Transitions in Central Bank Leadership*

Carlos Carvalho
PUC-Rio

Tiago Flórido
PUC-Rio

Eduardo Zilberman
PUC-Rio

June 2015
Preliminary Draft

Abstract

We assemble a novel dataset on transitions in central bank leadership in several countries, and study how monetary policy is conducted around those events. We find that policy is tight both at the last meetings of departing governors and first meetings of incoming leaders. This finding cannot be explained by endogenous transitions, fiscal policy, the effects of the zero lower bound, electoral cycles, nor uncertainty. Moreover, results are stronger when the central bank has less independence, is less transparent, the country's regulatory quality is lower, or the governor has more power. We offer an explanation for these results based on a simple signalling story. The incoming governor has incentives to tighten monetary policy to signal that she is a Hawk, and thus face more favorable inflation expectations going forward. In turn, the departing governor tightens policy in order to make it easier for an incoming Hawk to signal its type and separate from a Dove.

*For helpful comments and suggestions, we thank Tiago Berriel, Bernardo Guimarães, Leonardo Rezende, and seminar participants at IBMEC, LAMES 2014, PUC-Rio and UFPE. All errors are ours. Carvalho: cvianac@econ.puc-rio.br. Flórido: florido.tiago@gmail.com. Zilberman: zilberman@econ.puc-rio.br.

1 Introduction

Owing to the importance of monetary policy, a lot of attention is given to central banks as institutions, and to the people in their leadership positions. Recently, this fact came into particular prominence due to the end of Fed Chairman Ben Bernanke's term and the consequent speculations about possible successors. Before it was confirmed that vice-chairwoman Janet Yellen would be the next Fed chair, there were several articles in the media discussing pros and cons of different 'candidates'.[1] This reflects the importance that is assigned to the identity of a central bank's leader.

This view also finds resonance in the academic literature. Romer and Romer (2004), for example, analyze historical Fed transcripts and past speeches by Fed officials and find support to the idea that a central banker's views about the economy are a key determinant of monetary policy. In terms of theory, Rogoff (1985) flashes out the importance of central bankers, by showing that, in order to address the time inconsistency problem, the government should appoint a central banker who is more conservative than society as a whole.

Despite the importance assigned to central bank leaders, the literature has all but overlooked transitions in central bank leadership. In this paper, we study monetary policy around those episodes. Our contributions are twofold. First, to the best of our knowledge, we document a novel empirical fact. Namely, transition periods are associated with a tight monetary policy stance. Second, after failing to find empirical support for several possible explanations, we propose a simple signalling model that can explain that empirical regularity. In the model, an incoming governor has incentives to tighten monetary policy to signal that she is a Hawk (less tolerant with inflation), and thus face more favorable inflation expectations going forward. In turn, the departing governor tightens policy in order to make it easier for an incoming Hawk to signal its type and separate from a Dove (more tolerant with inflation).

Sections 2 and 3 establish the empirical fact. To that end, we assemble a novel dataset containing transitions in central bank leadership in 35 countries. In particular, our (unbalanced) panel has information on nearly 70 transitions in central bank leadership. After controlling for factors that affect monetary policy decisions through a standard Taylor rule, we can establish how central bank behavior differs in the first monetary policy meetings under a new central banker and in the last meetings of a departing one, relative to the other "usual" meetings.

We find that both first and last meetings – i.e. the transition period – are associated

[1] For a particularly stark opinion piece, see: "Why Janet Yellen, Not Larry Summers, Should Lead the Fed" by Joshph Stiglitz in the New York Times at September 6, 2013.

with a tight monetary policy stance. In our preferred specification, interest rates exceed the level predicted by a simple Taylor rule that accounts for policy inertia, inflation and activity by 0.076 (0.075) percentage points in the last (first) meetings, on average. To put this in context, note that over 50% of interest changes are of 0.25 percentage points.

We show that these results are not driven by two immediate endogeneity concerns. Namely, the timing of the transition and the choice of the new governor. Also, we fail to find empirical support for other possible straightforward explanations. In particular, our results cannot be explained by electoral cycles that might coincide with transitions in CB leadership, nor by fiscal developments around those times. Moreover, results are not driven by the zero lower bound constraint, and cannot be explained by unusually high uncertainty around CB transitions.

We then entertain heterogeneous effects that might speak to the signaling explanation that we put forward. In particular, we find that results are stronger when the central bank has less independence, is less transparent, and when the country's regulatory quality is lower. Also, results are stronger when the central bank's governor has more power. We claim that these heterogeneities are consistent with a simple signalling theory, in which monetary policy affects the public's beliefs about the extent to which the central banker dislikes inflation. Indeed, more independence, more transparency, better regulatory quality, and more powerful governors reduce the need for signalling the central bank's type (Hawk or Dove).

Altogether, we see our empirical results as suggestive of the presence of signalling in monetary policy. Of course, as it is impossible to exhaust all possibilities, other explanations may be consistent with our evidence.

To our knowledge, Hansen and McMahon (2014) is the only paper to offer empirical evidence on signalling in monetary policy. They use data from the Bank of England's Monetary Policy Committee to show that new members tend to be tougher on inflation initially to signal they are not dovish.[2] However, their paper is silent on the signalling incentives of departing central bankers. Another related paper is Johnson et al. (2012), who document that, as the end of their mandates approaches, regional Federal Reserve Bank presidents grow hawkish relative to members of the Board of Governors and continuing presidents.[3] In any case, as the compositions of different monetary policy committees overlap, incentives to adopt a more or less hawkish stance during one's' mandates should not translate

[2]Hansen and McMahon (2014) cite other references that contain empirical evidence that is somewhat consistent with signalling – although signalling was not the focus of those papers.

[3]The authors interpret this finding as evidence that consensus building occurs by conforming preferences rather than convincing arguments.

into systematic monetary policy tightening during transitions in main leadership.[4]

Section 4 presents our simple model of signalling that rationalizes monetary policy tightening during transitions in CB leadership. The model is built on Kydland and Prescott (1977) and Barro and Gordon (1983b) classical framework developed to study the time inconsistency problem of monetary policy.[5]

As in Vickers (1986), we assume that agents are uncertain about central bank's preferences regarding inflation. Moreover, the central bank is not able to commit to a given inflation plan. In contrast, other papers assume that uncertainty regards central banks' ability to commit to a given inflation plan (e.g., Cukierman and Liviatan (1991)). In any case, a standard result arises. The incoming central bank, whose type is uncertain, wants to signal that it dislikes inflation in order to face lower inflation expectations in the future. This is accomplished by tightening monetary policy in the first periods. In other words, the incoming central bank wants to build reputation.[6] Since this result is widely emphasized in the literature, we focus our exposition on the model's implications for the choice of the departing central bank.

The main mechanism at play behind a monetary contraction in the last meetings is the following. Suppose that the incoming central bank can be either a Dove or a Hawk, in the sense that it assigns, respectively, a lower or higher weight to inflation in its loss function. The public is uncertain about the incoming central bank's type. By tightening monetary policy, the departing central bank makes it specially costly for a Dove to pretend to be a Hawk, as this action would require a further increase in interest rates. In other words, lower inherited inflation makes separation of types "more likely". Hence, our empirical finding can be rationalized if the departing governor wishes to help a Hawk successor build its reputation.

In order to formalize this simple idea, we add two ingredients to a standard version of Barro and Gordon (1983b)'s model with uncertain CB types. First, we allow for some inflation inertia by assuming that current inflation is a convex combination of previous inflation and the component of current inflation that is under the CB's control. Besides being empirically plausible, this assumption is necessary to connect the decisions of different

[4]For a theory on signalling incentives in a monetary policy committee, see Sibert (2003).

[5]There is some literature on signalling in the Barro and Gordon (1983b) monetary policy framework. Contributions include Backus and Driffill (1985a), Backus and Driffill (1985b), Barro (1986), Vickers (1986), Cukierman and Liviatan (1991), Ball (1995), Walsh (2000), Sibert (2002), and King et al. (2008). For a review of this literature, see Walsh (2010), chapter 7.

[6]Throughout the paper, reputation means the public's perceived likelihood that the central bank will fight inflation. It does not refer to the concept used in the repeated games literature, in which a deviation from a low-inflation solution may trigger a punishment from the public, such as in Barro and Gordon (1983a).

central bankers through time. Second, we assume that CB's losses depend not only on (total) inflation, but also on the component of current inflation under its control. This extra term in the loss function captures the fact that changing inflation is costly. One reason may be that controlling inflation requires costly changes in the interest rate. Another reason may be that a high level of "controlled inflation" would lead to a general assessment that the central bank is incompetent, which may hurt teh central banker's career prospects.

Assume that a new central bank takes office today, so that the public does not know its type – Dove or Hawk. We show that a reduction in inherited inflation, which is treated as a parameter in the model, but interpreted as the choice of the outgoing central bank, increases the likelihood that both Dove and Hawk central banks separate their actions in the first meeting.[7] Hence, by tightening monetary policy, the departing central bank helps the Hawk build a reputation. In contrast, higher inherited inflation makes pooling of the incoming central bank's types more likely.

The model developed provides a simple framework to study transitions in monetary policy when there is uncertainty regarding the central bank's type. Few papers have considered transitions and signalling in a single framework. Ball (1995) and King et al. (2008), for example, allow transitions between central banks that are and that are not able to commit. As in Cukierman and Liviatan (1991), the ability to commit is private information. Ball (1995) aims to explain inflation dynamics in the postwar United States, whereas King et al. (2008) aim to provide a comprehensive framework to discuss key concepts regarding expectations management. None of these papers consider inertial inflation, so the intertemporal linkage between central banks, if any, is due solely to beliefs and expectations.

In contrast with Ball (1995) and King et al. (2008), we assume that type is the weight assigned to inflation rather than ability to commit. Other papers in the literature also consider transitions in preferences as recently exemplified by Debortoli and Nunes (2014).[8] However, we are not aware of a paper that allows for both transitions in preferences and signalling dynamics. The simple framework developed in this paper may serve as an ingredient in a more comprehensive macroeconomic model, that could be used to study other aspects of signalling in monetary policy.

[7] We treat previous inflation as a parameter for simplicity and ease of exposition. One could model the old central bank behavior by endowing it with similar preferences to the new central bank. In this case, under some extra assumptions, it can be shown that its optimal to increase the likelihood of separating.

[8] Debortoli and Nunes (2014) consider an environment that combines a standard central bank's loss function with a New Keynesian Phillips curve. They argue that, whenever transitions in preferences are not properly modeled, allowing transitions in interest rate rules to capture policy changes might lead to misleading results.

This paper is organized as follows. Section 2 describes the data and the empirical strategy. Section 3 presents the empirical results and provides an explanation for them. Section 4 develops the model. Finally, Section 5 concludes.

2 Data and empirical strategy

The aim of the paper is to establish if monetary policy differs at the end or the beginning of the mandate, i.e. during transitions in Central Bank (CB) leadership, from other periods. To do so, we assemble a novel dataset and design an empirical strategy to estimate the effects of transitions in monetary policy.

2.1 Data

The dataset is a panel composed of 35 countries, where each observation c, m consists of a country c and a monetary policy meeting m. One should note that m does not correspond to the same time period. After all, the m-th meeting we have for, say, the United States FOMC is not at the same date as the m-th meeting of the UK monetary policy committee. In fact, they do not even have the same periodicity: countries vary in the number of meetings held per year - spanning from monthly to quarterly meetings. In addition, countries enter the sample at different years, starting in 1984 with the US while Georgia is the last country to enter the sample in 2008. Table 13 in the Appendix A.1.2 lists all countries and their number of meetings and governors.

The panel is unbalanced because we only consider observations for which the instrument target is the interest rate. Moreover, we only consider countries where there is a meeting calendar, or we could track the date of every monetary policy decision. For instance, until the late 1990s many countries simply announce when there was change in policy. In the absence of a meeting calendar, we cannot track when the monetary committee actively decided to keep interest rates constant. Finally, we drop the financial crisis period - 2008 and 2009 - since this would confound our results on transition effects. Indeed, monetary policy during this period was conducted in an unconventional way, so it would be hard to establish whether monetary policy changed due to the transitions per se in comparison to regular policy.

The average number of meetings per country is 111 and the median is 102. The total number of transitions between governors in the sample is around 70. The main variables used in this study are:[9]

[9]For the data which were not seasonally adjusted from the source, we use ARIMA X12 procedure to

- Policy interest rate decisions(in %): $i_{c,m}$
- Inflation (YoY): $\pi_{c,m}$
- Activity level : $y_{c,m}$ (mostly unemployment when available, output growth otherwise).
- Dummy for the first meetings of a CB governor: $FM_{c,m}$
- Dummy for he last meetings of a CB governor: $LM_{c,m}$

The data come from four main sources: the OECD database, *Bloomberg* terminals, *Datastream - Thomson Reuters* terminals and individual central banks' websites. Since specific meeting dates are typically irregular, policy rate decisions and governors' transitions were obtained at each central bank's website. The macroeconomic series were mainly obtained from the OECD database and data terminals. However, there are countries whose time series are too short or not available on terminals. For these countries, we complement the macroeconomic series with data from central banks' websites and national data bureaus.

