

Original citation:

Hansen, Stephen and McMahon, Michael F.. (2016) First impressions matter : signalling as a source of policy dynamics. The Review of Economic Studies, 83 (4). pp. 1645-1672.

Permanent WRAP URL:

<http://wrap.warwick.ac.uk/85622>

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions. Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher's statement:

This is a pre-copyedited, author-produced PDF of an article accepted for publication in The Review of Economic Studies following peer review. The version of record Hansen, Stephen and McMahon, Michael F.. (2016) First impressions matter : signalling as a source of policy dynamics. The Review of Economic Studies, 83 (4). pp. 1645-1672. is available online at: <http://dx.doi.org/10.1093/restud/rdw007>

A note on versions:

The version presented here may differ from the published version or, version of record, if you wish to cite this item you are advised to consult the publisher's version. Please see the 'permanent WRAP URL' above for details on accessing the published version and note that access may require a subscription.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk

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

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

Michael McMahon
University of Warwick
CEPR, CAGE, CEP, CfM, and CAMA
m.mcmahon@warwick.ac.uk

July 5, 2015

Abstract

We provide the first direct empirical support for the importance of signalling in monetary policy by testing two key predictions from a novel structural model. First, all policymaker types should become less tough on inflation over time, and, second, types that weigh output more should have a more pronounced shift. Voting data from the Bank of England's Monetary Policy Committee strongly supports both predictions. Counterfactual results indicate signalling has a substantial impact on interest rates over the business cycle, and improves the committee designer's welfare. Implications for committee design include allowing regular member turnover and transparency regarding publishing individual votes.

Keywords: Signalling, Monetary Policy, Committees

JEL Codes: E52, E58, D78

*We are very grateful to the editor, Philipp Kircher, and three anonymous referees for their insightful suggestions. We would also 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, the Money-Macro Finance Group annual conference 2012, Barcelona GSE, Imperial College London, OFCE at Sciences Po and University of Essex.

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 aim to manage inflation expectations (King, Lu, and Pastén 2008). Examples include the establishment of independent central banks and recent forward guidance policies.

A particularly important time for controlling inflation expectations is during the replacement of senior policymakers, such as the Chairperson or Governor. In these transitions, the public naturally speculates about the preferences of incoming policymakers, about which there is typically substantially more uncertainty than their predecessors'. For example, Cottle (2012) contemplated how inflation averse incoming Bank of England Governor Mark Carney would be, and even before Chairman Bernanke announced his stepping down, *The Economist* (2013) and Appelbaum (2013) speculated about who, and how inflation averse, his replacement would be.

An idea with a long history in monetary policy is that new policymakers can take advantage of the uncertainty surrounding their preferences to signal inflation aversion to the public, thereby anchoring future expectations.¹ A central prediction of this literature is that policymakers act aggressively against inflation early in their careers but then adopt looser policies over time. This idea is also important in practice with Flanders (2011), for example, suggesting that Mario Draghi would go out of his way to rebut national stereotypes by being especially tough on inflation immediately following his appointment, and refrain from adopting unconventional, expansionary policies used by other central banks.²

Despite the importance of the signalling idea, there is surprisingly little evidence in the literature assessing whether it affects actual policy choices.³ The primary contribution of this paper is to provide evidence from a structural model that strongly supports the relevance of signalling for individual voting behavior on the Bank of England's Monetary Policy Committee (MPC), and to quantify its impact on policy outcomes and welfare.

To begin, we construct a model that captures the important details of the monetary

¹Papers in this tradition include 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).

²She 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."

³There is 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) and Erceg and Levin (2003) are examples. In signalling models, central bankers react strategically to the public's learning to affect inflation expectations, and their policy stance is determined endogenously.

policy decision-making process. Members serve for a finite time on an infinitely-lived committee that sets policy. In each period, the economy is in one of two states, an inflationary state which requires a high rate, or a non-inflationary state which requires a low rate. Members receive common information about the economy, representing economic data and staff forecasts, which they combine with heterogeneous private assessments to form beliefs on the state. Preferences are standard; all policymakers dislike deviations of inflation from a target rate, but differ in the weights they put on the output gap. We refer to these weights as preference types. A type who puts more (less) weight on output is inflation tolerant (averse).⁴

The model gives rise to cutoff voting rules in which all members vote for high rates when there is sufficient evidence the state is high. The level of evidence a member requires is called her cutoff. Members who place a higher weight on the output gap use a higher cutoff and so, *ceteris paribus*, vote low more often. Equilibrium multiplicity in models of signalling in monetary policy is generally problematic (Persson and Tabellini 2000, p. 407), and even more so for empirical work. Importantly, our model yields unique equilibrium dynamics.⁵ Since more inflation-averse members vote high more often in all periods, whenever the public observes a high vote by rookie members they reduce future inflation expectations. This gives every member, regardless of type, an additional incentive to vote for high rates when a rookie. This incentive is absent when a veteran, and therefore yields the first empirical prediction that *all policymakers' cutoffs increase with tenure*.

While the overall message of declining toughness on inflation with tenure is present in the existing literature, the model also yields a second, more novel prediction. The extent to which members care about reducing future inflation expectations relates to the weight they put on the output gap.⁶ So, the incentive to use a higher cutoff early in one's tenure is increasing in type, yielding the difference-in-differences prediction that *the magnitude of the increase in a banker's cutoff is increasing in her type*. These two

⁴Each of the N committee members draws one of K types, and all members behave strategically, whereas in the previous theoretical literature on signalling in monetary policy committees (Sibert 2003) there are two members with one of two types, a mechanistic hawk that always votes for zero inflation and a strategic dove that tries to build a reputation.

⁵In pure signalling models (Backus and Driffill 1985a,b, Vickers 1986, Sibert 2002), there are at least as many levels of inflation as there are preference types, so policy choices can perfectly reveal the banker's types. These models typically admit both separating and pooling equilibria. Papers such as Cukierman and Meltzer (1986), Faust and Svensson (2001), and Sibert (2009) address this problem by introducing exogenous noise in the mapping between the policymakers' intended policy and the actual policy. By contrast, in our model, all types have a positive probability of choosing high and low rates in every period, but different types differ in these probabilities.

⁶Papers 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 policymakers 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.

predictions together form a counterpoint to the view of Flanders (2011) quoted above. While it is true that members who are more inherently inflation tolerant (“Italians”) will signal more than the inherently inflation averse (“Germans”), we predict that both types will act tougher against inflation early in their tenures than later to establish a reputation for toughness.

The empirical analysis begins by describing reduced-form evidence broadly consistent with the model. In line with the first prediction, members are significantly less likely to vote for high interest rates after serving an initial period on the MPC. This is in line with a similar finding in Hansen and McMahon (2008), and also with Gerlach-Kristen (2003), who notes that MPC members seem less likely to dissent with time based on vote tabulations. Moreover, using two measures of preference type, we also find evidence that more inflation-tolerant members have a higher fall in the probability of voting high. The only other paper that examines voting dynamics on the MPC is Berk, Bierut, and Meade (2010).⁷ They estimate interest rate rules for internal and external members on the MPC and find that internal members become relatively more hawkish. This finding is consistent with our difference-in-differences prediction given that internals tend to be less inflation tolerant than external members.

Although none of these papers claim to provide a test of signalling, we argue that their evidence cannot be interpreted as such because they do not control for changes in the parameters that determine beliefs on the state—the common prior and the precision of private assessments—in order to pin down the evolution of the cutoff—the parameter that signalling affects. We instead use a structural estimator that separately estimates all these theoretical objects, allowing us to verify the two main dynamic predictions of the model. First, the average veteran member on the MPC uses a significantly higher cutoff than the average rookie member. Second, we show that this increase in the cutoff is significantly higher for more inflation-tolerant MPC members. To the best of our knowledge, these results represent the first empirical validation of models of signalling in monetary policy.

Our final contribution is to examine the extent to which rookies’ signalling actually translates into the committee’s choosing higher rates, and how this impacts welfare, by comparing simulated interest rate paths chosen by a signalling committee with those chosen by a non-signalling committee. We find that signalling induces at least a 25bps higher interest rate after a five-year business cycle with probability 0.93, while over the same period the likelihood that the signalling path is at least 50bps higher is 0.74. For a

⁷There are a number of other papers using MPC voting data that focus on static voting differences, in particular those between internal and external members (Gerlach-Kristen 2004, Bhattacharjee and Holly 2005, Spencer 2006, Besley, Meads, and Surico 2008, Harris and Spencer 2008, Hix, Hoyland, and Vivyan 2010, Hansen, McMahon, and Velasco 2014). All are reduced form with the exception of Hansen, McMahon, and Velasco (2014), who use a different estimator than the one we use in the main text.

social planner who weighs all mistakes equally, the overall effect of signalling is to improve welfare; while over five years the signalling committee would, on average, make 4.0 errors, a non-signalling committee would make 6.3. These results have a number of implications for committee design with the most obvious two being that there should be a reasonable amount of committee member turnover (rather than having members serve for too long) and that the committee environment should be transparent given that the public's being able to directly observe individual votes facilitates signalling.

The fundamental take-away 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” central banker often invoked to discuss the behavior of central banks. Instead, our estimates imply that a model in which preferences are heterogeneous and establishing credibility is crucial fits voting data very well in an institutional context admired for its independence.

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. Outside monetary policy, the mechanism we identify should be relevant in any context in which the policymaker'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.

The paper proceeds as follows. Section 2 lays out a theory of signalling. Section 3 explains the institutional setting of the MPC and the data we use. Section 4 presents reduced-form evidence on dynamics, before we turn to a structural analysis in section 5. Section 6 examines the robustness of the empirical results. Section 7 then uses the estimated structural parameters to quantify the impact of signalling on policy choices and welfare. Section 8 concludes.

2 A Model of Reputation and Policy Dynamics

A monetary policy committee of N (odd) policymakers chooses either a higher or lower interest rate $r_t \in \{0, 1\}$ in each of $t = 1, \dots, \infty$ periods.⁸ Members serve for two periods with staggered appointments, and in odd (even) periods there are N_1 (N_0) new members. A member in his first period is a *rookie* and in his second a *veteran*. Each member i in period t chooses a vote $v_{it} \in \{0, 1\}$, with r_t chosen by majority rule.

