_t.[14] Young member $i \in Y(t)$ votes high if and only if

$$\frac{\widehat{\omega}_{it}}{1-\widehat{\omega}_{it}} \geq \frac{\mu(0) + \chi(\theta_i)\left\{[\pi(0,0) - \pi(1,0)] - \delta\overline{D}\left(0, x(t)\right)\right\}}{\mu(1) - \chi(\theta_i)\left\{[\pi(0,1) - \pi(1,1)] - \delta\overline{D}\left(1, x(t)\right)\right\}} = B_Y[\theta_i, x(\tau(i))]. \tag{10}$$

where $\overline{D}\left(\omega_t, x(t)\right)$ is again an endogenous equilibrium quantity. While no single old member unilaterally determines policy, the probability that the period $t+1$ committee chooses $r_t = 1$ is increasing in the hawkishness of any single old member. At the same time, more hawkish young members are more likely to vote high. So, the public associates high votes by young members with lower future inflation, meaning the equilibrium value of $\overline{D}\left(\omega_t, x(t)\right)$ is positive. The one new theoretical message is that the strength of the signalling incentive in each period t can vary according to the composition of the committee.

2.2.2 Strategic voting

With strategic voting, members do not derive utility from their votes *per se*, but care only about the final policy decision, so that

$$u_{it} = u\left[r_t, \omega_t, \theta_i\right] = M\left[\pi(r_t, \omega_t)\right] + \chi(\theta_i)\left[\pi(r_t, \omega_t) - \pi^E_t\right]. \tag{11}$$

Here strategic refers to members' conditioning their votes on being pivotal (i.e. changing the committee decision) when choosing v_{it}; in both the strategic and sincere voting cases members behave strategically vis-à-vis the public's inflation expectation formation rule.

Let $\Pr\left[\,\mathrm{PIV}_{it} \mid \omega_t\,\right]$ be the expected probability that member i is pivotal in state ω_t. A *responsive* equilibrium is one in which all members vote high for some signal realizations and low for others. We focus on responsive equilibria in which members adopt threshold

[14] Symmetric strategies ensure that this change is independent of the identity of voter i, but the heterogenous committee composition in terms of new and old members means the change varies over time.

voting rules, so old members vote high if and only if[15]

$$\frac{\widehat{\omega}_{it}}{1-\widehat{\omega}_{it}}\frac{\Pr\left[\,\mathrm{PIV}_{it}\mid 1\,\right]}{\Pr\left[\,\mathrm{PIV}_{it}\mid 0\,\right]}\geq\frac{\mu(0)+\chi(\theta_i)\left[\pi(0,0)-\pi(1,0)\right]}{\mu(1)-\chi(\theta_i)\left[\pi(0,1)-\pi(1,1)\right]}. \tag{12}$$

and young members vote high if and only if

$$\frac{\widehat{\omega}_{it}}{1-\widehat{\omega}_{it}}\frac{\Pr\left[\,\mathrm{PIV}_{it}\mid 1\,\right]}{\Pr\left[\,\mathrm{PIV}_{it}\mid 0\,\right]}\geq\frac{\mu(0)+\chi(\theta_i)\left\{\left[\pi(0,0)-\pi(1,0)\right]-\frac{\delta\overline{D}(0,x(t))}{\Pr\left[\mathrm{PIV}_{it}\mid 0\right]}\right\}}{\mu(1)-\chi(\theta_i)\left\{\left[\pi(0,1)-\pi(1,1)\right]-\frac{\delta\overline{D}(1,x(t))}{\Pr\left[\mathrm{PIV}_{it}\mid 1\right]}\right\}}. \tag{13}$$

where $\overline{D}\left(\omega_t, x(t)\right)$ is the same as in (9), but computed under whatever beliefs the public has on the voting strategies employed with strategic behavior. We will refer to the right hand sides of (12) and (13) as the biases used by old and young members under strategic voting, thus modifying slightly the definition in (2). These are also the quantities we will estimate in our empirical exercise when we consider strategic voting.

The logic of the dynamics discussed above ultimately relied on the probability of low interest rates increasing in experienced bankers' dovishness. With strategic voting, it remains true that old banker i's dovish bias is increasing in θ_i. But whether this translates into the committee choosing low rates more often is complicated by two factors. First, θ_i affects other members' voting strategies, so there are second-order effects on colleagues' voting behavior of a shift in θ_i. Second, this reaction in turn impacts the probability i is pivotal, and so the probability he chooses low rates in equilibrium. Characterizing how the probabilities of being pivotal are affected by θ_i is difficult in general, so instead we state strategic dynamics given the empirically testable condition that committees composed of more dovish types vote low more often (which we empirically confirm in our data in appendix B).

Proposition 4 *Under strategic voting, in any responsive equilibrium in which the probability that $r_t = 0$ is strictly increasing in θ_i $\forall$ $i \in Y(t-1)$, $\boldsymbol{\theta}_{-i,\tau(i)+1}$, and q_t:*

1. *$B_x[\theta_i, \cdot]$ is strictly increasing in θ_i for $x = Y, O$.*
2. *$B_O[\theta_i, \cdot] - B_Y[\theta_i, \cdot] > 0$ for all $\theta_i, \boldsymbol{\theta}^O_{\tau(i)}, \boldsymbol{\theta}^O_{-i,\tau(i)+1}$*
3. *$B_O[\theta_i, \cdot] - B_Y[\theta_i, \cdot]$ is strictly increasing in θ_i for all $\boldsymbol{\theta}^O_{\tau(i)}, \boldsymbol{\theta}^O_{-i,\tau(i)+1}$*

One important question is whether the condition on equilibrium strategies in proposition 4 is in principle consistent with strategic behavior. To answer this, we consider a special case of our OLG model that corresponds to that of Sibert (2003), the only theory paper in the literature (of which we are aware) that explores reputation on a

[15] Duggan and Martinelli (2001) show in a simpler environment that such responsive equilibria are the only equilibria.

monetary policy committee rather than a single policy maker. Suppose that $N = 2$ and that each committee has one young and one old member; moreover, suppose that in case one member votes high and other low, r_t is chosen with probability α.

Example 2 [16] *When $\alpha = \frac{1}{2}$, the probability that $r_t = 0$ is strictly increasing in the dovishness of the old member.*

When the disagreement rule is symmetric (or, by continuity, close to symmetric), we can pin down unique equilibrium dynamics. This is because the probability that any member is pivotal is symmetric across states, so that strategic reasoning does not provide additional information beyond the private signal realization.

2.3 Discussion of preferences

While the assumption that utility is linear in the output gap is common in the literature, it has also been criticized. For example, Faust and Svensson (2001) point out that this specification implies that policy makers are indifferent to output volatility, and instead propose a quadratic loss term in the output gap. While this specification indeed yields a preference for a smooth output path, it sits somewhat awkwardly with our suspicion that most policy makers would prefer to see output higher if inflation remained at its target level. One way of reconciling these observations is to assume that utility is concave in the output gap by generalizing (1) to

$$u\left[r_t, \omega_t, \theta\right] = M\left[\pi(r_t, \omega_t)\right] + \chi(\theta) H\left[\pi(r_t, \omega_t) - \pi_t^E\right]. \tag{14}$$

where $H' > 0$ and $H'' < 0$. Now the single central banker in the second period or old members in committees will use the bias

$$\frac{u\left[0,0,\theta\right] - u\left[1,0,\theta\right]}{u\left[1,1,\theta\right] - u\left[0,1,\theta\right]} = \frac{\mu(0) + \chi(\theta)\left\{H\left[\pi(0,0) - \pi_t^E\right] - H\left[\pi(1,0) - \pi_t^E\right]\right\}}{\mu(1) - \chi(\theta)\left\{H\left[\pi(0,1) - \pi_t^E\right] - H\left[\pi(1,1) - \pi_t^E\right]\right\}}. \tag{15}$$

Unlike with linear utility, the bias now depends on the current level of inflation expectations, which can introduce a new motivation for signalling hawkishness. One can show that if $\pi(0,1) - \pi(1,1) = \pi(0,0) - \pi(1,0)$ and $H''' \leq 0$, then (15) is strictly decreasing in π_t^E. Thus reducing future inflation expectations when young can endogenously commit a banker to being tougher on inflation when old. Hence we conjecture that the introduction of non-linear utility can reënforce the value of signalling that our baseline model identifies under natural conditions.

Another assumption in the model, which follows the existing literature, is that the policy makers are looking to anchor inflation expectations against a tendency to rise.

[16] See details of the calculation in appendix A.

Of course, there may be periods, such as when the economy is in a liquidity trap, when policy makers wish to raise inflation expectations. We do not explore such cases as they do not apply in the sample period that we consider below, but in such circumstances our model would predict the reverse dynamics - members would be more dovish at the start of their tenure and, with experience, become more hawkish. The recent appointment of a new, much more dovish, Governor of the Bank of Japan might be well-captured by our model in such a situation.

3 Data and Reduced-Form Evidence

In this section we first outline the institutional setting and data from the Bank of England MPC and then turn to a reduced-form empirical analysis that suggests that there are significant policy dynamics.

3.1 Data

The MPC has met once a month since June 1997 to set UK interest rates. It has nine standing members (five Bank executives, or internal members, and four external members) who are required to vote independently.[17] Plurality rule determines the interest rate, with the Governor deciding in the case of a tie. Disagreements between members are the rule rather than the exception; 64% of the meetings in the sample have at least one deviation from the committee majority and there are many meetings decided by a vote of 5-4 or 6-3.[18]

Its remit, as defined in the Bank of England Act (1998) is to "maintain price stability, and subject to that, to support the economic policy of Her Majesty's government, including its objectives for growth and employment." In practice, the committee seeks to achieve a target inflation rate of 2%, based on the Consumer Price Index.[19] If inflation is greater than 3% or less than 1%, the Governor of the Bank of England must write an

[17]The standard term of office, which can be renewed, is three years (36 meetings) and the average served by members in our sample is 46 meetings; the current Governor, Mervyn King, is present in all 142 of our meetings, while Howard Davies served in only the first two meetings. According 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.

[18]For more institutional details, see Lambert (2006).

[19]This 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%.

open letter to the Chancellor explaining why.

We analyze the MPC voting record up to March 2009, when the interest rate reached its effective zero lower bound and a period of quantitative easing began; from this time, the main MPC decision concerned the additional policy of how many assets purchases to make. This sample yields a total of 142 meetings, and 1246 individual votes.[20]

In our model, individual votes (v_{it}) are restricted to take on one of two values that we denote 0 (a lower rate) and 1 (a higher rate). This assumption is not restrictive in the context of the Bank of England since in the overwhelming majority of meetings (135 of 142), all members either vote for the same rate, or one of two interest rates. We denote by $\widehat{v}_{it}$ the empirical counterpart of v_{it}. To construct it, we need to know which were the two interest rate decisions on the agenda in a given meeting; to identify these policy options, we follow the following procedure:

1. In periods with two unique votes by MPC members (84 of the 142 meetings), the votes define the two possible decisions under consideration.