We match the macroeconomic series with each central bank meeting of a given country c according to the following algorithm. First, we identify the calendar month of each meeting m. For instance, a meeting in the 17th of April counts as April. Then we match with the inflation and unemployment referring to that calendar month. However, some countries do not report unemployment monthly. In these cases we check the availability of quarterly data for unemployment and GDP growth. We use the quarterly value for the three months of the corresponding quarter, as if it was a monthly variable. For instance, if the rate of unemployment was 7% for the second quarter of a given year, we input 7% in the cells referring to April, May and June. Then we proceed as before matching the quarterly rate to meeting in the corresponding month. Thus, when there is no monthly unemployment rate, we use the quarterly unemployment rate and, when even such periodicity is lacking, we use GDP growth as the activity level variable.

Data on first and last meetings of governors are found in each central bank website. Normally, there is a webpage reporting the list of former governors with the initial and final dates of their mandates. If, for some governor in a given country, this information is ambiguous at the webpage, we checked the minutes of the relevant meetings in order to locate when the transition took place. In some transitions, the final meeting of a governor is not the one exactly before the first meeting of his successor, i.e. $LM_{c,m}$ does not necessarily lag $FM_{c,m}$. Sometimes a governor's tenure ends before the appointment of his successor. In

adjust for seasonality.

between the mandates, there may appear an acting governor for a couple meetings. In the Appendix A.1.1 we discuss in detail how these transitions are coded, but in any case our results vary little within reasonable code changes.

Finally, it is important to assuage a possible concern regarding our dataset. That is, the possibility that most transitions are clustered around a couple of years. Figure 1 shows that the transitions are scattered, with most of them happening after the late 1990s. In fact, most countries enters in the sample after the 1990s as shown in Figure 2.

Figure 1

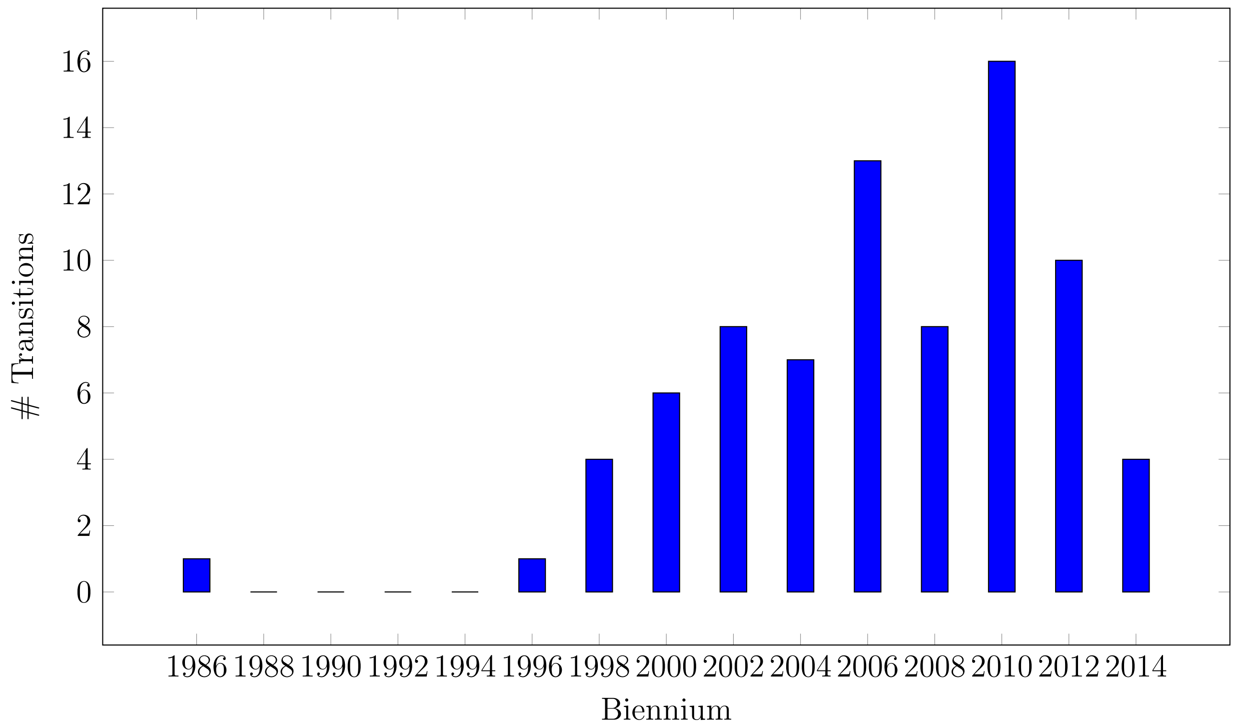

2.2 Empirical strategy

Recall that $FM_{c,m}$ ($LM_{c,m}$) is a dummy which takes value one when m is among the first (last) meetings of a given CB governor c. In order to estimate the effects of transitions in monetary policy, we add $FM_{c,m}$ and $LM_{c,m}$ to a simple Taylor rule in which inflation $\pi_{c,m}$, economic activity $y_{c,m}$ (either unemployment or GDP growth as explained above) and lagged interest rate $i_{c,m-1}$ are accounted for. In particular, we pool all observations and allow the coefficients on each of these variables to vary with countries. Moreover, we allow the intercept to vary with countries c and years $t = 1984, ..., 2014$. Hence, we estimate the

Figure 2

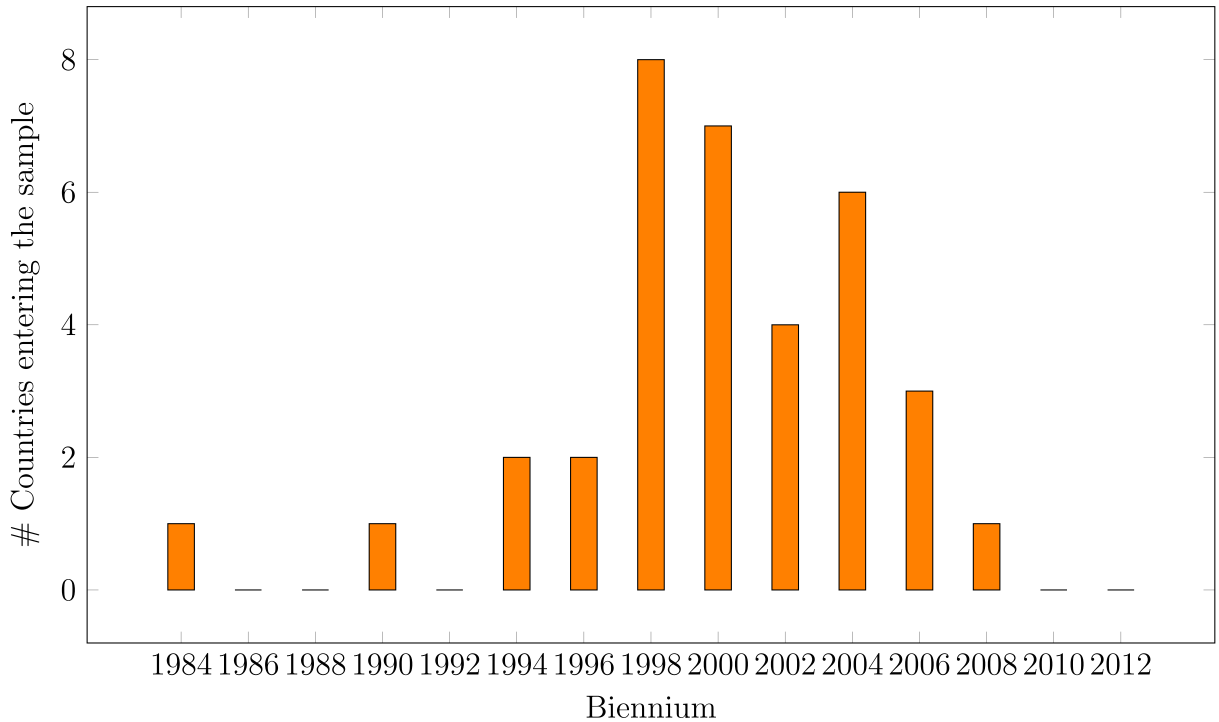

following Taylor rule by OLS:

$$i_{c,m} = \rho_c i_{c,m-1} + \alpha_{\pi,c}\pi_{c,m} + \alpha_{y,c}y_{c,m} + \beta_F FM_{c,m} + \beta_L LM_{c,m} + \delta_c + \delta_t + e_{c,m} \tag{1}$$

Our coefficients of interest are β_F and β_L. The idea is that, once changes in monetary policy warranted by macroeconomic factors are accounted for, β_F and β_L capture the effect of transitions in the interest rate $i_{c,m}$. In order to the exercise be meaningful, the bulk of variation in monetary policy due to macroeconomic factors must be accounted for. It is reassuring that, despite the simplicity of the functional form above, the R^2 of our baseline specification is 99.6%. In fact, the smoothing term improves a lot the fit of the regression - R^2 would be 92% otherwise. Finally, in the Appendix A.3, we report and discuss how the results would change were two lags included in the Taylor rule, or lagged values of inflation and activity were considered instead.

3 Results

In the last section, we underscore that the transition incentives faced by departing and incoming governors do not necessarily have to be limited to only the first and last meetings. For instance, a departing governor could influence his successor by changing policy at the penultimate meeting and not making any changes at the very last meeting. Hence, we report results from regression (1) for different specifications. For example, the variable $FM_{c,m}$ ($LM_{c,m}$) may include the first (last) n meetings. If $n = 2$, for instance, the specification consider the first (last) two meetings.

Before reporting the results, one word on inference: throughout the empirical section we use robust standard errors as usual. Even so, in the Appendix A.2 we report how our main results remain essentially the same when we use Driscoll and Kraay (1998) errors, which are resistant to many criticisms against robust errors.

3.1 Baseline

At Table 1, we report the results from regression (1) for specifications with n (the number of meetings) varying from one to four.

Table 1: Main Regression: $i_{c,m}$ is the dependent variable

# Meetings	1	2	3	4
$FirstMeet(\beta_F)$	0.052*	0.075***	0.061**	0.036
	[0.072]	[0.002]	[0.048]	[0.159]
$LastMeet(\beta_L)$	0.088*	0.076**	0.088***	0.089***
	[0.098]	[0.025]	[0.001]	[0.001]
Country FE	Y	Y	Y	Y
Year Dummy	Y	Y	Y	Y
# Obs	3881	3881	3881	3881

P-value between [ ], calculated with robust standard errors.

Table 1 shows that both the first and last few meetings are associated with higher interest rates than those prescribed by the Taylor rule in comparison to regular meetings. These results are statistically significant and robust across specifications. Moreover, they are economically relevant: take column 2, it is an increase of around 0.075 percentage points in the first meetings and 0.076 in final ones. As reference for this magnitude, we note that over 50% of interest changes are of 0.25 percentage points.

Besides being an important aspect of empirical Taylor rules, the interest rate smoothing addresses a concern that the new central banker might not be tightening policy. Assume that the departing banker increases interest rates above the level prescribed by the Taylor rule, generating a positive residual. Even if the new governor did not change policy, it is likely that macroeconomic conditions would have little changed from one meeting to the other so that the residual of the Taylor rule would remain positive. However, smoothing prevents this to be the case. As the smoothing coefficient is quite high (almost always above 0.9), most of the interest rate hike engendered by the departing banker is absorbed by the Taylor rule. Consequently, a positive coefficient of similar magnitude means that there was indeed a further tightening during the first few meetings.

In the following subsections, we show that in fact, the transition effects we estimate diminish as meetings distance themselves form actual change in leadership. Then, in the remainder of the empirical section, we report results considering $n = 2$ and $n = 3$ meetings. In this way we are not limited to a too short time interval nor we are allowing the transition to last too long.[10]

Finally, we argue that results are not driven by the endogeneity of the transition timing or governor's choice. After ruling out straightforward possible explanations that may rationalize results above, such as fiscal policy or political cycles, we propose an explanation based on signalling dynamics. The we use the heterogeneity of transitions to argue that these results are consistent with the proposed explanation.

3.1.1 Decaying effect

We argue above that a departing governor does not have to act precisely on the last meeting; he could tighten monetary policy a bit before and affect his successor in the same way. In particular, we consider different specifications that encompass transitions happening up to four meetings before, and after, the change in leadership.

This argument looses strength as meetings grow more distant from the actual governor's change. After all, policy should return to normal. The objective of this section is to show that as the meetings distance themselves from actual change, the transition effects diminish. In particular, we drop $FM_{c,m}$ and $LM_{c,m}$ from the specification in (1), but add other two dummy variables across seven different specifications. The j-th specification includes one

[10] Results for $n = 1$ and $n = 4$ are available upon request.

dummy variable that accounts for the j-th and $(j+1)$-th first meetings and another one that accounts for the last j-th and $(j+1)$-th meetings. For example, in the first specification, there are two dummies variables accounting for the first and last two meetings, respectively. Similarly, the second specification considers one dummy variable that accounts for the second and third meetings, as well as another dummy to account for the the penultimate and anti-penultimate meetings. The same logic applies for the subsequent specifications.

Figure 3 plots the coefficients for the dummy variables in these seven specifications. The upper (bottom) part of the graph plots the value of the coefficients associated with the first (last) meetings.

