CESifo Working Papers

www.cesifo.org/wp

First Impressions Matter: Signalling as a Source of Policy Dynamics

Stephen Hansen
Michael McMahon

CESIFO WORKING PAPER NO. 3782
CATEGORY 7: MONETARY POLICY AND INTERNATIONAL FINANCE
APRIL 2012

PRESENTED AT CESIFO AREA CONFERENCE ON APPLIED MICROECONOMICS, MARCH 2012

CESifo *Working Paper No. 3782*

First Impressions Matter: Signalling as a Source of Policy Dynamics

Abstract

We first establish that policymakers on the Bank of England's Monetary Policy Committee choose lower interest rates with experience. We then reject increasing confidence in private information or learning about the structure of the macroeconomy as explanations for this shift. Instead, a model in which voters signal their hawkishness to observers better fits the data. The motivation for signalling is consistent with wanting to control inflation expectations, but not career concerns or pleasing colleagues. There is also no evidence of capture by industry. The paper suggests that policy-motivated reputation building may be important for explaining dynamics in experts' policy choices.

JEL-Code: D780, E520.

Keywords: signalling, learning, monetary policy.

Stephen Hansen
University Pompeu Fabra
Barcelona GSE / Spain
stephen.hansen@upf.edu

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

October 2, 2011
We would like to thank Jan Eeckhout and Ronny Razin for helpful comments.

1 Introduction

In many, if not most, public policy settings, policymakers act repeatedly. Competition commissioners, utility regulators, and ombudsmen are some of the many public officials who serve in their positions for several years, during which time they repeatedly make similar technical policy decisions. While the dynamic behavior of politicians has received a relatively large amount of attention in the economics literature,[1] the dynamic behavior of unelected policymaking experts is less well understood. There are two broad reasons why behavior might change over time: first, experts might receive additional information as they accumulate experience, which in turn changes their beliefs about the payoffs of alternative policy choices; second, incentives might change over time, so that experts begin to favor policy stances they initially avoided, even given the same information. Understanding the extent to which either of these two forces plays a role in policy dynamics is clearly crucial for designing the institutional environment in which policymakers operate.

This paper explores these issues in the particular context of monetary policymaking on the Bank of England's Monetary Policy Committee (MPC). It begins by establishing a surprising fact: the average member of the MPC votes for significantly lower rates after accumulating a year and a half of experience. In the remainder of the paper, we explore a number of hypotheses that relate to both learning and incentives to understand what drives this effect.

Before this, we present a model of policymaking as a Bayesian decision problem in which MPC members receive private signals on whether the economy is inflationary or not. If these signals show sufficient evidence of inflationary pressures, then a member chooses a high interest rate; otherwise, she chooses a low interest rate. The amount of evidence that a member needs in order to vote for the high rate is interpreted as a measure of toughness on inflation, or hawkishness. The first hypothesis we test with the model is whether, as they gain experience on the MPC, members also gain expertise that allows them to better judge whether inflationary shocks have hit. We show that such learning-by-doing would manifest itself as an increasing tendency to vote against the interest rate that the public expects to be chosen. We then take this prediction to the data, and show there is no evidence for it. Instead, the parameter that measures hawkishness shifts significantly after 18 months of service. The remaining hypotheses we explore all concern why measured hawkishness shifts.

The second hypothesis we test is whether learning about the structure of the macroeconomy is an important explanatory factor in the shift in hawkishness. Three findings make this unlikely: (1) new members vote for significantly higher rates than their other, more experienced, colleagues; (2) the shift occurs for nearly all members; and (3) the

[1] See, for example, Drazen (2001), who reviews the literature on the political business cycle.

variance of individual estimates of the hawkishness parameter do not increase with time. While we do not construct a formal model of learning, we argue that this combination of results is difficult to reconcile with either a common or private learning model.

Since changes in information do not appear to explain MPC dynamics, we next explore hypotheses related to incentives. An established idea in the monetary literature (Backus and Driffill 1985a,b, Vickers 1986) is that monetary policymakers signal hawkishness in order to anchor the public's expectations of future inflation and improve policy outcomes. We adapt our voting model to allow for such signalling, and generate dynamics using a two-step logic. First, in equilibrium a policymaker that cares about signalling a hawkish type will vote for high rates more often than his innate policy preferences would otherwise dictate. Second, the strength of the signalling incentive declines over time as the future horizon over which reputation pays off shortens. So, experienced members vote more according to their innate preferences, and therefore vote for low rates more often than their earlier selves. Predictions (1)-(3) above come out of this model. An additional prediction is also generated by this model: the degree to which hawkish and dovish types shift in their toughness on inflation should be the same. We test this hypothesis and find supportive evidence.

Our model is also consistent (and observationally equivalent to) models in which members signal hawkishness due to other motives, but we can examine whether the data supports these. The first potential motive is career concerns: members might want to signal to potential future employers. We find little correlation between the observed shift in hawkishness and term length or occupational background, which points against the relevance of career concerns. The second potential motive is a desire to "fit in": new members may conform to the norms of the MPC, which stress the importance of controlling inflation. While this story has a more behavioral flavor, it fits within the framework of standard signalling models. We see that even members with extensive experience at the Bank prior to their appointment to the MPC display a shift in hawkishness, which indicates that conformity is unlikely to be the driving force behind the voting dynamics. This leaves policy motivations as the main candidate for explaining signalling.

The final incentive story we explore is regulatory capture. More specifically, members might vote for lower rates over time because of persuasion efforts of the financial sector. We compare the votes of a subset of members who have served on both the MPC and a shadow MPC assembled by the Times of London newspaper, and show that they are generally less hawkish on the latter, which is inconsistent with the capture hypothesis.

Our paper thus makes two contributions. First, it shows that signalling better fits the observed voting dynamics on the MPC than learning or capture. Not only does the signalling model do a better job of rationalizing the observed shift in hawkishness, but it also generates an independent prediction that we validate. This finding is important for committee design; for example, it implies that first-year hawkishness will still arise if members serve an apprenticeship period in the Bank prior to their MPC service.[2] Our second contribution is to narrow down why voters engage in signalling by rejecting career concerns and fitting in. Instead, policy-motivated signalling is the main alternative that emerges from our analysis.

Policy-motivated signalling is relevant in any context in which the outcomes about which the policymaker cares depend on the public's expectations about her actions. For example, the willingness of a company to engage in anticompetitive practices presumably depends on its belief that competition authorities will investigate and prosecute violations of competition law. By taking tough stances at the beginning of her career, an industry regulator can signal her intention to crack down on malfeasance, thereby discouraging bad behavior by firms in the future and further achieving her policy objective. Our paper indicates that exploring policy-motivated signalling could contribute to a better understanding of experts' policy choices in environments beyond monetary policy.

1.1 Literature Review

Our work is related to the existing literature on macroeconomic learning, signalling in monetary policy, and the career concerns of policy makers. Since at least the work of Brainard (1967), monetary economists have understood that monetary policymakers do not have full information about the structure of the economy, and that this has implications for optimal policy choices. A large literature has subsequently developed that examines how policymakers update their beliefs as more information becomes available; Evans and Honkapohja (2001) is the seminal reference.[3] Papers have examined the effects of policymakers learning about the behavior of inflation (Sargent 1999, Cho, Williams, and Sargent 2002, Primiceri 2006), as well as the supply-side of the economy, such as the natural rate of unemployment,(Orphanides and Williams 2005) and the level (and growth) of potential output (Bullard and Eusepi 2005). While we do not dispute that learning about the macroeconomy influences policy, our findings suggest that it does not generate short-run dynamics, which is not to say that it is not important over longer time periods.

The theoretical literature on policy-motivated signalling in monetary policy builds on the work of Barro and Gordon (1983a) and Barro and Gordon (1983b), who were

[2] In 2002 Michael Howard, then Treasury Spokesperson for the Conservative Party in the UK, argued that external members' terms should be lengthened from three years to four using the following logic: "(since) the first year is spent learning how the MPC works, the present arrangements leave only two years in which independent members can make an active contribution" (Select Committee on Economic Affairs 2003). While we find that first year behavior is indeed different from later behavior, our results suggest that Howard's intuition for why was incorrect.

[3] This literature also explores learning by households about the monetary policy regime; see, for example, Erceg and Levin (2003).

among the first to establish the importance for policymakers of establishing credibility as inflation fighters to keep actual inflation in check. These papers show that such credibility can emerge through an infinitely repeated game. Backus and Driffill (1985a), Backus and Driffill (1985b), and Vickers (1986), all borrowing from the signalling models of Kreps and Wilson (1982) and Milgrom and Roberts (1982), argue that credibility can also emerge in a finitely repeated game with type uncertainty.[4] In their equilibria, policymakers signal their toughness on inflation to the public early in their careers, but gradually become less tough on inflation over time as the end of the game nears and the value of reputation declines. While signalling is well understood as a theoretical device, we assess its empirical plausibility.

Finally, our paper relates to the literature on career concerns in political economy. Maskin and Tirole (2004) study the role of preference signalling in a career concerns model in which politicians get payoffs from re-election. In their model, signalling induces politicians to choose politically correct actions more often than they otherwise would early in their careers. Besley and Coate (2003) and Leaver (2009) provide evidence that career concerns are important for determining the policy choices of public utility commissioners in the United States (regarding prices and the incidence of rate reviews, respectively). Our paper makes the counter point that signalling for a policy motivate is, in some settings at least, potentially more important than signalling for a career motive.

2 The Basic Dynamic Fact

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) whom the Bank encourages to vote independently. 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.[5]

Throughout the paper, we analyze the MPC voting record up to March 2009, when the interest rate reached its effective zero lower bound and a period of quantitative easing began; from this time, the main MPC decision concerned the additional policy of how many assets purchases to make. This sample yields a total of 142 meetings, and 1246 individual votes.[6] All voting data is available from the Bank of England website. The

[4]Sibert (2003) uses a similar analysis to study the role of committees in monetary policy; our paper does not account for the influence of other committee members on an individual's vote, and instead treats each voter as an independent decision maker.

[5]For more institutional details, see Hansen and McMahon (2011).

[6]We drop the emergency meeting of September 2001.

variable vote_{it} gives the desired change in interest rate of member i in period t. So, for example, if a member voted for no change, $\text{vote}_{it} = 0$, and if he voted for an increase of 25 basis points, $\text{vote}_{it} = 0.25$. Most members serve three-year terms (36 meetings) although some members have served longer (the current Governor, Mervyn King, is present in all 142 of our meetings), and others have served less than 36 meetings.

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

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

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

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

$$\text{vote}_{it} = \alpha_i + \gamma \text{D(Experienced)}_{it} + \delta_t + \epsilon_{it} \tag{2}$$

This equation includes both member and time fixed effects (α_i, the member fixed effect, captures a member specific intercept while δ_t, the time fixed effect, captures the average vote in period t). The results, reported in column (1) of table 1, show that as members serve more time, they vote for lower interest rates on average.[7] Because we have included member fixed effects in (2), this does not reflect changing composition of the committee, but rather indicates that something at the individual level systematically shifts over time. The rest of the paper is concerned with what this "something" actually is.

We can also show that the reduction in average interest rates with experience at the individual level is robust to alternative definitions of experience. We create two alternative dummy variables called $\text{D(Experienced - 12M)}_{it}$ and $\text{D(Experienced - 24M)}_{it}$ along the lines of equation 1, except that these measure experience as any tenure over 12 and 24 meetings, respectively. The results with these alternative definitions are reported in columns (2) and (3) of table 1; again, we find that experienced members vote for lower rates on average. Since all three dummy variables give the same qualitative results, we will use the 18 meeting definition (as in equation (1)) simply because it represents half of a standard MPC member's term and therefore splits the sample into subsamples of roughly similar size.