2. In meetings with unanimous votes, we do not directly observe which alternative was under consideration. In such cases (51 of the 142 meetings) we make use of a survey of market economists which is conducted in the days leading up to the MPC meeting by Reuters.[21] The survey asks 30-50 market economists from financial institutions in London to predict the outcome of MPC voting by writing a probability distribution over possible interest rate 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 regarding inflationary pressures. As such, we can determine the two decisions over which the decision was made using the two most likely outcomes.[22]

3. In meetings where three unique interest rates are voted for (7 of the 142 meetings) we use a procedure which identifies the two most likely choices to make up the binary agenda. We then pool the votes for the third option into the most closely related option that is part of the binary agenda.

Once we have the two possible interest policies for each meeting, we simply set $\widehat{v}_{it} = 1$ if member i is observed to vote for the higher of the two rates.[23] Figure 2 displays the

[20] These data are available from the Bank of England (2010b). We use each regular MPC meetings in this period but we drop from the dataset the (unanimous) emergency meeting held after 9/11.

[21] These data are further discussed in the online data appendix.

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

[23] We observe votes, v_{it}, for nine unique interest rate changes in our data (-150bps, -100bps, -75bps, -50bps, -40bps, -25bps, no change, +25bps and +50bps). Depending on the set of votes we observe in a given meeting, a particular vote can be mapped into either $\widehat{v}_{it} = 0$ or $\widehat{v}_{it} = 1$. For example, if we observe

large variation across both internal and external members, across time, in their behavior measured by $\widehat{v}_{it}$.

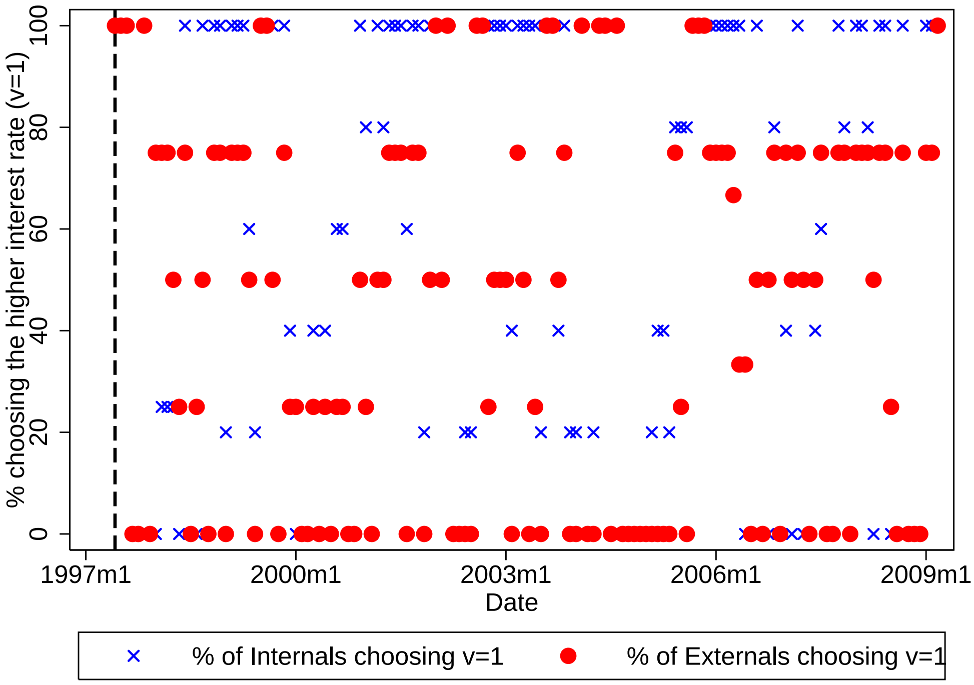

Figure 2: Percentage of Internal and External MPC Members choosing $\widehat{v}_{it} = 1$

Notes: This figure shows the percentage of MPC members, separately for internal (dots) and external (x's) members, who vote for the higher interest rate in each given meeting.

3.2 Reduced form evidence

As we are interested in voting dynamics, we begin by examining whether, in a reduced form sense, there is any behavior of interest. To do this, we define a dummy variable to indicate when a member has completed 18 meetings on the MPC:

$$\text{D(Experienced)}_{it} = \begin{cases} 0 & \text{if member } i \text{ has served in 18 or less meetings} \\ 1 & \text{if member } i \text{ has served in more than 18 meetings} \end{cases} \tag{16}$$

Accordingly, we define a member as *new* if $\text{D(Experienced)}_{it} = 0$ and *experienced* if $\text{D(Experienced)}_{it} = 1$.

As a first look at dynamic voting behavior, we estimate the following relationship:

$$\widehat{v}_{it} = \alpha_i + \gamma \text{D(Experienced)}_{it} + \delta_t + \epsilon_{it} \tag{17}$$

votes for no change and +25bps, then a vote for +25bps maps into $\widehat{v}_{it} = 1$, while if we observe votes for +25bps and +50bps, it maps into $\widehat{v}_{it} = 0$.

This equation includes both member and time fixed effects (α_i, the member fixed effect, captures a member specific intercept while δ_t, the time fixed effect, captures the average vote in period t).

Table 1: Reduced form evidence on the impact of experience

	(1) $\widehat{v}_{it}$	(2) $\widehat{v}_{it}$	(3) $\widehat{v}_{it}$	(4) $\widehat{v}_{it}$	(5) $\widehat{v}_{it}$
D(Experienced)	-0.092***			-0.091***	
	[0.001]			[0.001]	
D(Experienced - 12M)		-0.081***			
		[0.005]			
D(Experienced - 24M)			-0.053*		
			[0.070]		
D(Term End)					-0.0081
					[0.728]
Constant	0.98***	0.95***	0.93***	0.99***	1.01***
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
R-squared	0.704	0.703	0.702	0.708	0.701
Model	Panel LPM	Panel LPM	Panel LPM	Panel LPM	Panel LPM
Member effects?	YES	YES	YES	YES	YES
Time effects?	YES	YES	YES	YES	YES
Sample?	06/97-03/09	06/97-03/09	06/97-03/09	12/98-03/09	06/97-03/09
Obs?	1246	1246	1246	1106	1246

Notes: This regression presents OLS estimates of equation (17) with standard errors clustered by member. The dependent variable, $\widehat{v}_{it}$, is our measure of whether member i votes for the high interest rate in period t. The results show that, controlling for member and time fixed effects, members with experience (defined as 18 meetings experience in column (1), 12 meetings in column (2), and 24 meetings in column (3)) vote for lower interest rates than new members.

The results, reported in column (1) of table 1, show that as members serve more time, they vote for lower interest rates on average. We can also show that the reduction in average interest rates with experience at the individual level is robust to alternative definitions of experience. We create two alternative dummy variables called D(Experienced - 12M)$_{it}$ and D(Experienced - 24M)$_{it}$ along the lines of equation (16), except that these measure experience as any tenure over 12 and 24 meetings, respectively. The results with these alternative definitions are reported in columns (2) and (3) of table 1; again, we find that experienced members vote for lower rates on average. Since all three dummy variables give the same qualitative results, we will use the 18 meeting definition (as in equation (16)) simply because it represents half of a standard MPC member's term and therefore splits the sample into subsamples of roughly similar size. One might be concerned that the results of the analysis are somehow driven by the period at the start of the committee when all members were, by definition, new. In column (4) we drop the first 18 months of the committee and find that the results are unaffected. The final concern is that we are actually capturing some end of term effect as members attempt to get reappointed to their position. To examine this alternative signalling idea, we generate a dummy which measures whenever members are within 9 months of the end of their appointment.[24] There is no evidence of the end-of-term coinciding with a change in behavior.

The baseline size of the estimated effect of experience can broadly be taken to mean that an experienced member is 9pp less likely to vote for the higher rate on the agenda. If, for example, a new member had a 55% probability of voting for the high rate, an experienced member would have probability of 46%; a committee of all new members would, in such a circumstance, have a 62% probability of voting for the high rate compared to a 40% probability for a committee of all experienced members.

4 Structural Estimation: Methodology and Results

The reduced form results above suggest that there is "delayed dovishness" for the average MPC member when they serve on the committee. Because we have included member- and time- fixed effects in (17), this does not reflect changing composition of the committee, but rather indicates that something at the individual level systematically shifts over time. In this section we shed light on what this "something" actually is.

Primae facie it is unclear what evolves with time based on the reduced form evidence alone. One can well imagine that expertise might vary over tenure on the committee. For example, learning by doing might increase members' ability to perceive economic conditions, which would increase the informativeness of their signals. Also, as we will show in section 5, a shift in the bias need not necessarily lead to an observed shift in the actual interest rate chosen.

4.1 Estimation methodology

In our model, committee members vote for high interest rates if their signal exceeds a certain tenure-specific threshold:

$$s_{it} \geq \mathfrak{L}_Y^{-1} \left\{ \ln \left[\frac{1-q_t}{q_t} \Theta_i^x \right] \right\} \text{ for } x = \{Y, O\} \tag{18}$$

[24] We define:

$$\text{D(Term End)}_{it} = \begin{cases} 1 & \text{if member } i \text{ is serving in the last 9 meetings of their appointment} \\ 0 & \text{otherwise} \end{cases}$$

If a member's first term is renewed, D(Term End)_{it} resets to zero.

where $\Theta_i^Y \equiv B_Y(\theta_i)$ and $\Theta_i^O \equiv B_O(\theta_i)$ are the biases of young and old members as defined in the model.[25] Using this notation, we can summarize the three main predictions of our theoretical model as presented in figure 1 and propositions 3 and 4:

H1 $\Theta_i^Y < \Theta_i^O \; \forall \theta_i$.

H2 $\Theta_i^O - \Theta_i^Y$ is strictly increasing in θ_i.

H3 $\Theta_i^Y < \Theta_j^Y$ and $\Theta_i^O < \Theta_j^O \; \forall \theta_i < \theta_j$.

H1 says that the average member becomes less tough on inflation with experience, H2 states that this effect is greater for inherently more dovish members and H3 says that dovish members are always less tough on inflation than hawkish members. We now examine these predictions using a structural estimation of our model. In order to make our model empirically operational, we assume that private signals are unbiased and drawn from a normal distribution: $s_{it} \sim N(\omega_t, \sigma_i^2)$ where σ_i is a measure of expertise (lower σ_i indicates more expertise). The voting strategies (for $x = \{Y, O\}$) then become $v_{it} = 1$ if and only if

$$s_{it} \geq \underbrace{\frac{1}{2} - (\sigma_i^x)^2 \left[\ln\left(\frac{q_t}{1-q_t}\right) - \ln\left(\Theta_i^x\right) \right]}_{\equiv s_{it}^{\text{SIN}}\left(\Theta_i^x, \sigma_i^x, q_t\right)} \tag{19}$$

under sincere voting and

$$s_{it} \geq \underbrace{\frac{1}{2} - (\sigma_i^x)^2 \left[\ln\left(\frac{q_t}{1-q_t}\right) - \ln\left(\Theta_i^x\right) + \ln\left(\frac{\Pr\left[\text{PIV}_{it} \mid 1, \mathbf{s}_{-it}^{STR}\right]}{\Pr\left[\text{PIV}_{it} \mid 0, \mathbf{s}_{-it}^{STR}\right]}\right) \right]}_{\equiv s_{it}^{\text{STR}}\left(\Theta_i^x, \sigma_i^x, q_t\right)} \tag{20}$$

under strategic voting.