Figure 3: Decay: Rolling Transition (2 meetings)

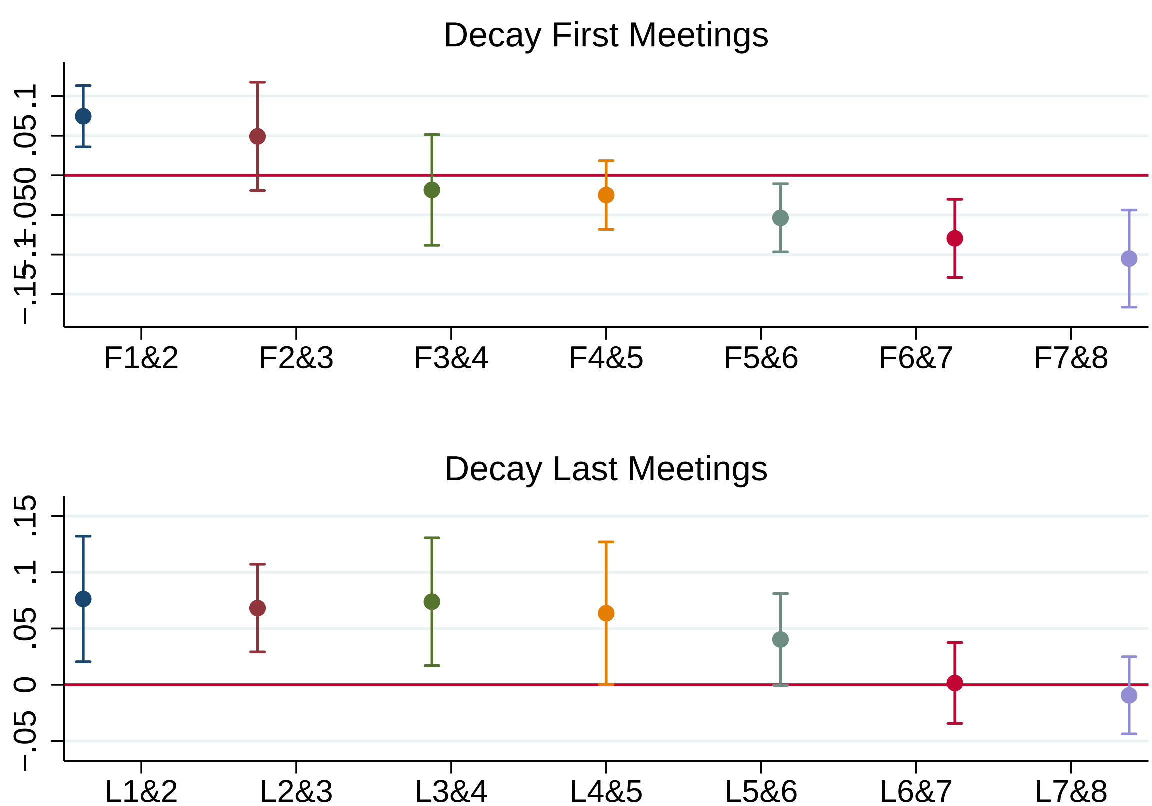

Notice that coefficients associated with last meetings fall as they grow distant from the actual leadership change. In particular, the coefficient becomes indistinct from zero around the 5th meeting. Similarly, coefficients associated with the first meetings fall and become zero around the 3rd meeting. In addition, they become statistically negative before returning to zero. We attribute this pattern to the lagged term included in the Taylor rule. In fact, as policy rate increases above the level prescribed by the Taylor rule in the first meetings, i_{m-1} is artificially higher in the next meetings. Hence, the Taylor rule prescribes an economically inadequate high level of interest rate in the next meetings. If the central bank understands this and does not follow this inadequate prescription of the Taylor rule, residuals are negative for a while until the excess tightening absorbed by the rule dissipates.

We find very similar results when we report the graph for a rolling window of three meetings instead, as shown by Figure 4.

Figure 4: Decay: Rolling Transition (3 meetings)

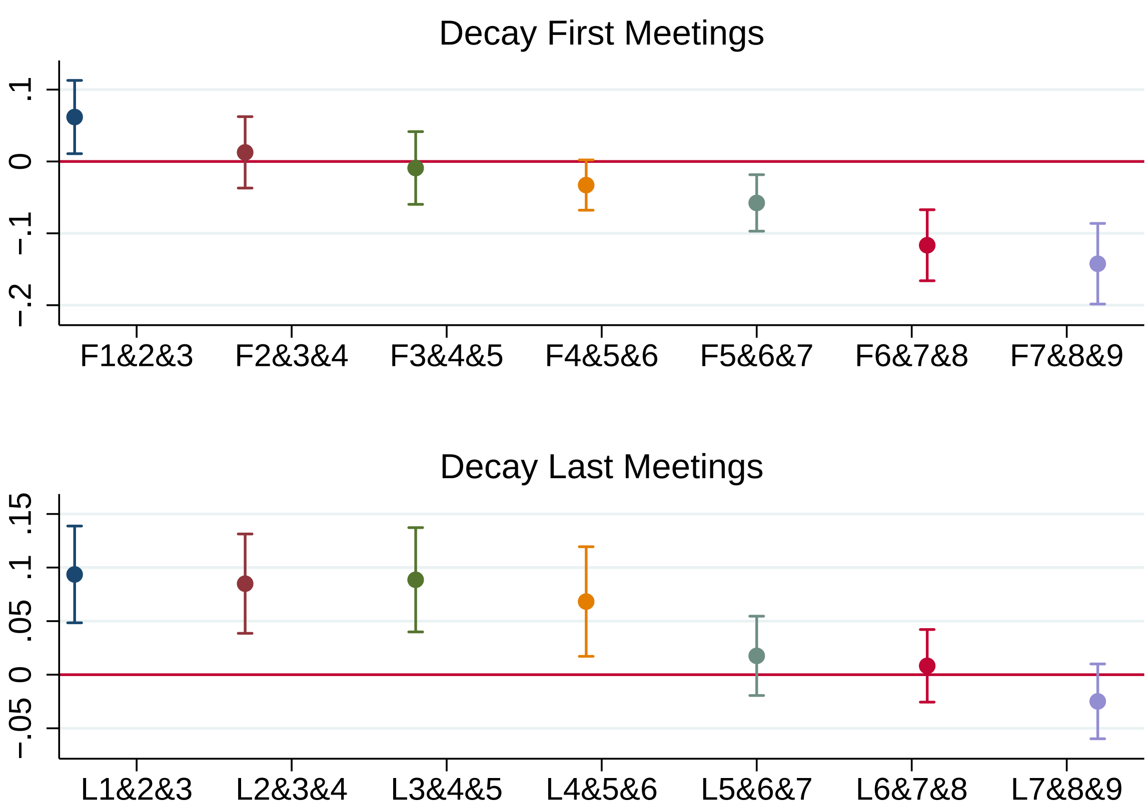

This decaying effect enhances the evidence discussed so far on leadership transition effects.

If there were other factors which do not spring from transitions driving our results, one might expect these factors to have influenced monetary policy for longer than 5 or 6 meetings before transitions. After all, these alleged confounders were strong enough to trigger transitions. In addition, in regard to first meetings effects, it is hard to come up with a confounder that initially is associated with tight monetary policy and then reverts to normal. For instance, a regime change that coincides with a transition should leave consequences for longer periods than the evidence reported in Figures 3 and 4 suggests.

3.1.2 Transition timing

One can argue that transitions are associated with tighter monetary policy stance because they are endogenous. Indeed, transitions could be more likely to occur at times when interest rates are above the level prescribed by the Taylor rule. We make sure that this is not driving our results by exploring transitions' timing. We analyze the heterogeneity between transitions where the governor mandate is fixed (e.g. US) and those where there is no such regularity (e.g. Brazil).

For fixed mandates, as their end is determined in advance, the transition should not have been caused by other factors that could also cause an increase in the interest rate. Consequently, we can rule out cases in which tighter monetary policy triggers leadership transitions. In the first column of Table 2, we report results once we interact $FM_{c,m}$ and $LM_{c,m}$ with $NotFix_c$, a dummy variable that takes value one if the mandate is not fixed, as described by the following equation.

$$i_{c,m} = \rho_c i_{c,m-1} + \alpha_{\pi,c}\pi_{c,m} + \alpha_{y,c}y_{c,m} + \beta_F FM_{c,m} + \beta_{FN} FM_{c,m} \times NotFix_c + \beta_L LM_{c,m} + \beta_{LN} LM \times NotFix_c + \delta_m + \delta_c + e_{c,m} \quad (2)$$

In the equation above, β_F and β_L capture the effects of the transitions where there is a fixed governor's mandate, whereas $\beta_F + \beta_{FN}$ and $\beta_L + \beta_{LN}$ are the effects of those transitions where the governor's mandate is not fixed. As before, β_F and β_L are the coefficients of interest.

The last row of the table separates the number of transitions in countries with and without fixed mandate: 61 and 10 transitions, respectively.

The first two columns of Table 2 show that β_F and β_L remain positive, statistically significant (except for one case), and with an economically relevant magnitude. Therefore the transition effects found in Table 1 do not appear to stem from endogenous transition timing.

Even in countries with fixed mandates, a governor could suddenly resign in a time of

Table 2: Fixed Regimes and Unannounced Resignation

# Meetings	2	3		2	3
FM	0.046**	0.039	FM	0.043*	0.036
	[0.036]	[0.254]		[0.069]	[0.330]
$FM \times NotFix$	0.179**	0.150*	$FM \times Unan$	0.130**	0.105*
	[0.043]	[0.075]		[0.044]	[0.096]
LM	0.048*	0.064***	LM	0.013	0.034*
	[0.062]	[0.003]		[0.577]	[0.083]
$LM \times NotFix$	0.190	0.156	$LM \times Unan$	0.221**	0.204**
	[0.267]	[0.203]		[0.025]	[0.012]
# Trans	$61-10$	$61-10$		55-16	55-16

P-value between [ ], calculated with robust standard errors.
Regressions still include FE and Year dummies.

crisis, which would trigger a transition with endogenous timing. In order to address this concern and be even more cautious, we create the variable $Unan$,[11] which includes not only transitions when the governor's mandate is fixed, but also transitions that followed an unannounced resignation. The coefficients of FM and LM capture the effects of fixed mandate transitions where there was no unannounced resignation. The right panel of Table 2 shows the results.

Notice that there are less transitions to estimate FM and LM, which makes it harder to obtain precise estimates. In fact, as there are only 51 transitions to estimate the effects of interest, the coefficients of FM and LM are not always statistically significant. Nevertheless, as the last two columns of Table 2 highlight, the coefficient of FM is significant using two meetings, whereas the coefficient associated with LM is if three meetings are considered instead. Moreover, the coefficients' magnitudes are broadly similar to those found using all fixed mandate transitions. We interpret the less robust statistical significance as a natural consequence of the reduced number of transitions.

Considering Table 2 as a whole, the empirical evidence suggests that causation direction goes from transitions to tighter monetary policy. There is something going on during transitions that leads to interest rates above the Taylor rule's predictions. It is not the case that results stem from transitions being more likely to occur during periods associated with monetary policy tightening.

[11] Details are in the Appendix A.1.3.

3.1.3 Governor's choice

Even if the leadership transitions are not endogenous, one could argue that the choice of the new governor is. After all, if for some reason, inflation around transitions is too high, the appointment of a hawkish governor is more likely. In this case, the Hawk could increase interest rates above those prescribed by the Taylor rule due to his preferences.

In this section we argue that though this criticism may play a role, it cannot account for the full effect. In order to control for different preferences across governors, we allow the intercept to also vary with governors. Hence, each governor may differ in the average interest rate chosen within his country in a given year. In other words, we add governor's fixed effects. A Hawk governor, for instance, should have a higher fixed effect than a Dove. In this case, β_F captures the difference of first meetings with other meetings for the same governor. A positive β_F implies that, on average, the same governor is more hawkish during his first meetings than throughout the rest of his tenure.

Table 3: Governor Fixed Effects

# Meetings	2	3
FM	0.098^{***}	0.084^{**}
	[0.002]	[0.021]
LM	0.039	0.053^{*}
	[0.289]	[0.071]
Gov FE	Y	Y
# Obs	3881	3881

P-value between [ ], calculated with robust standard errors.

Table 3 shows that results survive after we control for governor's fixed effects. As we add over one hundred governor dummies, of course, estimations loose some precision. Nevertheless, the coefficients less precisely estimated are those associated with LM, which are less threatened by the criticism of endogenous governors' appointments. Indeed, the departing governor was chosen well before the transition. As the aforementioned criticism concerns the coefficients associated with FM, it is reassuring that they remain positive, statistically significant and economically relevant. Consequently, different governors' preferences cannot explain our results.

3.2 Possible explanations

The previous section documents that monetary policy tightens around transitions, and rules out causality running from higher interest rates to transitions. In this section we rule out straightforward possible explanations for our findings. Namely, political transitions, fiscal policy and economic uncertainty. These variables might be confounding factors as they may correlate positively with both transitions in central bank leadership and hikes in interest rate.

We also show that results are not driven by either observations in which or periods when the zero lower bound is a concern. In these situations, the Taylor rule is not appropriate to describe monetary policy. Finally, we show that poor countries, where Taylor rules might not be adequate to describe monetary policy, are not driving our results.