In each period, an inflationary state variable $\omega_t \in \{0, 1\}$ is realized, with $\omega_t = 1$ indicating a more inflationary state. The state is drawn according to $\Pr[\omega_t = 1] = q_t \in (0, 1)$, where q_t is independent and identically distributed across periods with mean \bar{q} . Period t inflation $\pi_t \in \mathbb{R}$ depends both on r_t and ω_t through the relationship $\pi_t = \pi(r_t, \omega_t)$. Higher interest rates lower inflation for a given state of the economy, i.e. $\pi(1, \omega_t) < \pi(0, \omega_t) \forall \omega_t$, and a more inflationary state leads to higher inflation, i.e. $\pi(r_t, 1) > \pi(r_t, 0) \forall r_t$.

All central bankers discount the future by a factor δ , and member preferences within period t are given by

$$u_{it} = u(r_t, \omega_t, \theta_i) = -l(\pi_t) + \theta_i(\pi_t - \pi_t^E), \quad (1)$$

where l captures losses from deviations of inflation from its target level, and $\pi_t - \pi_t^E$ captures the gains from a positive output gap which is expressed using surprise inflation (as follows from an expectations-augmented Phillips Curve). A banker's type θ_i is the weight put on the output gap and hence a measure of inflation tolerance; $\theta_i = 0$ is an "inflation nutter" who only cares about inflation deviations (King 1997). θ_i is drawn independently from a prior distribution p_0 defined on Θ , a finite set with K non-negative elements and maximum value $\bar{\theta}$. Below we consider the cases in which θ_i is both public and private information. Period t inflation expectations π_t^E are the beliefs the public holds at time t on current inflation. In line with the literature, we consider the public to be a single, representative player and solve endogenously for π_t^E .

To ensure members do not trivially vote for either high ($v_{it} = 1$) or low ($v_{it} = 0$) interest rates in every period, members need to have a strict preference for matching the decision r_t to the state ω_t . Let $\mu_{\omega_t} \equiv l[\pi(1 - \omega_t, \omega_t)] - l[\pi(\omega_t, \omega_t)]$ be the net welfare loss of inflation from failing to match the decision to the state. The following assumptions are sufficient for ensuring this property:

A1 $\mu_{\omega_t} > 0$ for $\omega_t = 0, 1$

⁸The interest rates that form the binary agenda can change from meeting to meeting, and the higher of the two under consideration refers to the $r_t = 1$ policy choice. For example, in a meeting in which policymakers decide between not changing interest rates and raising by 25 basis points, $r_t = 1$ corresponds to raising rates. On other hand, in a meeting in which policymakers decide between cutting interest rates by 25 basis points and not changing interest rates, $r_t = 1$ corresponds to not changing.

$$\mathbf{A2} \quad \mu_1 - \bar{\theta} [\pi(0, 1) - \pi(1, 1)] > 0$$

$$\mathbf{A3} \quad [\pi(0, 1) - \pi(1, 1)] / [\pi(0, 0) - \pi(1, 0)] \in \left(\frac{1-\bar{q}}{1/\delta-\bar{q}}, \frac{1/\delta-(1-\bar{q})}{\bar{q}} \right)$$

A1 implies $r_t = \omega_t$ minimizes $l(\pi_t)$. A2 implies even the most inflation-tolerant type prefers $r_t = 1$ in state $\omega_t = 1$. A3 ensures that any effect of r_t on period $t+1$ expectations is never strong enough to overturn the motivation to match the policy to the state.⁹

Before voting, each member observes the (conditionally independent) private signal $s_{it} \sim \mathcal{N}(\omega_t, \sigma_{it}^2)$, where $\sigma_{it} = \sigma_R$ ($\sigma_{it} = \sigma_V$) if i is a rookie (veteran) in period t .¹⁰ We refer to σ as *expertise* since it measures the ability to perceive economic conditions. Allowing σ to evolve with time is important since, for example, members might become better able to observe economic conditions with experience. After observing s_{it} , member i uses Bayes' rule to update his belief on the state to $\hat{\omega}_{it} \equiv \Pr[\omega_t = 1 \mid s_{it}]$. The timing of each period t subgame is the following.

1. Rookie members join the committee, and each one draws a preference type θ_i .
2. q_t is observed by all players.
3. ω_t is drawn but not observed by any player.
4. The public forms π_t^E .
5. All committee members observe their private signals.
6. Committee members simultaneously choose votes, which all players observe.
7. π_t is observed by all players (which implicitly reveals ω_t).

2.1 Equilibrium concept

Let p_{it} be the *reputation* of member i at time t . This is the belief that i 's colleagues and the public hold on his preference type θ_i . All rookies' reputation is the prior p_0 , while veterans' reputations depend on their votes as rookies. Denote by \mathbf{p}_t the vector of reputations associated with members of the period t committee, and let N_t^R be the number of period t rookies. We limit attention to symmetric Markov strategies with respect to $I_t = (\mathbf{p}_t, q_t, N_t^R)$. Rookies use strategy $\bar{v}_R : \theta_i, s_{it}, I_t \rightarrow \{0, 1\}$ and veterans use strategy $\bar{v}_V : \theta_i, s_{it}, I_t \rightarrow \{0, 1\}$.¹¹

⁹This is proved in appendix A.2.

¹⁰The key property of normality is the monotone likelihood ratio. The theoretical results are robust to conditional signal distributions of the form $g(s \mid \omega_t)$ defined over the support (\underline{s}, \bar{s}) so long as $\frac{g(s|1)}{g(s|0)}$ is continuous and strictly increasing in s , with unbounded limits as s approaches \underline{s} and \bar{s} .

¹¹The restriction to pure strategies is innocuous. If one admitted mixed strategies, committee members would be indifferent between voting high and low only for a measure-0 set of signals.

The theory literature considers both sincere and strategic behavior in static voting games with private information (Austen-Smith and Banks 1996). Voters behave sincerely when they maximize their expected utility conditional on their signals. This is in contrast to strategic voting, in which voters additionally condition on actually being pivotal. Here we apply a concept similar to sincere voting by assuming that equilibrium strategies maximize expected utility given signals and given how the public forms period t inflation expectations π_t^E , which we denote by the mapping $\bar{\pi}^E : I_t \rightarrow \mathbb{R}$. In words, π_t^E is formed by taking the period t committee member's reputations \mathbf{p}_t , the distribution over states q_t , and the committee composition N_t^R , and using them to forecast π_t . In equilibrium this forecast must be consistent with voting strategies, as the following formalizes. (As we discuss below, we consider the case of strategic voting in an online appendix and the main messages of the paper are unchanged.)

Definition 1 *A Sincere Markov Perfect Bayesian Equilibrium is a pair of strategies $(\bar{v}_R^*, \bar{v}_V^*)$ and expectation formation rule $\bar{\pi}^{E*}$ where:*

1. $\bar{v}_V^* = \arg \max_{v_{it} \in \{0,1\}} \mathbb{E}[u(v_{it}, \omega_t, \theta_i) \mid s_{it}, I_t] \text{ given } \bar{\pi}^{E*}.$
2. $\bar{v}_R^* = \arg \max_{v_{it} \in \{0,1\}} \mathbb{E}[u(v_{it}, \omega_t, \theta_i) + \delta u(\bar{v}_V^*, \omega_{t+1}, \theta_i) \mid s_{it}, I_t] \text{ given } \bar{\pi}^{E*}.$
3. $\bar{\pi}^{E*}$ satisfies Bayes' rule given \bar{v}_R^* and \bar{v}_V^* .

When θ_i is private information, equilibrium inflation expectations are formed in the following way. First, the public uses the period $t - 1$ observed votes of rookies to update its beliefs on their preference types, which go into \mathbf{p}_t . Second, given I_t and the equilibrium voting strategies, the public computes the probability of $r_t = 1$. Third, given the relationship $\pi_t = \pi(r_t, \omega_t)$, the public computes the expected value of π_t . Hence there is a link between observed votes and π_t^E , which gives rise to the signalling channel.

2.2 Equilibrium policy dynamics

To begin the discussion of the solution, we follow Duggan and Martinelli (2001) and define *cutoff voting rules* for member i in meeting t as those in which he chooses $v_{it} = 1$ if and only if $\frac{\hat{\omega}_{it}}{1 - \hat{\omega}_{it}} \geq C_{it} \in (0, \infty)$, where C_{it} is i 's period t *cutoff*. The cutoff measures the amount of evidence for the inflationary state ($\omega_t = 1$) that is required in order to vote high. Loosely speaking, the cutoff captures the inclination to choose lower interest rates in a given meeting and so it can be related to the common classification of policymakers into “hawks” and “doves”. A voter with $C_{it} = 1$ is neutral in the sense of voting in the direction of whichever state is most likely. Members with $C_{it} < 1$ are then hawkish as they require weaker evidence of an inflationary state to vote for high rates, and, by similar logic, members with $C_{it} > 1$ are dovish.

Lemma 1 *Cutoff voting rules imply choosing $v_{it} = 1$ if and only if*

$$s_{it} \geq \frac{1}{2} - \sigma_{it}^2 \left[\ln \left(\frac{q_t}{1 - q_t} \right) - \ln(C_{it}) \right] \equiv s_{it}^*(C_{it}, \sigma_{it}, q_t). \quad (2)$$

In other words, a high vote is chosen if the signal reaches a critical threshold s_{it}^* . A feature of cutoff voting rules is that observing v_{it} never allows the public to perfectly infer θ_i because even though different types have different probabilities of choosing high rates, votes also depend on the realization of private signals. Nonetheless, the empirical consequences of a change in the cutoff is clear: when C_{it} increases, the probability of voting high decreases since the threshold the signal must reach increases.

2.2.1 Policymaking with no signalling

We first consider a game in which θ_i is public information. This is a useful benchmark because its equilibrium is equivalent to one of a game in which inflation expectations do not react to observed votes, and the motive to use one's vote to signal one's type is absent (NS = no signalling). We will use this benchmark in the counterfactual analysis in section 7.

Proposition 1 *With known θ_i , there is a unique equilibrium in which rookies and veterans use cutoff voting rules with cutoff*

$$C^{NS}(\theta_i) = \frac{\mu_0 + \theta_i [\pi(0, 0) - \pi(1, 0)]}{\mu_1 - \theta_i [\pi(0, 1) - \pi(1, 1)]}. \quad (3)$$

The voting rule in (3) derives from members' expected utility maximization problem treating π_t^E as fixed. The numerator is the cost of a wrong decision in state 0 (the utility that is lost from choosing $v_{it} = 1$ 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 $v_{it} = 1$ implies lower output, and similarly a lower benefit in state 1. This means that he adopts a higher cutoff, and votes high less often. Of course, in equilibrium inflation expectations π_t^E are not fixed but consistent with bankers' strategies given by (3). When the public knows that a member has a higher θ_i , it increases its inflation expectations in line with the higher cutoff.¹²

¹²One can show that bankers with $\theta_i > 0$ would like to commit to using a voting rule with a lower cutoff than the one used in equilibrium (proof available on request), and internalize the effect of the cutoff on π_t^E . In this sense, types that put positive weight on the output gap over-inflate the economy as in the Barro and Gordon model.