Using this notation, the likelihood L_t that we observe a vector of votes $\mathbf{v}_t$ is:

$$\begin{aligned} L_t = q_t \prod_i \left[1 - \Phi\left(\frac{s_{it}^*(\cdot) - 1}{\sigma_i}\right)\right]^{\widehat{v}_{it}} \left[\Phi\left(\frac{s_{it}^*(\cdot) - 1}{\sigma_i}\right)\right]^{1-\widehat{v}_{it}} + \\ (1-q_t) \prod_i \left[1 - \Phi\left(\frac{s_{it}^*(\cdot) - 1}{\sigma_i}\right)\right]^{\widehat{v}_{it}} \left[\Phi\left(\frac{s_{it}^*(\cdot)}{\sigma_i}\right)\right]^{1-\widehat{v}_{it}} \end{aligned} \tag{21}$$

where $\Phi(\cdot)$ is the normal cdf.

[25] Here we abstract away from the dependence of the bias on variables other than θ_i; our empirical estimates will measure the average value of the bias across different committee compositions and priors.

To structurally estimate Θ and σ variables for young and old members, we use a two-stage estimator. In this approach, we rewrite this likelihood function as

$$L_t^{ALT} = q_t \prod_i (\kappa_{1it})^{\widehat{v}_{it}} (1-\kappa_{1it})^{1-\widehat{v}_{it}} + (1-q_t) \prod_i (\kappa_{0it})^{\widehat{v}_{it}} (1-\kappa_{0it})^{1-\widehat{v}_{it}} \tag{22}$$

where $\kappa_{1it} \equiv 1 - \Phi\left(\frac{s_{it}^*(\cdot)-1}{\sigma_i}\right)$ and $\kappa_{0it} \equiv 1 - \Phi\left(\frac{s_{it}^*(\cdot)}{\sigma_i}\right)$ are the probabilities of voting high in states 1 and 0. We then model q_t and the κ terms as functions of observed covariates (proxies for the prior P_t; time-varying voter characteristics X_{it} which includes their experience; and meeting characteristics that might affect voters' willingness to vote high Z_t) as follows:

$$q_t = \frac{\exp(\alpha \cdot P_t)}{1+\exp(\alpha \cdot P_t)} \tag{23}$$

and

$$\kappa_{0it} = \frac{\exp(\beta_0 \cdot X_{it} + \beta_1 \cdot P_t + \beta_2 \cdot X_{it} \cdot P_t + \beta_3 \cdot Z_t)}{1+\exp(\beta_0 \cdot X_{it} + \beta_1 \cdot P_t + \beta_2 \cdot X_{it} \cdot P_t + \beta_3 \cdot Z_t)} \tag{24}$$

$$\kappa_{1it} = \frac{\kappa_{0it} + \exp(\gamma_0 \cdot X_{it} + \gamma_1 \cdot P_t + \gamma_2 \cdot X_{it} \cdot P_t + \gamma_3 \cdot Z_t)}{1+\exp(\gamma_0 \cdot X_{it} + \gamma_1 \cdot P_t + \gamma_2 \cdot X_{it} \cdot P_t + \gamma_3 \cdot Z_t)}. \tag{25}$$

Under either specification of the κ terms, the estimation of the structural parameters follows a two-step procedure:[26]

1. Estimate using maximum likelihood the α, β, and γ parameters of the mixture model and obtain fitted values $\widehat{q}_t$, $\widehat{\kappa}_{0it}$, and $\widehat{\kappa}_{1it}$.

2. Use these fitted values to recover the structural parameters from the theoretical voting probabilities. An estimate of period-t expertise comes via

$$\widehat{\sigma}_{it} = \frac{1}{\Phi^{-1}(1-\widehat{\kappa}_{0it}) - \Phi^{-1}(1-\widehat{\kappa}_{1it})}. \tag{26}$$

   We uncover Θ_{it} via

$$\widehat{s}_{it}^* = \frac{\Phi^{-1}(1-\widehat{\kappa}_{0it})}{\Phi^{-1}(1-\widehat{\kappa}_{0it}) + \Phi^{-1}(\widehat{\kappa}_{1it})} \tag{27}$$

   along with $\widehat{q}_t$ and (19) [(20)] under the assumption of sincere (strategic) voting.[27]

The second stage yields an estimate of preference and precision parameters for each voter for each unique value of $\widehat{q}_t$. We report for young and old members separately the median values of these estimates, $\widehat{\Theta}_i$ and $\widehat{\sigma}_i$, as point estimates.[28] In order to conduct hypothesis

[26]For further details on this two-step procedure, see Iaryczower and Shum (2012), Hansen, McMahon, and Velasco (2012) and Hansen and McMahon (2013).

[27]More precisely, (20) defines an NxN system of equations that must be solved to recover biases.

[28]The estimated dynamics using the mean of these estimates are qualitatively unchanged.

tests, we bootstrap on the first-stage estimates. Because we report estimates of Θ that are pooled across time periods and members, one can interpret them as uncovering the average dynamics on the committee.

In order to implement this approach, in P_t we use two proxies for the prior that we argue correlate with economic conditions. Full details of the construction of both are presented in the online data appendix, and here we limit ourselves to a brief discussion. The first proxy q_t^R is based on the same Reuters survey data we use to construct $\widehat{v}_{it}$ (see above). q_t^R is the average probability survey respondents place on the higher interest rate choice over the total average probability placed on both policy rates. A second proxy q_t^M is derived from the cross-section of prices for short sterling futures options on the first day (Wednesday) of the MPC meeting. Short sterling futures contracts are standardized futures contracts which settle on the 3-month London Interbank Offered Rate (LIBOR) on the contract delivery date. We use the distribution of these prices to construct a distribution over risk-neutral traders' beliefs on the expected value of the change in 3-month LIBOR. We then extract probabilities attached to discrete changes in the interest rate, and construct q_t^M as the average probability placed on the higher interest rate option over the total average probability placed on both policy options.

Certainly both q_t^R and q_t^M have weaknesses. q_t^R is not a perfect measure of the unobservable q_t because, for example, the data on which respondents form their beliefs is a subset of that available to the MPC since the committee is regularly given advance access to data that will only be released subsequently to the wider public. Also, q_t^R predicts the outcome of MPC voting, not the probability of the realization of the underlying state variable. For q_t^M the first concern is that the sterling options that go into the constructed probabilities are based on LIBOR rather than the interest rate that MPC members choose ("Bank Rate"). Second, it is based on beliefs about the 3-month interest rate and as such reflects expected changes over the next three meetings not simply the one immediately following. Third, the beliefs are associated with risk-neutral traders; to the extent that actual traders are risk averse the beliefs backed out from option price data will be biased by the presence of a risk-premium in the observed market data. Finally, as with the Reuters data, q_t^M captures predictions about LIBOR, not underlying inflation states. Despite these problems it is a useful measure because it aggregates the opinions of a large number of agents (all traders in the sterling options market) and, unlike with the Reuters data, these opinions are backed by real money and so potentially less subjective and manipulable.[29] However, it is worth emphasizing that we do not take q_t^R and q_t^M to be the actual q_t, but estimate their relationship with q_t.

[29]Though, following the recent LIBOR scandals, the irony of using market data related to LIBOR for the fact that it is not manipulable is not lost on us.

4.2 Estimation results

Before turning to the estimates of the structural parameters, table 2 below presents the results of the first-stage maximum likelihood estimation of equations (23), (24), and (25). In this baseline specification we control for whether a member is experienced or new using the D(Experienced) variable defined above to capture the effect of tenure. We also include a variable D(Internal) in order to control for whether the member is an internal or external MPC member (i.e. the member's type of appointment). We do this because Hansen, McMahon, and Velasco (2012) found that along the internal-external split members display different levels of expertise and biases, and both types of heterogeneity are useful for testing our model and assessing its implications. As a result of this expertise difference, it is important, as shown in Hansen and McMahon (2013), that we interact D(Internal) with the q_t^R and q_t^M variables to capture the differential effect of the prior across different levels of expertise.[30] In terms of meeting characteristics Z_t, we control for meetings in which members had at least one choice on the agenda to hike interest rates,[31] and for a measure of the committee composition that captures if the meeting has three or more new members serving.

Given the non-linearities in moving from first-stage estimates to the second-stage estimates, and the fact that the first-stage estimates for some of the main variables of interest contain interaction terms, the first-stage results are less interesting than the structural estimates to which we shortly turn and so we will not discuss each coefficient estimate in detail. However, it is worth pointing out that in column (1) of table 2 we report estimates for the coefficients on our proxies for the prior and find a large and significant relationship between the prior q_t and the Reuter's proxy q_t^R, while the relationship between the q_t and the proxy derived from market data q_t^M is insignificant. This indicates that our Reuter's survey data is a good predictor of members' prior beliefs to which market price data adds little. Also, the voting probabilities depend on the composition of the committee, as the signalling model predicts.

Using these first-stage results, the estimated structural parameters (the second-stage results) are presented in table 3. The first two rows of Table 3 show the estimates of Θ for new and experienced members assuming both sincere and strategic voting. We also report the difference between these two estimates, $\Delta\Theta = \Theta^O - \Theta^Y$. The table also reports, in brackets below the difference estimates, the p-value associated with the difference being zero or negative. The p-values are calculated using 1,000 boot-strapped draws for the

[30] We have also run the regression with interactions terms between the proxies and D(Experienced) without changing the results. Since we find no average difference in expertise between new and experienced members, there is no theoretical reason to include the interactions and we omit them for reasons of parsimony.

[31] The most common such meetings are those that have a choice of no change and a choice of raising by 25 basis points.

Table 2: Structural Model: First Stage Estimates

	(1) $\widehat{q}_t$	(2) $\widehat{\kappa}_{0t}$	(3) $\widehat{\kappa}_{1t}$
Reuter's q^R	7.10***	1.04	3.39***
	[0.001]	[0.348]	[0.001]
Market q^M	3.33	0.34	0.11
	[0.141]	[0.815]	[0.934]
D(Experienced)		-0.80**	-1.40***
		[0.058]	[0.001]
D(Internal)		-3.05***	1.71
		[0.01]	[0.467]
D(Internal) x D(Experienced)		0.57	0.05
		[0.314]	[0.969]
D(Internal) x Reuter's q^R		3.20***	-2.15
		[0.009]	[0.335]
D(Internal) x Market q^M		3.12	4.55
		[0.148]	[0.181]
D($\geq$ 3 New)		0.87***	-0.62**
		[0.008]	[0.042]
D(Hike)		1.40***	1.61***
		[0.008]	[0.01]
Constant	-5.43***	-3.73***	-0.65
	[0.001]	[0.001]	[0.496]

Notes: This table shows the estimated coefficients of equations (23), (24), and (25) for the first-stage maximum likelihood estimation of the baseline specification. Columns (1), (2) and (3) provide equations from which we get the fitted values $\widehat{q}_t$, $\widehat{\kappa}_{0t}$ and $\widehat{\kappa}_{1t}$ respectively. Significance of coefficient estimates is reported using p-values in brackets.

estimates.