3.2.1 Political transitions

In this section we discuss how transitions in central bank leadership interact with political transitions. The concern is that if the changes at the central bank coincide with changes in government, our results could be driven by monetary policy responding to some government variable we do not account for. In a rough manner, the analysis on transition timing at section 3.1.2 helped somewhat in avoiding this confounder by showing that the results survive even when we focus on countries with fixed central banker's mandate. Indeed, usually the goal of such mandates is precisely to make monetary policy less susceptible to the executive influence and the mandates's length is many times designed to not coincide with political cycles. Notwithstanding our previous efforts, we still feel that this concern warrants a thorough analysis. After all, it not difficult to imagine a scenario where even with fixed mandate, a central banker would schedule his resignation to coincide with a political cycle in order to smooth overall policy changes or even due to unobservable pressures from the government.

As a result, this section compares the effects of CB governor's changes when they coincide with political transitions with the effect of CB governor's changes during normal times. We will do so by interacting our main regressors of interest with dummies which mark election years and beginning of mandate years. Moreover, we distinguish between elections and reelections so that we create 4 dummies:

1. $ElecY_{ct}$: The meeting ct takes place in a election year when the winner is taking office for the first consecutive time.

2. $ReelecY_{ct}$: The meeting ct takes place in a election year when the incumbent wins the election.

3. $BegMandElecY_{ct}$: The meeting ct takes place in the year when a new head of government took office after a election.

4. $BegMandReelecY_{ct}$: The meeting ct takes place in the year when the incumbent head of government took office after a reelection.

Before we begin discussing the regressions, we clarify a few points regarding the data. First, the specific position of the head of government changes across countries. In presidential systems, it is naturally the president, who usually takes office in the year following the election (e.g. US, Brazil). In parliamentary systems, the head of government is the prime minister (even if the country does have a president) who is elected following a general election. In most cases, the prime minister takes office immediately after the election so that $ElecY_{ct}$ and $BegMandElecY_{ct}$ will coincide for most parliamentary systems. Second, it is also important to note that the dummies refer to calendar year and not the 12 month period before/after an election. For instance, if a presidential election in country c takes place in September 2014, all country c meetings in 2014 have $ElecY_{ct}$ or $ReelecY_{ct}$ equal 1. Finally, we note that, for convenience, reelection years refer to when a incumbent succeeded. If a president tries to get himself reelected and loses, the year in question will count as $ElecY_{ct}$ and not as $ReelecY_{ct}$.

Let $W_{ct} \in \{ElecY_{ct}, ReelecY_{ct}, BegMandElecY_{ct}, BegMandReelecY_{ct}\}$ be one of the four dummies described above. Results are reported in Table 4, where each column refers to a specification that considers a single political transition variable.

Notice that our empirical results are not driven by correlation to political cycles. Although we found that first election years are associated with higher interest rates, the coefficients of FM and LM in the first columns have similar magnitudes to the results of Table 1. We can also note that Table 4 coheres with a pattern appearing throughout our empirical analysis: FM is more precisely estimated using 2 meetings as transition whereas LM is more precise with 3 meetings (see Tables 1, 3, 8). Therefore Tables 4 presents evidence consistent with our main results: first and final meetings have tighter monetary policy; they are economically significant (coefficients' sizes are similar to Table 1); and they cannot be explained away by appealing to political cycles involving elections or beginning of mandates.

3.2.2 Fiscal policy

In this section we address the concern that fiscal policy could be driving our results. This could be the case if transitions were more likely to occur in times of fiscal build ups. Stories that justify

Table 4: Political Cycles

Dummy W	*Elec*	*Reelec*	TookOf	Retook
# Meetings	2	2	2	2
FM	0.077***	0.075***	0.075*	0.075***
	[0.001]	[0.001]	[0.001]	[0.001]
LM	0.076**	0.076**	0.076**	0.076**
	[0.022]	[0.024]	[0.023]	[0.023]
W	0.076***	0.013	0.004	-0.005
	[0.000]	[0.582]	[0.838]	[0.845]
# Meetings	3	3	3	3
FM	0.062**	0.061**	0.061**	0.061**
	[0.042]	[0.045]	[0.045]	[0.045]
LM	0.084***	0.086***	0.086***	0.086***
	[0.001]	[0.001]	[0.001]	[0.001]
W	0.076***	0.012	0.003	-0.006
	[0.000]	[0.606]	[0.856]	[0.826]
# Trans	59×12	65×6	56×15	66×5

P-value between [ ], calculated with robust standard errors.
Regressions still include FE and Year dummies.
2 First and 2 Last meetings per transition

this relation usually hinge on some change in government or policy, so we were already protected to some extent against this criticism by our results controlling for political transitions. However, we believe this to be a concern serious enough to warrant particular attention to fiscal developments and how they affect central bank transitions.

The variable we use to account for fiscal development is the ratio of government expenditures to GDP (henceforth GY). This is a quarterly measure which is available for most countries in our sample[12]. The matching of these quarterly data to each central meeting follows the algorithm detailed previously for quarterly unemployment rates. Then, we add $GY_{c,t}$ as an economic factor to the main equation (1). Importantly, we allow the coefficient on this variable to vary with countries. As we let GY to affect the interest rates, the fact that fiscal build ups make the central bank tighten monetary policy is controlled for.

Table 5 reports the results, which remain essentially the same as those of Table 1.

[12]See Appendix A.1.5 for the list of countries for which we have only yearly GY.

Table 5: Fiscal Policy Robustness

# Meetings	2	3
FM	0.087***	0.071**
	[0.000]	[0.022]
LM	0.092***	0.103***
	[0.009]	[0.001]
# Obs	3781	3781

P-value between [ ], calculated with robust standard errors.

3.2.3 Uncertainty

In this section we address the concern that uncertainty could be driving our results. The idea is that the transition in central bank leadership increases the public's uncertainty and that this could explain the monetary contractions. In common theoretical frameworks, the central bank would ease policy to calm markets. Nonetheless, it is possible to imagine a case (involving an emerging market economy, for instance) in which uncertainty increases capital flights, what could force the central bank to increase interest rates. In order to assuage this fear, we include two measures of uncertainty in our Taylor rule: standard deviation of log returns of the main stock market index of each country (Stock) and the standard deviation of log returns of the nominal exchange rate (Exchange). The motivation is straightforward: uncertain periods should be reflected in volatile stock and/or exchange markets. Finally, we allow the coefficients of these variables to vary with countries. Table 6 reports the transition coefficients of interests and shows that our results are not being caused by markets' uncertainty.

Table 6: Uncertainty

# Meetings	2	3
FM	0.085***	0.065**
	[0.005]	[0.034]
LM	0.091***	0.090***
	[0.008]	[0.001]
# Obs	3416	3416

P-value between [ ], calculated with robust standard errors.

3.2.4 Zero lower bound

Another concern is that the zero lower bound could be biasing our results. If the Taylor rules of many countries were predicting negative rates and if more transitions occurred in the countries at the zero lower bound, the residuals would be positive by construction and could make the coefficients pf FM and LM artificially positive. To make sure this is not the case, we change the sample we use to estimate our main regression in two ways. First, we drop every observation in which the interest rate is equal or lower to 0.5 % ($i > 0.5$). Second, we also report the more drastic approach of dropping every observation after the year 2007 ($Y < 2008$). This, however, makes us lose a big portion of the transitions we have to estimate the coefficients, as Figure 1 indicates. Third, we include the crisis years but drop interest rate is equal or lower to 0.5 % ($FS \setminus ZB$). Table 7 shows that most results change little and are unlikely to be caused by biases arising from the zero lower bound.

Table 7: Zero Lower Bound

	$i > 0.5$	$i > 0.5$	$Y < 2008$	$Y < 2008$	$FS \setminus ZB$	$FS \setminus ZB$
# Meet	2	3	2	3	2	3
FM	0.070***	0.056*	0.078*	0.070*	0.039	0.037
	[0.004]	[0.084]	[0.062]	[0.075]	[0.231]	[0.260]
LM	0.074**	0.085***	0.084	0.119***	0.064*	0.082***
	[0.035]	[0.002]	[0.136]	[0.005]	[0.059]	[0.002]
# Obs	3648	3648	2538	2538	4281	4281

P-value between [ ], calculated with robust standard errors.

3.2.5 Poor countries

In this section we show that our results do not rely on one particular group of countries - the poorest in our sample. As we use data from 35 countries, there are varying levels of income (e.g., Norway, Brazil and Kenya). It would decrease the interest of our results were they driven solely by the poor countries. After all, these countries have more idiosyncrasies and often Taylor rules describe inadequately their monetary policy. Hence we decompose our results between the 9 poorest countries, those with GDP per capita less than US$10000 PPP, and the rest.

Table 8 shows that results survive. Coefficients of FM and LM remain positive, statistically significant and economically relevant. These coefficients capture transition effects for countries in the sample that are not poor. In addition the coefficients of $FM \times Poor$ and $LM \times Poor$ show how large, but very noisy, the effects are in poor countries. Results regarding poor countries should be read with caution as the number of transitions considered is only fifteen.

Table 8: Poor countries (GDP per capita less than $10000 PPP)

# Meet	2	3
FM	0.063***	0.061*
	[0.008]	[0.067]
$FM \times Poor$	0.092	0.015
	[0.277]	[0.877]
LM	0.061*	0.069***
	[0.079]	[0.009]
$LM \times Poor$	0.113	0.135
	[0.310]	[0.128]
# Trans	56 − 15	56 − 15

P-value between [ ], calculated with robust standard errors.
Regressions still include FE and Year dummies.

3.3 Proposed explanation

After discarding many possible explanations, we propose one based on signalling dynamics for the results obtained in Section 3.2, i.e. $\beta_F > 0$ and $\beta_L > 0$. Of course, as it is impossible to exhaust all possibilities, other explanations may remain. In the next section, we report heterogenous effects that are consistent with the proposed explanation.

Consider a new governor who knows that the public is uncertain whether she is a Hawk or a Dove.[13] She has incentives to tighten monetary policy in order to signal she is a Hawk and face lower inflation expectations in the future. This mechanism at play translates into a positive β_F. As to the departing governor, if she wants to help a Hawk successor to signal her type, contracting monetary policy would make it harder for a Dove to pretend she is a Hawk. After all, Doves should find it even more costly to tighten monetary policy further after an interest rate increase. As this last result is less intuitive and, to our knowledge, novel, it is the focus of the signalling model

[13]There are plenty of anecdotal evidence that substantiate the idea that the public has great uncertainty about new central bankers. In fact, "So, Mr. Carney, Hawk or Dove" at the WSJ and "ECB: Clearing the way for an Italian hawk?" at the BBC demonstrate such uncertainty. Moreover, central bankers are aware of this special uncertainty, as the quote bellow illustrates.

> "I think any action we take – because we are certainly in the spotlight today – will be looked at very eagerly and there are psychological reactions coming from what we do." Winn, Cleveland Fed President, at Volcker's first meeting.

developed in Section 4. The model makes explicit how a monetary contraction ($\beta_L > 0$) helps to sustain a separating equilibrium.

In the next section, we argue that some heterogeneities in the estimated coefficients are consistent with the signalling dynamics described above. Altogether, we see this bulk of empirical results as suggestive of the presence of signalling in monetary policy.

3.4 Heterogeneities

In this section, we explore some heterogeneities to argue that its direction is consistent with the signalling dynamics described above, in which the new governor must prove she is a Hawk and the departing one wants to help this process. As we cannot exhaust other alternative interpretations, if the heterogeneity goes in the same direction one would expect were signalling true, this provides further evidence of this view and may help discard other confounding stories.

3.4.1 Independence, transparency and regulatory quality

The assumption underlying this section is that signalling incentives should be stronger where there is greater institutional uncertainty. After all, in places where the public trusts that no governor will try to exploit the short term benefit of inflating the economy, there would be little reason for a new governor to signal that she is a Hawk by distorting monetary policy. On the other hand, in places where every leadership change brings back fears of bad policy, a new governor has strong incentives to signal that she is committed to fighting inflation.

This raises the question of how to measure what we referred to as institutional uncertainty. We consider three indices: *Transparency* from Crowe and Meade (2007), *Independence* from Dincer and Eichengreen (2014) and *Regulatory Quality* from the World Bank Governance Indicators. Notice that the first and second indices refer specifically to the Central Banks, while the third refers to the whole country. Although every index has flaws, they should roughly capture this institutional uncertainty. As all indices were constructed by different authors, there is less risk of one methodology being the sole driver of results. Also, since the indices were not constructed to study leadership transitions, we do not worry about hindsight biasing our results.

We interact the variables of interest with these three indices. Results are reported in Table 9.

In fact, the heterogeneity goes in the direction we expect: the more independent, transparent or the better the regulatory quality, the smaller the tightening of monetary policy during first and last meetings is. The interaction coefficients are always negative and most are statistically significant. Even in the cases the coefficients are less precise, the size of the point estimates are very similar to their significant counterparts. In addition, note that the indices' coefficients are economically relevant (indices are normalized) and that transitions are still associated with tight monetary policy for countries with average institutional qualities.