2.2.2 Policymaking with signalling

We now turn to the more realistic case in which θ_i is private information. Inflation expectations now react to observed votes, and the signalling channel is present.

Proposition 2 *With unknown θ_i , in all equilibria rookies and veterans use cutoff voting rules with cutoffs $C_R^S(\theta_i, q_t, N_t^R)$ and $C_V^S(\theta_i)$, respectively, where*

1. $C_V^S(\theta_i) = C^{NS}(\theta_i)$.
2. $C_R^S(\theta_i, q_t, N_t^R)$ is strictly increasing in θ_i .
3. $C_V^S(\theta_i) - C_R^S(\theta_i, q_t, N_t^R) > 0 \forall \theta_i, q_t, N_t^R$.
4. $C_V^S(\theta_i) - C_R^S(\theta_i, q_t, N_t^R)$ is strictly increasing in $\theta_i \forall q_t, N_t^R$.

While equilibria may not be unique, all share the same qualitative features. First, veterans use the same cutoff as when their preferences are observed. Second, equilibrium cutoffs for both rookies and veterans increase in θ_i . Third, all types use a higher cutoff in the second period than the first. Finally, the difference in cutoffs between veterans and rookies is increasing in θ_i . Figure 1 illustrates the predicted dynamics in the cutoff for three different preference types ordered by inflation tolerance (θ_3 places the highest weight on output deviations and so is most tolerant of inflation; θ_1 places the lowest weight on output deviations and is least tolerant) with and without signalling incentives.




Figure 1: Predicted Dynamics in Equilibrium Cutoffs for Types $\theta_1 < \theta_2 < \theta_3$.

Notes: This figure shows how the equilibrium cutoff evolves for three representative preference types θ_1 , θ_2 , and θ_3 between their first and second periods on the committee. All types use the same cutoff when veterans whether or not types are public or private information. When types are known, each adopts the same cutoff when a rookie as when a veteran. When types are unknown, each adopts a lower cutoff when a rookie, with the extent of the difference with the veteran cutoff increasing in type.

To understand this result, it is first useful to understand in more detail how committee members' reputations \mathbf{p}_t are formed. All period t rookies share the same reputation p_0 , while every period t veteran has one of two reputations depending on whether he voted high or low as a rookie in period $t - 1$. (Recall that rookies use symmetric strategies). How high and low votes change a voter's reputation depends on q_{t-1} and N_{t-1}^R , which directly enter equilibrium strategies, and also ω_{t-1} , which is observed prior to period t and determines the probability of voting high given a threshold $s_{i,t-1}^*$. So, without loss of generality, one can express $\mathbf{p}_t = \mathbf{p}(V_{t-1}^R, q_{t-1}, N_{t-1}^R, \omega_{t-1})$ where V_{t-1}^R is the number of rookies who vote high in period $t - 1$.

In terms of equilibrium behavior, veterans in period t have the same utility maximization problem as when preferences are public information; they treat π_t^E as fixed, and their votes do not change future expectations. So in equilibrium they use the same cutoff as without signalling. On the other hand, rookies in period t must consider the effect of their vote on π_{t+1}^E , which we define as

$$\Delta(q_t, N_t^R \mid \omega_t) \equiv \mathbb{E} \left\{ \pi^{E*} \left[\mathbf{p}(V_{-i,t}^R, q_t, N_t^R, \omega_t), q_{t+1}, N_{t+1}^R \right] \right\} - \mathbb{E} \left\{ \pi^{E*} \left[\mathbf{p}(V_{-i,t}^R + 1, q_t, N_t^R, \omega_t), q_{t+1}, N_{t+1}^R \right] \right\}, \quad (4)$$

where $V_{-i,t}^R$ is the number of high rookie votes excluding rookie i 's. In words, $\Delta(q_t, N_t^R \mid \omega_t)$ is the expected change in future inflation expectations from voting low rather than high given current economic conditions.¹³

Rookies then maximize utility by using the cutoff

$$C_R^S(\theta_i, q_t, N_t^R) = \frac{\mu_0 + \theta_i \{ [\pi(0, 0) - \pi(1, 0)] - \delta \Delta(q_t, N_t^R \mid 0) \}}{\mu_1 - \theta_i \{ [\pi(0, 1) - \pi(1, 1)] - \delta \Delta(q_t, N_t^R \mid 1) \}}. \quad (5)$$

The key step in the proof is to show that the equilibrium sign of Δ must always be positive. The logic follows three steps. First, since veterans' cutoff is strictly increasing in θ_i , the public increases inflation expectations when it believes more inflation-tolerant veterans set policy. Second, given that (5) is increasing in θ_i (for any values of Δ , in or out of equilibrium), in every equilibrium it must be the case that more inflation-tolerant rookies are more likely to choose $v_{it} = 0$. Finally, combining these two observations means that when the public observes $v_{it} = 0$, it associates the rookie with a more inflation-tolerant type, which leads it to increase π_{t+1}^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 reputation for inflation aversion that anchors future inflation expectations at a lower level.

¹³The expectation is taken with respect to q_{t+1} , as well as the types of other rookies and the signals they draw, which, along with knowledge of the cutoffs they use, determines $V_{-i,t}^R$.

The model predicts that the increase in cutoff is greater the more inflation tolerant the type because the signalling incentive in (5) is directly linked to the weight the banker places on the future output gap. More inflation-tolerant types have a higher weight on this, and so, intuitively, care more about convincing the market they are inflation averse than inherently inflation-averse types do. This gives rise to the difference-in-differences on the evolution of the cutoff.

2.3 Discussion of modelling assumptions

The model abstracts from details such as the transmission from the interest rate to inflation in order to explicitly model the individual decisions of MPC members, and to set up a tight link between the theory and subsequent empirical exercise. While embedding it into a richer macroeconomic setting would be an important next step, the model as it stands provides a baseline framework through which to assess the importance of signalling in monetary policy. We now address some of its specific simplifications.

Three potentially restrictive assumptions are that we assume (i) members vote sincerely, (ii) members are ex-ante identical, and (iii) members have no career concerns. We shall address each of these limitations but we defer their discussion until section 6, after the empirical analysis of behavior, so that we can empirically explore the implications of alternative modelling assumptions.

Another feature of the model is that members serve for two periods. In the first section of the online appendix we explore a T -period signalling model. While this introduces additional effects of changes in reputation, the solution features the same qualitative dynamics between any two periods t and $t + 1$ as in the two period model: all types use a higher cutoff over time, with a more pronounced evolution for higher types.

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). A primary advantage of this is that cutoffs are independent of current inflation expectations. Otherwise, changes in future inflation expectations from current votes would also change future voting behavior via changes in cutoffs.

Next, as is typically the case in the literature, policymakers in our model wish to anchor inflation expectations against a tendency to rise. In an alternative model, such as a liquidity trap model, policymakers may wish to raise inflation expectations. While such cases do not apply to the sample period we consider, in such circumstances a similar signalling channel would predict a tendency towards higher cutoffs early in tenure.

Finally, we assume that members care only about inflation expectations when they serve on the committee, whereas one might imagine utility would depend on the present discounted value of all future periods' inflation expectations. Our formulation is without

loss of generality in the sense that we only write utility as a function of the variables that members' votes affect in equilibrium. For example, while period t veterans might care about π_{t+1}^E , its equilibrium value is independent of their period t votes since their types do not impact r_{t+1} . In section 6 we explore the alternative assumption that members care about the committee's reputation independently of their own.

3 The Bank of England MPC as a Test of the Model

The theoretical model has two key empirical predictions corresponding to the third and fourth properties, respectively, described in proposition 2. The first (empirical prediction 1) is that members increase their cutoffs as they transition from rookies to veterans. The second (empirical prediction 2) is that more inflation-tolerant members—those a higher θ —increase their cutoffs by more than more inflation-averse members. We now proceed to test the model using voting data from the Bank of England's MPC. This section provides the necessary background details and data sources, while the next two present reduced-form and structural evidence in support of the signalling model.

3.1 Institutional Background

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.¹⁴ The standard term of office during our sample, apart from for the Governor and the two Deputy Governors, is three years (36 meetings); Governor-level positions carry five year terms. The Act that created the MPC allows for the reappointment of all members, internal and external. The average served by members in our sample is 46 meetings; only former Governor Mervyn King is present in all 142 of our sample meetings.

Within the monthly meeting, and after a general discussion of economic and business trends, each member summarizes his or her view to the rest of the MPC and suggests which vote they favor (although, as Lambert (2006) notes, they can, if they wish, wait to hear others' views before committing to a vote). This process begins with the Deputy Governor for monetary policy and concludes with the Governor, but the order for the

¹⁴According 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.

others is not fixed. To formally conclude the meeting, the Governor proposes an interest rate decision that he believes will command a majority. Each member then chooses whether to agree with the Governor’s proposal, or dissent and state an alternative interest rate. 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.

The MPC’s 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. If inflation is greater than 3% or less than 1%, the Governor of the Bank of England must write an open letter to the Chancellor explaining why. Upside and downside misses are treated equally seriously.¹⁵

3.2 Data

The paper analyzes the MPC voting record through March 2009, when the interest rate reached its effective zero lower bound and a period of quantitative easing (QE) began; from then, the main MPC decision concerned how many assets to purchase. We use each regular MPC meeting in this period but drop from the dataset the (unanimous) emergency meeting held after 9/11. This sample yields a total of 142 meetings, and 1246 individual votes (all these data are available from Bank of England (2010b)).

The four main empirical counterparts to the theoretical objects from the model required for estimation are the binary voting agenda, proxies for the prior q_t , and, for each member, classification as rookie or veteran, and measures of the preference type θ . The first two of these are described in detail in Hansen, McMahon, and Velasco (2014) (and its online appendix), so here we provide only a brief description. Descriptive statistics of the final two covariates are in the appendix (table B.1).

In meetings with two observed votes, we set $v_{it} = 1$ if member i votes for the higher one. In meetings with one observed vote, we use a Reuters survey of City of London financial institutions conducted prior to each MPC meeting in which respondents are asked to submit a probability distribution over outcomes. We take the second alternative in the meeting to be the one that receives the highest average weight in the survey (the observed vote is always one of the two outcomes given highest average weight).¹⁶

¹⁵ “The remit is not to achieve the lowest possible inflation rate. Inflation below the target of 2% is judged to be just as bad as inflation above the target. The inflation target is therefore symmetrical” Bank of England (2015).