Table 3: Baseline Estimates of Structural Parameters

	New	**Experienced**	**Difference**
$\boldsymbol{\theta(SIN)}$	0.22	1.28	1.05
			[0.000]
$\boldsymbol{\theta(STR)}$	0.06	0.74	0.69
			[0.002]
$\boldsymbol{\sigma}$	0.40	0.38	-0.02
			[0.291]

Notes: This table shows the structural estimates for New (column 1) and Experienced (column 2) members, as well as the difference between them (column 3). The rows report the estimates for preferences under sincere voting ($\theta(SIN)$) and strategic voting ($\theta(STR)$), as well as the precision parameter (σ). We report, in brackets below the difference estimate, the p-value of a one-sided test that the difference is significantly non-zero; the test is calculated using a boot-strapped distribution of estimates.

The results concerning member preferences are consistent with the reduced form results while also guiding us as to their source. Consistent with H1 (the prediction that $\Theta^Y < \Theta^O$) we find evidence of a significant upward shift in the Θ parameter with experience. This is true whether we assume members vote sincerely or strategically. Table 3 also presents estimates for the behavior of expertise with experience. There is no evidence of any change, statistically speaking, in the expertise parameter (σ) parameter. This result is of independent interest since it suggests that members do not accumulate additional expertise with experience. Instead, voting dynamics are driven entirely by a shift in the average member's bias.

To conclude, the dynamics in the strategic case were stated in terms of a restriction on equilibrium strategies that required the probability that the committee chose low rates to be increasing in the dovishness of each experienced members. In appendix B we show that this restriction is consistent with the estimated thresholds.

4.3 Type dependence of the signalling effect

While the results above are evidence in support of the our model's first prediction on dynamics, we also want to test H2 (whether $\Theta^O - \Theta^Y$ is increasing in dovishness). The first way we do so is to decompose the experience effect by type of appointment (internal or external). Previous literature has emphasized that internal members have a more hawkish bias than externals, so we believe it represents a good proxy for hawkishness. In table 4, we use the baseline estimates reported above, but split the results by D(Internal); this table confirms that internal members are more hawkish (lower Θ). The overall dynamic behavior reported above is common to both internals (or hawks) and externals (or doves); both have a significant shift in their measured bias, but not in their expertise. At the same time, external members have a larger shift in their measured Θ than internals and externals are always more dovish than internals. In table 4 we report the difference in differences for all the structural parameters, along with p-values for the hypothesis test that the experience effect is greater for the external members. We find evidence that supports H2 and H3.

Table 4: Internals vs Externals

	Internal			External			
	New	**Experienced**	**Difference**	**New**	**Experienced**	**Difference**	**Diff-in-Diff**
$\boldsymbol{\theta(SIN)}$	0.04	0.69	0.65	0.65	5.96	5.32	-4.67
			[0.015]			*[0.000]*	[0.000]
$\boldsymbol{\theta(STR)}$	0.02	0.44	0.42	0.12	2.13	2.01	-1.59
			[0.036]			*[0.000]*	[0.002]
$\boldsymbol{\sigma}$	0.31	0.34	0.03	0.56	0.54	-0.02	0.05
			[0.308]			*[0.403]*	[0.336]

Notes: This table replicates table 3 for internal and external members separately (see that table for details). The final column compares the effect of experience between the two groups for preferences under sincere voting ($\theta(SIN)$) and strategic voting ($\theta(STR)$), as well as the precision parameter (σ). The terms reported in brackets are p-values, calculated using a boot-strapped distribution of estimates, for a one-sided test of difference from zero.

While table 4 presents evidence consistent with H2 and H3, it is also important to examine its robustness. First, internal and external members may differ in ways other than the weights they put on output (for example, their average term lengths are different and they may have different career concerns). Second, not everyone agrees that the internal-external split approximates hawk-dove differences (although our estimates explicitly indicate they do after controlling for informational variables). For example, Hix, Hoyland, and Vivyan (2010) do an analysis of individual members and find that internal members are generally in "the centrist group" while external members can be found in both the extremes of hawk and dove.

To explore their approach, we take the ideal point estimates from a recent working paper (Eijffinger, Mahieu, and Raes (2013), EMR hereafter) that updates the MPC ideal points ranking using the methodology of Hix, Hoyland, and Vivyan (2010). We use their estimates of members' ideal points to create a dummy variable D(Hawk) that approximately splits the sample; D(Hawk) and D(Internal) have a weak correlation of around 0.15. We then add D(Hawk) and its interaction with D(Experienced) in the variables that go into κ_{ω_t} in the first stage regression. Table 5 reports the point estimates for the structural parameters by D(Hawk) that this new model yields. The conclusion is identical: both groups have an increasing bias, but the more dovish members have a significantly greater increase.

As further robustness checks, we consider two alternative constructions for D(Hawk). The first takes the fixed effects from the reduced form regression estimated above in table

Table 5: Hawks vs Doves 1 (EMR Measure)

	Hawks			Doves			
	New	**Experienced**	**Difference**	**New**	**Experienced**	**Difference**	**Diff-in-Diff**
$\boldsymbol{\theta(SIN)}$	0.08	0.20	0.12	0.76	3.52	2.76	-2.64
			[0.029]			*[0.005]*	[0.005]
$\boldsymbol{\theta(STR)}$	0.02	0.13	0.10	0.45	1.92	1.47	-1.37
			[0.01]			*[0.014]*	[0.014]
$\boldsymbol{\sigma}$	0.41	0.34	-0.07	0.33	0.34	0.01	-0.08
			[0.021]			*[0.34]*	[0.052]

Notes: This table replicates table 3 for hawk and dove members separately (see that table for details). We identify the hawks and doves using the Eijffinger, Mahieu, and Raes (2013) results. The final column compares the effect of experience between the two groups for preferences under sincere voting ($\theta(SIN)$) and strategic voting ($\theta(STR)$), as well as the precision parameter (σ). The terms reported in brackets are p-values, calculated using a boot-strapped distribution of estimates, for a one-sided test of difference from zero.

1, and uses them to split the sample broadly into two, with those members with higher (lower) fixed effects labelled hawks (doves). The results with this definition are presented in table 6 and confirm the results found above. The second alternative splits the sample according to the raw percentage of times a member voted high during their tenure on the committee; those voting high most often are labelled hawks. This measure suffers from the fact that members who served at times when the prior was high would vote high more often and so be classed as hawks even though our model predicts that all types would vote high in these circumstances. Despite this potential noise in the measure, the results using this split (reported in table 7) shows the same qualitative story as for the other hawk-dove splits. We are therefore confident that our findings are not driven by one particular split of the data. One new empirical finding that these direct hawk-dove splits reveal is that hawks' estimated expertise increases with time, indicating that, for some subsets of voters at least, learning by doing might be relevant.

Table 6: Hawks vs Doves 2 (Fixed Effects Measure)

	Hawks (FE)			Doves (FE)			
	New	**Experienced**	**Difference**	**New**	**Experienced**	**Difference**	**Diff-in-Diff**
$\boldsymbol{\theta(SIN)}$	0.04	0.21	0.17	1.21	3.40	2.19	-2.01
			[0.005]			*[0.052]*	[0.067]
$\boldsymbol{\theta(STR)}$	0.01	0.16	0.14	0.36	2.99	2.63	-2.49
			[0.012]			*[0.007]*	[0.006]
$\boldsymbol{\sigma}$	0.37	0.32	-0.05	0.34	0.38	0.05	-0.09
			[0.097]			*[0.177]*	[0.076]

Notes: This table replicates table 3 for a second split of hawk and dove members separately (see that table for details). For this table, hawks are defined as those members with higher fixed effects as estimated in the reduced form regression reported in table 1. The final column compares the effect of experience between the two groups for preferences under sincere voting ($\theta(SIN)$) and strategic voting ($\theta(STR)$), as well as the precision parameter (σ). The terms reported in brackets are p-values, calculated using a boot-strapped distribution of estimates, for a one-sided test of difference from zero.

Table 7: Hawks vs Doves 3 (% of High Votes Measure)

	Hawks (percent)			Doves (Percent)			
	New	Experienced	Difference	New	Experienced	Difference	Diff-in-Diff
$\boldsymbol{\theta(SIN)}$	0.06	0.10	0.04	0.44	4.23	3.79	-3.75
			[0.315]			*[0.000]*	[0.000]
$\boldsymbol{\theta(STR)}$	0.03	0.05	0.02	0.04	0.58	0.54	-0.52
			[0.235]			*[0.000]*	[0.000]
$\boldsymbol{\sigma}$	0.38	0.29	-0.09	0.39	0.39	-0.01	-0.08
			[0.008]			*[0.522]*	[0.077]

Notes: This table replicates table 3 for a third classification of hawk and dove members separately (see that table for details). For this "% of high votes measure" hawks measure, a member is classified as a hawk if (s)he has a high percentage of (her)his total votes that are high. The final column compares the effect of experience between the two groups for preferences under sincere voting ($\theta(SIN)$) and strategic voting ($\theta(STR)$), as well as the precision parameter (σ). The terms reported in brackets are p-values, calculated using a boot-strapped distribution of estimates, for a one-sided test of difference from zero.

Figure 3 summarizes the results of the different splits of the committee into hawkish and dovish members (sincere case). It shows quite clearly that, as predicted by our model and plotted in figure 1, more dovish types shift more than hawkish types. This confirms H2. It also shows that the dovish bias is ordered by dovishness, which is consistent with H3.

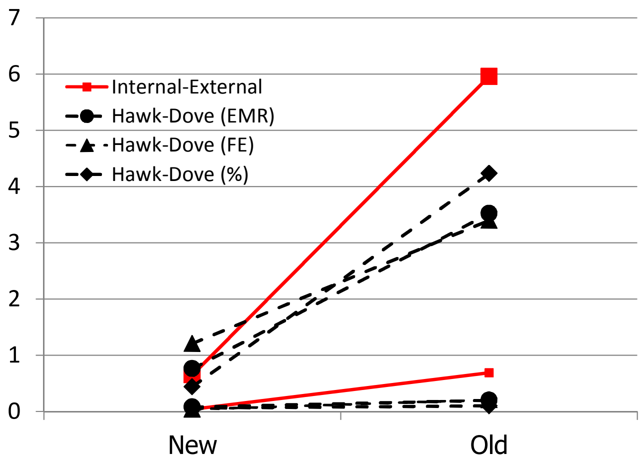

Figure 3: Change in Bias Across Different Measures of Hawk and Dove

Notes: This figure summarizes the empirical results of tables 4 to 7 for sincere voting. It shows the average effect of experience with the committee split by different measures of hawkishness. As predicted by our model, and plotted in figure 1, more dovish types shift more than hawkish types.