[7] This fact was previously established in Hansen and McMahon (2008).

Table 1: MPC voting and the impact of experience

	(1) Vote	(2) Vote	(3) Vote
D(Experienced)	-0.024*		
	[0.067]		
D(Experienced - 12M)		-0.019*	
		[0.093]	
D(Experienced - 24M)			-0.016*
			[0.092]
Constant	0.24***	0.24***	0.24***
	[0.000]	[0.000]	[0.000]
R-squared	0.903	0.903	0.903
Number of members	27	27	27
Model	Panel OLS	Panel OLS	Panel OLS
Member effects?	FE	FE	FE
Time effects?	YES	YES	YES

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

3 Distinguishing Changes in Inflation Toughness from Changes in Confidence

To understand why members' voting behavior changes over time, one must first consider the determinants of interest rate decisions and how to empirically measure them. The first major factor that we consider is how tough a given member is on inflation (or, in the language of the monetary literature, his "hawkishness"). The second is his private forecast of inflation conditions, and in particular the weight he puts on it (or, as we will term it, "confidence"). Hansen and McMahon (2011) develop an empirical methodology to separately identify these, and show that each is important for explaining the MPC voting record. This section first briefly reviews this approach, and then uses it to examine the extent to which changes in each contribute to explaining the basic dynamic fact established in the previous section.

3.1 Empirical methodology

The first step in separating hawkishness from confidence is to construct a simple theoretical voting model. In the overwhelming majority of meetings (135 of 142), all members either vote for the same rate, or one of two interest rates. As such, member i's vote in period t can be modelled as a choice from $v_{it} \in \{0, 1\}$, where 0 corresponds to the lower of two rates and 1 to the higher. Voting for the higher rate should be interpreted as decreasing expected future inflation. In period t an unknown state variable $\omega_t \in \{0, 1\}$

is realized, with $\omega_t = 0$ corresponding to a "low-inflation state" and $\omega_t = 1$ to a "high-inflation state". Because of the one-member, one-vote ethos of the MPC, and the relative unimportance of strategic voting effects under majority rule (Goeree and Yariv 2010), one can model member i's period t vote as a Bayesian decision problem in which she chooses the high rate ($v_{it} = 1$) if and only if

$$\widehat{\omega}_{it} \geq 1 - \theta_i \tag{3}$$

where $\widehat{\omega}_{it}$ is her belief that the economy is in the high-inflation state and $\theta_i \in (0, 1)$ is the threshold that this belief must reach, or the "burden-of-proof" needed,[8] to justify voting for the high rate. θ_i essentially captures how tough member i is on inflation: the larger it is, the more evidence she needs of high inflation to vote high. We will refer to a member with a higher (lower) θ parameter than another as more *hawkish* (*dovish*). The voting rule in (3) can arise from models in which θ_i represents preferences over the rate of inflation,[9] or member-specific aversions to type I and type II errors of inflation outcomes away from the inflation target; alternatively, it can arise from models in which θ_i represents a belief about expected future inflation, or a belief about some structural parameter of the macroeconomy like the supply gap or natural rate of interest.

Confidence instead relates to the formation of $\widehat{\omega}_{it}$. Let q_t be a common public prior belief on the probability that the high-inflation state $\omega_t = 1$ has arisen. In addition, suppose that member i receives a private signal $s_{it} \sim N(\omega_t, \sigma_i^2)$ (where these signals are i.i.d. across members and periods). Here σ_i^2 literally represents the precision of the private signal, but behaviorally speaking it measures the extent to which the member is swayed by the conventional wisdom as embodied in q_t in his assessment of inflation conditions; a straightforward implication of Bayes' rule is that a member with a lower σ_i^2 will put more weight on his private signal and less on the public prior. Thus, σ_i^2 measures the extent to which member i is willing to ignore public forecasts; this could either be because he has additional information on inflation shocks, or another factor, such as a desire to stand out.

Because the normal distribution satisfies the monotone likelihood ratio property, the voting rule in (3) is equivalent to a cutoff rule with respect to the signal: member i chooses $v_{it} = 1$ if and only if $s_{it} \geq s_{it}^* (\theta_i, \sigma_i, q_t)$, the precise expression for which is presented in appendix A as equation (A.1).

Heterogeneity in the θ and σ parameters is separately identifiable given the prior q_t. To understand why, it is useful to consider figure 1, which plots the probability that a member will choose the higher rate ($v_{it} = 1$) as a function of q_t. A member who is tougher

[8]This is the terminology used by Feddersen and Pesendorfer (1998), who use a variant of this setup to model voting behavior on juries.

[9]Section B.1 provides such a derivation.

on inflation than another will vote for high rates more often, independently of q_t. This is represented by the horizontal shift in the curve in figure 1a. By contrast, a member with a higher σ_i than another will be more influenced by the value of the prior. So, when q_t is low (high) and the public view is that rates should be low (high), the high-σ member will vote for low (high) rates more often. This is represented by the rotation of the curve in figure 1b.

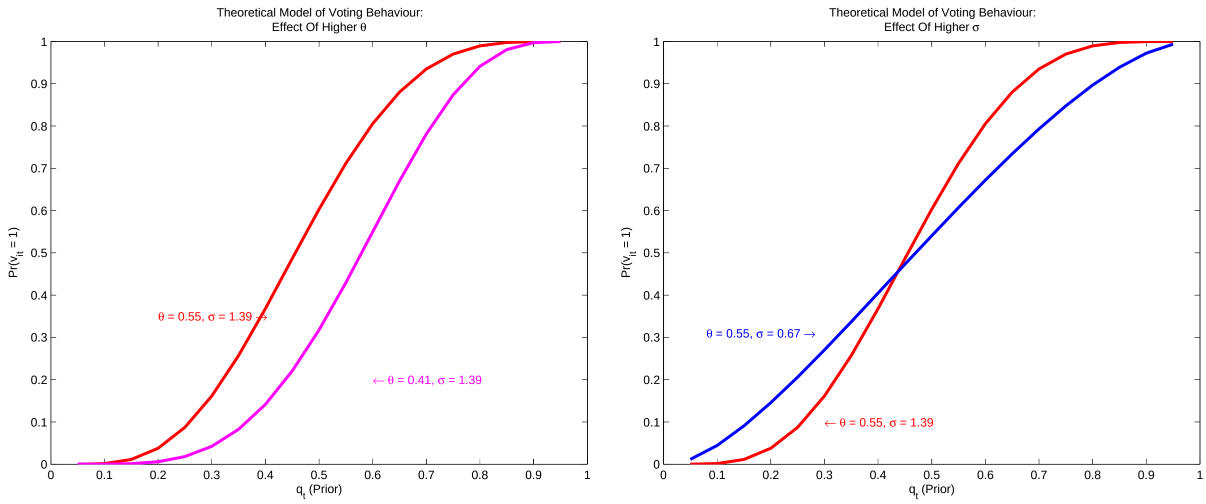

(a) Change in Inflation Toughness

(b) Change in Signal Distribution

Figure 1: Distinguishing Information and Preferences

Notes: These figures show the theoretical probability that a member votes for the high interest rate ($\Pr(v_{it} = 1)$) in a meeting as a function of the prior belief that the economy is in an inflationary state (q_t). The different curves represent different combinations of confidence and hawkishness to show how the probability of voting high changes.

To actually use this model as an empirical tool requires knowing the public prior and v_{it}; although votes are observable and we can construct a high and low vote in any non-unanimous meeting, the rate which corresponds to the high rate (and so $v_{it} = 1$) is unknown in unanimous meetings since the MPC voting record does not indicate the other vote under consideration. In order to construct a proxy $\widehat{v}_{it}$ for v_{it}, Hansen and McMahon (2011) use a monthly Reuters survey of City of London economists in which respondents write probability distributions over voting outcomes at the next MPC meeting. In unanimous meeting, the two rates under consideration are taken as the two rates on which the market places highest probability.[10] In periods with two observed votes, $\widehat{v}_{it}$ is set equal to 1 if and only if member i votes for the higher observed rate in period t.

In order to construct a proxy for the prior, Hansen and McMahon (2011) combine

[10] Incidentally, the observed rate in the voting record always lies in this set.

the Reuters data with options price data and the voting record of a shadow MPC, and extract a single common factor $\widehat{q}_t$ that is used as a proxy of q_t.[11] The main concern with this proxy is that two of the three inputs (the Reuters survey and the options price data) are predictions of what the MPC will do, not necessarily beliefs about inflationary conditions. Hansen and McMahon (2011) explore this issue in some detail, and conclude that the common factor analysis appears to successfully purge any prediction bias.[12]

One can then plug the $\widehat{q}_t$ and $\widehat{v}_{it}$ proxies into the likelihood function (derived in appendix A) and structurally estimate θ and σ parameters for different subgroups of MPC members.[13] For the purposes of this paper, the two subgroups under consideration are the population of new and experienced members, as defined in the previous section.

3.2 Application to voting dynamics

Primae facie it is unclear which parameters evolve with time. One can well imagine that confidence levels vary over tenure on the committee. For example, learning by doing might increase members' ability to perceive economic conditions, which would decrease the estimated value of σ_i for experienced members. On the other hand, members may be overconfident in their initial views and listen more to public opinion over time, which would increase the estimated value of σ_i. At the same time, they may adjust how tough they are on inflation for a variety of reasons (explored in detail in later sections), which would lead to changes in the estimated value of θ for new and experienced members.

Before proceeding to structural estimation, we first employ a reduced form approach to attempt to shed light on whether θ or σ—or both—evolve over time. To do this we introduce our proxy $\widehat{q}_t$ into a probit regression in which the left hand side variable is the proxy $\widehat{v}_{it}$:

$$\Pr\left(\widehat{v}_{it} = 1\right) = \alpha_0 + \alpha_1 \mathrm{D(Experienced)}_{it} + \alpha_2 \widehat{q}_t + \alpha_3 \mathrm{D(Experienced)}_{it} \times \widehat{q}_t + \epsilon_{it}. \quad (4)$$

One can use the parameter estimates from equation (4)—reported in table 2—to predict the probability that new and experienced members vote for high rates for different value of $\widehat{q}_t$. Figure 2 plots these predicted probabilities, which are reduced-form analogues

[11] This common factor correlates with actual votes as would the theoretical prior: it strongly and positively predicts the probability a member votes high and has a concave relationship with the probability that a member dissents from the majority.

[12] In fact, the data from the shadow MPC, which we describe and use in section 6.3, only go back to 2002. One can also construct a second proxy for q_t based on just the Reuters and options data to cover the entire sample. This proxy also appears to purge prediction bias, and we use it in all subsequent regressions.

[13] A potential concern is that within-group heterogeneity contaminates the estimates, but Monte Carlo simulations presented in Hansen and McMahon (2011) indicate that the estimated θ and σ parameters are unbiased estimators of the average θ and σ of the group members. In some tests below we will use individual estimates, although we cannot obtain them for all members.

of the theoretical probabilities in figure 1. Two notable features stand out. First, the probability that experienced members vote for high rates is lower than the probability that new members do for almost all values of the prior. This indicates that as members gain experience, they become more dovish (i.e. their θ falls). Second, the slopes of the curves in figure 2 are similar, indicating that confidence remains more or less constant over time (i.e. σ does not change with experience).