¹⁶ In the seven meetings with three unique votes, we identify the two most likely choices to make up

The first proxy for the prior—denoted q_t^R —uses the Reuter’s survey and is the average probability placed on the higher of the two outcomes on the agenda over the average probability of observing either outcome on the agenda. The second—denoted q_t^M —relies on backing out the probabilities of different outcomes using the price distributions of LIBOR futures described in Bank of England (2011). It is again the average probability placed on the higher agenda outcome over the probabilities of either agenda outcome.¹⁷

To classify members as rookies and veterans we use the indicator variable

$$D(\text{Vet})_{it} = \begin{cases} 0 & \text{if member } i \text{ has served in 18 or fewer meetings} \\ 1 & \text{if member } i \text{ has served in more than 18 meetings} \end{cases} \quad (6)$$

where 18 meetings represents half the median term length. Below we explore the effect of using different splits by committee tenure.

We classify member types using two different but related measures (and explore robustness to other measures in section 6). While θ_i is by definition private information for member i , individual voting histories provide information that allows the public and econometrician to update their beliefs—in every period members with a higher θ_i are more likely to choose $v_{it} = 0$. Our approach assigns members to an inflation-averse or inflation-tolerant group based on the percentage of votes cast that are high.¹⁸ In the model all preference types adopt relatively similar cutoffs when rookies, which suggests veteran votes are more informative for inferring preference types. Therefore, our first measure $D(\theta^{PCTExp})_i$ splits the sample of members evenly based on the percentage of total votes cast that are high while a veteran. $D(\theta^{PCTExp})_i = 1$ ($= 0$) denotes a more inflation-averse (inflation-tolerant) member. However, as five of the 27 members in the sample never served as veterans, this measure has some missing values. Moreover, rookie votes might actually be more informative than veteran votes if expertise increases with time, although our empirical results below show this is not the case. We therefore define our second measure $D(\theta^{PCT})$ as the percentage of all votes cast by member i that are high. Table B.1 shows member classification for both proxies.

the binary agenda, then pool the votes for the third option into the nearest option that is part of the binary agenda. See online appendix E of Hansen, McMahon, and Velasco (2014) for full details.

¹⁷For five meetings the Reuter’s survey and price distributions are not available, so when using estimators that rely on q_t^R and q_t^M the sample size is 1201 rather than 1246.

¹⁸The difference-in-differences prediction depends on the true type θ not the public’s belief on θ at time t . We can therefore use the entire voting history to proxy θ , and apply it to all voting periods even though the public has not yet seen all the votes that go into constructing the proxy.

4 Reduced-Form Evidence

We begin the analysis of voting dynamics with reduced-form evidence by estimating a linear probability model on the binary vote variable:

$$v_{it} = \mu_i + \delta_t + \lambda_0 D(\text{Vet})_{it} + \epsilon_{it}, \quad (7)$$

where μ_i is a member fixed effect, δ_t is a time fixed effect that captures the average vote in meeting t , and ϵ_{it} is an error term. The results are reported in column (1) of table 1, and show that veteran members are significantly less likely to vote for high rates than rookies, with a gap of 9.2 percentage points. The inclusion of member and time effects indicates this change is driven by some systematic shift at the individual level rather than, for example, changing economic conditions.

Column (2) of table 1 reports estimates of a regression that replaces $D(\text{Vet})_{it}$ with separate indicators for having served on the MPC more than twelve, 24, and 36 months. The main finding is that after members have served in their first twelve months, they are significantly less likely to choose high rates, but there is no significant additional change in voting probabilities as members serve longer and longer.¹⁹ Finally, column (3) reports results from replacing $D(\text{Vet})_{it}$ with an indicator for whether a member is within nine months of the end of a given (renewable) term of office. Being near the end of a term has no effect on the probability the average member chooses high rates, indicating that the effect of being a veteran is not driven by approaching the end of one's current term.

The fact that the probability of voting high declines after some time on the MPC is consistent with members' cutoffs increasing over time (empirical prediction 1). To test empirical prediction 2, that there is a difference-in-differences between high and low types in the change in the cutoff, we estimate the following relationship:

$$v_{it} = \mu_i + \delta_t + \lambda_0 D(\text{Vet})_{it} + \lambda_1 D(\theta^X)_i \times D(\text{Vet})_{it} + \epsilon_{it} \quad (8)$$

where $D(\theta^X)_i$ is either of the two measures of type discussed above. Columns (4) and (5) present estimates for both measures. In each case, there is a significant difference-in-differences (the inflation tolerant have a higher fall in the probability of choosing high rates), but the fall in the probability for more inflation-averse members is not significant.

¹⁹This is consistent with the T -period model we solve in the online appendix in which the signalling effect is strongest in initial periods before declining steadily over time.

Table 1: Reduced-form evidence on the impact of experience

Main Regressors	(1) v_{it}	(2) v_{it}	(3) v_{it}	(4) v_{it}	(5) v_{it}
D(Vet)	-0.092*** [0.001]			-0.031 [0.351]	-0.038 [0.314]
D(12+ meetings)		-0.077** [0.010]			
D(24+ meetings)		-0.012 [0.704]			
D(36+ meetings)		-0.0057 [0.880]			
D(Term End)			-0.0081 [0.728]		
D(θ^{PCTExp}) x D(Vet)				-0.13*** [0.003]	
D(θ^{PCT}) x D(Vet)					-0.080* [0.055]
Constant	0.98*** [0.000]	0.97*** [0.000]	1.01*** [0.000]	1.22*** [0.000]	0.95*** [0.000]
R-squared	0.704	0.703	0.701	0.707	0.705
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	06/97-03/09	06/97-03/09
Obs?	1246	1246	1246	1204	1246


Notes: This regression presents OLS estimates of equation (7) with standard errors clustered by member. The dependent variable, v_{it} , is our measure of whether member i votes for the high interest rate in period t .

5 Structural Estimation

One reason against interpreting the reduced-form evidence as a strict test of the signalling model (and why one should not jump to any positive or negative conclusions based on the results) is that it concerns changes in voting probabilities, whereas the theoretical predictions concern changes in the cutoff C , only one component of the voting probability.

For example, if σ is constant over a member's tenure, a rise in C unambiguously lowers the probability of voting high. However, if σ changes over time, one cannot interpret a fall in the probability as incontrovertible evidence of C 's rising. To demonstrate this, we examine theoretically how voting probabilities change when members gain expertise with tenure without changing their cutoff. Suppose the average committee member is a hawk, with $C = 0.4$, and suppose that members gain expertise over time, with $\sigma = 0.7$ for rookies and $\sigma = 0.3$ for veterans. Figure 2a plots the theoretical difference in the probability of voting high between veterans and rookies as a function of q_t . For most values, the difference is negative, and the average value is around -0.09 , the measured

probability change in column (1) of table 1. The intuition for why greater expertise generates a decline in the probability of voting high is that, since all members want to get the correct decision ($r_t = \omega_t$), a more expert (veteran) member who sees the true state of the economy more clearly will tend to vote with the state and is less affected by his cutoff. That the cutoff has less (more) influence on the decision of a member with a low (high) σ is shown formally in equation (2) where σ_{it} directly interacts with C_{it} . So, in this example, the rookie will be more inclined to vote for high rates than the veteran even though the cutoff does not change.



(a) Change in $\Pr[v_{it} = 1]$ with increasing expertise. (b) Change in $\Pr[v_{it} = 1]$ for a hawk and a dove.

Figure 2: Change in probability of voting high when different characteristics change.

Notes: These figures show the theoretical difference between veteran and rookie members in the probability of voting for the high interest rate as a function of q_t . In the left panel, the expertise of the individual is increased (σ falls from 0.7 to 0.3). The right panel shows how the probability of voting high will be affected if an initial hawk (dove) becomes more dovish.

Even if σ were constant over time, using reduced-form estimates to test the difference-in-differences prediction would be difficult. The mapping from changes in C to changes in the probability of voting high is non-linear and depends on how extreme preferences are. For example, if the change in C moves a hawk from being very hawkish to slightly hawkish, it will have a bigger effect than a shift that moves a dove from being close to neutral, to slightly dovish. Figure 2b illustrates this problem. We assume that a hawk and dove share $\sigma = 0.5$ when rookies and veterans, and both have changes in cutoffs consistent with the signalling model: the hawk's cutoff rises from 0.15 to 0.55 and the dove's from 1.1 to 4.0. Nevertheless, the fall in the probability of voting high is equivalent for both of them further highlighting the difficulty of interpreting the reduced form regressions.

5.1 Structural estimation methodology

These problems with interpreting reduced-form evidence necessitate the use of structural estimation in order to test the signalling model. The advantage of the structural estimation is that we can estimate the evolution of both members' cutoff and expertise. Another advantage of structural estimation is that it allows us to conduct counterfactual analysis.

By lemma 1, policymakers vote high whenever their private signals exceed $s_{it}^*(C_{it}, \sigma_{it}, q_t)$ which is given by equation (2). This means that we can write the probability member i in meeting t votes high as $\kappa_{0it} \equiv 1 - \Phi\left(\frac{s_{it}^*}{\sigma_{it}}\right)$ when $\omega_t = 0$ and $\kappa_{1it} \equiv 1 - \Phi\left(\frac{s_{it}^* - 1}{\sigma_{it}}\right)$ when $\omega_t = 1$, where Φ is the standard normal cdf. Using these conditional probabilities, the probability of observing the voting data is

$$\prod_t \left[q_t \prod_i (\kappa_{1it})^{v_{it}} (1 - \kappa_{1it})^{1-v_{it}} + (1 - q_t) \prod_i (\kappa_{0it})^{v_{it}} (1 - \kappa_{0it})^{1-v_{it}} \right]. \quad (9)$$

To estimate the evolution of the parameters of the model—the prior q_t , cutoff C_{it} , and expertise σ_{it} —we need to write them as functions of observed covariates.²⁰ We choose functional forms to constrain $q_t \in [0, 1]$ and $C_{it}, \sigma_{it} \geq 0$. The specification for q_t is:

$$\ln\left(\frac{q_t}{1 - q_t}\right) = \alpha_0 + \alpha_1 \cdot q_t^R + \alpha_2 \cdot q_t^M. \quad (10)$$

The baseline specification for the cutoff to test empirical prediction 1 is:

$$\ln(C_{it}) = \beta_0 + \beta_1 \cdot \text{D(Vet)}_{it} + \beta_2 \cdot \text{D}(N^R)_t + \beta_3 \cdot \text{D(Int)}_i + \beta_4 \cdot \text{D(Hike)}_t, \quad (11)$$

where $\text{D}(N^R)_t$ indicates whether the period t committee composition includes at least three rookies, D(Int)_i indicates whether member i is internal, and D(Hike)_t indicates whether the agenda includes at least one option to raise rates. The inclusion of D(Vet)_{it} , which captures the effect of tenure, and $\text{D}(N^R)_t$, which controls for the balance of the committee between rookies and veterans, comes directly from the theoretical model. We also include D(Int)_i to control for ex-ante heterogeneity between internal and external members; Hansen, McMahon, and Velasco (2014) show that the key observable dimension along which MPC members vary in terms of the structural voting parameters is internal-external. The inclusion of D(Hike)_t controls for agenda-specific preferences.²¹ In section

²⁰An alternative estimation approach is to use a two-step procedure along the lines of Iaryczower and Shum (2012) and Hansen, McMahon, and Velasco (2014). Its main disadvantage is that the mapping between the theoretical and empirical models is less direct, but its advantage is that it can accommodate estimation of strategic voting models. In the online appendix we use this approach to estimate both the sincere and strategic voting models, and find the results are qualitatively identical.