5 The Impact of Signalling on Policy Choices

In this section, we examine how signalling affects individual policy choices.[32] As one can see from equation (19) above, the probability a member votes high is not determined solely by the bias; both the prior and a member's expertise will determine the effect that a given shift in Θ has on the likelihood that he votes for the higher interest rate.[33] As the prior moves to the extreme values ($q_t \to 0$ or $q_t \to 1$), the shift in bias makes no difference. This is because our model assumes, contrary to most of the literature on signalling in monetary policy, that all members want to get the decision right first and foremost. Hence, when the prior is clear on the choice to make, a member's bias doesn't influence his choice. The effect of a shift in Θ on voting behavior is also influenced by expertise: (19) makes clear that a lower σ reduces the responsiveness of s^* to Θ. The intuition is that when a banker sees the state clearly, he is more likely to simply follow what his signal tells him is the right decision and ignore the relative utility loss across states of the economy.

To illustrate the influence of q_t and σ, we examine the probability of voting for the high interest rate by internal and external members separately. We choose this split because internal and external members differ in terms of the shift of the bias, but also in their expertise, which allows us to illustrate the point most clearly. Figure 4 plots the unconditional probability that an average new and old member votes for the high interest rate for different values of q_t based on the estimates in table 4. For any value of q_t, the (signalling) new central banker is more likely to vote high compared with the (non-signalling) old central banker. At the same time, as q_t becomes extreme, the probabilities for both internal and external members converge to zero or one. Also, the change in probability for the average external member is markedly higher. This could either be because external members have less expertise, or because their shift in Θ is higher.

To work out what drives externals to be more likely to choose lower rates with time, figure 5 is helpful. Figure 5a simply replicates figure 4 but plots the difference between new and experienced bankers' probabilities of voting high (the *experience effect*) for different values of the prior; this shows even more clearly how the effect for externals is much larger. In figure 5b, we plot the same gap but under the counterfactual assumption that external members have internal members' expertise levels (both when new and old) and vice versa. As can be seen, although external members still have a larger shift in Θ, the fact they now have more expertise than internal members essentially reverses the

[32]Throughout this section we perform counterfactual exercises under the sincere voting assumption. The analysis of the strategic voting case yields even larger effects of signalling on individual decisions but we maintain the sincere assumption so as to remain consistent with the existing literature on monetary policy committees.

[33]The effect of signalling on the MPC in any given period also depends trivially on the committee composition (how many members are young and seeking to signal).

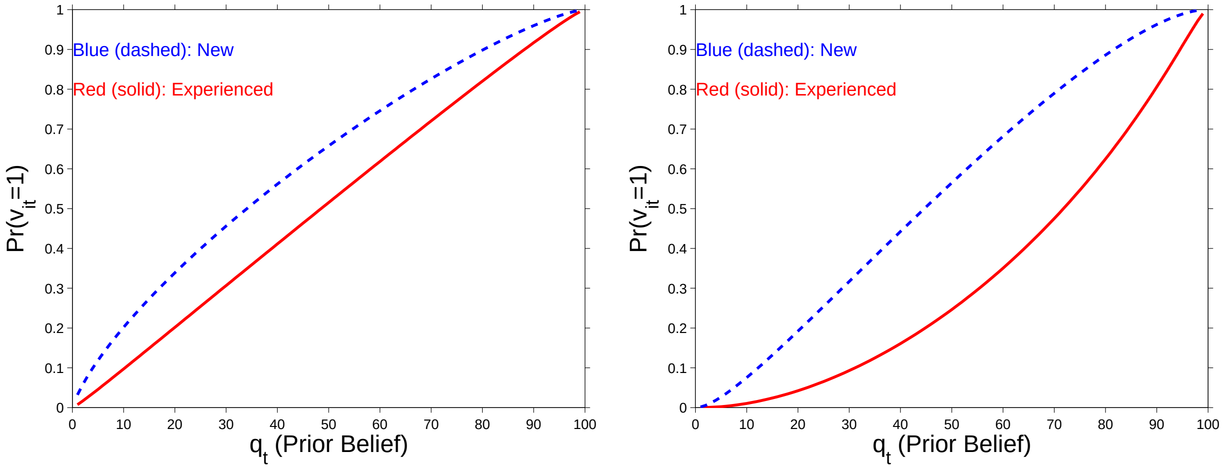

(a) Internal Members **(b)** External Members

Figure 4: Probability new and experienced members choose $v_{it} = 1$ over q_t

Notes: The top two figures show the theoretical probability that an individual committee member chooses the higher interest rate as a function of the prior belief that $\omega_t = 1$ (q_t). The different curves represent new members (dashed) and experienced members (solid).

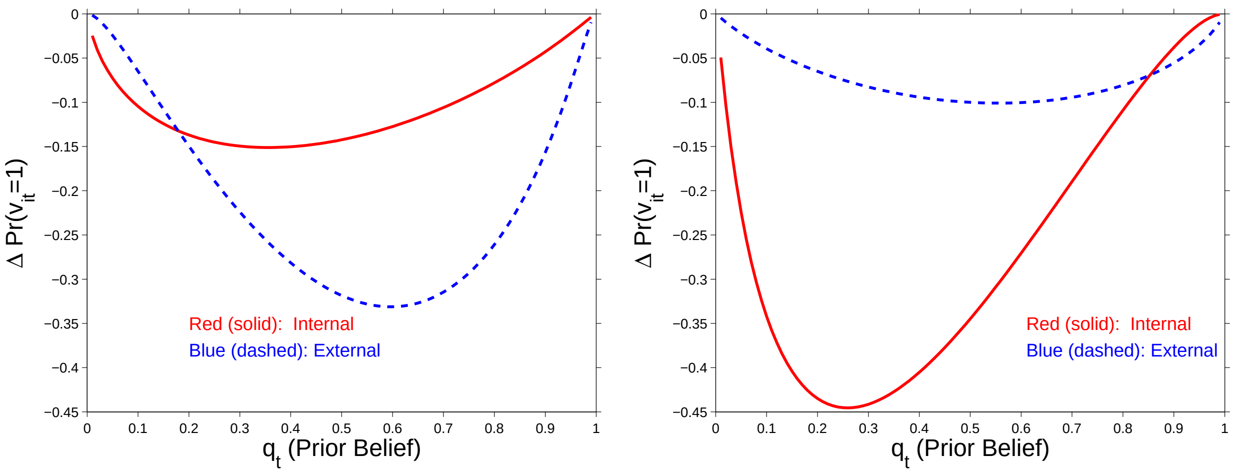

(a) Experience Effect over q_t **(b)** Experience Effect over q_t if switch expertise

Figure 5: Probability new and experienced members vote high over q_t

Notes: The left figure shows the difference between experienced and new members (the experience effect) for each type of appointment; this comes from the difference between the two lines in figure 4. The right figure draws an alternative experience effect plot but under the assumption that external members have the estimated expertise of internal members, and vice versa.

magnitude of the fall in the probability of choosing high rates. Thus the level of expertise is crucial for determining the quantitative impact of signalling.

Table 8: Reduced form evidence on the impact of experience by internal and external

	(1) $\widehat{v}_{it}$
D(Experienced)	-0.15***
	[0.000]
D(Internal) x D(Experienced)	0.14***
	[0.001]
Constant	0.92***
	[0.000]
R-squared	0.707
Model	Panel LPM
Member effects?	YES
Time effects?	YES
Sample?	06/97-03/09
Obs?	1246

Notes: This regression replicates the regression reported in table 1 but presents OLS estimates of equation (17) with standard errors clustered by member. They show that, controlling for member and time fixed effects, members with experience (defined as 18 meetings experience in column (1), 12 meetings in column (2), and 24 meetings in column (3)) vote for lower interest rates than new members.

These observations underscore the difficulty of using reduced-form evidence to ascertain whether members of a committee are signalling. Imagine that we wished to test our model's implications that all preference types were subject to signalling incentives, and we chose to look at internals compared with externals. One approach would be to perform an OLS regression of equation (17) but augmented with an interaction between D(Internal) and D(Experienced); table 8 reports the estimation results. Since this regression is effectively identifying the coefficients off the unconditional probabilities we plot in figure 4, it is not surprising that these results suggest little signalling by internal members (and would lead us to reject the predictions of our theoretical model). But, in fact, that conclusion would not be correct and this finding is simply a manifestation of the fact that it is not possible to use reduced form policy decisions to identify the effect of signalling when members might differ in terms of their expertise; a structural analysis as used in this paper is necessary.

6 Learning as an explanation of our results

Of course, other explanations for the empirical findings that we uncover might be possible. For example, macroeconomists have long considered the possibility that policy makers will learn over time. In this situation one could think of Θ as a belief about some unknown structural parameter of the macroeconomy. Imagine that members all start period t with Θ_t, an imperfect estimate of some true parameter Θ_{TRUE}. Suppose further that after the implementation of r_t, some new data (such as actual inflation data) arrived that allowed members to update their beliefs to a new value Θ_{t+1}. In this case, one might well imagine that our estimates of Θ could change over time.

But any learning story would need to be consistent with all three patterns in the data. First, in terms of the prediction on the average rise in Θ, the arrival of iid shocks should generate no shift on average since beliefs formed using Bayes' Rule are martingales. To get this effect, one would need to assume that the initial average belief on Θ were biased downwards, or that Θ itself were a convex function of some uncertain parameter about which learning occurred, in which case beliefs would form a submartingale.

Second, to capture the prediction that doves' average Θ rises more than hawks', one would need assumptions along the lines of doves' having initial beliefs that were more biased than hawks' or having Θ functions that were "more convex" in the variable about which learning occurred.

Finally, a fairly robust finding in the literature on learning (Kalai and Lehrer 1994) is that, as rational agents are exposed to increasing amounts of information on a parameter, their beliefs tend to converge even if they begin with non-common priors. This directly contradicts our estimates that show the difference between Θ for hawkish and dovish members *increases* over time.

While we do not wish to claim learning plays no role in monetary policy making, we do think the recent literature has underplayed the idea that signalling actively influences policy decisions. As our paper shows that in some settings this role may be of potentially first-order importance, we feel future work should take signalling more seriously as a generator of dynamic behavior.

7 Conclusion

This paper argues that one should take seriously the idea that independent monetary policy makers care about their reputation for hawkishness. It does so by building a new model of signalling in monetary policy that generates a constellation of predictions on dynamic behavior that correspond very well to observed voting on an important independent monetary policy committee (the MPC). The introduction of private assessments of the inflationary state means the prior and expertise crucially affect how the signalling incentive maps into policy choices, making a structural approach necessary for uncovering the evolution of the underlying bias.