Table 2: Estimated coefficients for equation (4)

	(1) D($\widehat{v}_{it} = 1$)
D(Experienced)	-0.99***
	(0.000)
$\widehat{q}_t$	3.30***
	(0.000)
D(Experienced) × $\widehat{q}_t$	1.29***
	(0.005)
Constant	-1.39***
	(0.000)

Robust p-values in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Notes: This table shows the estimated coefficients for equation (4) using $\widehat{v}_{it}$ as the dependent (dummy) variable and D(Experienced) (as defined in (1)) and $\widehat{q}_t$, the proxy for the common prior, as explanatory variables.

Of course, these remarks are based only on informal observation rather than on statistical testing. To formalize the conclusions, we structurally estimate θ and σ parameters for new and experienced members using maximum likelihood estimation on the likelihood function generated by our model (the details of which are in appendix A). The results are presented in table 3, and are consistent with the intuitions from figure 2. First, there is a strong downward shift in the θ parameter (from 0.57 to 0.48) that is statistically significant at the 1% level. Second, there is no movement, statistically speaking, in the σ parameter. This result is of independent interest, since it suggests that members neither accumulate additional expertise with experience nor adjust the weight they attach to public beliefs for any other reason. Instead, voting dynamics are driven entirely by a shift in the average member's hawkishness. One can use the estimated parameters to predict the probability that new and experienced members vote for high rates; these are plotted in figure 3. These curves are quite similar to the reduced form estimates in figure 2, which provides reassurance that the distributional assumptions behind our structural estimates are not imposing unreasonable patterns on the raw data.

The goal of this section was to decompose the observed shift in MPC voting behavior from section 2 into a term representing hawkishness and a term representing confidence, and we have found that the former explains nearly everything. For the rest of the paper,

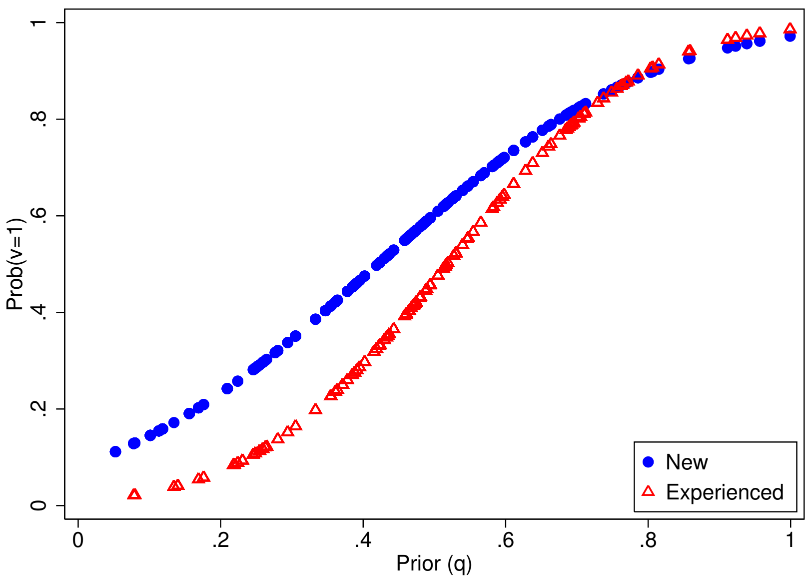

Figure 2: Predicted probabilities of voting high derived from equation (4)

Notes: This figure shows the empirical probability that a member votes for the high interest rate ($\Pr(\widehat{v}_{it} = 1)$) in a meeting as a function of the proxy for the prior belief that the economy is in an inflationary state ($\widehat{q}_t$). The estimated probabilities for new and experienced members are derived from the estimates of equation (4) presented in table 2. New members are more likely to vote for the high interest rate for almost all values of the prior.

we will use the term *experience effect* to refer to the difference in the estimated θ parameters for experienced and new members. So, for example, this section has shown that, on average, the experience effect is negative. While the results of this section get us some way towards understanding why experienced members vote for lower rates, the question remains: why does the average member get softer on inflation over time?

4 Evaluating Learning

The first explanation we propose for the experience effect, inspired by the macro literature on model uncertainty, is that θ is not a fixed preference parameter, but an uncertain structural parameter of the macroeconomy, the average belief about which changes over time. Since the properties of Bayesian learning are well understood, we will not formalize or explicitly solve a learning model. Instead we will informally discuss what we mean by learning, and sketch out the properties we believe would be consistent with it.

Consider the simplest possible dynamic model of MPC voting, in which member i votes in two periods — $t = 1$ and $t = 2$. Let θ_1^i be his prior belief on the value of θ in period 1. Now suppose that some new information arrives (such as actual inflation data after the implementation of the period 1 interest rate) the realization of which is

Table 3: Structural estimates of the effect of experience on MPC voting

	(1) θ	(2) σ
New	0.57***	0.97***
	[0.53 - 0.62]	[0.69 - 1.24]
Experienced	0.48***	1.04***
	[0.45 - 0.51]	[0.84 - 1.24]
Difference	-0.090***	0.073
	[-0.14 - -0.037]	[-0.27 - 0.42]

Clustered, by member, confidence interval in brackets
*** p<0.01, ** p<0.05, * p<0.1

Notes: This table shows structural estimates of θ and σ for new and experienced members. We model θ_i and σ_i in the likelihood function in appendix A as $\theta_i = \alpha_{0\theta} + \alpha_{1\theta}\mathrm{D}(\mathrm{Experienced})_{it}$ and $\sigma_i = \alpha_{0\sigma} + \alpha_{1\sigma}\mathrm{D}(\mathrm{Experienced})_{it}$, respectively, and use maximum likelihood to estimate the α parameters. The last row shows the difference of row (2) less row (1).

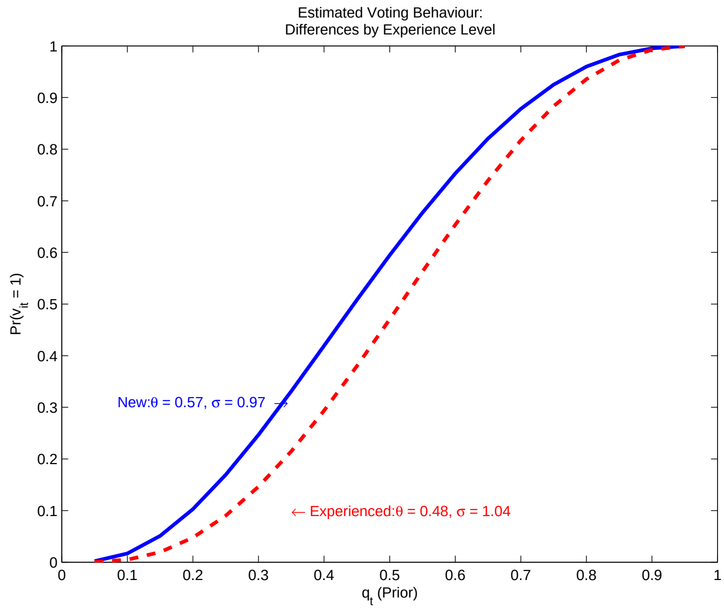

Figure 3: Predicted probability functions based on structural estimates of θ and σ

Notes: This figure plots the predicted probability that a member votes for the high interest rate ($\Pr(v_{it} = 1)$) as a function of the prior belief that the economy is in an inflationary state (q_t) under the assumption that our model, with θ and σ for new and experienced members set equal to their estimated values in table 2, generates the data.

correlated with the true value of θ. This will cause him to update his belief to some new value θ_2^i in period 2. While beliefs formed using Bayes' Rule are martingales, the average value of θ_2^i across members could be lower than the average value of θ_1^i if the initial average belief on θ were biased upwards, or if θ itself were a concave function of some uncertain parameter about which learning occurred, in which case beliefs about θ would form a supermartingale.

The empirical predictions of learning depend on whether it is common (all members receive the same information about θ) or private (members receive private signals on θ). Due to how we conceive of the θ parameter, we believe the most natural assumption is common learning. We model learning about ω_t as in part private because we think of it as a frequent, transitory shock whose magnitude is not immediately discernible from available public data. In contrast, learning about a structural parameter like the supply gap generally occurs infrequently and only after a long series of observations of historical interest rate decisions and economic outcomes, both of which are public information.

One way to test for common learning is to see if new members consistently have a higher estimated θ than experienced members who are serving at the same time. It is consistent with common learning for all members to share the same belief on θ which gradually drifts down over time. By contrast, it is not consistent with common learning for new members to systematically anchor their beliefs on θ above concurrently serving experienced members, because this indicates that new members are not using the information of experienced members to adjust their estimates of the θ parameter.

While we do not have enough data to estimate the difference between new and experienced members for each year of the sample, we can break down the sample into smaller time windows and estimate the experience effect in each of them. To do this, we estimate the levels of θ and σ for new and experienced members just as in table 3, but use rolling five-year windows instead of the entire sample. Specifically, we first estimate them using the first 5 years of the MPC from June 1997 to May 2002, and store the saved results as May 2002 (the end of the window). We then advance the window forward by one month (i.e. one meeting) and repeat the estimates for the period from July 1997 until June 2002. As this procedure progresses, for each month until March 2009, meetings enter and leave the sample such that every meeting enters in at least one five-year window, and most enter many times. The resulting estimates are shown in figure 4.[14]

The top of the figure traces out the estimated θ for new (green dash) and experienced (red solid) members, while the bottom measures their difference (experienced minus new) and plots the 95% confidence intervals around it. As one can see, new members consistently have a higher estimated θ than experienced ones. In those windows towards the end of the sample period the gap closes and, although the the difference remains nega-

[14] We do not show the σ estimates, but in all cases our estimates of θ control for σ.

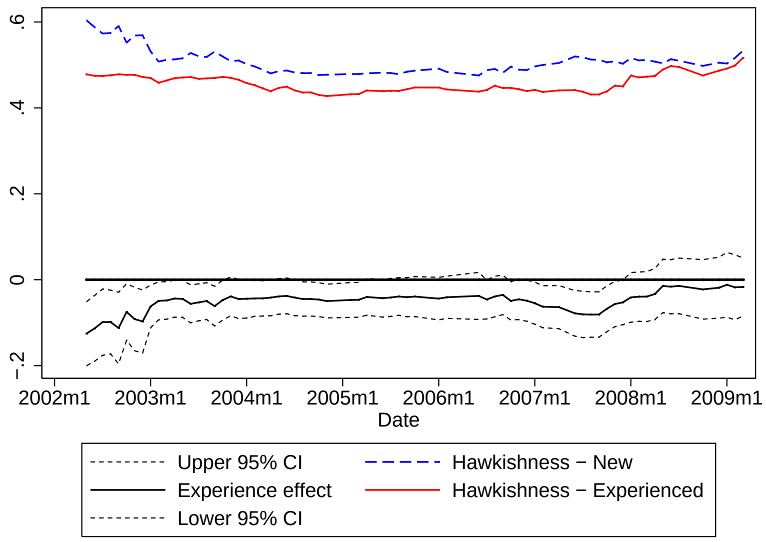

Figure 4: Rolling window estimates of θ for new and experienced members

Notes: This figure shows the structural estimates of θ for new and experienced members, as well as their difference. The estimates are calculated in the same way as the estimates reported in table 3 except that in this figure we use a rolling 5-year window to derive the estimates. This means that the estimate for January 2005 (2005m1 on the figure) is based on a sample of data between February 2000 and January 2005. The first estimate, reported as 2002m5, is based a window of data in the window from June 1997 to May 2002; the last estimate is based on data between April 2004 and March 2009.

tive, it is no longer statistically different from zero at reasonable significance levels; this is perhaps related to the onset of the global financial crisis. Moreover, the estimated θ for experienced members appears fairly constant over the time. In a learning model, this means that, after 18 months, members consistently form the same estimate of θ. If learning about θ were common, one would expect experienced members to simply tell new members what the "correct" belief was, or for new members to work out for themselves what the correct θ was from looking at the history of previous voters. The bottom line here is that a common learning explanation for the experience effect appears unlikely.