Table 9: CB Independence[a], CB Transparency[b] and Country Regulatory Quality[c]

# Meet	2		2		2
FM	0.073^{***}	FM	0.068^{***}	FM	0.059^{***}
	[0.003]		[0.004]		[0.007]
$FM \times Ind$	-0.038^{*}	$FM \times Trp$	-0.059^{***}	$FM \times RQ$	-0.068^{***}
	[0.070]		[0.004]		[0.002]
LM	0.081^{**}	LM	0.078^{**}	LM	0.071^{**}
	[0.024]		[0.024]		[0.020]
$LM \times Ind$	-0.051^{**}	$LM \times Trp$	-0.040	$LM \times RQ$	-0.043
	[0.016]		[0.155]		[0.180]
# Meet	**3**		**3**		**3**
FM	0.062^{*}	FM	0.062^{**}	FM	0.053^{*}
	[0.056]		[0.060]		[0.065]
$FM \times Ind$	-0.020	$FM \times Trp$	-0.018	$FM \times RQ$	-0.057^{**}
	[0.523]		[0.487]		[0.028]
LM	0.091^{***}	LM	0.089^{***}	LM	0.083^{***}
	[0.001]		[0.001]		[0.001]
$LM \times Ind$	-0.034^{**}	$LM \times Trp$	-0.042^{*}	$LM \times RQ$	-0.045^{*}
	[0.049]		[0.077]		[0.084]

P-value between [ ], calculated with robust standard errors. Indices are normalized
[a] Independence (Crowe and Meade 2007)
[b] Transparency (Dincer and Eichengreen 2013)
[c] Regulatory Quality (World Bank Governance Indicators).

3.4.2 Monetary policy committee

Nowadays most monetary policy decisions are made by committees. Without repudiating the working assumption that leadership transitions are the most relevant, we can use the committee structure to exploit heterogeneities and, consequently, assess if they are consistent with the signalling interpretation offered for our results.

Governor's strength This paper focuses only on leadership transitions, i.e. the change of governor. Hence, wherever the governor is stronger, transition effects are expected to be larger. One can evaluate the strength of the governor according to the characteristics of the committee he belongs. Blinder (2007) proposes the following typology, in increasing order of governor's strength.

1. Individualistic Committee.

2. Genuinely Collegial Committee.

3. Autocratically Collegial Committee.

4. Individual Governor.

According to Blinder, one is characterized by all members being exhorted to vote their own mind, with the governor often on the losing side of the vote (e.g. UK); two is the case in which there is an atmosphere that strives for consensus and thus the governor plays a role in forging this consensus (e.g. ECB or Bernanke); in three the governor plays the dominant role and can shift the board to his preferred policy (e.g. Volcker or Greenspan);[14] four is obviously the case with the strongest governor as he is the sole determiner of policy. We create a variable called *Blin* that classifies central banks according to this topology.

The caveat with this typology is that it is inevitably subjective. For instance, both the UK and the US have similar committee structures on the surface - one vote per member, which is released to the public - but Blinder argues that tradition gives the US governor a much greater sway over the board than the UK one. Despite this caveat, we use Blinder's opinion for the countries he did categorize; search in central bank's staff papers of each country how they categorize their own central bank; and, as a last resort, take our best guess based on the committee structure and minutes. The Appendix A.1.4 discusses in details how this index is constructed.

Table 10 shows that the coefficient of $LM \times Blin$ goes in the direction consistent with signalling. The stronger the governor (higher $Blin$), the stronger the monetary contraction at the last meetings. On the other hand, we did not find, as predicted by signalling dynamics, that the first meeting's effect is greater with stronger governors. In this sense, Blinder's typology provides partial evidence in favor of the signaling interpretation.

Governor was previously part of the committee The assumption behind this exercise is that the public should have a better idea of the type of a new governor if he was already part of the monetary policy committee before he held office. Hence, we create a dummy variable $PrevBoard$ indicating whether the governor was part of the previous board. If this is the case, there is less need of signalling by both old and new governors. Hence we expect a smaller policy tightening at the first meeting (smaller incentives to prove she is a Hawk) and at the last meeting (smaller incentives for the departing to help the signalling process.)

Overall, Table 11 shows that this exercise is inconclusive. We cannot reject the hypothesis that effects at first and last meetings are different different from zero in both specifications. However,

[14]Blinder tells an anecdote in which Greenspan started on the losing vote, but managed to persuade the committee to vote unanimously in favor of his choice.

Table 10: Governor Power - Blinder Index

# Meet	2	3
FM	0.074***	0.064**
	[0.002]	[0.043]
$FM \times Blin$	0.023	0.019
	[0.159]	[0.256]
LM	0.078**	0.088***
	[0.024]	[0.001]
$LM \times Blin$	0.059***	0.057***
	[0.002]	[0.000]

P-value between [ ], calculated with robust standard errors.
It includes country FE and year dummies.

there is some very weak evidence that policy tightening at the first meetings is smaller when the new governor was part of the board. In fact, there is a close to significant effect, with p-value of 13%, regarding the two but not three meetings per transition dummy. Given the substantive decrease of the number of transitions used to estimate the effect of *PrevBoard*, we believe that this finding provides a slight support for our interpretation that signalling dynamics are driving our results.

To sum up, with Blinder's typology, we find strong evidence regarding the last meetings but no significant effect regarding the first meetings. By considering transitions in which governors belonged to the board, we find weak evidence regarding the first meetings. Altogether, Tables 10 and 11, by exploring the structure of monetary policy committees, provide some evidence consistent with the existence of signaling dynamics during transitions.

4 Model

We consider a model built on Kydland and Prescott (1977) and Barro and Gordon (1983b) important contributions.[15] Both papers study the time inconsistency of policy. The objective is to study monetary policy during transitions in central bank leadership. Since contractions in monetary policy in the first meetings are consistent with standard results in the literature on signalling and monetary policy, e.g. Barro (1986), we use the model to justify why a more contractionary policy stance takes place in the last meetings. In particular, we argue that, by tightening monetary

[15] This section is extremely preliminary. It may change substantially in future drafts.

Table 11: Mon. Pol. Committees

# Meet	2	3
FM	0.098^{***}	0.043
	[0.002]	[0.158]
$FM \times PrevBoard$	-0.065	0.053
	[0.128]	[0.474]
LM	0.063	0.065^{*}
	[0.188]	[0.074]
$LM \times PrevBoard$	0.033	0.060
	[0.593]	[0.222]
# Trans	46×25	46×25

P-value between [ ], calculated with robust standard errors.
It includes country FE and year dummies.

policy, the departing central bank facilitates the distinction between Dove and Hawk incoming central banks. Before adapting their models for that purpose, we briefly summarize their main contribution through a basic setup which serves as a benchmark for the rest of the analysis.

4.1 Basic setup

Time is discrete and the horizon is finite, i.e. $t = 1, ..., T$. The relation between output y_t and inflation π_t is given by the following Phillips curve:

$$y_t = y_t^n + a(\pi_t - E[\pi_t]), \tag{3}$$

where y_t^n is the natural level of output, $E[\pi_t]$ is the expected inflation, and $a > 0$ measures the output response to inflation surprises.

For each period t, taking $E[\pi_t]$ as given, the central bank (CB) chooses π_t in order to minimize the current loss function,

$$\frac{\pi_t^2}{2} - \lambda(y_t - y_t^n), \tag{4}$$

subject to the Phillips curve (3).

In equilibrium, rational expectations require that $\pi_t = E[\pi_t]$. The classic result of inconsistency arises. In particular, the desire to stimulate output leads to positive inflation, $\pi_t = E[\pi_t] = a\lambda > 0$, without output gains, i.e. $y_t = y_t^n$. In contrast, if in $t = 1$, the CB could credibly commit to $\pi_t = 0$ for $t = 1, ..., T$, then society would be better as $\pi_t = E[\pi_t] = 0$ and $y_t = y_t^n$ arise in equilibrium.

We define $\kappa \equiv a\lambda$, which is the inflationary bias that arises in this basic setup.

4.2 Novel elements

In order to study monetary policy decisions during transitions, we add two ingredients to the basic setup.

First, inflation π_t comprises the sum of two components,

$$\pi_t = \gamma\pi_{t-1} + (1-\gamma)\pi_t^c, \quad \text{for} \quad t = 1, ..., T, \tag{5}$$

where $\gamma \in (0, 1)$ measures the degree of inertia in the economy and π_t^c is the inflation under control of the CB. Hence, π_{t-1} is the state variable and π_t^c is the control variable.

This extension is necessary to connect the decisions of different central bankers through time. Indeed, π_{t-1}^c chosen by the previous CB would affect current inflation π_t and, thus, the current CB's choice of π_t^c.

Second, the CB not only cares about current inflation π_t but also about inflation under control π_t^c. In particular, the current loss function reads

$$\frac{\theta(\pi_t^c)^2}{2} + \frac{\pi_t^2}{2} - \lambda(y_t - y_t^n), \tag{6}$$

where $\theta > 0$ measures the weight attributed to the controllable part of inflation. If θ is low (high), we say that the CB is dove (hawk).

This extension is necessary to generate non-trivial dynamics. Otherwise, if there is no cost to change inflation π_t^c (i.e. $\theta = 0$), then the CB could simply adjust π_t^c to set total inflation π_t at its preferred level. As a result, previous inflation π_{t-1} becomes irrelevant.

These two ingredients, inertial inflation and losses from changing π_t^c, allow us to transform the basic setup, inspired by Kydland and Prescott (1977) and Barro and Gordon (1983b), into a dynamic model. Inertial inflation is an intuitive assumption, easily motivated by some degree of nominal rigidity built in contracts or some kind of indexation. In contrast, the assumption that, apart from total inflation π_t, inflation under control also enters the loss function merits some digression.

We offer two interpretations for $\theta > 0$. The first is that it is costly to change inflation. In practice, the CB does not control inflation directly. Instead, it controls policy instruments, such as the interest rate, that affects inflation. One finds many reasons in the literature to avoid abrupt changes in the interest rate: to avoid financial stress (Cukierman, 1991); better control over long-term interest rates (Woodford, 2003); politico-economic costs associated with committee decision making (Riboni and Ruge-Murcia, 2010). If the central banker cares about any of these reasons, it will find costly to change the part of inflation under control today from its optimum level.

The second interpretation is that θ can capture career concerns. The public may consider the inherited state of the economy when evaluating the competence of a CB. Hence, central bankers that deliver the same inflation rate, but inherit different ones, should be perceived differently. If the CB cares about how competent it is perceived to bring inflation close to zero, there is an extra cost associated with generating inflation under its control, π_t^c.

In order to show that signalling dynamics are consistent with the empirical results above, we assume that θ is private information. In particular, $\theta \in \{\theta^H, \theta^D\}$, with $\theta^H > \theta^D$, where H and D stand for Hawk and Dove, respectively. Thus the Hawk CB finds inflation under control π_t^c more costly. As the public tries to infer CB's type from its actions, the model becomes a signaling game: there may be separating, pooling or mixed equilibria depending on the parameters.

4.3 Central bankers' problems

We assume that at $t = 1$, a new central bank (NCB) takes office inheriting inflation π_0 from the old central bank (OCB). Although π_0 is treated as a parameter in the model, we interpret it as a choice variable of the OCB. Indeed, by doing comparative statics in π_0, one may inspect the properties of the equilibrium that the OCB can induce. In particular, π_0 affects the existence conditions of separating and pooling actions and, as a result, how the beliefs update process unfolds. We show that a reduction in π_0 is warranted if the OCB wishes to foster type revelation, thereby substantiating the tight monetary stance in last meetings found in the data.

We stress that we do not model the OCB's decision process explicitly for simplicity and easy of exposition. The argument above implicitly assumes that, for some reason, facilitating type revelation yields utility to the OCB. Importantly, it is possible to model an OCB with similar preferences to the NCB so that the OCB finds optimal to increase the likelihood of separating. The main assumption we need to add in this case is that the OCB knows the NCB's type and cares enough about future periods after its tenure.

After taking office at $t = 1$, the NCB stays in office for T periods and discounts the future with $\beta \geq 0$. Its loss function at $t = j$ reads:

$$L_j = \sum_{t=j}^{T} \beta^{t-j} \left[\frac{\theta(\pi_t^c)^2}{2} + \frac{\pi_t^2}{2} - \lambda(y_t - y_t^n) \right]. \tag{7}$$

Notice that the NCB knows its decision at t influences its later decision at $t+1$ through the state variable π_t. This kind of mechanism is found elsewhere in the literature as in Alesina and Tabellini (1990) and in Debortoli and Nunes (2013).

In order to be consistent with the basic setup in Section 4.1, we assume that expected inflation $E[\pi_t]$ is set before the NCB chooses its control variable π_t^c, but $E[\pi_{t+1}]$ is set after NCB's choice at t. Hence, the NCB takes current inflation expectations as given but recognizes that, in equilibrium,

$E[\pi_{t+1}] = \pi_{t+1}$ and, thus, $y_{t+1} = y^n_{t+1}$. In other words, the NCB knows it cannot stimulate output in the following period.

Finally, for the rest of the paper, we assume that $T = 2$. It is the smallest number of periods that allows the model to capture an important incentive faced by the NCB. In particular, the Dove NCB may want to mimic the Hawk NCB in order to face more favorable inflation expectations in the next period.

In the Appendix B.1, we solve the full information case. The rest of section 4 deals directly with the incomplete information case.