More generally, the results can be taken as a starting point for re-assessing the role of reputation in the design of institutions for monetary policymaking. Since the emergence of a consensus in the 1990s that independent central banks should set policy to establish credibility, the literature has, in our view, tended to de-emphasize the relevance of reputation. By showing that it is important for understanding the behavior of independent experts, our work opens the door to new research on optimal committee structures. For example, a preliminary conclusion is that there is a trade-off between rotating members relatively frequently (which maintains uncertainty on preferences and so the strength of the signalling incentive) and human capital accumulation (which our estimates for hawkish members show might be important in some cases). We leave a thorough exploration of this and other design ideas for future work, but it highlights that a theoretical framework with both bias and expertise can yield important and non-obvious insights.

In terms of the contemporary policy debate, our results are useful for clarifying how one should expect policy makers to behave. Consider again the suggestion by Flanders that a hawkish German central banker would have come to the ECB job wishing to show off his inner dove, and that this behavior may have been more desirable than a dovish Italian (Draghi) coming in trying to show off his toughness on inflation. Our analysis suggests that while the latter is true, the former is not likely to have happened. That is, so long as both members are concerned about keeping inflation expectations contained, both types will enter the job and adopt a more hawkish bias than their later selves. While the dove might initially be further from the voting rule he would use without signalling, it is worth remembering that the inherent differences between the types mean that the hawk will be tougher on inflation than the dove.

References

Appelbaum, B. (2013): "Possible Fed Successor Has Admirers and Foes," http://www.nytimes.com/2013/04/25/business/janet-l-yellen-possible-fed-successor-has-admirers-and-foes.html?emc=eta1&_r=0, 25th April 2013, last accessed 25 April 2013.

Austen-Smith, D., and J. S. Banks (1996): "Information Aggregation, Rationality, and the Condorcet Jury Theorem," *American Political Science Review*, 90(1), 34–45.

Backus, D., and J. Driffill (1985a): "Inflation and Reputation," *American Economic Review*, 75(3), 530–38.

——— (1985b): "Rational Expectations and Policy Credibility Following a Change in Regime," *Review of Economic Studies*, 52(2), 211–21.

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 04 April 2013.

Bank of England Act (1998): http://www.bankofengland.co.uk/about/legislation/1998act.pdf, last accessed 04 April 2013.

Barro, R. J. (1986): "Reputation in a model of monetary policy with incomplete information," *Journal of Monetary Economics*, 17(1), 3–20.

Barro, R. J., and D. B. Gordon (1983a): "A Positive Theory of Monetary Policy in a Natural Rate Model," *Journal of Political Economy*, 91(4), 589–610.

Barro, R. J., and D. B. Gordon (1983b): "Rules, discretion and reputation in a model of monetary policy," *Journal of Monetary Economics*, 12(1), 101–121.

Besley, T. (2006): *Principled Agents?: The Political Economy of Good Government*, The Lindahl Lectures. OUP Oxford.

Bianchi, F., and L. Melosi (2012): "Constrained Discretion and Central Bank Transparency," Mimeograph.

Bullard, J., and S. Eusepi (2005): "Did the great inflation occur despite policymaker commitment to a Taylor rule?," *Review of Economic Dynamics*, 8, 324–359.

Cho, I. K., N. Williams, and T. Sargent (2002): "Escaping Nash inflation," *The Review of Economic Studies*, 69, 1–40.

Cottle, D. (2012): "So, Mr. Carney, Hawk or Dove?," http://blogs.wsj.com/source/2012/11/27/so-mr-carney-hawk-or-dove/, 27 November 2012, last accessed 04 April 2013.

Cukierman, A., and A. H. Meltzer (1986): "A Theory of Ambiguity, Credibility, and Inflation under Discretion and Asymmetric Information," *Econometrica*, 54(5), 1099–1128.

Drazen, A. (2001): "The Political Business Cycle After 25 Years," in *NBER Macroeconomics Annual 2001*. MIT Press, Cambridge, Massachusetts.

Duggan, J., and C. Martinelli (2001): "A Bayesian Model of Voting in Juries," *Games and Economic Behavior*, 37(2), 259–294.

Eijffinger, S. C. W., R. J. Mahieu, and L. Raes (2013): "Inferring Hawks and Doves from Voting Records," Discussion Paper DP9418, CEPR.

Erceg, C. J., and A. T. Levin (2003): "Imperfect credibility and inflation persistence," *Journal of Monetary Economics*, 50(4), 915–944.

Evans, G. W., and S. Honkapohja (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press.

Faust, J., and L. E. O. Svensson (2001): "Transparency and Credibility: Monetary Policy with Unobservable Goals," *International Economic Review*, 42(2), 369–97.

Flanders, S. (2011): "ECB: Clearing the way for an Italian hawk?," `http://www.bbc.co.uk/blogs/thereporters/stephanieflanders/2011/05/ecb_clearing_the_way_for_an_it.html`, 5th May 2011, last accessed 04 April 2013.

Gerlach-Kristen, P. (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. McMahon (2013): "Estimating Bayesian Decision Problems with Heterogeneous Priors," manuscript.

Hansen, S., M. McMahon, and C. Velasco (2012): "How Experts Decide: Preferences or Private Assessments on a Monetary Policy Committee?," manuscript.

Hix, S., B. Hoyland, and N. Vivyan (2010): "From Doves to Hawks: A Spatial Analysis of Voting in the Monetary Policy Committee of the Bank of England," *European Journal of Political Research*, 49(6), 731–758.

Iaryczower, M., and M. Shum (2012): "The Value of Information in the Court: Get It Right, Keep It Tight," *American Economic Review*, 102(1), 202–37.

Kalai, E., and E. Lehrer (1994): "Weak and strong merging of opinions," *Journal of Mathematical Economics*, 23(1), 73–86.

King, R. G., Y. K. Lu, and E. S. Pastén (2008): "Managing Expectations," *Journal of Money, Credit and Banking*, 40(8), 1625–1666.

Kreps, D. M., and R. Wilson (1982): "Reputation and imperfect information," *Journal of Economic Theory*, 27(2), 253–279.

Lambert, R. (2006): "Who sets our interest rates?," `http://www.bankofengland.co.uk/publications/other/monetary/inside_mpc.pdf`, last accessed 04 April 2013.

Levy, G. (2007): "Decision Making in Committees: Transparency, Reputation, and Voting Rules," *American Economic Review*, 97(1), 150–168.

Milgrom, P., and J. Roberts (1982): "Predation, reputation, and entry deterrence," *Journal of Economic Theory*, 27(2), 280–312.

Orphanides, A., and J. Williams (2005): "The decline of activist stabilization policy: natural rate misperceptions, learning and expectations," *Journal of Economic Dynamics and Control*, 29, 1927–1950.

Ottaviani, M., and P. Sorensen (2000): "Herd Behavior and Investment: Comment," *American Economic Review*, 90(3), 695–704.

Primiceri, G. E. (2006): "Why Inflation Rose and Fell: Policy-Makers' Beliefs and U. S. Postwar Stabilization Policy," *The Quarterly Journal of Economics*, 121(3), 867–901.

Sargent, T. J. (1999): *The Conquest of American Inflation*. Princeton University Press.

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.

——— (2009): "Is Transparency about Central Bank Plans Desirable?," *Journal of the European Economic Association*, 7, 831–857.

The Economist (2013): "Ben. And then?," `http://www.economist.com/news/finance-and-economics/21574521-who-will-run-fed-next-year-ben-and-then`, 30th March 2013, last accessed 04 April 2013.

Vickers, J. (1986): "Signalling in a Model of Monetary Policy with Incomplete Information," *Oxford Economic Papers*, 38(3), 443–55.

Visser, B., and O. H. Swank (2007): "On Committees of Experts," *The Quarterly Journal of Economics*, 122(1), 337–372.

A Proofs

A.1 Proof of Proposition 1

Proof. Consider first period 2 The expected utility of choosing r_t is

$$\begin{aligned} U(r_t,\theta) &= \widehat{\omega}_2 \left\{ M\left[\pi(r_2,1)\right] + \chi(\theta)\left[\pi(r_2,1) - \pi_2^E\right]\right\} + \\ &\quad (1-\widehat{\omega}_2)\left\{ M\left[\pi(r_2,0)\right] + \chi(\theta)\left[\pi(r_2,0) - \pi_2^E\right]\right\}. \end{aligned}$$

Thus C chooses $r_t = 1$ whenever $U(1,\theta) > U(0,\theta)$, which yields the bias in the proposition. Since θ is observable, r_1 does not impact on π_2^E, and so the decision rule is identical in the first period as in the second. ■

A.2 Proof of Proposition 2

Proof. The voting rule in (5) translates into the equilibrium strategy $\rho_2^* = 1$ whenever

$$s_2 \geq \mathfrak{L}_2^{-1}\left\{ \ln\left[\frac{1-q_2}{q_2} B_2(\theta)\right]\right\} = s_2^*\left(\theta, q_2\right) \tag{A.1}$$

where $s_2^*\left(\theta, q_2\right)$ is strictly increasing in θ. For the rest of the proof we take $\widehat{\rho}_2 = \rho_2^*$.

We first show that assumption A1 implies

$$[\pi(0,\omega_t) - \pi(1,\omega_t)] - \delta[\overline{\pi}_2^E(0,\omega_t) - \overline{\pi}_2^E(1,\omega_t)] > 0 \ \forall \omega_t.$$

π_2^E can be written

$$\begin{aligned} \pi_2^E &= \Pr\left[\omega_2 = 1\right]\left\{ \begin{array}{c} \mathbb{E}_\theta[\,1 - G_2\left(s_2^* \mid \omega_2 = 1\right) \mid p_1\,]\pi(1,1) + \\ \mathbb{E}_\theta[\,G_2\left(s_2^* \mid \omega_2 = 1\right) \mid p_1\,]\pi(0,1) \end{array}\right\} + \\ &\quad \Pr\left[\omega_2 = 0\right]\left\{ \begin{array}{c} \mathbb{E}_\theta[\,1 - G_2\left(s_2^* \mid \omega_2 = 0\right) \mid p_1\,]\pi(1,0) + \\ \mathbb{E}_\theta[\,G_2\left(s_2^* \mid \omega_2 = 0\right) \mid p_1\,]\pi(0,0) \end{array}\right\}. \end{aligned}$$

So $\overline{\pi}_2^E(0,\omega_t) - \overline{\pi}_2^E(1,\omega_t)$ is strictly smaller than

$$\begin{aligned} &\mathbb{E}\{\, q_2[\pi(0,1) - \pi(1,1)] + (1-q_2)[\pi(0,0) - \pi(1,0)] \mid q_1 \,\} = \\ &\overline{q}[\pi(0,1) - \pi(1,1)] + (1-\overline{q})[\pi(0,0) - \pi(1,0)]. \end{aligned}$$

Assumption A1 guarantees the following two inequalities are satisfied:

$$\begin{aligned} \pi(0,0) - \pi(1,0) &> \delta\left\{\overline{q}[\pi(0,1) - \pi(1,1)] + (1-\overline{q})[\pi(0,0) - \pi(1,0)]\right\} \\ \pi(0,1) - \pi(1,1) &> \delta\left\{\overline{q}[\pi(0,1) - \pi(1,1)] + (1-\overline{q})[\pi(0,0) - \pi(1,0)]\right\}. \end{aligned}$$