²¹To be completely in line with the model, we could also model C_{it} as a function of q_t , but identifying this dependence separately from the independent effect of q_t on the threshold in lemma 1 is impossible. So our estimates of C_{it} for rookies should be interpreted as the average effect of signalling over different

6, we show that the estimates are robust to running the more parsimonious specification excluding the additional controls.

The specification for expertise is:

$$\ln(\sigma_{it}) = \gamma_0 + \gamma_1 \cdot D(\text{Vet})_{it} + \gamma_2 \cdot D(\text{Int})_i. \quad (12)$$

We include the $D(\text{Vet})_{it}$ term to measure the effect of tenure, and, for similar reasons as above, include $D(\text{Int})_i$ in order to account for ex-ante heterogeneity between internal and external members in terms of expertise (Hansen, McMahon, and Velasco 2014). We leave the inclusion of $D(\text{Hike})_t$, which would suggest agenda-specific expertise, for a robustness check in section 6.

Using (9)-(12) and the MPC data, we estimate the α , β , and γ parameters via maximum likelihood. Given that our main interest is how the cutoff and expertise evolve, we use the estimated versions of (11) and (12) to obtain estimates of C_{it} and σ_{it} . In the tables below, we report the sample average of the cutoff and expertise for rookies and veterans separately and report bootstrapped standard errors.²²

To test empirical prediction 2, we expand the specification in (11) to include one of the member type indicators [$D(\theta^{PCTExp})_i$ or $D(\theta^{PCTExp})_i$] and its interaction with $D(\text{Vet})_{it}$. Using a similar approach, we can then recover estimates of the average C and average σ for rookies and veterans, separately for inflation-averse and inflation-tolerant members.

5.2 Results

Table 2 shows the estimated values for the structural parameters from the baseline specification.²³ The estimated member cutoffs explain the reduced-form results while also guiding us as to their source. Consistent with the theoretical model, we find evidence of a significant upward shift in the cutoff with experience. There is no statistically significant change in the expertise parameter. This result is of independent interest since it suggests that members do not accumulate additional expertise with experience. Instead, voting dynamics are driven by a shift in the average member's cutoff in line with empirical prediction 1.

Table 3 reports the estimated cutoffs for our two distinct measures of inflation aversion. Each measure produces qualitatively identical results. All types have a significant increase in their cutoffs, but the increase is significantly greater for more inflation-tolerant

values of q_t . Also, the key predictions of signalling are independent of the value of q_t .

²²We use standard parametric bootstrapping. For each simulation, we plug the maximum likelihood estimates of α , β , and γ into the likelihood function (9) and use it to draw new votes. We then re-estimate α , β , and γ and re-recover estimates of the cutoff and expertise for rookies and veterans. We perform 1,000 such simulations, which yields a distribution from which one can compute confidence intervals in the usual way.

²³Table B.2 in appendix B reports the maximum likelihood coefficient estimates for the baseline case.

Table 2: Baseline Estimates of Structural Parameters

	Baseline		
	Rookie	Veteran	Difference
$C(\theta)$	0.87	3.17	2.30 [0.003]
σ	0.40	0.41	0.01 [0.406]

Notes: This table shows the structural estimates for rookie and veteran members, as well as the difference between them. 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 bootstrapped distribution of estimates.

members. This difference-in-differences is perhaps the strongest evidence supporting the relevance of signalling on the MPC.

Table 3: Evolution of Cutoffs by Member Type

	Inflation Averse			Inflation Tolerant			Diff-in-Diff
	Rookie	Veteran	Difference	Rookie	Veteran	Difference	
$D(\theta^{PCTExp})$	0.33	1.58	1.26 [0.007]	1.52	11.22	9.69 [0.006]	-8.43 [0.008]
$D(\theta^{PCT})$	0.30	0.68	0.38 [0.048]	2.45	18.28	15.84 [0.000]	-15.46 [0.000]

Notes: This table reports the estimated cutoffs for high and low θ_i types based on our two indicator measures described in the text (calculated using (i) the percentage of high votes when a veteran and (ii) the percentage of high votes over entire tenure). The final column compares the effect of experience on the increase in the cutoff between the two groups. The terms reported in brackets are p-values, calculated using a bootstrapped distribution of estimates, for a one-sided test of difference from zero.

The fact that cutoffs diverge over time suggests that veterans should be more likely to dissent from the MPC decision than rookies. As a consistency check, table 4 presents estimates of an OLS regression similar to equation (7) but with whether member i dissents in meeting t as the dependent variable. Consistent with the structural results, veterans are found to be significantly more likely to dissent. This is essentially a more formal econometric test (e.g. with time and member fixed effects) with a longer sample of a similar reduced-form finding in Gerlach-Kristen (2003).²⁴

Our two empirical findings are also important because they distinguish the signalling model from an alternative learning model in which the cutoff is a belief about some unknown structural parameter of the macroeconomy that changes in response to new

²⁴Riboni and Ruge-Murcia (2014), and references therein, provide a more comprehensive discussion of the literature of monetary policy dissents.

Table 4: Reduced-form evidence on the impact of experience on dissent likelihood

Main Regressors	(1) D(Dissent)	(2) D(Dissent)
D(Vet)	0.16*** [0.000]	
D(12+ meetings)		0.086** [0.012]
D(24+ meetings)		0.071** [0.050]
D(36+ meetings)		-0.068 [0.101]
Constant	-0.073 [0.519]	0.047 [0.715]
R-squared	0.302	0.296
Model	Panel LPM	Panel LPM
Member effects?	YES	YES
Time effects?	YES	YES
Sample?	06/97-03/09	06/97-03/09
Obs?	1246	1246

Notes: This regression presents OLS estimates of a variant of equation (7) with standard errors clustered by member. The difference in this regression is that the dependent variable is $D(\text{Dissent})_{it}$ —a dummy variable capturing whether member i dissented relative to the committee decision in period t . We also present in column (2) the marginal effect of each year of service (when the baseline is a member’s first year). The results show that, controlling for member and time fixed effects, veteran members are more likely than rookies to dissent from the MPC decision, and consistent with earlier regressions we find the differential effect is strongest at the start of a member’s tenure.

data.²⁵ We are not aware of any paper in the macro learning literature that yields the dynamic patterns that we identify and we are skeptical that straightforward extensions of the current generation of learning models could do so. In particular, 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. The difference-in-differences finding for the evolution of cutoffs directly contradicts this prediction. 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.

Another alternative explanation concerns a growing willingness to express views that differ from the majority. Johnson, Ellis, and Kotenko (2012) study the policy preferences of members of the US Federal Reserve’s Federal Open Market Committee (FOMC). They find that Regional Fed Presidents adopt a more hawkish stance towards the end of their

²⁵This literature has examined the effects of policymakers’ 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 potential output (Bullard and Eusepi 2005).

tenure. They interpret this as showing that the Chair’s push for consensus simply makes Presidents less likely to express their true preferences and they initially express the dovish preferences of the chair. One might wonder if such early conformism is responsible for our MPC findings but there are a number of reasons that make this explanation unlikely as the main driver of the behavior we uncover, though such a channel may still somewhat complement the channel explored in this paper.

First, the FOMC has a very different structure and protocols to the MPC. In particular, it has a strong norm for conforming to the Chair’s view that is explicitly absent in the MPC. Second, contrary to what one might expect if the explanation for a change in behavior were related to a more carefree attitude as members approached the end of their time on the MPC, in column (3) of Table 1 we find no change in behavior as MPC members approach the end of their terms in a similar reduced-form regression to that in Johnson, Ellis, and Kotenko (2012). Third, we find that all MPC members change their cutoffs over time, whereas on the FOMC only one group appears to modify their behavior away from the Chair. In other words, the MPC does not seem to have a focal, stable cutoff from which a certain group diverges over time.

6 Robustness

In section 2.3 we deferred the discussion of a number of theoretical simplifications. Here we return to those simplifications while also discussing a number of empirical robustness exercises. The tables associated with this section are contained in the online appendix.

In the model, if MPC members behaved strategically rather than sincerely, they would condition their votes on being pivotal (i.e. changing the decision). Such behavior introduces more meaningful interactions among committee members than the sincere model, in which members essentially behave independently. The main new theoretical effect is that rookies’ votes also signal to their colleagues as well as the public. The online appendix describes the strategic model and its implications. We also describe how to estimate the model under strategic voting using the two-step estimator of Iaryczower and Shum (2012). We show that our findings are robust to the assumption of strategic voting.

Although our empirical specifications allow for differences between internal and external members, our theory model treats all policymakers as ex-ante identical. In fact, the model could easily be re-written to accommodate external and internal members’ types being drawn from different distributions. Our main predictions would then hold conditional on being an internal or external member: all bankers with the same appointment status would have an increasing cutoff, with more inflation-tolerant members’ cutoffs increasing more. The next section reports the results of table 2 disaggregated by internals and externals, and shows the differences prediction is satisfied within appointment status.

In the appendix, we report the results in table 3 separately by appointment status, and show this stricter difference-in-differences prediction is consistent with our results.