One can also evaluate a private learning explanation for the experience effect. A model of private learning would generate two empirical predictions that we can explore with the data. First, while individual beliefs can *on average* decline over time, there should be at least some individuals for whom the belief on θ goes up (or, to use the notation above, some members i for whom $\theta_i^2 > \theta_i^1$). The most direct way to test this prediction is to simply estimate the experience effect for each individual separately. In general, our estimator requires more data than is present in the voting record of just one member to converge and yield precise estimates. Nevertheless, we run it on all 27 members separately, with results presented in table 4. As one can see, only 17 of the

27 regressions converge (those that do not converge are indicated with ·), and of those that do converge, the estimates of the experience effect are not terribly precise. Still, the results illustrate the predominant pattern of declining hawkishness. In 12 of the 17 cases that converge, the effect is negative, while for the remaining 5 members, two show no change and three have a positive effect. The only experience effects that are measured as significant at the 10% level are negative, of which there are five.

Table 4: Experience Effect by individual member

Member	θ^{New}	θ^{Exp}	Experience effect	p-value
Allsop	0.41	0.27	-0.15	0.054
Barker	0.42	0.49	0.08	0.251
Bean	0.50	0.47	-0.03	0.752
Bell	0.40	0.35	-0.05	0.475
Besley	.	.	.	.
Blanchflower	.	.	.	.
Budd	.	.	.	.
Buiter	.	.	.	.
Clementi	0.62	0.51	-0.11	0.030
Dale	.	.	.	.
Davies	.	.	.	.
Fisher	.	.	.	.
George	0.63	0.52	-0.11	0.046
Gieve	0.54	0.54	-0.00	0.997
Goodhart	0.81	0.54	-0.26	0.227
Julius	0.55	0.32	-0.23	0.067
King	0.81	0.57	-0.24	0.262
Lambert	.	.	.	.
Large	0.66	0.66	0.00	0.987
Lomax	0.44	0.49	0.05	0.555
Nickell	0.64	0.33	-0.31	0.138
Plenderleith	0.63	0.52	-0.11	0.042
Sentence	0.67	0.76	0.09	0.746
Tucker	0.58	0.56	-0.02	0.835
Vickers	.	.	.	.
Wadhwani	0.45	0.20	-0.25	0.002
Walton	.	.	.	.

Notes: This table shows structural estimates of the hawkishness parameter (θ) for individual members when new and experienced. We model θ_i and σ_i in the likelihood function in appendix A as $\theta_i = \alpha_{i\theta} + \alpha_{i\theta}\mathrm{D(Experienced)}_{it}$ and $\sigma_i = \alpha_{i\sigma} + \alpha_{i\sigma}\mathrm{D(Experienced)}_{it}$, respectively, and use maximum likelihood to estimate the α parameters. Columns (1) and (2) reports the appropriate linear combinations of the α parameters. Column (3) reports the difference of columns (2) and (1), and column (4) the p-value for the null hypothesis that the experience effect is different from zero. A "·" rather than an entry signifies that the corresponding regression did not converge, or converged without producing meaningful estimates.

In order to obtain more precise results of the experience effect at the individual level, we measure it for different randomly selected subgroups, which we refer to as *pseudo-committees*. More specifically, we randomly draw nine members from our sample 4,000 times (with replacement), and estimate the experience effect for each pseudo-committee. Figure 5 displays a kernel density of the estimated experience effects and table 5 tabulates the estimates according to their magnitude and estimated significance. In 98% of cases we estimate a negative experience effect, and in 38% of cases this estimate is significant; in no case do we find a positive experience effect at a $p - value$ less than 0.5. This evidence, combined with that in table 4, suggests that upward movements in θ are relatively rare on the MPC, and that the overall average fall in θ is actually composed of falls in θ of various magnitude at the individual level. A private learning model in which there were only a small probability that a member responded to new information by increasing his belief would require some unusual assumptions on the environment and/or behavior.

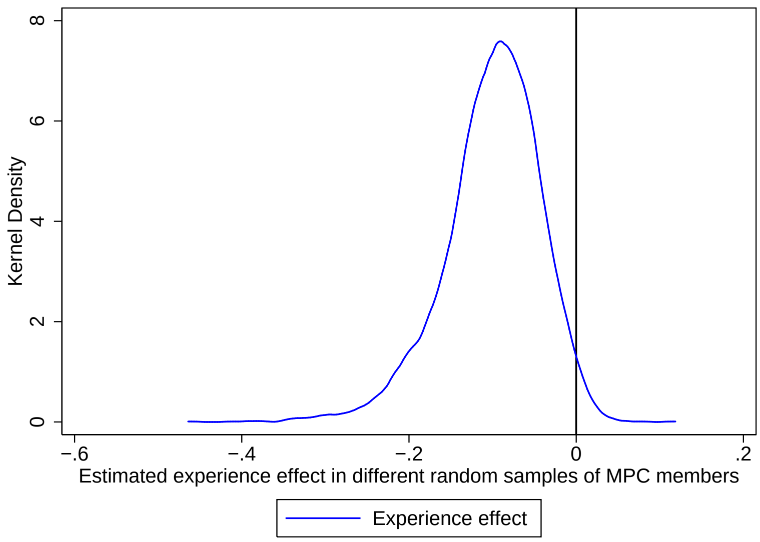

Figure 5: Kernel Density plots of estimated experience effect using random combinations of members

Notes: We construct 4,000 randomly selected pseudo-committees of 9 members. For each pseudo-committee, we model θ_i and σ_i as described in table 3 and estimate θ for new and experienced members. We compute the experience effect as the difference in these estimates. This figure plots the kernel density of the 4,000 experience effect estimates. We use the Epanechnikov kernel, and set the half-width to the value that would minimize the mean integrated squared error if the underlying distribution were Gaussian.

The second prediction of private learning is that the variance of the distribution of individual beliefs should be higher for experienced members (the variance of θ_2^i should be higher than the variance of θ_1^i). Conditional on θ_1^i, the realization of θ_2^i is a random variable, and, in a private learning model, the individual values of θ_2^i are not perfectly correlated. As such, at the population level, the dispersion in θ_2^i should be higher than the dispersion in θ_1^i. In order to determine whether this is the case, one can simply compare the kernel densities for new and experienced members that the above 4,000 replications generate. These are displayed in figure 6. In fact, the dispersion in the population of experienced members is *less* than the dispersion in the new population (the variance of

Table 5: The experience effect across 4,000 random pseudo-committees

Item	Number	Per cent
Negative, p-value < 0.01	337	8
Negative, 0.05> p-value > 0.01	633	16
Negative, 0.1 > p-value>0.05	544	14
Negative, 0.5 > p-value > 0.1	1,955	49
Negative, p-value > 0.5	450	11
Positive, p-value > 0.5	78	2
Positive, 0.5 > p-value > 0.1	3	0
Total	4,000	100

Notes: We construct 4,000 randomly selected pseudo-committees of 9 members. For each pseudo-committee, we model θ_i and σ_i as described in table 3 and estimate θ for new and experienced members. We compute the experience effect as the difference in these estimates. This table breaks the 4,000 estimates into groups according their signs and p-values for the statistical test of difference from zero.

the former is 0.00340 and of the latter 0.00376). Thus, apart from a small portion of pseudo-committees in which experienced members begin voting for much lower rates (the skewness of the experienced population is -1.163, while for the new it is 0.176), individual estimates of θ become more concentrated over time.

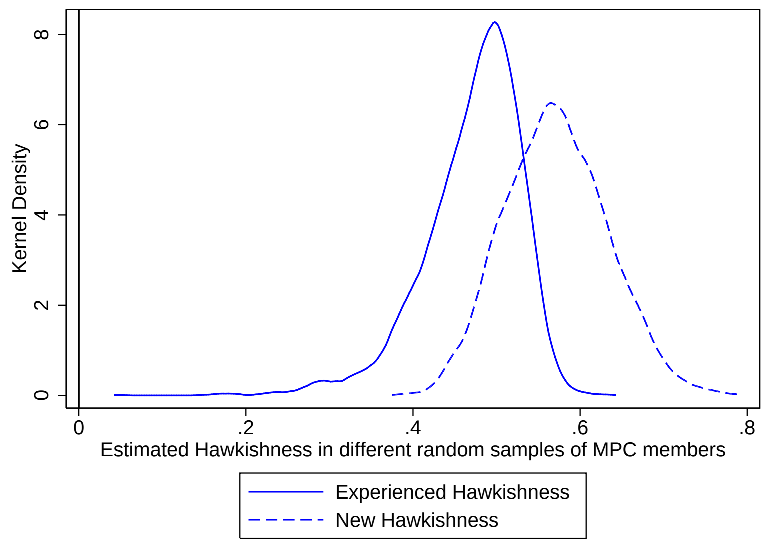

Figure 6: Kernel Density plots of estimated θ for new and experienced members using random combinations of members

Notes: We construct 4,000 randomly selected pseudo-committees of 9 members. For each pseudo-committee, we model θ_i and σ_i as described in table 3 and estimate θ for new and experienced members. This figure plots the kernel densities of the 4,000 new θ and experienced θ estimates. We use the Epanechnikov kernel, and set the half-width to the value that would minimize the mean integrated squared error if the underlying distribution were Gaussian.

In summary, we have shown that if a learning model explains the experience effect, it would have to be a private learning model. We have then shown that θ estimates generally declines at the individual level, not just the overall population level, and that the dispersion individual θ estimates does not increase over time. While we do not rule out the possibility that MPC members are uncertain about some structural parameters in the economy about which they learn, we do claim that the weight of evidence presented in this section points away from learning as of first-order importance for explaining the experience effect. The next section presents an alternative idea—signalling—that seems to better account for the results of this section.

5 Signalling and Policy Dynamics

Another possibility is that θ_i is a fixed preference parameter known only to member i, and, in line with the papers cited in section 1.1, that new monetary policymakers want to signal that they are hawkish in order to anchor inflation expectations. We formulate a simple signalling model consistent with these ideas to illustrate how signalling generates dynamics,[15] and show the predicted dynamics are consistent with the evidence presented in the previous section. We then test an additional prediction of this signalling model: the experience effect should be independent of whether a member is innately hawkish or dovish. The voting record of the MPC is broadly consistent with this.

5.1 Model

Suppose member i votes for two periods $t = 1$ and $t = 2$ and that $\theta_i \in \{\theta_L, \theta_H\}$, where $\theta_H > \theta_L$ is a fixed constant known by member i but not an evaluator, who one can think of as a representative member of the public. The evaluator instead places probability r on $\theta_i = \theta_H$. We will refer to type θ_L as a Dove and type θ_H as a Hawk. Let $R(v_{i1})$ be the evaluator's updated belief that member i is a Hawk after observing the first period vote v_{i1}, and $\Delta = R(1) - R(0)$ be the difference in this belief when the evaluator observes a high vote rather than a low vote in the first period. In other words, Δ is the net reputational reward to voting high in the first period. To introduce the idea that members care about their reputation, we assume that v_{i1} is chosen if and only if

$$\widehat{\omega}_{i1} \geq 1 - \theta_i - \beta\Delta \tag{5}$$

[15]In our voting model, a policymaker chooses a discrete (one of two) interest rate, whereas in the monetary literature he chooses a continuous inflation rate. This means we must make a different equilibrium construction compared with other signalling papers in the monetary literature, but we do not claim the underlying economic intuition is original.

while v_{i2} is chosen if an only if

$$\widehat{\omega}_{i2} \geq 1 - \theta_i; \tag{6}$$

that is, we modify the voting rule in (3) by introducing reputational concerns in the first period but not the second. This captures the fact that, at the end of voters' tenures, there is no future period in which reputation pays off, so that only the current policy payoff of a particular interest rate matters. Here β measures the strength of the signalling incentive in the first period: when $\beta = 0$ the model is the standard one, and, as β increases away from 0, the voter cares increasingly about how outsiders perceive his hawkishness. Appendix B.1 explicitly derives the voting rule in (5) from a model in which the member places a weight λ on the policy payoff of his vote (a quadratic loss function in the realized inflation rate about the inflation target, with the Hawk using a higher target than the Dove) and a weight $1 - \lambda$ on the evaluator's belief he is tough on inflation.