Now, π_2^E can be more compactly expressed as

$$\pi_2^E \;=\; \mathbb{E}_{\omega_2}\{\, \pi(1,\omega_2) + [\pi(0,\omega_2) - \pi(1,\omega_2)]\mathbb{E}_\theta[\, G_2\,(s_2^* \mid \omega_2) \mid p_1\,]\,\} \qquad \text{(A.2)}$$

The voting rule in (6) implies that C follows a cutoff rule in which he chooses $v_1 = 1$ whenever $s_1 \geq s_1^*(\theta, q_1)$, which is strictly increasing in θ. Let $\widehat{s}_1(\theta, q_1)$ denote the public's belief on this threshold for type θ and $\widehat{\mathbf{s}}_1(q_1)$ the vector of such beliefs. Then

$$p_1(1) = \frac{\{1 - G_1\,[\widehat{s}_1(\theta, q_1) \mid \omega_1]\}\, p_0^\theta}{\sum\limits_\theta \{1 - G_1\,[\widehat{s}_1(\theta, q_1) \mid \omega_1]\}\, p_0^\theta};\ \text{and}\ p_1(0) = \frac{G_1\,[\widehat{s}_1(\theta, q_1) \mid \omega_1]\, p_0^\theta}{\sum\limits_\theta G_1\,[\widehat{s}_1(\theta, q_1) \mid \omega_1]\, p_0^\theta}.$$

$\frac{p_1(0)}{p_1(1)}$ is proportional to $\frac{G_1[\widehat{s}_1(\theta,q_1)|\omega_1]}{1-G_1[\widehat{s}_1(\theta,q_1)|\omega_1]}$ which is increasing in θ, implying $p_1(0)$ first order stochastically dominates $p_1(1)$. Since $G_2\,(s_2^* \mid \omega_2)$ is strictly increasing in θ

$$D\,[\widehat{\mathbf{s}}_1(q_1), \omega_1, q_2] \;\equiv\; \mathbb{E}_{\omega_2}\left\{\, [\pi(0,\omega_2) - \pi(1,\omega_2)] \begin{pmatrix} \mathbb{E}_\theta[\, G_2\,(s_2^* \mid \omega_2) \mid p_1(0)\,] - \\ \mathbb{E}_\theta[\, G_2\,(s_2^* \mid \omega_2) \mid p_1(1)\,] \end{pmatrix} \right\} > 0$$

for all ω_1, q_2 and that, a fortiori, $\overline{D}\,[\widehat{\mathbf{s}}_1(q_1), \omega_1] = \mathbb{E}_{q_2}\{\, D\,[\widehat{\mathbf{s}}_1(q_1), \omega_1, q_2]\,\} > 0$. From these arguments, C in the first period votes high whenever

$$s_1^*(\theta, q_1) = \mathfrak{L}_1^{-1}\left\{ \ln\left[\frac{1 - q_1}{q_1} B_1(\theta, q_1) \right] \right\}$$

where

$$B_1(\theta, q_1) = \frac{\mu(0) + \chi(\theta)\,\{[\pi(0,0) - \pi(1,0)] - \delta\overline{D}\,[\widehat{\mathbf{s}}_1(q_1), 0]\}}{\mu(1) - \chi(\theta)\,\{[\pi(0,1) - \pi(1,1)] - \delta\overline{D}\,[\widehat{\mathbf{s}}_1(q_1), 1]\}}.$$

It remains to be shown that an equilibrium exists. Let $\mathbf{s}_1^*(q_1)$ denote the vector of equilibrium thresholds. We need to show that the K equations

$$\mathbf{s}_1^*(q_1) = \mathfrak{L}_1^{-1}\left\{ \ln\left[\frac{1 - q_1}{q_1} \frac{\mu(0) + \chi(\theta)\,\{[\pi(0,0) - \pi(1,0)] - \delta\overline{D}\,[\mathbf{s}_1^*(q_1), 0]\}}{\mu(1) - \chi(\theta)\,\{[\pi(0,1) - \pi(1,1)] - \delta\overline{D}\,[\mathbf{s}_1^*(q_1), 1]\}} \right] \right\} \qquad \text{(A.3)}$$

have a solution. We have shown above that $\delta\overline{D}\,[\widehat{\mathbf{s}}_1(q_1), \omega_1] \in [0, 1]$. This implies that the RHS of (A.3) maps the compact set

$$E = \{\, \mathbf{s}_1^*(q_1) \mid s_H \geq s_1^*(\theta_H, q_1) \geq \ldots \geq s_1^*(\theta_L, q_1) \geq s_L \,\}$$

where

$$s_H = \mathfrak{L}_1^{-1}\left\{ \ln\left[\frac{1 - q_1}{q_1} \frac{\mu(0) + \chi(\theta_H)\,(1 - \delta)}{\mu(1) - \chi(\theta_H)\,(1 - \delta)} \right] \right\} \text{ and } s_L = \mathfrak{L}_1^{-1}\left\{ \ln\left[\frac{1 - q_1}{q_1} \frac{\mu(0)}{\mu(1)} \right] \right\}$$

into itself. Because the RHS of (A.3) is also continuous in $\mathbf{s}_1^*(q_1)$ we can conclude by Brouwer's Fixed Point Theorem that a solution exists.

We now turn to proving the second claim. Let $I_{\omega_t} = \pi(0, \omega_t) - \pi(1, \omega_t)$.

$$B_2(\theta) - B_1(\theta) = \frac{\mu(0) + \chi(\theta) I_0}{\mu(1) - \chi(\theta) I_1} - \frac{\mu(0) + \chi(\theta) \left(I_0 - \delta \overline{D}\left[\cdot, 0\right]\right)}{\mu(1) - \chi(\theta) \left(I_1 - \delta \overline{D}\left[\cdot, 1\right]\right)}$$

Differentiating with respect to θ gives

$$\frac{\chi'(\theta) I_0[\mu(1) - \chi(\theta) I_1] + \chi'(\theta) I_1[\mu(0) + \chi(\theta) I_0]}{[\mu(1) - \chi(\theta) I_1]^2} - \frac{\begin{pmatrix} \chi'(\theta)\left(I_0 - \delta \overline{D}\left[\cdot, 0\right]\right)[\mu(1) - \chi(\theta)\left(I_1 - \delta \overline{D}\left[\cdot, 1\right]\right)] + \\ \chi'(\theta)\left(I_1 - \delta \overline{D}\left[\cdot, 1\right]\right)[\mu(0) + \chi(\theta)\left(I_0 - \delta \overline{D}\left[\cdot, 0\right]\right)] \end{pmatrix}}{[\mu(1) - \chi(\theta)\left(I_1 - \delta \overline{D}\left[\cdot, 1\right]\right)]^2}$$

which is positive whenever

$$\frac{I_0 \mu(1) + I_1 \mu(0)}{[\mu(1) - \chi(\theta) I_1]^2} - \frac{\left(I_0 - \delta \overline{D}\left[\cdot, 0\right]\right)\mu(1) + \left(I_1 - \delta \overline{D}\left[\cdot, 1\right]\right)\mu(0)}{[\mu(1) - \chi(\theta)\left(I_1 - \delta \overline{D}\left[\cdot, 1\right]\right)]^2} > 0.$$

Since we showed above that $\delta \overline{D}\left[\cdot, \omega_1\right] > 0$, this condition clearly holds. ■

A.3 Proof of Proposition 3

Proof. Members $i \in Y(t-1)$ in period $t = \tau(i) + 1$ have a dominant strategy and choose $v_{it} = 1$ whenever

$$s_{it} \geq s_O^*\left(\theta_i, q_t\right) \equiv \mathfrak{L}_O^{-1}\left\{\ln\left[\frac{1 - q_t}{q_t} B_O(\theta_i)\right]\right\}. \tag{A.4}$$

Therefore set $\widehat{\rho}_O = \rho_O^*$.

Since young members' equilibrium strategies are independent of their beliefs on old members' types, we can without loss of generality consider $\rho_Y^*[\theta_i, s_{it}, x(t)]$. By the same arguments as those in proposition 2, assumption A1 guarantees that in any equilibrium young voters choose $v_{it} = 1$ whenever $s_{it} \geq s_Y^*[\theta_i, q_t, x(t)]$ where s_Y^* is increasing in θ_i. Let $\widehat{\mathbf{s}}_Y(q_t, x(t))$ denote the public's belief on these cutoffs. Now let

$$D\left[\widehat{\mathbf{s}}_Y(q_t, x(t)), \widehat{\mathbf{s}}_Y(q_{t+1}, x(t+1)), \omega_t, \mathbf{v}_{-it}^Y, q_{t+1}\right] \equiv$$
$$\mathbb{E}_{\omega_{t+1}}\begin{pmatrix} [\pi(0, \omega_{t+1}) - \pi(1, \omega_{t+1})] \times \\ \mathbb{E}_{\boldsymbol{\theta}_{-i}}\left\{\Pr\left[\mathrm{PIV}_{i,t+1} \mid \widehat{\mathbf{s}}_Y(q_{t+1}, x(t+1)), \mathbf{s}_O^*(q_{t+1}), \boldsymbol{\theta}_{-i}, \omega_{t+1}\right] \mid \mathbf{p}_{-i,t+1}\left(\mathbf{v}_{-it}^Y\right)\right\} \times \\ \{\mathbb{E}_{\theta_i}[G_O\left(s_O^* \mid \omega_{t+1}\right) \mid p_{i,t+1}(0)] - \mathbb{E}_{\theta_i}[G_O\left(s_O^* \mid \omega_{t+1}\right) \mid p_{i,t+1}(1)]\} \end{pmatrix}$$

be the change in period $t+1$ inflation expectations that voter $i \in Y(t)$ induces by choosing

$v_{it} = 0$ over $v_{it} = 1$ given the public's period $t+1$ information set. $\text{PIV}_{i,t+1}$ denotes the set of events in which voter i is pivotal in period $t+1$—this corresponds to the union of all events in which exactly $\frac{N-1}{2}$ others vote 1.

We obtain

$$p_{i,t+1}(1) = \frac{\{1 - G_Y[\widehat{s}_Y(\theta_i, q_t, x(t)) \mid \omega_t]\} p_0^\theta}{\sum_{\theta_i} \{1 - G_Y[\widehat{s}_Y(\theta_i, q_t, x(t)) \mid \omega_t]\} p_0^\theta}; \text{ and}$$

$$p_{i,t+1}(0) = \frac{G_Y[\widehat{s}_Y(\theta_i, q_t, x(t)) \mid \omega_t] p_0^\theta}{\sum_{\theta_i} G_Y[\widehat{s}_Y(\theta_i, q_t, x(t)) \mid \omega_t] p_0^\theta}.$$

By the same arguments as those in proposition 2, $p_{i,t+1}(0)$ first order stochastically dominates $p_{i,t+1}(1)$ so that $\mathbb{E}_{\theta_i}[G_O(s_O^* \mid \omega_{t+1}) \mid p_{i,t+1}(0)] > \mathbb{E}_{\theta_i}[G_O(s_O^* \mid \omega_{t+1}) \mid p_{i,t+1}(1)]$ for $\omega_{t+1} = 0, 1$. Thus we can conclude that the equilibrium value of D must be positive.