One might also be concerned that other drivers give rise to the dynamic behavior of MPC members. First, we consider that the behavior might be driven by members' desire to pursue different career paths rather than anchor inflation expectations. We check to see whether the estimated evolution of cutoffs differs depending on members' career background as well as future career. All subgroups display a significant increase in their estimated cutoff, suggesting our results above are not driven by any particular group seeking out specific career goals.

We also consider whether a desire for reappointment could be driving results given that both internals and externals can be reappointed after the end of their three-year terms (five-years for Governor-level members). One might especially worry if reappointment endogenously follows from more dovish interest rate choices towards the end of the first term as this might be an alternative explanation for our empirical findings. In the online appendix, we examine the relationship between the type proxies and reappointment, and show that, if anything, more hawkish (internal) members are more likely to be reappointed to serve a second term. While not conclusive evidence against any changes in behavior driven by a desire to be reappointed, these findings at least reassure us that such effects are not the main drivers of the dynamic behavior.

The final alternative source of dynamic behavior we consider is that members might care about the committee's reputation independently of their own, for example because they care about the public's perception of the Bank. Although veterans do not change inflation expectations in the model, the presence of such concerns would mean that one should treat the estimates of veterans' cutoffs not as their true preferences, but as a mix of preferences and the desire to protect the committee's reputation. If this were the case, it is likely that such concerns would have been greatest immediately following the establishment of the MPC when there was the most uncertainty about how it would operate. During these initial meetings, moreover, all members were, by our definition, rookies. We therefore remove the first 18 meetings from the sample and reestimate the structural exercise. We find that even after the MPC matured as an institution, the predictions of our model are confirmed. Again, this does not rule out any concern about committee reputation, but rather reinforces that it is not the main driver of our empirical results.

As described in section 3, our two proxies for inflation aversion derive from simple statistics of how often each member votes for high rates. We now consider some alternative measures of the member type. The first alternative measure comes from Eijffinger, Mahieu, and Raes (2013) (EMR hereafter). They update the ideal point estimates of MPC members using the Bayesian simulation methodology of Hix, Hoyland, and Vivyan

(2010). These ideal points are estimates from an item response model and they measure MPC members' preferences for higher or lower interest rates. We rank members based on EMR's reported ideal points, and use the ranking to create an indicator variable $D(\theta^{EMR})_i$ that roughly splits the sample equally.

The second measure simply uses the fixed effect estimated from equation (7). The main advantage of this measure over our simple percentage of high votes measure is that it controls for common drivers of voting behavior through the inclusion of time effects. We split members into those with high and low fixed effects and define a dummy variable, $D(\theta^{FE})_i$, similar to the other three measures. The results in the online appendix show that the difference-in-differences prediction is robust to all our measures of inflation aversion.

Finally, we conduct a number of checks on the robustness of the specification used for the structural analysis. First we examine a more flexible specification for equation (12) in which we allow the expertise to vary with the agenda by including the $D(\text{Hike})_t$ variable. Second, we include only the terms predicted by our model, $D(\text{Int})_i$ and $D(N^R)_t$, in equation (11) and otherwise leave the baseline specification as it is. Finally, since Mervyn King was present in all the committee meetings in the sample, and hence is not consistent with our turnover assumption, we show that he is not driving the results by excluding him from the sample and re-estimating the structural parameters. In all cases, our results remain robust to the changes in the specification.

7 Signalling, Policy Rates, and Welfare

Given the structural estimates, we can quantify the effect of signalling on policy. First, by proposition 2, the structural estimates of cutoffs for veterans are the same ones veterans would use without signalling. Furthermore, by proposition 1, rookies would also use these cutoffs without signalling. So, to assess the effect of signalling, one can compare rates chosen by a signalling committee in which rookies and veterans use the cutoffs we estimate (C_{it}^S) with those chosen by a non-signalling committee in which both rookie and veteran members adopt veteran cutoffs (C_{it}^{NS}).²⁶

Our counterfactual exercise uses a committee of five internal and four external members to match the structure of the MPC. In odd (even) periods there are three (two) rookie internals, three (one) rookie externals, two (three) veteran internals, and one (three) veteran externals.²⁷ In order to mimic the actual MPC, we need to allow for cutoffs and expertise to vary across internals and externals, whose behavior several studies have found to differ. Table 5 reports the estimated structural parameters from the baseline specifica-

²⁶If concerns for the reputation of the committee also affect behavior (see section 2.3), then this exercise will understate the effect of signalling on policy.

²⁷The exact composition of rookies and veterans in odd and even periods does not greatly affect the results. We choose this composition since it matches the $D(N^R)_t$ control used in (11).

tion separately for internal and external members. In line with Hansen, McMahon, and Velasco (2014), internal members have greater expertise and are more inflation averse. Also, the external cutoff shifts more than the internal cutoff, which is consistent with the difference-in-differences prediction. But ex ante heterogeneity might also explain this result. For example, there may be less uncertainty about internals' types, reducing their incentive to signal.

Table 5: Structural Parameters for Internals and Externals

	Internal			External			Diff-in-Diff
	Rookie	Veteran	Difference	Rookie	Veteran	Difference	
$C(\theta)$	0.17	0.78	0.61 [0.000]	1.42	7.18	5.76 [0.000]	-5.14 [0.000]
σ	0.30	0.33	0.04 [0.055]	0.48	0.54	0.06 [0.055]	-0.02 [0.055]

Notes: This table replicates table 2 for internal and external members separately (see that table for details). The final column compares the effect of experience between the two groups for the cutoff C , as well as the precision parameter σ . The terms reported in brackets are p-values, calculated using a bootstrapped distribution of estimates, for a one-sided test of difference from zero.

The other estimated object important for the counterfactual exercise is the distribution of q_t . When the prior is extreme, one would expect the impact of signalling on policy to be small since everyone has a clear view on the right decision. However, when the prior is nearer to 0.5, there is more uncertainty and signalling should play a larger role. The structural model produces 137 different estimates of q_t , one for each meeting for which we have q_t^R and q_t^M data. These are roughly uniform on the $[0, 1]$ interval.

We consider the behavior of the committee over 5 years, or 60 meetings. For each period $t = 1, \dots, 60$, we generate two policy choices— r_t^S for the signalling (actual) committee and r_t^{NS} for the non-signalling (counterfactual) committee—using the following procedure:

1. Draw q_t from one of the 137 fitted values (with replacement).
2. Draw ω_t given q_t .
3. Draw a signal for each member from $\mathcal{N}(\omega_t, \sigma_{it}^2)$.
4. Draw a vote for each member using member-specific parameters:²⁸
 - (a) For the actual committee (S), $v_{it} = 1 \Leftrightarrow s_{it} \geq s_{it}^* (C_{it}^S, \sigma_{it}, q_t)$.
 - (b) For the counterfactual committee (NS), $v_{it} = 1 \Leftrightarrow s_{it} \geq s_{it}^* (C_{it}^{NS}, \sigma_{it}, q_t)$.

²⁸To be clear, C_{it}^S and σ_{it} take on any of the four estimated values in table 5 depending on whether i is internal or external and whether he is a rookie or veteran in period t . By contrast, C_{it}^{NS} only takes on two values, the estimated veteran cutoffs for internals or externals.

5. Compute r_t^S and r_t^{NS} by majority rule given individual votes.

Since $s_{it}^*(C_{it}^{NS}, \sigma_{it}, q_t) \geq s_{it}^*(C_{it}^S, \sigma_{it}, q_t)$, a member who votes for low rates when signalling will never vote for high rates when not. So it is either the case that $r_t^S = r_t^{NS}$, or that $r_t^S = 1 > 0 = r_t^{NS}$ and we can measure the overall effect of signalling, over a given horizon, as the number of times the actual committee chooses a high rate while the counterfactual committee chooses a low one. One can express this impact in basis point terms by multiplying the number of rate differences by the standard interest rate increment, 25.

As the votes depend on draws of signals for each member, a policy rate path derived from this exercise is a single draw of a random variable, and the ultimate objects of interest are the distributions of interest rate paths. We simulate these by drawing 10,000 different signalling and non-signalling paths. In the first three rows of table 6 we report cumulative probabilities of differences in the interest rate paths. In a twelve-month period, signalling induces a 25bps difference with probability 0.41. After 36 months this probability grows to 0.80 and after 60 months to 0.93. Larger effects also have large probabilities. The odds are nearly even that signalling generates a 50bps difference after 36 months and a 75bps difference after 60 months. Given that, according to Joyce, Lasasosa, Stevens, and Tong (2011), the Bank of England's QE program reduced yields on five- to 25-year gilts by around 100bps on average, the effects we estimate are economically substantial.²⁹

Table 6: Counterfactual analysis comparing a signalling and non-signalling committee

	After X Months				
	12	24	36	48	60
Pr(At least 25bps extra cut)	0.41	0.66	0.80	0.88	0.93
Pr(At least 50bps extra cut)	0.09	0.28	0.46	0.63	0.74
Pr(At least 75bps extra cut)	0.01	0.08	0.21	0.34	0.49
Mean Errors (Signalling Committee)	0.8	1.6	2.4	3.2	4.0
Mean Errors (Non-Signalling Committee)	1.2	2.5	3.8	5.1	6.3
Pr(Fewer errors)	0.01	0.02	0.01	0.01	0.01
Pr(Equal errors)	0.63	0.38	0.23	0.14	0.09
Pr(More errors)	0.35	0.61	0.76	0.85	0.90

Notes: The upper panel of this table summarises the distribution of differences between the signalling and non-signalling interest rate paths from 10,000 simulations of committee behavior. The lower panel shows the mean number of errors made on each committee, as well as well various cumulative probabilities of the differences between the two committees.

Perhaps an even more important economic question is the effect of signalling on welfare. As discussed in section 3.1, the Bank of England treats upside and downside misses

²⁹Given the nature of the counterfactual exercise, the difference between the two committees would continue to grow with longer time horizons. Interpreting effects at horizons longer than a business cycle are problematic, however, because variables treated as exogenous like the agenda and the distribution of q_t would begin to differ in economies whose interest rates diverged by as much as 100bps.

equally, so we can assess welfare by comparing whether the actual or counterfactual paths generate more errors. Signalling can both help welfare (by inducing high rates when $\omega_t = 1$) and hurt it (by doing so when $\omega_t = 0$). Nevertheless, we find that signalling clearly improves welfare on average. Moreover, the probability that the non-signalling committee makes fewer errors is consistently around 1%. But once the time horizon reaches two years, the signalling committee has a greater than 60% probability of performing strictly better and by five years this is 90%.