4.4 Analysis

One shortcoming of signalling models is that different beliefs can sustain multiple equilibria for a given set of parameter values. Multiple equilibria hinder the analysis of the mechanisms at play. In order to circumvent this problem, we consider a specific set of beliefs in line with Cukierman and Liviatan (1991) and Walsh (2000). In particular, agents always expect a Hawk NCB to choose its preferred action as if it did not fear being mistaken for a Dove NCB. Thus, whether actions are pooled or separated depends on the Dove NCB's choice. If it prefers to mimic the Hawk NCB's choice of inflation, actions are pooled. If, instead, it prefers to reveal its type by choosing its preferred level of inflation, then actions are separated. This "refinement criterion" guarantees uniqueness of equilibrium.

Let π^{cH}_{1S} and π^{cD}_{1S} be the Hawk and Dove, respectively, NCBs' preferred choice of π^c_1 when actions are separated at $t = 1$. Similarly, let π^c_{1P} be chosen by the Hawk NBC when it is expected that the Dove NCB pool its actions at $t = 1$.

Let $\mu \in (0, 1)$ be the prior probability that the NCB's type is θ^H. The public has the following expectations in a separating equilibrium: $E[\pi^c_{1S}] = \mu\pi^{cH}_{1S} + (1 - \mu)\pi^{cD}_{1S}$. Of course, in a pooling equilibrium, the expected controlled inflation is the chosen one, π^c_{1P}. Notice that expectations' formation embed results from the optimization problems of each central banker, which is close to the spirit of Kydland and Prescott (1977). In the Appendix B.3 we discuss how different refinement criteria may alter our results.

In the Appendix B.2, we characterize the optimal levels of inflation at $t = 1$ and $t = 2$ for each possible equilibrium path. There are two possible equilibrium paths with pure strategies. The first entails NCBs separating their actions at $t = 1$ and $t = 2$, whereas the second considers NCBs pooling their actions at $t = 1$ but separating them $t = 2$. Although actions cannot be pooled at $t = 2$ as both Dove and Hawk NCBs prefer to separate, we call the first (second) separating (pooling) equilibrium due to what happens in the first period. We also consider the case with mixed strategy. For each possible equilibrium, given the relevant set of beliefs, we characterize the unique path of inflation levels for each type of the NCB. The uniqueness follows from our refinement criterion as the Hawk NBC is expected to play its preferred action without fear of being mistaken for a Dove NBC. Hence, the Hawk NBC always chooses its single preferred action in the sense that it minimizes its loss function.

4.4.1 Equilibria existence

In this section, we discuss conditions that determine the existence of a particular equilibrium. In the Appendix B.2, we characterize the candidates for an equilibrium, but it remains to show that it in fact exists.

The definition of equilibrium has three requirements. First, each type of the NCB minimizes its loss function taking current expectations and beliefs as given, but accounting for the effect of its choice on future expectations and beliefs. Second, expectations are rational; that is, expectations must reflect the weighted average, with beliefs determining the weights, of the Hawk and Dove NBCs' equilibrium strategies. Third, beliefs are updated following a Bayes' rule on the equilibrium path.

Recall that we name an equilibrium after what happens in first period.

A - Separating equilibrium. Let μ_{2S} be the belief in a separating equilibrium that the NCB is a Hawk at $t = 2$. Bayes' rule and our refinement criterion imply that beliefs are updated according to:

$$\mu_{2S} = \begin{cases} 1, & \text{if } \pi_1 = \pi_{1S}^H \\ 0, & \text{if } \pi_1 = \pi_{1S}^D \\ 0, & \text{if } \pi_1 \neq \pi_{1S}^H \text{ or } \pi_1 \neq \pi_{1S}^D \end{cases}.$$

In words, if observed inflation is different from the equilibrium one chosen by a Hawk NCB, agents believe the NCB is Dove, updating their beliefs to zero. Notice that we rely on our refinement criterion to make explicit how beliefs are updated off the equilibrium path.[16] If they observe the equilibrium inflation chosen by a Hawk NCB, then beliefs are updated to one. The precise formula for π_{1S}^i, $i \in \{D, H\}$, can be found in equation (18) of Appendix B.2.

In order to confirm that the separating equilibrium exists, one must check whether the Dove or Hawk NCB has incentives to deviate from its equilibrium strategy, i.e. its choice for inflation, given expectations and beliefs. Consider the Hawk NCB. If it chooses an inflation rate different from π_{1S}^{cH}, agents will think it is a Dove NCB in the second period and, thus, expected inflation will be higher. This outcome not only worsens welfare in the second period, but also in the first

[16] Recall that the refinement criterion requires the Hawk NCB to always choose its preferred actions π_{tS}^H at $t = 1, 2$, which are its choices when it does not fear being mistaken by a Dove NCB.

period. In fact, π_{1S}^{cH} was found by minimizing the Hawk NCB's loss function, taking expectations as given. Hence, if the Hawk NCB deviates, it harms itself in every period.

Alternately, the Dove NCB could potentially improve its welfare by pretending to be a Hawk, i.e. by choosing π_{1S}^{cH}, in order to generate lower expected inflation in the second period. Hence, the Dove NCB faces a tradeoff: it can choose its preferred level of inflation at $t = 1$ and reveal its type; or it can pretend to be the Hawk NCB at $t = 1$ and improve its welfare at $t = 2$.

Let L_S^D be the loss of the Dove NCB associated with a separating equilibrium. Define L_{SD}^D as the loss associated with deviating from the prescribed equilibrium strategy and trying to pass itself for a Hawk , i.e. to choose π_{1S}^{cH}. For the separating equilibrium to exist, it cannot be profitable for a Dove NCB to pretend to be the Hawk NCB, given beliefs and expectations. Hence, it is required that:

$$L_S^D \leq L_{SD}^D.$$

Proposition 1. *For γ small enough, there exists $\beta^S \geq 0$ such that $L_S^D \leq L_{SD}^D$ for all $\beta \in [0, \beta^S]$, and that $L_S^D > L_{SD}^D$ for all $\beta \in (\beta^S, \infty)$.*

The intuition of this proposition is straightforward. For β small, the Dove NCB cares less about the second period, choosing its preferred inflation level and, thus, engendering a separating equilibrium. Alternatively, for β large, the Dove NCB cares more about the second period and, thus, the benefits accrued from lower expected inflation at $t = 2$ surpass the costs of pretending to be Hawk at $t = 1$.

B - Pooling equilibrium. Let μ_{2P} be the belief in a pooling equilibrium that the NCB is a Hawk at $t = 2$. Recall that π_{1P}^c is the inflation rate chosen by a Hawk NCB in a pooling equilibrium, whose precise formula is equation (19) in Appendix B.2. Bayes' rule and our refinement criterion imply that beliefs are updated according to:

$$\mu_{2P} = \begin{cases} \mu, & \text{if } \pi_1 = \pi_{1P} \\ 0, & \text{if } \pi_1 \neq \pi_{1P} \end{cases}$$

If agents observe anything other than the equilibrium inflation of a Hawk NCB expecting to be imitated, they will revise their beliefs zero, so to be sure the NCB is Dove. Otherwise, beliefs are not updated. Notice that the Hawk NCB has no incentives to deviate as it would be worse off in both periods. Indeed, it would incur a cost at $t = 1$ and face higher expectations at $t = 2$. In contrast, the Dove NCB may wish to deviate from the pooling equilibrium as π_{1P} is not the Dove NCB's preferred inflation at $t = 1$.

Let L_P^D be the loss of the Dove NCB associated with a pooling equilibrium. Define L_{PD}^D as the loss associated with deviating from the prescribed strategy in a pooling equilibrium. For the pooling equilibrium to exist, it cannot be profitable for a Dove NCB to deviate and reveal its type, given beliefs and expectations. Hence, it is required that:

$$L_P^D \leq L_{PD}^D.$$

Proposition 2. *For γ small enough, there exists $\beta_P \geq 0$ such that $L_P^D \leq L_{PD}^D$ for all $\beta \in (\beta^P, \infty)$, and that $L_P^D > L_{PD}^D$ for all $\beta \in [0, \beta^P]$.*

The intuition behind this proposition is quite straightforward. A Dove NCB mimics a Hawk NCB's strategy at $t = 1$ in order to face lower expectations at $t = 2$. Consequently, if the Dove NCB does not care much about the future, i.e. β is low, it never plays the pooling strategy and the equilibrium collapses. Alternatively, if β is large, any extra loss borne at $t = 1$ is acceptable because of the welfare gain at $t = 2$.

In summary, whether a separating or pooling equilibrium prevails depends mainly on the discounting factor β: for β low, it is not worth to mask oneself as Hawk in order to improve future expectations – separating equilibrium prevails; for β large, it is worth to sacrifice one's favorite choice at $t = 1$ for more favorable expectations at $t = 2$ – pooling equilibrium prevails. For values of β that cannot sustain either a separating or a pooling equilibrium, a mixed strategy equilibrium arises, in which the Dove NCB randomizes between pooling and separating actions (we discuss this case further ahead and in Appendix B.2.3). Provided that the conditions of Propositions 1 and 2 are satisfied, the dependence on β can be depicted on the Figure 5 below. Notice also that $\beta^S \leq \beta^P$ as both pooling and separating equilibria cannot coexist due to our refinement criterion.

Figure 5: Types of Equilibrium Depending on β

4.4.2 Comparative Statics in π_0

In the previous section, we show how the kind of prevailing equilibrium depends on the discount factor β. However, it also hinges on the inherited inflation π_0, which the NCB treats as exogenous. If we interpret a decrease in π_0 as a monetary contraction in the last meetings of the OCB, one may inspect, through the lens of this model, theoretical reasons that corroborate our empirical finding.

In this section, we show that a reduction in π_0 increases the likelihood of separating actions, and thus, helps the Hawk NCB build reputation.

Pure Strategies Equilibria

In the context of pure strategies equilibria, we show that a reduction in π_0, makes it easier (harder) to sustain a separating (pooling) equilibrium. In other words, by tightening policy, the OCB helps the public to discover whether its successor is a Hawk or not. This is captured by the Proposition 3, which is the relevant theoretical result when looking at pure strategies equilibria. Define $\Delta^S(\pi_0) \equiv L^D_{DS} - L^D_S$ and $\Delta^P(\pi_0) \equiv L^D_{DP} - L^D_P$, which are the gains a Dove NCB gets by deviating from a separating and pooling equilibriums, respectively. Let $\bar{\pi}_0$, which depends on the parameters of the model, be an upper bound on π_0.[17]

Proposition 3. *For $\pi_0 < \bar{\pi}_0$, $\Delta^S(\pi_0)$ decreases in π_0 and $\Delta^P(\pi_0)$ increases in π_0.*

Proposition 3 states that a reduction in π_0 leads to an increase (decrease) in the loss difference from deviating from a separating (pooling) equilibrium, thereby making this deviation less (more) attractive for the Dove NCB. In other words, a contraction in monetary policy makes separating more attractive and pooling less attractive.

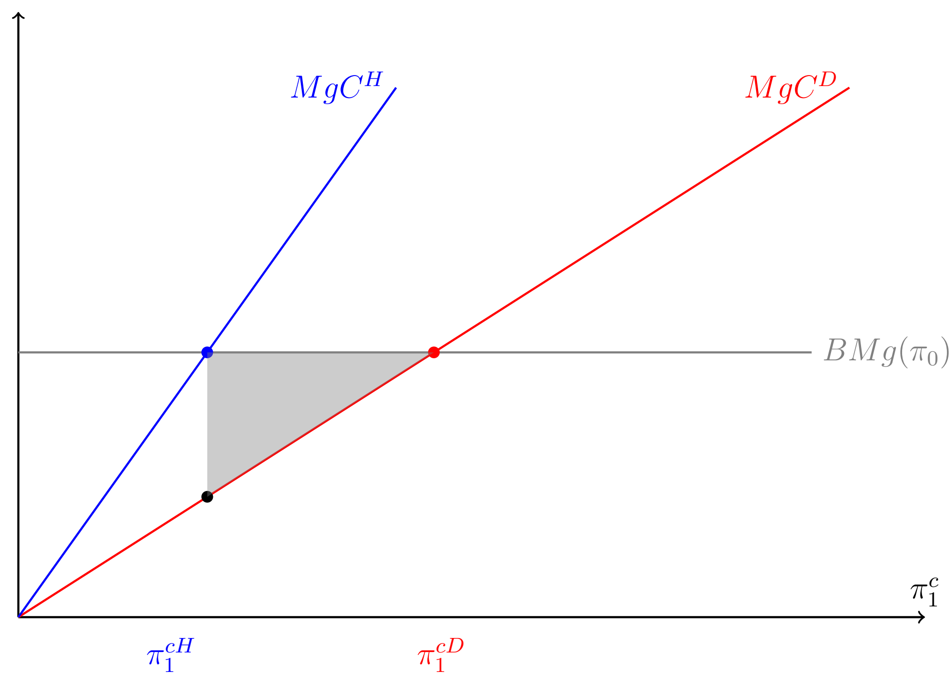

Figure 6: Intuition for Proposition 3 - Part I

The intuition behind Proposition 3 can be represented in Figures 6 and 7. The preferred actions for Hawk (blue circle) and Dove (red circle) NCBs are located at the points that the marginal cost

[17] We discuss the intuition for this upper bound in the proof of Proposition 3.

of choosing controlled inflation at $t = 1$ equalizes its benefit. The linearity of the marginal costs is a consequence of the typical quadratic loss specification. Moreover, $\theta^H > \theta^D$ implies that MgC^H is steeper than MgC^D and, as a consequence, $\pi_1^{cH} < \pi_1^{cD}$. The area of the shaded gray triangle represents the cost a Dove NCB incurs for mimicking the Hawk's preferred action - it is the area where the marginal benefit of π^c is above the marginal cost for the Dove NCB.