We use Perfect Bayesian Equilibrium as the solution concept for the model. One possibility is that β takes on sufficiently high values to lead members to always vote for high rates in equilibrium; however, the MPC voting record indicates that such equilibria are not reasonable: both new and experienced members sometimes vote for high and sometimes for low rates. We call a *responsive equilibrium* one in which both types select $v_{it} = 0$ and $v_{it} = 1$ with positive probability; i.e., the member is responsive to the realization of his signal whether he is a Hawk or Dove. The following proposition establishes the existence of such equilibria, along with their main empirical implication. The proof is in appendix B.2.

Proposition 1 *Responsive equilibria have the following properties:*

1. *For all $\beta \in (0, \theta_L)$ a responsive equilibrium exists.*

2. *The equilibrium value of Δ in a responsive equilibrium is positive.*

3. *Let $\Delta(\beta)$ denote the equilibrium value of Δ given β. Then there exists an $\alpha^* > 0$ such that $\alpha\beta\Delta(\alpha\beta) < \beta\Delta(\beta)$ for all $\alpha \in [0, \alpha^*]$.*

Although simple, proposition 1 immediately provides a theoretical rationale for the average experience effect uncovered in section 3. Suppose that members are either Hawks or Doves; that, in their first 18 months, they wish to signal their hawkishness; and that experienced members simply follow their innate policy preferences. Then, if new members are in a responsive equilibrium, our estimate of their average θ is an estimate of

$$r\theta_H + (1 - r)\theta_L + \beta\Delta \tag{7}$$

or the average value of θ in the population plus a reputation term that is positive. On the other hand, our estimate of the average θ for experienced members is

$$r\theta_H + (1-r)\theta_L. \tag{8}$$

The difference between (7) and (8) is $\beta\Delta$, which would explain the experience effect. In words, because new members care about establishing a reputation as Hawks, they vote as if they were more hawkish than they actually are,[16] while experienced members revert to their innate policy preferences, and vote for high rates less often.

It is not signalling *per se* that generates dynamics in the model, but the different weight that members attach to their reputation in periods 1 and 2. One might worry that the particularly extreme assumption that we have used in supposing that reputation has zero weight for experienced members is crucial. The third result in the proposition alleviates these: as long as the weight on reputation falls sufficiently far for experienced compared to new members, all of the results continue to hold.[17]

Proposition 1 also provides a rationale for the evidence uncovered in section 4. Consider the estimates we would recover if we estimated θ for an individual who was a Hawk. Then equations (7) and (8) would become

$$\theta_H + \beta\Delta \tag{9}$$

and

$$\theta_H. \tag{10}$$

Thus the experienced effect estimated at the individual level for Hawks would be $\beta\Delta$. But by the exact same logic, the experience effect at the individual level for Doves would also be $\beta\Delta$. This observation rationalizes the three main findings from the previous section. First, because all new members engage in signalling and experienced members do not, if one were to take a snapshot of the committee at any given point in time, new members would have a higher estimated θ than experienced ones. Second, if one were to look at the experience effects at the individual level, they would all fall. Third, the average distance between new and experienced members is the same: the distribution of estimated individual θ parameters in the new population is simply a mean-shift of the distribution in the experienced population. All three findings were, we argued, inconsistent with various forms of learning, while they are perfectly consistent with a signalling model.

This discussion not only rationalizes the previous empirical findings, but yields a new

[16]Note that even though in equilibrium the evaluator knows that this behavior is taking place, and adjusts his belief updating accordingly, he still increases (decreases) his assessment that member i is a Hawk after observing $v_{it} = 1$ ($v_{it} = 0$).

[17]Any reputation concerns that operate across all periods would not be separately identifiable from the average θ.

prediction: if signalling is relevant on the MPC, we would expect the experience effect to be type-independent. This is important to test because it provides some independent verification for the model's plausibility. It also contrasts with learning yet again. Suppose that θ_L and θ_H represented initial beliefs on an unknown θ parameter rather than preferences. Learning would predict a gradual convergence of these non-common priors over time (Kalai and Lehrer 1994). The prediction also contrasts with a story in which Hawks coexist with Doves in the first year before the former "give-in" to the latter and mimic their behavior.

5.2 Testing type-independence of the experience effect

The first step in testing the type-independence of the experience effect is to break the sample into Hawks and Doves. This requires ranking members on the basis of individual θ estimates. To generate this ranking we use three methods:[18]

1. Rank members on the basis of individual θ_i estimates, which we denote $\widehat{\theta_i}$. While this is the ideal measure, we are unfortunately only able to estimate individual θ values for 20 of the 27 members in our sample;[19] for the other members, our maximum likelihood estimator does not converge.

2. Reuse the 4,000 draws underlying figure 5 described in section 4, and, for each one, estimate the average overall θ on the pseudo-committee rather than the average θ for new and experienced members. Then rank members according to the average of the θ estimates on all pseudo-committees in which they are present, which we denote $\widehat{\theta}_i^{\text{Alt}}$. To see why this ranking is meaningful, consider some member A with associated θ_A and another member B with associated θ_B, where $\theta_A > \theta_B$. With a large numbers of replications, members A and B will on average be drawn with the representative committee member. So, the average of the θ estimates on the pseudo-committees on which A serves should be greater than the average of the θ estimates on which B serves.

3. Rank by total fraction of votes cast that are high.

Table 6 shows the value of each hawkishness metric for each member (the $\cdot$ represents a value we were unable to estimate). All three metrics produce rankings that are very close. The Spearman rank correlation between the rankings produced by methods 1 and

[18] Although the methods presented below rank the members based on their average θ, the groupings we arrive at are little changed, and there is no change in the qualitative results, if we instead used experienced θ. The advantage of using the average θ is that we can compare across the three rankings for more members.

[19] We are able to get individual estimates for three extra members compared with the estimates presented in table 4 because here we only require an average θ and not a separate θ for new and experienced.

2 (2 and 3) is 0.97 (0.93), which is highly significantly different from zero (ditto). The similarity can also be seen visually in figure 7, which contains two scatter plots of method 2 against methods 1 and 3. Since the ranking of method 2 is highly correlated with those of the other methods, and because it ranks 26 of the 27 members (Paul Fisher is excluded as he is present for only the March 2009 meeting) based on estimated θ values, we use it as our preferred measure. We create a dummy variable D(Hawk)_i equal one if $\widehat{\theta}_i^{\text{Alt}}$ is above the median value (0.509). Figure 7 codes members with $\text{D(Hawk)}_i = 1$ as blue and members with $\text{D(Hawk)}_i = 0$ as red.

Table 6: Measures to rank individual hawkishness

Member	$\widehat{\theta}_i$	$\widehat{\theta}_i^{\text{Alt}}$	% High Votes	D(Hawk)_i
Allsop	0.310	0.487	33.3	0
Barker	0.475	0.503	52.1	0
Bean	0.469	0.501	51.0	0
Bell	0.382	0.496	36.1	0
Besley	0.740	0.527	67.7	1
Blanchflower	0.124	0.481	2.9	0
Budd	.	0.519	66.7	1
Buiter	.	0.515	50.0	1
Clementi	0.539	0.514	48.3	1
Dale	.	0.508	44.4	0
Davies	.	0.512	100.0	1
Fisher	.	.	100.0	.
George	0.539	0.515	53.4	1
Gieve	0.539	0.509	48.6	0
Goodhart	0.732	0.523	52.8	1
Julius	0.313	0.494	20.0	0
King	0.602	0.537	61.3	1
Lambert	0.442	0.504	50.0	0
Large	0.671	0.526	77.5	1
Lomax	0.475	0.504	48.3	0
Nickell	0.388	0.492	44.4	0
Plenderleith	0.548	0.516	50.0	1
Sentence	0.701	0.525	66.7	1
Tucker	0.584	0.522	59.8	1
Vickers	.	0.518	60.7	1
Wadhwani	0.202	0.477	13.9	0
Walton	.	0.509	75.0	0

Notes: We construct three measures of individual hawkishness. $\widehat{\theta}_i$ are direct structural estimates; $\widehat{\theta}_i^{\text{Alt}}$ is the average of the θ estimates on all pseudo-committees on which individual i served; and "% High Vote" corresponds to the fraction of all the members' votes for which $\widehat{v}_{it} = 1$. D(Hawk)_i is a dummy variable which equals 1 whenever $\widehat{\theta}_i^{\text{Alt}}$ is above the median for all 26 members for whom we can estimate it—we do not have an estimate for Paul Fisher as he serves only 1 meeting in our sample. Missing values are represented by ".".

Broadly speaking, members with $\text{D(Hawk)}_i = 1$ correspond to Hawks in our theoretical model, and those with $\text{D(Hawk)}_i = 0$ correspond to Doves. In order to test how the experience effect varies across these types, we now structurally estimate θ for new and old members separately by D(Hawk)_i, with results presented in table 7. The test of whether

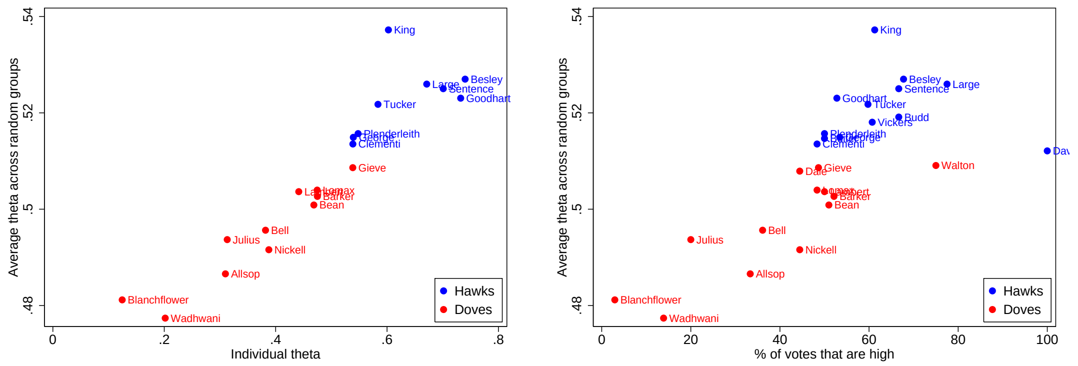

(a) Relationship between two measures of individual hawkishness

(b) Relationship between hawkishness and % of votes that are high

Figure 7: Ranking individual MPC members based on individual hawkishness

Notes: These figures show the relationship among different ways to measure individual hawkishness. $\widehat{\theta}_i$ are direct structural estimates; $\widehat{\theta}_i^{\text{Alt}}$ is the average of the θ estimates on all pseudo-committees on which the individual i served; and "% High Vote" corresponds to the fraction of all the members' votes for which $\widehat{v}_{it} = 1$. Both figures plot $\widehat{\theta}_i^{\text{Alt}}$ on the vertical axis; figure 7a plots $\widehat{\theta}_i$ on the horizontal axis and figure 7b plots "% High Vote" on the horizontal axis. An individual is a Hawk if $\widehat{\theta}_i^{\text{Alt}} \geq 0.509$ and a Dove otherwise.

the experience effect differs across the types is whether the quantity in the lower-right corner of the table is different than zero. This is not rejected at standard levels of statistical significance (and would not be even at the 35% significance level). Moving away from statistical testing and simply looking at the magnitudes of the estimates does show some convergence between Hawks and Doves when new and experienced: the difference in the estimated θ values is 0.24 when new and 0.19 when experienced. Our message is not that learning does not take place at all on the MPC, but that the dynamics predicted by a signalling model better explain the experience effect than the dynamics predicted by a learning model. Table 7 shows two groups of voters, one more hawkish than another, in disagreement both during the beginning of their tenures and at the end, with a tendency to vote for lower rates over time but little tendency towards convergence. We interpret this pattern as more consistent with signalling than learning.