Signalling creates a benefit since dovish externals behave more neutrally when rookies, but also a cost since already hawkish internals behave even more hawkishly. Why then does signalling improve welfare? The effect of the cutoff on voting behavior declines with expertise: a member with a very low value of σ simply follows his very precise signal and disregards the cutoff.³⁰ Since internals are more expert than externals, the beneficial effect of signalling on externals' cutoffs is more relevant for welfare than the detrimental effect on internals'.

8 Conclusion

This paper argues that one should take seriously the idea that independent monetary policy makers care about their reputation for inflation aversion. It does so by building a new model of signalling in monetary policy with two predictions on dynamic behavior that can be tested with structural estimation. Although the Bank of England is admired for its independence, the voting behavior of its MPC members fits very well a model in which preferences are heterogeneous and establishing credibility is crucial. This suggests that independent central banks do not automatically replicate the policies of the metaphorical “hard-nosed” banker often invoked to discuss the behavior of central banks.

Beyond showing that signalling is important for explaining voting behavior, the paper's welfare results have clear implications for committee design, and particularly those features which can affect the incentive to signal. Since signalling appears to affect behavior most at the beginning of tenure and fade over time, our results provide a rationale for regular committee member turnover to generate uncertainty on policymakers' types and maintain the strength of the signalling incentive. Second, more subtly, the cutoffs in table 5 suggest that reappointing internal members and replacing external ones might dominate replacing the whole committee.³¹ Finally, to the extent that individual signalling

³⁰As one can see from lemma 1, as $\sigma \rightarrow 0$ the threshold that the signal must cross to vote high converges to $\frac{1}{2}$, which is independent of the cutoff C .

³¹We cannot simply run a counterfactual with this alternative committee structure and the cutoffs in table 5 because the equilibrium strength of the signalling incentive for external members will not be the same as in the committee with rookie internals. Quantifying this difference would require estimates of additional model parameters.

incentives are positively related to the public's ability to observe policymakers' choices (as in Cukierman and Meltzer (1986), Faust and Svensson (2001), and Sibert (2009)), the results support a transparent regime (in line with Bank of England policy) in which the public can directly observe individual votes as opposed to just the committee decision.

Of course, our conclusions are based on the voting behavior of the MPC, so one might wonder the extent to which they are externally valid. An initial point is that numerous other central banks, such as the Swedish Riksbank, have very similar institutional frameworks that combine a committee of experts with an inflation target. Nevertheless, it is important to consider how the particular structure of the MPC might affect the results.

First, the MPC has nine members. In countries with smaller (larger) committees, one would expect signalling to have a stronger (weaker) effect on voting behavior. A rational public changes its inflation expectations after observing a single member vote high depending on whether that member will affect future policy. Since members are more likely to change policy on smaller committees, the reaction of expectations to individual votes should be correspondingly higher.

Second, the MPC has an inflation target, which should partially solve the credibility problem and provide some discipline for its members to be tough on inflation. For example, Cukierman and Muscatelli (2008) find evidence that UK monetary policy makers became less inflation tolerant (a lower θ_i in our model) after the adoption of an inflation target. In central banks with greater policy discretion, time consistency problems might be more serious and signalling may be more important.

Third, the MPC encourages members to vote for their preferred policies even if these contradict the Chair's. Thus, individual votes can be taken as informative signals of true preferences. In committees in which the Chair plays a more dominant role, like the Federal Open Market Committee (FOMC), members typically do not dissent in vote against the Chair's proposal, making it much more difficult for the public to infer preferences based on voting records. In this case, other signals, such as member speeches, may be more important communication channels.

Ultimately the impact of different characteristics of institutional design on signalling incentives is an empirical question which we leave for future research.

References

- APPELBAUM, B. (2013): “Possible Fed Successor Has Admirers and Foes,” <http://www.nytimes.com/2013/04/25/business/janet-1-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.
- (2011): “Notes On The Bank Of England Option-Implied Probability Density Functions,” <http://www.bankofengland.co.uk/statistics/IMPLIEDPDFS/BofEOptionImpliedProbabilityDensityFunctions.pdf>, last accessed 04 April 2013.
- (2015): <http://www.bankofengland.co.uk/monetarypolicy/Pages/framework/framework.aspx>, last accessed 8 June 2015.
- 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.
- BERK, J. M., B. BIERUT, AND E. MEADE (2010): “The Dynamic Voting Patterns of the Bank of England’s MPC,” DNB Working Papers 261, Netherlands Central Bank, Research Department.
- BESLEY, T., N. MEADS, AND P. SURICO (2008): “Insiders versus Outsiders in Monetary Policymaking,” *American Economic Review*, 98(2), 218–23.
- BHATTACHARJEE, A., AND S. HOLLY (2005): “Inflation Targeting, Committee Decision Making and Uncertainty: The Case of the Bank of England’s MPC,” Discussion paper, Department of Applied Economics, University of Cambridge.

- 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.
- CUKIERMAN, A., AND A. MUSCATELLI (2008): “Nonlinear Taylor Rules and Asymmetric Preferences in Central Banking: Evidence from the United Kingdom and the United States,” *The B.E. Journal of Macroeconomics*, 8(1), 1–31.
- 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.
- 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. (2003): “Insiders and Outsiders at the Bank of England,” *Central Banking*, XIV(1), 96–102.
- (2004): “Is the MPC’s voting record informative about future UK monetary policy?,” *Scandinavian Journal of Economics*, 106(3), 299–313.
- HANSEN, S., AND M. MCMAHON (2008): “Delayed Doves: MPC Voting Behaviour of Externals,” CEP Discussion Papers dp0862, Centre for Economic Performance, LSE.
- HANSEN, S., M. MCMAHON, AND C. VELASCO (2014): “How Experts Decide: Preferences or Private Assessments on a Monetary Policy Committee?,” *Journal of Monetary Economics*, 67(October), 16–32.

- HARRIS, M., AND C. SPENCER (2008): “Decade of dissent: explaining the dissent voting behavior of Bank of England MPC members,” MPRA Paper 9100, University Library of Munich, Germany.
- 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.
- JOHNSON, E. D., M. A. ELLIS, AND D. KOTENKO (2012): “Consensus building on the FOMC: An analysis of end of tenure policy preferences,” *Economics Letters*, 117(1), 368 – 371.
- JOYCE, M. A. S., A. LASAOSA, I. STEVENS, AND M. TONG (2011): “The Financial Market Impact of Quantitative Easing in the United Kingdom,” *International Journal of Central Banking*, 7(3), 113–161.
- KALAI, E., AND E. LEHRER (1994): “Weak and strong merging of opinions,” *Journal of Mathematical Economics*, 23(1), 73–86.
- KING, M. (1997): “Changes in UK monetary policy: Rules and discretion in practice,” *Journal of Monetary Economics*, 39(1), 81–97.
- 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.
- PERSSON, T., AND G. TABELLINI (2000): *Political Economics*. MIT Press.
- 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.
- RIBONI, A., AND F. RUGE-MURCIA (2014): “Dissent in monetary policy decisions,” *Journal of Monetary Economics*, 66(C), 137–154.

- 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.
- SPENCER, C. (2006): “The Dissent Voting Behaviour of Bank of England MPC Members,” Department of Economics Discussion Papers 0306, Department of Economics, University of Surrey.
- 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 Omitted Proofs

A.1 Proof of Lemma 1

Proof. The pdf of $s_{it} \mid \omega_t$ is proportional to $\exp \left[\frac{1}{2\sigma_{it}^2} (s_{it} - \omega_t)^2 \right]$, so

$$\hat{\omega}_{it} = \frac{\exp \left[-\frac{1}{2\sigma_{it}^2} (s_{it} - 1)^2 \right] q_t}{\exp \left[-\frac{1}{2\sigma_{it}^2} (s_{it} - 1)^2 \right] q_t + \exp \left[-\frac{1}{2\sigma_{it}^2} s_{it}^2 \right] (1 - q_t)}$$

and

$$\frac{\hat{\omega}_{it}}{1 - \hat{\omega}_{it}} = \frac{q_t}{1 - q_t} \frac{\exp \left[-\frac{1}{2\sigma_{it}^2} (s_{it} - 1)^2 \right]}{\exp \left[-\frac{1}{2\sigma_{it}^2} s_{it}^2 \right]} = \frac{q_t}{1 - q_t} \exp \left[-\frac{1}{2\sigma_{it}^2} (1 - 2s_{it}) \right]$$

Following a cutoff voting rule is equivalent to voting high whenever $\ln \left(\frac{\hat{\omega}_{it}}{1 - \hat{\omega}_{it}} \right) \geq \ln(C_{it})$, which is equivalent to

$$\ln \left(\frac{q_t}{1 - q_t} \right) - \frac{1}{2\sigma_{it}^2} (1 - 2s_{it}) \geq \ln(C_{it}) \Rightarrow s_{it} \geq \frac{1}{2} - \sigma_{it}^2 \left[\ln \left(\frac{q_t}{1 - q_t} \right) - \ln(C_{it}) \right].$$

■

A.2 Proof of Proposition 1

Proof. A veteran's expected utility of choosing v_{it} is

$$U(v_{it}, \theta_i) = \hat{\omega}_{it} \{ -l[\pi(v_{it}, 1)] + \theta_i \pi(v_{it}, 1) \} + (1 - \hat{\omega}_{it}) \{ -l[\pi(v_{it}, 0)] + \theta_i \pi(v_{it}, 0) \} - \theta_i \pi_t^E.$$

$U(1, \theta_i) \geq U(0, \theta_i)$ is equivalent to using a cutoff voting rule with cutoff $C^{NS}(\theta_i)$. Moreover, since π_{t+1}^E is independent of observed period t votes, rookies face the identical utility maximization problem as veterans. ■

A.3 Proof of Proposition 2

Proof. As explained in the text, a veteran in the signalling model adopts the same voting rule as with no signalling, so we focus on the equilibrium behavior of rookies. Their expected utility of choosing v_{it} is

$$U_R(v_{it}, \theta_i) = \hat{\omega}_{it} \{ -l[\pi(v_{it}, 1)] + \theta_i \pi(v_{it}, 1) \} + (1 - \hat{\omega}_{it}) \{ -l[\pi(v_{it}, 0)] + \theta_i \pi(v_{it}, 0) \} - \theta_i \pi_t^E - \delta \theta_i \mathbb{E} \{ \pi^{E*} [\mathbf{p}[(V_{-it}^R + v_{it}, q_t, N_t^R), q_{t+1}, N_{t+1}^R, \omega_{t+1}]] \},$$

from which one derives the cutoff in (5).