The fall in π_0 falls causes an increase in the marginal benefit for both types (or, equivalently, a shift downwards in the marginal cost lines). As it can be seen in Figure 7, this increase the cost of a Dove passing himself as Hawk, which is represented by the larger dark gray triangle. This is the main mechanism by which a tighter monetary stance of OCB helps to sustain a separating equilibrium at the same time it hinders the sustainment of a pooling equilibrium.

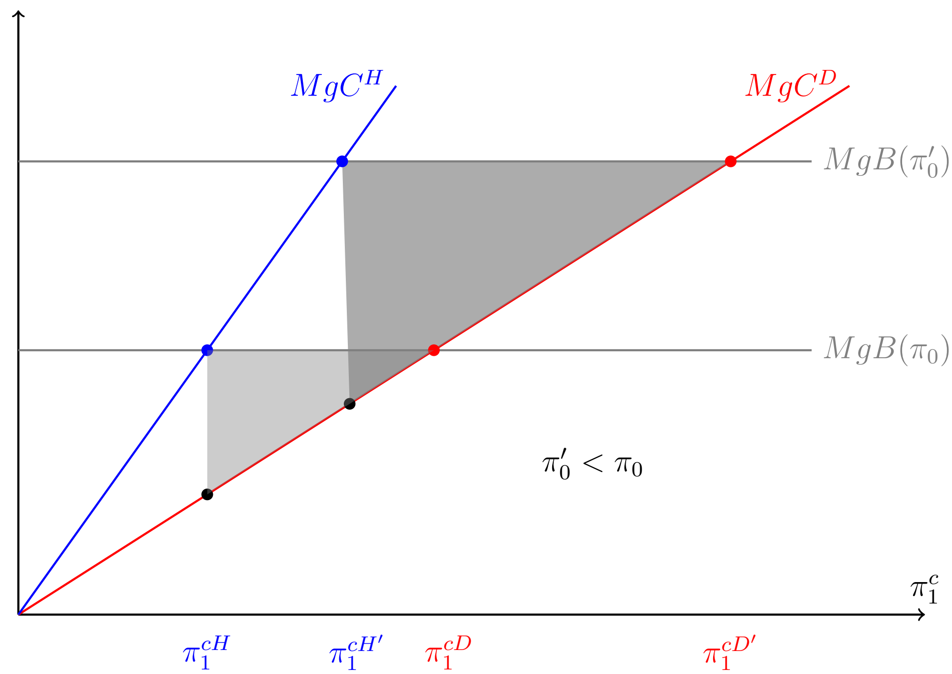

Figure 7: Intuition for Proposition 3 - Part II

Proposition 3 therefore helps to substantiate our empirical findings regarding central bank leadership transitions. If the OCB wishes to reveal whether NCB is a Hawk or a Dove, Proposition 3 prescribes a tighter policy stance in the last meetings. After all, this tight policy stance will make it harder for a Dove to pretend to be a Hawk. Mapping directly with interest rate choice, only a true Hawk would raise interest rates on top of the recent increase conducted in the last meetings.

Having discussed the intuition behind Proposition 3, as well as its mapping with the empirical part, we now provide further details linking π_0 to β_S and β_P, the threshold values of discounting that trigger the existence of equilibrium (see Propositions 1 and 2). Then we will discuss the analogous result to Proposition 3 for mixed strategies equilibrium.

Changes in π_0 alter the values of β that sustain a type of pure strategy equilibrium. The tighter the monetary policy is, for instance, the more patient the NCBs can be without destroying the separating equilibrium. Conversely, they must have an even higher discount factor if a pooling equilibrium is to be sustained (see Figure 8). Naturally, π_0 also alters the values of other parameters of the model, e.g. γ, that help sustain each equilibrium. We focus, however, on β because it has the most intuitive effect on the type equilibria: patience fosters the pooling equilibrium and undermines the separating one.

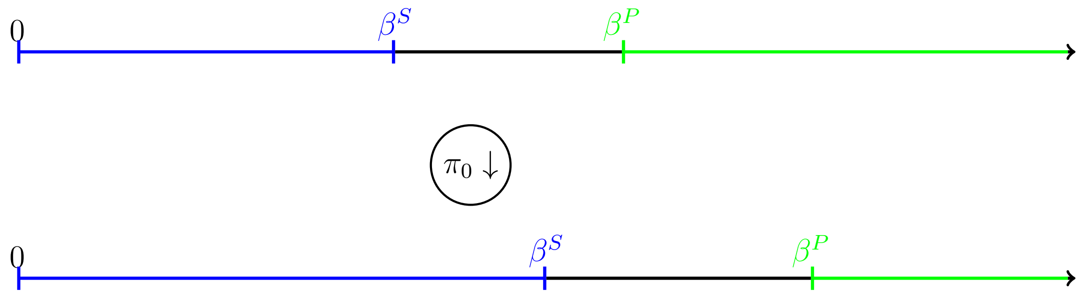

Figure 8: Parameter Space

Mixed Strategies Equilibrium

As mentioned before, for certain intermediate values of β, it is possible that neither a separating nor a pooling equilibrium in pure strategies is sustained. That is, the Dove NCB is not patient enough too pool its action, but also not impatient enough to separate it. In other words, if both NCBs play pure strategies, the Dove NCB will always decrease its losses by deviating if the Hawk NCB expects the equilibrium to be either pooling or separating.

Recall that our refinement criterion implies that the Hawk NCB plays its preferred action without fear of being mistaken for a Dove. Hence, in a mixed strategy equilibrium, only the Dove NCB mix its strategies. In particular, it separates its action with probability α and pools it with the Hawk's choice with probability $1-\alpha$. If actions are separated at $t=1$, beliefs that the NCB is a Hawk are updated to zero. If actions are pooled instead, Bayes' rule implies that beliefs at $t=2$ are given by $\mu_{post} = \frac{\mu}{\mu+(1-\mu)(1-\alpha)}$.

By minimizing the Hawk NCB's loss function properly accounting for the evolution of beliefs, one obtains Hawk's choice of inflation at $t=1$, π_{1M}^{c} (see equation (20) in the Appendix B.2.3). The Dove NCB mixes with probabilities α and $1-\alpha$ its separating and pooling actions, π_{1M}^{c} and π_{1S}^{cD}, respectively.

The value of α is pinned down by equalizing the loss functions associated with the Dove NCB's pooling and separating strategies. In other words, the Dove NCB must be indifferent between separating from and pooling with the Hawk NCB's inflation choice.

The following proposition reinforces Proposition 3 by stating that a contraction in monetary policy, i.e. a reduction in π_0, increases the probability α that actions are separated. Again, if the OCB wishes to reveal the NCB's type by separating actions, Proposition 4 prescribes a contraction in monetary policy in the last meetings, a result that is in line with the empirical finding above.

Proposition 4. *For $\pi_0 < \bar{\bar{\pi}}_0$ and $\mu > \underline{\mu}$, α decreases in π_0.*

Figure 9, which plots α as a function of π_0 in a calibrated version of the model, illustrates Proposition 4. The intuition is the same behind Proposition 3. A lower π_0 decreases the cost of deviating from a separating equilibrium. In a mixed strategy equilibrium, this translates into higher probability α the separating action is played. In addition, for the Dove NCB to be indifferent between both actions, the pooling equilibrium must be more attractive at $t = 2$ to compensate the costs of mimicking incurred at $t = 1$. The higher value of α ensures this by increasing μ_{post}, which implies a lower inflation expectations at $t = 2$.

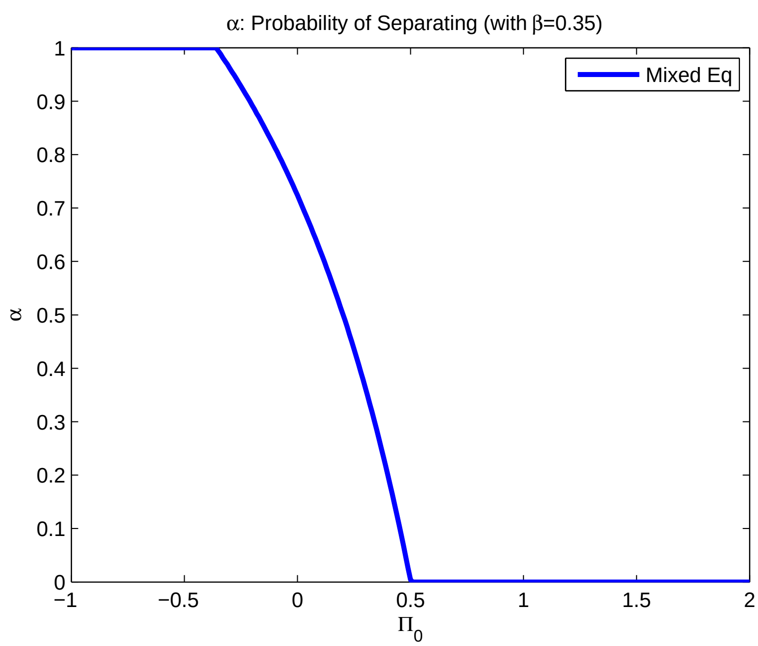

Figure 9: Mixed Strategies Equilibrium: α as a function of π_0

As μ_{post} increases with α, a monetary policy contraction also raises μ_{post}. Intuitively, as the probability α of the Dove NCB choosing the separating action increases, if agents observe the pooling action, π^c_{1M}, then they attribute a higher probability μ_{post} that the NCB is a Hawk. Thus, the OCB affects the belief updating process.

The mixed strategies case has the advantage of showing how the fall in π_0 affects an endogenous variable of the model α instead of altering the parameter space that sustains one equilibrium or the other. Indeed, the effect on α has the obvious interpretation of an increase in the likelihood of "separating actions" arising. This is precisely the explanation we offer for our empirical finding that, on average, monetary policy is tight in the last meetings of departing central bank leaders.

4.5 Discussion

As discussed above, the model was designed to understand the empirical result of policy tightening in the last meetings, which cannot be explained by standard models in the literature on monetary policy and signalling. A shortcoming of the model is that equilibrium outcomes in the first meetings of a NCB have a less clear mapping with our empirical results. In fact, whenever a Dove NCB separates and reveals its type, it sets a higher level of controlled inflation – i.e., it loosens policy – contrasting with the empirical result of tight policy in the first meetings. Other papers in the literature that might be invoked to explain this empirical fact would be subject to the same criticism. However, in our case, this criticism is more accute, since we argue that monetary policy becomes more contractionary in the last meeting precisely to induce a separating equilibrium.

One way to reconcile model and data is to argue that most transitions in the data regard incoming Hawks (relative to the public's beliefs) rather than Doves. However, this cannot be easily checked. Instead, we conjecture that a slight modification to the model could circumvent this shortcoming. What seems to be needed is a more gradual belief updating process, such that types are not revealed immediately. One way to do so is to add some noise to inflation, so that agents can never be sure whether the NCB is a Dove or a Hawk – even in a separating equilibrium. In particular,

$$\pi_1 = \gamma\pi_0 + \pi^c_1(1-\gamma) + u, \text{ where } u \text{ i.i.d } N(0, \sigma^2).$$

In this case, agents still update using Bayes' rule, but it is a gradual process that preserves uncertainty regarding the NCB's type. Hence, even if a Dove NCB does not profit from mimicking the Hawk NCB completely, it still has incentives to choose a lower inflation than it would otherwise. Indeed, there is an incentive to approach the Hawk NCB's choice slightly, to induce more favorable posterior beliefs on the part of the public. Without this noise u, small reductions in inflation produced no benefit for a Dove NCB. It had to choose between mimicking a Hawk or revealing itself as a Dove. Once the noise u is added to inflation, the Dove NCB may not desire a pooling equilibrium, but since it does not fully reveal its type, it may have incentives to tighten monetary policy a bit. Results regarding first meetings and the model would be thus reconciled. In addition, as we argue above, early tightening due to signalling concerns has already been explored by the literature on monetary policy and signalling.

5 Conclusion

We document a novel fact. Namely, transition periods in central bank leadership are robustly associated with tighter monetary policy. We argue that this result is unlikely to stem from endogenous transitions; this argument is built upon the timing of the transitions and the government's choice. Furthermore, after assessing and discarding other possible explanations for this result, we offer an interpretation based on signaling dynamics. A new governor tightens policy to signal she is a Hawk, whereas the departing governor tightens policy to foster type revelation.

In order to explore our explanation in the data, we consider several heterogeneities between transitions to assess whether the results go in the direction one would expect were the proposed mechanism true. Indeed, we show that results are stronger when the central bank has less independence, is less transparent, and when the country's regulatory quality is lower. In addition, they are stronger when the outgoing governor has more power. Although we cannot affirm that our explanation is the only acceptable one, the fact that these heterogeneities go in the expected direction suggests that signaling dynamics may be an important part of the story.