A second way of testing symmetry is to repeat the kernel density analysis from section 4, but on Hawks and Doves. We again carry out 4,000 random draws from the 27 MPC voters, but now we draw 7 Hawks and 7 Doves (with replacement), compute the experience effect for each type separately, and plot the kernel densities of the estimates in figure 8. The vertical dashed lines in the figure mark the overall group experience effects measured in table 7; as one can see, the pseudo-committee estimates are concentrated on these group averages. Table 8 tabulates the estimates.

Table 7: The experience effect for Hawks and Doves

	(1) Dove	(2) Hawk	(3) Difference
New	0.44***	0.68***	-0.24***
	[0.39 - 0.48]	[0.63 - 0.74]	(0.000)
Experienced	0.37***	0.56***	-0.19***
	[0.27 - 0.47]	[0.53 - 0.59]	(0.000)
Difference	-0.066	-0.12***	0.054
	[-0.16 - 0.031]	[-0.18 - -0.062]	(0.351)

Clustered, by member, 95% confidence interval in brackets
*** p<0.01, ** p<0.05, * p<0.1

Notes: This table shows structural estimates of θ for new and experienced members according to whether they are classified as Hawks or Doves, as described in the text. We model θ and σ in the likelihood function in appendix A as

$$\theta = \alpha_{0\theta} + \alpha_{1\theta}\mathrm{D(Experienced)}_{it} + \alpha_{2\theta}\mathrm{D(Hawk)}_{i} + \alpha_{3\theta}\mathrm{D(Experienced)}_{it}\cdot\mathrm{D(Hawk)}_{i}$$
$$\sigma = \alpha_{0\sigma} + \alpha_{1\sigma}\mathrm{D(Experienced)}_{it} + \alpha_{2\sigma}\mathrm{D(Hawk)}_{i} + \alpha_{3\sigma}\mathrm{D(Experienced)}_{it}\cdot\mathrm{D(Hawk)}_{i},$$

and use maximum likelihood to estimate the α parameters. The last row shows the difference between experienced members (row 2) and new members (row 1). The last column shows the difference between new and experienced Hawks (column 2) and Doves (column 1), as well as the difference-in-differences (bottom right cell).

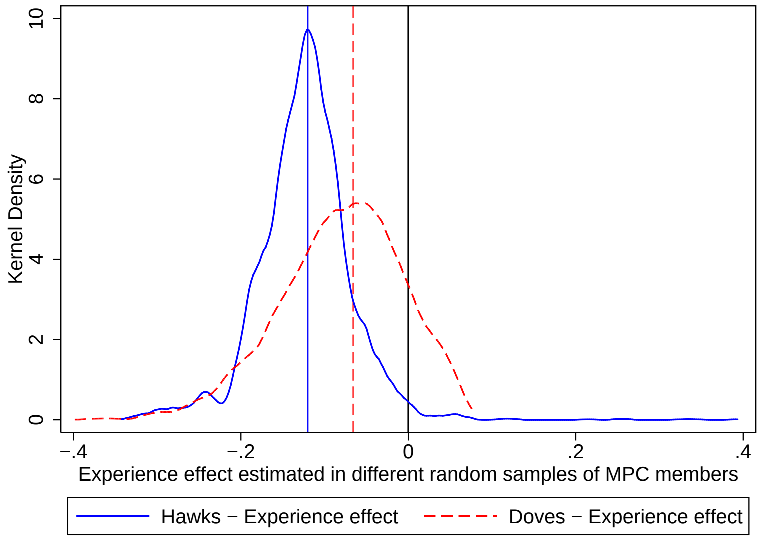

Figure 8: Kernel density plots of experience effect on random pseudo-committees

Notes: We construct 4,000 randomly selected pseudo-committees of 7 Hawks and 7 Doves. For each pseudo-committee, we model θ_i and σ_i as described in table 7 and estimate θ for new and experienced Hawks and Doves. We compute the experience effect as the difference in these estimates. This figure plots the kernel density of the 4,000 experience effect estimates; the thin vertical lines plot the structural estimates of the group experience effect contained in table 7. We use the Epanechnikov kernel, and set the half-width to the value that would minimize the mean integrated squared error if the underlying distribution were Gaussian.

Table 8: The experience effect across 4,000 random pseudo-committees
Split by Hawks and Doves

(a) Hawks

Item	Number	Per cent
Negative, p-value < 0.01	1,304	33
Negative, 0.05> p-value > 0.01	988	25
Negative, 0.1 > p-value>0.05	583	15
Negative, 0.5 > p-value > 0.1	886	22
Negative, p-value > 0.5	190	5
Positive, p-value > 0.5	39	1
Positive, 0.5 > p-value > 0.1	2	0
Positive, 0.1 > p-value > 0.05	0	0
Positive, 0.05 > p-value > 0.01	0	0
Positive, p-value < 0.01	8	0
Total	4,000	100

(b) Doves

Item	Number	Per cent
Negative, p-value < 0.01	521	13
Negative, 0.05> p-value > 0.01	364	9
Negative, 0.1 > p-value>0.05	213	5
Negative, 0.5 > p-value > 0.1	1,357	34
Negative, p-value > 0.5	1,009	25
Positive, p-value > 0.5	199	5
Positive, 0.5 > p-value > 0.1	174	4
Positive, 0.1 > p-value > 0.05	38	1
Positive, 0.05 > p-value > 0.01	78	2
Positive, p-value < 0.01	47	1
Total	4,000	100

Notes: We construct 4,000 randomly selected pseudo-committees of 7 Hawks and 7 Doves. For each pseudo-committee, we model θ_i and σ_i as described in table 7 and estimate θ for new and experienced Hawks and Doves. We compute the experience effect as the difference in these estimates. This table breaks the 4,000 estimates into groups according their signs and p-values for the statistical test of difference from zero, as well as by Hawks and Doves.

There is in fact some slight evidence of convergence between Hawks and Doves. Some portion of Doves have a positive estimated experience effect, while the Hawk distribution is clearly below the Dove. One could interpret Hawks and Doves as partially "meeting-in-the-middle" after an initial period of more extreme disagreement. We emphasize, however, that this really is only partial because individuals in both groups fall substantially. Across the replications, 99% of Hawk subgroups and 87% of Doves subgroups have a negative experience effect. Furthermore, 73% of Hawk subgroups have a negative experience effect significant at the 10% level, while only 1% has a positive significant experience effect; the corresponding numbers for Doves are 27% and 4%. Thus, while Hawks appear to contribute more to the experience than Doves overall (although, as table 7 shows, not significantly so), Doves are by no means free of the negative experience effect. The important intuition from the signalling model is that both types care about their reputation equally, so both should become less tough on inflation as their reputation concerns decline. Since the downward shift in θ for both groups is present in the data, we again claim the signalling model contributes to understanding voting dynamics in a way that learning cannot.

5.3 The timing of the experience effect

In models with multi-period signalling, the influence of reputation is generally declining. An additional check on the signalling model is that the estimated θ declines not just before and after some threshold, but year-on-year. Accordingly, we re-introduce the D(Experienced - 12M) and D(Experienced - 24M) dummy variables from section 2 that coded members as experienced if they have served more than 12 and 24 months, and include them both in the structural equations for θ and σ. We also define a new dummy variable D(Experienced - 36M) that measures experience past the 36 month cutoff and include it along with D(Experienced - 12M) and D(Experienced - 24M) in a second regression. The output from both regressions is contained in table 9.

Table 9: Effect of each year of experience on θ

	(1) θ		(2) θ
First 12 months	0.61***	First 12 months	0.62***
	[0.54 - 0.67]		[0.55 - 0.69]
12 to 24 Months	0.51***	12 to 24 Months	0.53***
	[0.46 - 0.56]		[0.49 - 0.57]
24+ Months	0.48***	24 to 36 Months	0.49***
	[0.45 - 0.51]		[0.39 - 0.59]
		36+ Months	0.50***
			[0.47 - 0.53]
Experience Effects		Experience Effects	
Year 2 - Year 1	-0.100**	Year 2 - Year 1	-0.089**
	[-0.18 - -0.020]		[-0.17 - -0.0035]
Year 3 - Year 1	-0.13***	Year 3 - Year 1	-0.13**
	[-0.20 - -0.061]		[-0.25 - -0.0080]
		Year 4 - Year 1	-0.12***
			[-0.20 - -0.040]

Clustered, by member, confidence interval in brackets
*** p<0.01, ** p<0.05, * p<0.1

Notes: This table shows structural estimates of θ for members in different years of their service. For the estimates in column (1), we model θ and σ in the likelihood function in appendix A as

$$\theta = \alpha_{0\theta} + \alpha_{1\theta}\text{D(Experienced - 12M)}_i + \alpha_{2\theta}\text{D(Experienced - 24M)}_i$$
$$\sigma = \alpha_{0\sigma} + \alpha_{1\sigma}\text{D(Experienced - 12M)}_i + \alpha_{2\sigma}\text{D(Experienced - 24M)}_i$$

and for the estimates in column (2) we include a third dummy variable D(Experienced - 36M) in each equation. These variables indicate whether a member has served the corresponding number of months on the MPC. The bottom rows of the table show the experience effect for each year of service by subtracting the θ estimate for the year of service from the θ estimate for the first year of service.

The first column shows that the estimated θ coefficient declines substantially from the first year to the second, and less so from the second to the third; the second shows that after three years of experience, the θ estimate stabilizes completely. This is consistent with a steady decline of reputational incentives. Not only does the θ estimate decrease over time, it decreases at a decreasing rate. This means that the experience effect is much more a beginning-of-term effect than an end-of-term effect.

6 Other Dynamic Incentives

We motivated the above signalling model as one in which members signal to control inflation expectations. Of course, as well as policy-motivated signalling, our model is also consistent with other motivations for signalling hawkishness. This section first evaluates two of these alternatives. First we examine whether members signal to advance their careers. Various measures of career heterogeneity do not correlate in any significant way with the experience effect, so we conclude that career concerns do not drive the effect. Second, we explore if members signal to each other, and again conclude this is unlikely.

We conclude the section by examining another dynamic incentive story other than signalling—member capture by the financial industry. Again, we find little supportive evidence. This leaves the policy motivation emphasized in the monetary literature as a principal candidate for rationalizing the dynamic incentive story.

6.1 Career concerns

The career concerns literature, initiated by Holmström (1999), maintains that agents seek to build a good reputation in order to receive better employment opportunities or a higher wage in the future. The political economy literature has interpreted career concerns as the desire for reelection. In the MPC context, reappointment in equivalent to reelection. So, to the extent that a member's reappointment depends on selectors' beliefs about his hawkishness, he may want to vote for high rates in order to increase the chance of reappointment.