According to lemma 1, cutoff voting rules can equivalently be expressed in terms of signal thresholds. Accordingly, define

$$s_R^*(\theta, q_t, N_t^R) \equiv s^*[C_R^S(\theta, q_t, N_t^R), \sigma_R, q_t] \text{ and } s_V^*(\theta) \equiv s^*[C_V^S(\theta), \sigma_V, q_t]$$

as the equilibrium thresholds adopted by rookies and veterans.³² Also define

$$H(\omega_{t+1}, I_{t+1}) \equiv \mathbb{E} \left\{ \Pr \left[r_{t+1} = 0 \mid s_R^*(\theta, q_{t+1}, N_{t+1}^R), s_V^*(\theta), \omega_{t+1} \right] \mid \mathbf{p}_{t+1} \right\},$$

One can then express $\pi^{E*}(I_{t+1})$ as

$$q_{t+1} \{ \pi(1, 1) + [\pi(0, 1) - \pi(1, 1)]H(1, I_{t+1}) \} + (1 - q_{t+1}) \{ \pi(1, 0) + [\pi(0, 0) - \pi(1, 0)]H(0, I_{t+1}) \}.$$

Note that the conditional probability inside the expectation in $H(\omega_{t+1}, I_{t+1})$ is strictly increasing in any veteran's type since the probability any veteran i votes $v_{it} = 0$ is increasing in s_V^* , which is itself increasing in θ_i .

The main proof proceeds in three steps. First, assumption A3 implies $[\pi(0, \omega_t) - \pi(1, \omega_t)] - \delta \Delta(q_t, N_t^R \mid \omega_t) > 0 \forall \omega_t$ by guaranteeing that $\pi(0, 0) - \pi(1, 0) > \delta \bar{\Delta}$ and $\pi(0, 1) - \pi(1, 1) > \delta \bar{\Delta}$, where $\bar{\Delta} = \bar{q}[\pi(0, 1) - \pi(1, 1)] + (1 - \bar{q})[\pi(0, 0) - \pi(1, 0)]$ is an upper bound on the change in period $t + 1$ inflation expectations from observing an additional high vote by rookies.

Second, in every equilibrium $\Delta(\cdot \mid \omega_t) \geq 0$. Let p_0^θ be the prior probability attached to a rookie i having type θ . After observing v_{it} and ω_t , the public forms the posterior belief on θ_i given by

$$p_i(v_{it}, \omega_t) = \frac{\{1 - \Phi[(s_R^*(\theta, q_t, N_t^R) - \omega_t) / \sigma_R]\}^{v_R} \Phi[(s_R^*(\theta, q_t, N_t^R) - \omega_t) / \sigma_R]^{1-v_R} p_0^\theta}{\sum_{\theta} [1 - \Phi[(s_R^*(\theta, q_t, N_t^R) - \omega_t) / \sigma_R]]^{v_R} \Phi[(s_R^*(\theta, q_t, N_t^R) - \omega_t) / \sigma_R]^{1-v_R} p_0^\theta},$$

where Φ is the standard normal cdf. Now observe that

$$p_i(0)/p_i(1) \propto \frac{\Phi[(s_R^*(\theta, q_t, N_t^R) - \omega_t) / \sigma_R]}{1 - \Phi[(s_R^*(\theta, q_t, N_t^R) - \omega_t) / \sigma_R]},$$

which is monotonically increasing in θ . Since $H(\omega_{t+1})$ is also monotonically increasing in members' types, it is immediate that

$$\pi^{E*}[\mathbf{p}(V_{-i,t}^R + 1, q_t, N_t^R), q_{t+1}, N_{t+1}^R] > \pi^{E*}[\mathbf{p}(V_{-i,t}^R, q_t, N_t^R), q_{t+1}, N_{t+1}^R].$$

Third, we must show an equilibrium exists. Since N_t^R only depends on whether t is an odd or even period, this is a solution to a system of $2K$ equations. Let $C_{R0}^S(\theta, q_t)$ [$C_{R1}^S(\theta, q_t)$] be the K equilibrium cutoffs used in even (odd) periods. Moreover, let $\Delta_0(q_t, N_0 \mid \omega_t)$ and $\Delta_1(q_t, N_1 \mid \omega_t)$ be the equilibrium strengths of the signalling incentive in even and odd periods, respectively. These each depend continuously on both sets of cutoffs $C_{R0}^S(\theta, q_t)$ and $C_{R1}^S(\theta, q_t)$. The equilibrium system is then

$$C_{Rj}^S(\theta, q_t) = \frac{\mu_0 + \theta \{ [\pi(0, 0) - \pi(1, 0)] - \delta \Delta_j(q_t, N_j \mid 0) \}}{\mu_1 - \theta \{ [\pi(0, 1) - \pi(1, 1)] - \delta \Delta_j(q_t, N_j \mid 1) \}}$$

³²Here we drop the i index in the arguments of the equilibrium thresholds because rookies and veterans are assumed to use symmetric strategies.

for $\theta = \underline{\theta}, \dots, \bar{\theta}$ and $j = 0, 1$. By the arguments above, the right-hand side of this system maps

$$C_L \leq C_{Rj}^S(\theta, q_t) \leq \dots \leq C_{Rj}^S(\bar{\theta}, q_t) \leq C_H \text{ for } j = 0, 1$$

into itself, where $C_L = \frac{\mu_0 + \theta \{[\pi(0,0) - \pi(1,0)] - \delta \bar{\Delta}\}}{\mu_1 - \theta \{[\pi(0,1) - \pi(1,1)] - \delta \bar{\Delta}\}}$ and $C_H = \frac{\mu_0 + \theta [\pi(0,0) - \pi(1,0)]}{\mu_1 - \theta [\pi(0,1) - \pi(1,1)]}$. Hence the conditions of Brouwer's fixed point theorem are met, and a solution to the system exists.

It remains to be shown that $C_V^S(\theta) - C_{Rj}^S(\theta, q_t)$ is increasing in θ . Letting $a_{\omega_t} \equiv \pi(0, \omega_t) - \pi(1, \omega_t)$ and $b_{\omega_t}^j(q_t) \equiv \delta \Delta_j(q_t, N_j | \omega_t)$, then the following hold:

$$\begin{aligned} \frac{\partial C_V^S(\theta)}{\partial \theta} &= \frac{a_0[\mu_1 - \theta a_1] + a_1[\mu_0 + \theta a_0]}{[\mu_1 - \theta a_1]^2} = \frac{a_0\mu_1 + a_1\mu_0}{[\mu_1 - \theta a_1]^2} \\ \frac{\partial C_{Rj}^S(\theta, q_t)}{\partial \theta} &= \frac{[a_0 - b_0^j(q_t)][\mu_1 - \theta a_1 + \theta b_1^j(q_t)] + [a_1 - b_1^j(q_t)][\mu_0 + \theta a_0 - \theta b_0^j(q_t)]}{[\mu_1 - \theta a_1 + \theta b_1^j(q_t)]^2} = \\ &\quad \frac{[a_0 - b_0^j(q_t)]\mu_1 + [a_1 - b_1^j(q_t)]\mu_0}{[\mu_1 - \theta a_1 + \theta b_1^j(q_t)]^2} \end{aligned}$$

Since $\delta b_{\omega_t}^j(q_t) > 0$, we conclude that $\frac{\partial C_V^S(\theta)}{\partial \theta} - \frac{\partial C_{Rj}^S(\theta, q_t)}{\partial \theta} > 0$. ■

B Appendix Data Tables

Table B.1: MPC Members in the Sample

Member	D(Int) _i	First Meeting	Last Meeting	Total Meetings	% as Veteran	D(θ^{PCT}) _i	D(θ^{PCTExp}) _i
Davies	Internal	Jun 1997	Jul 1997	2	0.0	0	
George	Internal	Jun 1997	Jun 2003	73	75.3	0	0
King	Internal	Jun 1997		142	87.3	0	0
Plenderleith	Internal	Jun 1997	May 2002	60	70.0	1	0
Clementi	Internal	Sep 1997	Aug 2002	60	70.0	1	0
Vickers	Internal	Jun 1998	Sep 2000	28	35.7	0	1
Bean	Internal	Oct 2000		102	82.4	1	0
Tucker	Internal	Jun 2002		82	78.0	0	0
Large	Internal	Oct 2002	Jan 2006	40	55.0	0	0
Lomax	Internal	Jul 2003	Jun 2008	60	70.0	1	0
Gieve	Internal	Feb 2006	Feb 2009	37	51.4	1	1
Dale	Internal	Jul 2008		9	0.0	1	
Fisher	Internal	Mar 2009		1	0.0	0	
Buiter	External	Jun 1997	May 2000	36	50.0	1	1
Goodhart	External	Jun 1997	May 2000	36	50.0	0	1
Julius	External	Sep 1997	May 2001	45	60.0	1	0
Budd	External	Dec 1997	May 1999	18	0.0	0	
Wadhvani	External	Jun 1999	May 2002	36	50.0	1	1
Allsop	External	Jun 2000	May 2003	36	50.0	1	0
Nickell	External	Jun 2000	May 2006	72	75.0	1	0
Barker	External	Jun 2001		94	80.9	0	1
Bell	External	Jul 2002	Jun 2005	36	50.0	1	1
Lambert	External	Jun 2003	Mar 2006	34	47.1	1	0
Walton	External	Jul 2005	Jun 2006	12	0.0	0	
Blanchflower	External	Jun 2006		34	47.1	1	0
Besley	External	Sep 2006		31	41.9	0	1
Sentence	External	Oct 2006		30	40.0	0	0

Notes: This table provides summary statistics concerning the MPC members in our sample.

Table B.2: MLE Estimates for the Baseline Model

	Equation (10) $\ln \left(\frac{q_t}{1-q_t} \right)$	Equation (11) $\ln (C_{it})$	Equation (12) $\ln (\sigma_{it})$
q_t^R	4.64*** [0.000]		
q_t^M	1.21 [0.155]		
$D(\text{Vet})_{it}$		1.28*** [0.000]	0.11 [0.106]
$D(\text{Int})_i$		-2.14*** [0.000]	-0.49*** [0.000]
$D(\text{Hike})_t$		-2.69*** [0.000]	
$D(N^R)_t$		-0.18 [0.496]	
Constant	-2.99*** [0.000]	1.22*** [0.009]	-0.72*** [0.000]

Notes: This table provides maximum likelihood estimates of (10), (11) and (12) to maximize the likelihood (9). P-values are reported in the brackets below the estimated coefficient.