Now let

$$\begin{aligned}&\overline{D}[\widehat{\mathbf{s}}_Y(q_t, x(t)), \widehat{\mathbf{s}}_Y(Q, x(t+1)), \omega_t] = \\ &\mathbb{E}_{\mathbf{v}_{-it}^Y, q_{t+1}}\left\{ D\left[\widehat{\mathbf{s}}_Y(q_t, x(t)), \widehat{\mathbf{s}}_Y(q_{t+1}, x(t+1)), \omega_t, \mathbf{v}_{-it}^Y, q_{t+1}\right] \,\middle|\, \mathbf{s}_Y^*(q_t, x(t)) \right\}.\end{aligned}$$

be young voter i's expected impact on period $t+1$ inflation expectations from voting low rather than high in period t (where $q_{t+1} \sim Q$). Here the expectation over young colleagues' votes is computed using their equilibrium strategies in period t. An equilibrium of the model is a solution to the $2K$ equations

$$s_Y^*[q_t, 0] = \mathfrak{L}_Y^{-1}\left\{\ln\left[\frac{1-q_t}{q_t} B_Y(q_t, 0)\right]\right\} \text{ and } s_Y^*[q_t, 1] = \mathfrak{L}_Y^{-1}\left\{\ln\left[\frac{1-q_t}{q_t} B_Y(q_t, 1)\right]\right\}$$

where

$$B_Y(q_t, 0) = \frac{\mu(0) + \chi(\theta)\left\{[\pi(0,0) - \pi(1,0)] - \delta\overline{D}[\mathbf{s}_Y^*(q_t, 0), \mathbf{s}_Y^*(Q, 1), 0]\right\}}{\mu(1) - \chi(\theta)\left\{[\pi(0,1) - \pi(1,1)] - \delta\overline{D}[\mathbf{s}_Y^*(q_t, 0), \mathbf{s}_Y^*(Q, 1), 1]\right\}}$$

and

$$B_Y(q_t, 1) = \frac{\mu(0) + \chi(\theta)\left\{[\pi(0,0) - \pi(1,0)] - \delta\overline{D}[\mathbf{s}_Y^*(q_t, 1), \mathbf{s}_Y^*(Q, 0), 0]\right\}}{\mu(1) - \chi(\theta)\left\{[\pi(0,1) - \pi(1,1)] - \delta\overline{D}[\mathbf{s}_Y^*(q_t, 1), \mathbf{s}_Y^*(Q, 0), 1]\right\}}.$$

Following similar reasoning as in proposition 3, one can obtain existence. Point 3. also follows identical reasoning as in the single banker case. ■

A.4 Proof of Proposition 4

Proof. First note that the probability a young member votes high depends on θ_i only via $B_Y(\theta_i, \cdot)$. Through differentiating $B_Y(\theta_i, \cdot)$ with respect to θ_i, one obtains that $B_Y(\theta_i, \cdot)$ is increasing if and only if

$$\begin{aligned} V &= \mu(1)\left\{[\pi(0,0)-\pi(1,0)] - \frac{\delta\overline{D}(0,x(t))}{\Pr\left[\,\mathrm{PIV}_{it}\mid 0\,\right]}\right\} + \\ &\quad \mu(0)\left\{[\pi(0,1)-\pi(1,1)] - \frac{\delta\overline{D}(1,x(t))}{\Pr\left[\,\mathrm{PIV}_{it}\mid 1\,\right]}\right\} > 0 \end{aligned}$$

Since this is independent of θ_i, one can conclude that $B_Y(\theta_i, \cdot)$ is either strictly increasing in θ_i (if $V > 0$), independent of θ_i (if $V = 0$), or decreasing in θ_i (if $V < 0$). Suppose it is independent of θ_i. Then the probability young member i votes $v_{i\tau(i)} = 1$ is independent of θ_i, meaning period $t+1$ inflation expectations are independent of $v_{i\tau(i)}$. But then $\delta\overline{D}(\omega_t, x(t)) = 0$ and $V > 0$, a contradiction.

Now suppose $B_Y(\theta_i, \cdot)$ is decreasing in θ_i. Then as explained in arguments in propositions 2 and 3, $p_{i,\tau(i)+1}(1)$ first order stochastically dominates $p_{i,\tau(i)+1}(0)$ and, because $\Pr\left[\,r_{\tau(i)+1} = 0\,\right]$ is strictly increasing in θ_i, inflation expectations in period $\tau(i)+1$ computed under $p_{i,\tau(i)+1}(1)$ must be higher than those computed under $p_{i,\tau(i)+1}(0)$. But then $\delta\overline{D}(\omega_t, x(t)) < 0$, meaning that $V > 0$, a contradiction.

Hence we conclude that $B_Y(\theta_i, \cdot)$ must be increasing in θ_i in any responsive equilibrium in which the hypothesis is satisfied. This implies that $\delta\overline{D}(\omega_t, x(t)) > 0$, from which the three properties derive as in proposition 2. ■

A.5 Derivation of Example 2

Equations (12) and (13) write as (replacing i with Y or O to indicate young and old)

$$\begin{gathered} \frac{\widehat{\omega}_{Ot}}{1-\widehat{\omega}_{Ot}} \times \\ \frac{(1-\alpha)\sum_{\theta_Y}\Pr\left[\,s_Y \geq s_Y^*(\theta_Y,\theta_O)\mid \omega=1\,\right]p_0^{\theta_Y} + \alpha\sum_{\theta_i}\Pr\left[\,s_Y < s_Y^*(\theta_Y,\theta_O)\mid \omega=1\,\right]p_0^{\theta_Y}}{(1-\alpha)\sum_{\theta_Y}\Pr\left[\,s_Y \geq s_Y^*(\theta_Y,\theta_O)\mid \omega=0\,\right]p_0^{\theta_Y} + \alpha\sum_{\theta_i}\Pr\left[\,s_Y < s_Y^*(\theta_Y,\theta_O)\mid \omega=0\,\right]p_0^{\theta_Y}} \geq \\ \frac{\mu(0)+\chi(\theta_O)[\pi(0,0)-\pi(1,0)]}{\mu(1)-\chi(\theta_O)[\pi(0,1)-\pi(1,1)]}. \end{gathered} \tag{A.5}$$

and

$$\frac{\widehat{\omega}_{Yt}}{1-\widehat{\omega}_{Yt}}\frac{(1-\alpha)\Pr[\,s_O \geq s_O^*(\theta_O) \mid \omega=1\,]+\alpha\Pr[\,s_O < s_O^*(\theta_O) \mid \omega=1\,]}{(1-\alpha)\Pr[\,s_O \geq s_O^*(\theta_O) \mid \omega=0\,]+\alpha\Pr[\,s_O < s_O^*(\theta_O) \mid \omega=0\,]} \geq$$

$$\frac{\mu(0)+\chi(\theta_Y)\left\{[\pi(0,0)-\pi(1,0)]-\frac{\delta\overline{D}(0)}{(1-\alpha)\Pr\left[s_O\geq s_O^*(\theta_O)\,\middle|\,\omega=0\right]+\alpha\Pr\left[s_O<s_O^*(\theta_O)\,\middle|\,\omega=0\right]}\right\}}{\mu(1)-\chi(\theta_Y)\left\{[\pi(0,1)-\pi(1,1)]-\frac{\delta\overline{D}(1)}{(1-\alpha)\Pr\left[s_O\geq s_O^*(\theta_O)\,\middle|\,\omega=1\right]+\alpha\Pr\left[s_O<s_O^*(\theta_O)\,\middle|\,\omega=1\right]}\right\}}. \quad \text{(A.6)}$$

where $s_Y^*(\theta_Y,\theta_O)$ and $s_O^*(\theta_O)$ are the cutoffs that young and old members' signals must reach to induce them to vote high. Plugging in $\alpha=\frac{1}{2}$ yields

$$\frac{\widehat{\omega}_{Ot}}{1-\widehat{\omega}_{Ot}} \geq \frac{\mu(0)+\chi(\theta_O)[\pi(0,0)-\pi(1,0)]}{\mu(1)-\chi(\theta_O)[\pi(0,1)-\pi(1,1)]}. \quad \text{(A.7)}$$

and

$$\frac{\widehat{\omega}_{Yt}}{1-\widehat{\omega}_{Yt}} \geq \frac{\mu(0)+\chi(\theta_Y)\left\{[\pi(0,0)-\pi(1,0)]-2\delta\overline{D}\,(0)\right\}}{\mu(1)-\chi(\theta_Y)\left\{[\pi(0,1)-\pi(1,1)]-2\delta\overline{D}\,(1)\right\}}. \quad \text{(A.8)}$$

In this case, $s_O^*(\theta_O)$ is immediately solvable from (A.7) and is strictly decreasing in θ_O, while s_Y^* does not depend on θ_O. So the hypothesis of proposition 4 holds.

B The responsiveness of member thresholds to more hawkish experienced members

The equilibrium dynamics under strategic voting rely on the fact that the probability that the period t committee chooses the high rate is increasing in old members' hawkishness (see proposition 4). While it may be theoretically possible that a committee composed of more hawkish types could opt for low rates more often, the important question is how the committee reacts in our empirical setting when we add an additional old hawk controlling for the number of old members. In table 9 we report the results of a regression of the estimated cutoffs, $\widehat{s}_{it}$, on the number of experienced hawks as measured by Eijffinger, Mahieu, and Raes (2013) and described in section 4.3.[34]

In column (1) of table 9, we use all members' cutoffs in meetings with at least one old member as the left hand side variable, and find the average member adopts a lower cutoff (and so votes high more often) when there is an additional hawk. This is consistent with the hypothesis of proposition 4. In column (2), we use just the cutoffs of new members in these meetings, and find no significant effect of old members' hawkishness on their voting behavior, although the direction of the point estimate is consistent with new members' voting high more often.

Table 9: Effect of hawkish experienced members on voting cutoffs

	(1) $\widehat{s}_{it}$	(2) $\widehat{s}_{it}$
Number of Old Hawks	-0.012*	-0.026
	[0.077]	[0.192]
Common Prior	-0.27***	-0.20***
	[0.000]	[0.000]
D(Experienced)	0.26***	
	[0.000]	
Total Old Members	0.025***	0.038***
	[0.000]	[0.003]
Constant	0.32***	0.33***
	[0.000]	[0.000]
R-squared	0.398	0.188
Number of mem_num	26	21
Model	Panel	Panel
Member effects?	YES	YES
Time effects?	No	No
Obs?	1061	293

Notes: This regression presents estimates of the impact on member cutoffs of an additional experienced member. They show that, controlling for member fixed effects, the average member on a committee with more experienced hawks votes low more often.

[34] We also explored the use of internal-external as a measure of hawkishness and the results are similar.