If career concerns indeed existed on the MPC, one would imagine that their strength would vary among members. One group for whom they might be more salient are external members. Internal members are often drawn from the world of central banking and are appointed to the MPC as part of their more permanent positions at the Bank of England. As such, they can expect to continue to serve the MPC while they remain as Bank of England staff. External members come on more as unknown quantities and, if not reappointed, must find new employment outside the Bank. Also, external members serve shorter terms on average,[20] and term length has been used as a measure of career concerns in the past (Leaver 2009). To test this idea, we estimate the experience effect separately for internal and external members; the results are in table 10. If signalling arose because because of career concerns, and external members had more career concerns than internal members, we would expect the experience effect to have a large magnitude for the former. While the qualitative results of table 10 go in the right direction (the experience effect for external members is over three times that of internals), the difference in difference is not significant. We interpret this as showing that the joint movement down with experience

[20]External members serve three year terms, while three (of five) internal members serve five year terms.

is more robust than any differential effect of experience between internal and external members.

Table 10: Internal and external members θ Estimates

	(1) $\text{D(External)}_i = 0$	(2) $\text{D(External)}_i = 1$	(3) Difference
New	0.60***	0.54***	0.059
	[0.54 - 0.66]	[0.38 - 0.70]	(0.496)
Experienced	0.54***	0.34***	0.19**
	[0.51 - 0.57]	[0.19 - 0.50]	(0.016)
Difference	-0.059***	-0.20**	0.14
	[-0.10 - -0.015]	[-0.38 - -0.013]	(0.155)

Robust 95% confidence interval in brackets
*** p<0.01, ** p<0.05, * p<0.1
Column (1)-(3), H0: Estimate = 0

Notes: This table shows structural estimates of θ for new and experienced members split by internal ($\text{D(External)}_i = 0$) and external members ($\text{D(External)}_i = 1$); those in the latter group are Goodhart, Buiter, Julius, Budd, Wadhwani, Allsop, Nickell, Barker, Bell, Lambert, Walton, Blanchflower, Besley, and Sentence. We model θ and σ in the likelihood function in appendix A just as described in table 7, but replace D(Hawk)_i with D(External)_i, and use maximum likelihood to estimate the α parameters. The last row shows the difference between experienced members (row 2) and new members (row 1). The final column shows the difference within new or experienced members between external (column 2) and internal members (column 1), as well as the difference-in-differences (bottom right cell).

There are also other divisions on the MPC that one can explore; for example, whether a member held an academic position prior to joining the MPC. Since the most important activity that academic research economists carry out is the production of basic research, they may not care as much about reappointment as non-academics since recommencing an active research career is increasingly difficult with time spent exclusively in the policy world. The fact that tenured academics do not face future occupational uncertainty is a further reason why their career concerns may be lower. Another division along the same lines is between members who joined the MPC from private sector positions and those who did not. Members from the private sector are in some sense in the opposite position to academics: the ability to progress in their careers does not seem to decline with MPC tenure (in fact, due to the exposure that MPC service offers, it might increase) and their future job prospects, while not poor, are also not certain. Tables 11 and 12 show the results. In both cases the experience effect associated with the group we have said should have less career concerns is less; however, just as for the internal-external division, there is little statistical evidence of a difference-in-difference.[21]

[21] Besley, Meads, and Surico (2008) have a related finding that monetary policy reaction functions are homogenous across similar groups to the ones we examine here; they also interpret this as a lack of evidence for career concerns.

Table 11: Academic members θ Estimates

	(1) D(Academic)$_i$ = 0	(2) D(Academic)$_i$ = 1	(3) Difference
New	0.57***	0.60***	-0.038
	[0.49 - 0.64]	[0.35 - 0.86]	(0.779)
Experienced	0.51***	0.39***	0.12
	[0.45 - 0.56]	[0.25 - 0.53]	(0.115)
Difference	-0.056**	-0.22*	0.16
	[-0.11 - -0.0034]	[-0.44 - 0.012]	(0.181)

Robust 95% confidence interval in brackets
*** p<0.01, ** p<0.05, * p<0.1
Column (1)-(3), H0: Estimate = 0

Notes: This table shows structural estimates of θ for new and experienced members split by those members who joined the MPC directly from an academic position (D(Academic)$_i$ = 1); this group includes Goodhart, Buiter, Vickers, Allsop, Nickell, Bean, Blanchflower, and Besley. We model θ and σ in the likelihood function in appendix A just as described in table 7, but replace D(Hawk)$_i$ with D(Academic)$_i$, and use maximum likelihood to estimate the α parameters. The last row shows the difference between experienced members (row 2) and new members (row 1). The final column shows the difference within new or experienced members between non-academic (column 2) and academic members (column 1), as well as the difference-in-differences (bottom right cell).

Table 12: Private Sector θ Estimates

	(1) D(Private Sector)$_i$ = 0	(2) D(Private Sector)$_i$ = 1	(3) Difference
New	0.62***	0.49***	0.13*
	[0.52 - 0.71]	[0.39 - 0.58]	(0.052)
Experienced	0.51***	0.42***	0.087
	[0.44 - 0.57]	[0.28 - 0.55]	(0.256)
Difference	-0.11***	-0.069	-0.043
	[-0.20 - -0.030]	[-0.20 - 0.061]	(0.582)

Robust 95% confidence interval in brackets
*** p<0.01, ** p<0.05, * p<0.1
Column (1)-(3), H0: Estimate = 0

Notes: This table shows structural estimates of θ for new and experienced members split by those members who joined the MPC directly from a private sector appointment (D(Private Sector)$_i$ = 1); this group includes Clementi, Julius, Wadhwani, Barker, Bell, Lambert, Walton, and Sentence. We model θ and σ in the likelihood function in appendix A just as described in table 7, but replace D(Hawk)$_i$ with D(Private Sector)$_i$, and use maximum likelihood to estimate the α parameters. The last row shows the difference between experienced members (row 2) and new members (row 1). The final column shows the difference within new or experienced members between non-private sector (column 2) and private sector members (column 1), as well as the difference-in-differences (bottom right cell).

To some extent, the career concerns stories we have told for tables 10 - 12 are slightly arbitrary, and should be taken more as conjecture than fully-fledged hypotheses. What seems less ad hoc to us is the idea that non-academics, internals, and career bureaucrats have very different professional backgrounds and career trajectories. If career concerns were the underlying driving force behind the experience effect, one would expect the experience effect to correlate *in some way* across the divisions we have explored. The fact that they do not, combined with the fact that the experience effect in general is significant and negative across most specifications, leads us to the conclusion that career concerns do not drive the majority of the signalling on the MPC.

6.2 Members justifying their appointment to other members

The next story we explore is more behavioral. We ask whether members may try to signal their type to other members of the committee in order "fit in". It may be that within the MPC there is a norm that tough stances on inflation are best, and recently appointed members vote for high rates to show that they belong in the group before gradually reverting to their true preferences. While we do not construct a formal test of this idea, informal observation is not consistent with it. Table 4 in section 4 displays estimates of the experience effect at the individual level. Two of the members with the largest negative experience effect are Professor Charles Goodhart and Sir Mervyn King. Prior to their appointments to the MPC, both men accumulated vast experience in the UK monetary policymaking world. Goodhart had spent 20 years as a senior advisor at the Bank of England,[22] while King, after leaving his position as Professor at the London School of Economics, had been the Bank of England's chief economist since 1991 and had been instrumental in shaping the analysis of the UK economy as carried out by the Bank staff. It seems that these two members would not have been very concerned about proving their credentials to their colleagues on the MPC. On the other hand, since both were members of the inaugural MPC in 1997, they may have wanted to show that the new monetary policymaking institution was serious about inflation.

6.3 Regulatory capture

A separate idea is that MPC members do not signal at all, but are instead directly influenced by the policy choices that outsiders want to see implemented. The theory of regulatory capture, articulated by Stigler (1971) and Posner (1974), maintains that government regulators inevitably end up serving the interests of the industries they regulate, either because they have higher incentives to lobby and pressure the regulators, or be-

[22]He did not go straight from his advisory position to the MPC, but rather had a post as Professor at the London School of Economics from 1985 to 1997.

cause members of the industry that a regulator regulates end up eventually becoming regulators themselves and vice versa. Commentators both in the media (Frank 2009, Persaud 2009) and in academia (Baxter 2011) have argued that one of the causes of the recent financial crisis was the capture of financial regulators by the financial industry.

While the Bank of England is not a financial regulator (or was not during our sample period), its policies clearly have a substantial impact on financial markets. One might imagine, therefore, that the financial industry applied systematic pressure on MPC members to pursue a "cheap money" approach in order to stimulate lending and borrowing. During their tenures, MPC members regularly go to official dinners, give speeches, and generally interact with members of the financial industry. Also, as we have already mentioned, some MPC members have come directly from positions in the financial industry. So, an alternative story to explain our results is that during their first 18 months on the MPC, members pursue what they genuinely believe to be the optimal monetary policy, but that due to the continuous lobbying effort of the financial industry, they eventually give up and begin to vote for lower rates with a higher probability.

Fortunately we have a test of this hypothesis. The Times of London newspaper has, since 2002, formed its own shadow MPC of monetary policy experts, some of whom are former, and in some cases future, members of the real MPC. Members of the Times MPC vote each month prior to the MPC meeting on what they believe to be the best interest rate, and their votes are subsequently published in the Times newspaper. Since there are few implications of Times MPC votes for anyone other than the members themselves, one would imagine that they simply express their genuine opinion about interest rates. If it were the case that actual MPC members faced industry pressure to keep rates lower than they genuinely believed to be social-welfare maximizing, then it should be the case that a member of the MPC votes for systematically lower rates than he would if he were voting on the Times MPC.

There are five MPC members (Budd, Goodhart, Lambert, Sentence, and Wadhwani) who have also served on the Times MPC. As an initial test of regulatory capture, we compute the frequency with which these members chose the high vote during the first 18 months on the MPC, their subsequent time on the MPC, and their tenure on the Times MPC. The results are in table 13. All members apart from Lambert choose the high vote less often while on the Times MPC than they do as either new or experienced members of the MPC, which goes against the capture hypothesis. The fact that members vote for high rates less frequently on the Times MPC than experienced MPC members might also indicate that some signalling incentives also remain in the final years of MPC service.

We then estimate the θ parameter for the pooled votes of the five members across the three periods covered in table 13. Our estimator does not converge when we include Wadhwani, so we discard him and proceed with estimation of the other four; table 14

Table 13: Voting on the MPC and the Times MPC:
Fraction of High Votes (%) and count (n) of total votes on different committees

Member	Times 1st		MPC New		MPC Exp		Times 2nd	
	(%)	(n)	(%)	(n)	(%)	(n)	(%)	(n)
Budd	.	.	67	18	.	.	56	77
Goodhart	.	.	56	18	50	18	33	9
Lambert	75	4	44	18	56	16	.	.
Sentence	38	8	72	18	58	12	.	.
Wadhwani	.	.	17	18	11	18	4	46

Notes: This table shows the % of total votes cast that were high, and the total number of votes, for the five members who have served on both the Times (shadow) MPC and the MPC. Two of the five members (Lambert and Sentence) served on the Times MPC before their appointment to the MPC while the other three served following their appointment.

presents the results. The estimated θ coefficient for the Times MPC is slightly below the coefficient for experienced periods on the MPC. This again goes against the capture hypothesis, which would predict the Times MPC coefficient to be closer to the new MPC coefficient than the experienced.

Table 14: Structural estimates of θ on the MPC and the Times MPC

	θ
New	0.69***
	[0.50 - 0.88]
Experienced	0.58***
	[0.47 - 0.70]
Times	0.52***
	[0.47 - 0.56]
Difference	-0.067
	[-0.19 - 0.053]

Clustered, by member, 95% confidence interval in brackets
*** p<0.01, ** p<0.05, * p<0.1

Notes: This table shows structural estimates of the average θ for members when new to the MPC ("New"), experienced on the MPC as defined by greater than 18 months service ("Experienced") and during their service on the Times MPC ("Times"). The final row reports the difference between their service on the Times MPC (row 3) and the experienced MPC members (row 2). The estimates exclude Sushil Wadhwani.

7 Conclusion

This paper's first contribution is to show that significant dynamics are present on the MPC. After rejecting two separate learning explanations for this (increasing confidence in private information and learning about the macroeconomy), we show that a signalling model accounts for the dynamics. We rule out two potential motivations for this signalling (career concerns and fitting in), and a separate dynamic incentive story (regulatory capture), leaving policy-motivated signalling as a plausible driver of MPC voting dynamics.

Our paper presents two avenues for future research. Specifically in the monetary literature, while we show that the balance of evidence favors a policy-motivated signalling explanation for voting dynamics, formulating and testing more detailed hypotheses using richer datasets (for example, analyzing how the behavior of inflation expectations derived from asset prices reacts to observed voting behavior) is a logical next step. Second, within macroeconomics specifically and political economy generally, much recent research has used the stories we have rejected to understand policymakers' behavior. While we do not argue that learning, career concerns, regulatory capture, or other behavioral considerations are not important for behavior, we do think the recent existing literature has underplayed the idea that signalling preferences for policy motivations has a role to play. As our paper shows that in some settings this role may be the most important, future work should take such policy-motivated signalling more seriously as a generator of dynamic behavior.

References

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.

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.

Baxter, L. G. (2011): "Capture in Financial Regulation: Can We Channel it Toward the Common Good?," *Cornell Journal of Law and Public Policy*, forthcoming.

Besley, T., and S. Coate (2003): "Elected Versus Appointed Regulators: Theory and Evidence," *Journal of the European Economic Association*, 1(5), 1176–1206.

Besley, T., N. Meads, and P. Surico (2008): "Insiders versus Outsiders in Monetary Policymaking," *American Economic Review*, 98(2), 218–23.

Brainard, W. C. (1967): "Uncertainty and the Effectiveness of Policy," *The American Economic Review, Papers and Proceedings of the Seventy-ninth Annual Meeting of the American Economic Association*, 57(2), pp. 411–425.

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.

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

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

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

Feddersen, T., and W. Pesendorfer (1998): "Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts under Strategic Voting," *American Political Science Review*, 92(1).

Frank, T. (2009): "Obama and 'Regulatory Capture'," Wall Street Journal, 24 June.

Goeree, J. K., and L. Yariv (2010): "An Experimental Study of Collective Deliberation," *Econometrica*, forthcoming.

Hansen, S., and M. F. McMahon (2008): "Delayed Doves: MPC Voting Behaviour of Externals," Discussion Paper 862, Centre for Economic Performance, London School of Economics.

——— (2011): "How Experts Decide : Identifying Preferences versus Signals from Policy Decisions," Discussion Paper 963, Economics Department, University of Warwick.

Holmström, B. (1999): "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies*, 66(1), 169–82.

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

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

Leaver, C. (2009): "Bureaucratic Minimal Squawk Behavior: Theory and Evidence from Regulatory Agencies," *American Economic Review*, 99(3), 572–607.

Maskin, E., and J. Tirole (2004): "The Politician and the Judge: Accountability in Government," *American Economic Review*, 94(4), 1034–1054.

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.

Persaud, A. (2009): "Boomtime politicians will not rein in the bankers," Financial Times, 26 November.

Posner, R. A. (1974): "Theories of Economic Regulation," *Bell Journal of Economics*, 5(2), 335–358.

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

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

Select Committee on Economic Affairs (2003): "The MPC and Recent Developments in Monetary Policy," Second Report.

Sibert, A. (2003): "Monetary Policy Committees: Individual and Collective Reputations," *Review of Economic Studies*, 70(3), 649–665.

Stigler, G. J. (1971): "The Theory of Economic Regulation," *Bell Journal of Economics*, 2(1), 3–21.

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

A Likelihood Function Derivation

Let $\widehat{w}_{it}\left(s_{it}\right)$ be the updated belief after observing s_{it}. Then

$$\ln\left(\frac{\widehat{\omega}(s_{it})}{1-\widehat{\omega}(s_{it})}\right)=\ln\left(\frac{q_t}{1-q_t}\right)+\ln\left(\frac{f_1}{f_0}\right)=\ln\left(\frac{q_t}{1-q_t}\right)+\frac{2s_{it}-1}{2\sigma_i^2}.$$

where $f_1 \sim (1,\sigma_i^2)$ is the distribution of s_{it} conditional on $\omega_t=1$ and $f_0 \sim (0,\sigma_i^2)$ is the distribution of s_{it} conditional on $\omega_t=0$. The first equality is implied by Bayes' rule and the second by the distributional assumptions on f_0 and f_1. The voting rule in (3) can be rewritten as choosing $v_{it}=1$ whenever $\ln\left(\frac{\widehat{\omega}(s_{it})}{1-\widehat{\omega}(s_{it})}\right)\geq\ln\left(\frac{1-\theta_i}{\theta_i}\right)$, which, when combined with the above expression, implies choosing $v_{it}=1$ whenever

$$s_{it}\geq s_{it}^*=\frac{1}{2}-\sigma_i^2\ln\left(\frac{\theta_i}{1-\theta_i}\frac{q_t}{1-q_t}\right). \tag{A.1}$$

This yields the likelihood function

$$L_{it}=\begin{cases}q_t\left(1-\Phi\left(\frac{s_{it}^*-1}{\sigma_i}\right)\right)+(1-q_t)\left(1-\Phi\left(\frac{s_{it}^*}{\sigma_i}\right)\right) & \text{if } v_{it}=1\\ q_t\Phi\left(\frac{s_{it}^*-1}{\sigma_i}\right)+(1-q_t)\Phi\left(\frac{s_{it}^*}{\sigma_i}\right) & \text{if } v_{it}=0.\end{cases}$$

We model $\theta_i=\alpha_\theta+\beta_\theta\mathrm{D(Experience)}_i$ and $\sigma_i=\alpha_\sigma+\beta_\sigma\mathrm{D(Experience)}_i$.

B Signalling Model

B.1 Derivation of Voting Rule with Reputational Concerns

Here we derive the voting rule in (5) by adapting the baseline model from Hansen and McMahon (2011). Suppose that member i has the utility function

$$\lambda\mathbb{E}\left[-(\pi_t-\pi_i^*)^2\right]+(1-\lambda)R(v_{i1})$$

where $\pi_i^*\in\{\pi_L^*,\pi_H^*\}$ is his preferred inflation target. The inflation rate π_t is distributed according to $F\left[\pi_t\mid v_{i1}\right]$, whose expected value is ω_t-v_{i1}, and whose variance σ^2 is finite and independent of v_{i1}. Using the fact that

$$\mathbb{E}\left[(\pi_t-\pi_i^*)^2\right]=V\left[\pi_t\right]+(\mathbb{E}[\pi_t]-\pi_i^*)^2$$

one can express the payoff of choosing $v_{i1}=1$ as

$$-\lambda\left(-1+\widehat{\omega}_{it}-\pi_i^*\right)^2+(1-\lambda)R(1) \tag{B.1}$$

and that of choosing $v_{i1} = 0$ as

$$-\lambda\left(\widehat{\omega}_{it} - \pi_i^*\right)^2 + (1-\lambda)R(0). \tag{B.2}$$

Straightforward algebra reveals that (B.1) is larger than (B.2) whenever

$$\widehat{\omega}_{it} \geq 1 - (0.5 - \pi_i^*) - \frac{1-\lambda}{2\lambda}\left(R(1) - R(0)\right) \tag{B.3}$$

Defining $\theta_L = 0.5 - \pi_L^*$, $\theta_H = 0.5 - \pi_H^*$, and $\beta = \frac{1-\lambda}{2\lambda}$ reduces this model to the one in the main text.

B.2 Proof of Proposition 1

Proof. In any responsive equilibrium, both types must vote high if and only if the realized signal exceeds a finite cutoff. Suppose the evaluator believes that type θ_L uses cutoff s_L^* and that type θ_H uses cutoff s_H^*. This implies that

$$\begin{aligned}\Delta(s_L^*, s_H^*) =& \frac{r\Pr\left[\, s_H \geq s_H^* \,\right]}{r\Pr\left[\, s_H \geq s_H^* \,\right] + (1-r)\Pr\left[\, s_L \geq s_L^* \,\right]} - \\ & \frac{r\Pr\left[\, s_H < s_H^* \,\right]}{r\Pr\left[\, s_H < s_H^* \,\right] + (1-r)\Pr\left[\, s_L < s_L^* \,\right]}.\end{aligned} \tag{B.4}$$

It is straightforward to show that $\Delta(s_L^*, s_H^*) \gtreqless 0$ as $s_L^* \gtreqless s_H^*$. Given the evaluator's beliefs s_L^* and s_H^*, the optimal thresholds for both types s_L' and s_H' are, following the expression (A.1) in appendix A,

$$s_L' = \frac{1}{2} - \sigma_i^2 \ln\left(\frac{\theta_L - \beta\Delta(s_L^*, s_H^*)}{1 - \theta_L + \beta\Delta(s_L^*, s_H^*)}\frac{q_1}{1-q_1}\right) \tag{B.5}$$

$$s_H' = \frac{1}{2} - \sigma_i^2 \ln\left(\frac{\theta_H - \beta\Delta(s_L^*, s_H^*)}{1 - \theta_H + \beta\Delta(s_L^*, s_H^*)}\frac{q_1}{1-q_1}\right). \tag{B.6}$$

Clearly $s_L' > s_H'$, so that when one imposes the equilibrium conditions $s_L' = s_L^*$ and $s_H' = s_H^*$, one obtains $\Delta > 0$. This establishes the second statement in the proposition.

Proving the second statement is equivalent to showing that equations (B.5) and (B.6) have a fixed point. Since $\Delta \in [0,1]$, $s_x' \in [\underline{s}_x, \overline{s}_x]$ where

$$\underline{s}_x = \frac{1}{2} - \sigma_i^2 \ln\left(\frac{\theta_x - \beta}{1 - \theta_x + \beta}\frac{q_1}{1-q_1}\right) \tag{B.7}$$

$$\overline{s}_x = \frac{1}{2} - \sigma_i^2 \ln\left(\frac{\theta_x}{1 - \theta_x}\frac{q_1}{1-q_1}\right) \tag{B.8}$$

for $x \in \{L, H\}$. $\beta < \theta_L$ is a sufficient condition for (B.7) and (B.8) to be finite. Under this condition, (B.5) and (B.6) map $\underline{s}_x \leq s_x^* \leq \overline{s}_x$ into $\underline{s}_x \leq s_x' \leq \overline{s}_x$ for $x \in \{L, H\}$. Moreover, since (B.5) and (B.6) are continuous in s_L^* and s_H^*, one can apply Brouwer's Theorem to establish the existence of a fixed point.

To prove the third point, denote by $s_x^*(\beta)$ the equilibrium cutoff that type $x \in \{L, H\}$ uses with weight β. Setting α^* such that

$$s_x^*(\beta) \leq \frac{1}{2} - \sigma_i^2 \ln\left(\frac{\theta_x - \alpha^*\beta}{1 - \theta_x + \alpha^*\beta}\frac{q_1}{1-q_1}\right)$$

holds for $x \in \{L, H\}$ and using the logic from the proof of the second statement completes the proof. ■