Finally, we build a simple model to formalize this idea. Our model shows how a monetary tightening carried out by the departing central bank can foster type revelation. In a pure strategies equilibrium, a contraction in monetary policy can engender a separating equilibrium. In a mixed strategies equilibrium, a contraction in monetary policy affects the belief updating process, leading agents' posterior that the new central bank is a Hawk to increase. If the departing governor wants to foster type revelation, or help the incoming Hawk to build reputation, these results rationalize the tight monetary stance in the last meetings that are borne out in the data.

References

Alberto Alesina and Guido Tabellini. A Positive Theory of Fiscal Deficits and Government Debt. *Review of Economic Studies*, 57(3):403–14, July 1990.

David Backus and John Driffill. Inflation and reputation. *American Economic Review*, 75(3):pp. 530–538, 1985a.

David Backus and John Driffill. Rational expectations and policy credibility following a change in regime. *Review of Economic Studies*, 52(2):pp. 211–221, 1985b.

Laurence Ball. Time-consistent policy and persistent changes in inflation. *Journal of Monetary Economics*, 36(2):329 – 350, 1995.

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

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

Robert J. Barro and David B. Gordon. A positive theory of monetary policy in a natural rate model. *Journal of Political Economy*, 91(4):589–610, August 1983b.

Alan S. Blinder. Monetary policy by committee: Why and how? *European Journal of Political Economy*, 23(1):106–123, March 2007.

Olivier Coibion and Yuriy Gorodnichenko. Why Are Target Interest Rate Changes So Persistent? *American Economic Journal: Macroeconomics*, 4(4):126–62, October 2012.

Christopher Crowe and Ellen E. Meade. The Evolution of Central Bank Governance around the World. *Journal of Economic Perspectives*, 21(4):69–90, Fall 2007.

Alex Cukierman. Why does the fed smooth interest rates? In Michael T. Belongia, editor, *Monetary Policy on the 75th Anniversary of the Federal Reserve System*, pages 111–147. Springer Netherlands, 1991.

Alex Cukierman and Nissan Liviatan. Optimal accommodation by strong policymakers under incomplete information. *Journal of Monetary Economics*, 27(1):99–127, February 1991.

Davide Debortoli and Ricardo Nunes. Lack of commitment and the level of debt. *Journal of the European Economic Association*, 11(5):1053–1078, October 2013.

Davide Debortoli and Ricardo Nunes. Monetary Regime Switches and Central Bank Preferences. *Journal of Money, Credit and Banking*, 46(8):1591–1626, December 2014.

N. Nergiz Dincer and Barry Eichengreen. Central Bank Transparency and Independence: Updates and New Measures. *International Journal of Central Banking*, 10(1):189–259, March 2014.

John C. Driscoll and Aart C. Kraay. Consistent Covariance Matrix Estimation With Spatially Dependent Panel Data. *The Review of Economics and Statistics*, 80(4):549–560, November 1998.

Stephen Hansen and Michael McMahon. First impressions matter: Signalling as a source of policy dynamics. Working paper, University of Warwick, 2014.

Eric D. Johnson, Michael A. Ellis, and Diana Kotenko. Consensus building on the fomc: An analysis of end of tenure policy preferences. *Economics Letters*, 117(1):368 – 371, 2012.

Robert G. King, Yang K. Lu, and Ernesto S. Pastén. Managing expectations. *Journal of Money, Credit and Banking*, 40(8):1625–1666, 2008.

Finn E Kydland and Edward C Prescott. Rules rather than discretion: The inconsistency of optimal plans. *Journal of Political Economy*, 85(3):473–91, June 1977.

Alessandro Riboni and Francisco J. Ruge-Murcia. Monetary Policy by Committee: Consensus, Chairman Dominance, or Simple Majority? *The Quarterly Journal of Economics*, 125(1):363–416, February 2010.

Kenneth Rogoff. The Optimal Degree of Commitment to an Intermediate Monetary Target. *The Quarterly Journal of Economics*, 100(4):1169–89, November 1985.

Christina D. Romer and David H. Romer. Choosing the Federal Reserve Chair: Lessons from History. *Journal of Economic Perspectives*, 18(1):129–162, Winter 2004.

Anne Sibert. Monetary policy with uncertain central bank preferences. *European Economic Review*, 46(6):1093 – 1109, 2002.

Anne Sibert. Monetary policy committees: Individual and collective reputations. *The Review of Economic Studies*, 70(3), 2003.

John Vickers. Signalling in a model of monetary policy with incomplete information. *Oxford Economic Papers*, 38(3):pp. 443–455, 1986.

Carl E Walsh. Market Discipline and Monetary Policy. *Oxford Economic Papers*, 52(2):249–71, April 2000.

Carl E. Walsh. *Monetary Policy and Theory.* MIT Press, 3rd edition, 2010.

Michael Woodford. Optimal Interest-Rate Smoothing. *Review of Economic Studies*, 70(4):861–886, October 2003.

Appendix

A Data and empirical robustness

A.1 Data

A.1.1 Transition coding

In this section we detail different ways to code the transitions in the data, and explain our choice. Table 12 reports six possibilities, A through F. The following example illustrates some of the challenges in our coding decisions. Assume that one governor's term expires, but the body of government responsible for nominating the new governor has not yet announced its decision. In such case an acting governor must be conducting the meetings. This acting governor may be later appointed to office, and thus, become the official governor. When should the first meeting be labeled? As soon as she becomes the acting governor (possibilities D or E)? Or after she is officially appointed (possibility C)? Similarly, when should the last meeting be labeled? These are the choices one must make. In this section we explain our options.

Table 12: Coding of different transitions

Transition	Assumed Office	Will Become Official	Was Acting Before	Acting Predecessor
A	Y	-	N	N
B	Y	-	N	Y
C	Y	-	Y	-
D	N	Y	-	Y
E	N	Y	-	N
F	N	N	-	-

The dummy variable $FM_{c,m}$ (first meetings) considers transitions of types A, B and C in Table 12. Therefore we consider the first meeting as the one right after the new governor took office. We rule out transitions in which the new central banker was an acting governor. Our view is that an acting governor might not have power to change policy much so to print his own mark. In other words, there is a stand-by until the leadership appointment is settled. Similarly, the dummy variable $LM_{c,m}$ (last meetings) considers transitions of types A, B, E and F in Table 12. This is consistent with the aforementioned argument. Indeed, we discard cases in which the final meeting would happen before the governor is officially appointed but the same person was the acting governor before. After all, the arguments presented throughout the paper consider different people in a transition period.

Notice that our definition of transition implies that a departing governor may not know the identity of its successor. This is not a problem as the theoretical results hold whether the old central bank (OCB) knows its successor type or not.

A.1.2 Countries

Table 13 reports the list of countries in the sample; the total number of meetings per country; the first and last year each countries appear in the sample; and the number of different governors (excluding changes during the financial crisis) per country.

A.1.3 Fixed regimes and unannounced resignation

We determine whether a country has a fixed regime for central bankers by checking whether an appointment also specifies how long a governor remains in charge. Countries without fixed mandates for central bankers (at some point of the sample) are Brazil, Colombia, Peru and Thailand. There are of course some caveats. For instance, Tunisia has fixed mandate for central bankers, but in practise none of them stay until the end of their term. Tunisian government seems to have enough power to change the central bank's leader at will. These caveats do not undermine our results because we also analyze the heterogeneity of unannounced resignations. Places where a fixed mandate is not respected in practice have many unannounced resignations. In contrast, resignations in countries such as Norway are announced in advance, many times to match the calendar year. This antecedence addresses the problem of endogenous timing. If the mandate was fulfilled or the resignation was announced in advance, it is unlikely that transitions are happening due to tighter monetary policy. We consider the resignation to be announced when the public acknowledges it at least two months before. Any choice of months is inevitably arbitrary, it is reassuring that results change little when we consider one or three months as the cutoff for announced resignations.

Table 13: Countries

	Country	# Meetings	First Year	Last Year	# Governors**
1	Albania	103	2001	2014	2
2	Australia	270	1990	2014	3
3	Brazil	156	1999	2014	3
4	Chile	169	2000	2013	4
5	Colombia	242	1995	2014	2
6	Czech Rep	172	1998	2014	4
7	ECB	214	1999	2014	3
8	Georgia	61	2008	2014	1
9	Ghana	60	2002	2014	2
10	Guatemala	82	2005	2014	3
11	Hungary	141	2002	2014	3
12	India	60	2005	2014	2
13	Indonesia	110	2005	2014	3
14	Israel	231	1995	2014	4
15	Japan*	162	1998	2013	2
16	Kenya	48	2006	2014	2
17	Mexico	94	2005	2014	2
18	New Zealand	121	1999	2014	3
19	Nigeria	60	2003	2014	3
20	Norway	128	1999	2013	2
21	Pakistan	40	2005	2014	5
22	Peru	162	2001	2014	5
23	Philippines	124	2002	2014	2
24	Poland	182	1999	2014	4
25	Serbia	124	2007	2014	3
26	South Africa	78	2001	2014	1
27	South Korea	183	1999	2014	5
28	Sweden	174	1994	2014	3
29	Switzerland	62	2000	2014	4
30	Thailand	112	2001	2014	4
31	Tunisia	175	2000	2014	5
32	Turkey	115	2005	2014	3
33	United Kingdom	202	1997	2014	3
34	United States	300	1984	2013	3
35	Uruguay	25	2007	2013	1

* Between March 2001 and February 2006, Japan's monetary target was money growth. We drop these meetings from the sample

** This is the number of governors ignoring changes during 2007-2008.

A.1.4 Governor's strength Index

The goal of this section is to make as clear as possible how we extended the typology proposed by Blinder (2007) to the set of countries present in our sample. Before we begin, we are the first to point out that this is a tentative extension which used more heuristics than ideal.

The procedure used in the paper was: first we checked whether Blinder himself had classified some of the countries; second we searched for papers (usually from central bank staff) where the authors applied Blinder's typology to their own country; third, lacking the previous options, we assessed the committee structure and its minutes and made our best guess regarding which of the 4 types is the best fit for the country in question. We assign number from 1 to 4 according to:

1. Individualistic Committee.
2. Genuinely Collegial Committee.
3. Autocratically Collegial Committee.
4. Individual Governor.

As there is no clear cut classification in some countries, we allow the index to vary in 0.5 increments to reflect such uncertainty. In addition, we allow different governors within a country to be classified differently, though we only do that for a couple of countries where there are strong reasons to do so: United States (following Blinder), Israel and South Korea.

Blinder (2007) classified 9 countries of our sample. In order of governor's strength: New Zealand, Canada, Australia, US, Japan, Switzerland, Euro Area, Sweden and UK. He also admits that his classification is a subjective one.

> "I have ranked the same nine banks on their degree of "democracy" in making monetary policy decisions - ranging from the individual governor in New Zealand to the Bank of England's highly-democratic Monetary Policy Committee. This ranking is admittedly subjective, but I checked it with several colleagues and made some modifications of my original views - an ersatz Delphi method."

In Table 14 we report the classification for each country following the proceeding outlined above. In the cases the classification derived from a paper/staff report, we also document the webpage of the paper in question. In the cases Blinder and the reports were silent, our best guess was based on the committee structure discussed in a central bank webpage (e.g. decomposition of nominal votes seems to indicate less governor's strength).

Table 14: Countries

	Country	Blinder Index	Webpage
1	Albania	2.5	
2	Australia	3	
3	Brazil	2.5	
4	Chile	2	
5	Colombia	2	
6	Czech Rep	1.5	
7	ECB	2	
8	Georgia	3.5	https://www.nbg.gov.ge/index.php?m=553
9	Ghana	3	
10	Guatemala	2.5	
11	Hungary	1	
12	India	3	http://rbi.org.in/scripts/BS_SpeechesView.aspx?Id=395
13	Indonesia	2	http://www.bis.org/publ/work262.pdf
14	Israel*	3.5 /2	
15	Japan	2.5	
16	Kenya	2	
17	Mexico	2.5	
18	New Zealand	4	
19	Nigeria	2.5	http://www.bis.org/events/fmda07.pdf
20	Norway	3	http://www.bis.org/publ/work274.pdf
21	Pakistan	2.5	
22	Peru	2.5	
23	Philippines	2	http://www.bsp.gov.ph/downloads/EcoNews/EN12-05.pdf
24	Poland	1.5	http://www.suerf.org/download/collmay11/ppt_/1sirchenko.pdf
25	Serbia	2.5	
26	South Africa	3	http://www.scielo.br/pdf/rep/v31n4/06.pdf
27	South Korea**	3/1	http://www.kmfa.or.kr/paper/econo/2008/12.pdf
28	Sweden	1	
29	Switzerland	2.2	
30	Thailand	2	http://www.bis.org/publ/work262.pdf
31	Tunisia	3	
32	Turkey	2	
33	United Kingdom	1	
34	United States***	3/2	
35	Uruguay	2	

* Israel changed from 3.5 to 2 in 2013 following a big change in how the committee was organized.
**South Korea's classification is 3 until 2002 and 1 starting in 2013 as explained in the paper cited in the webpage column.
*** US's classification is 3 for the Volcker and Greenspan period and 2 for the Bernanke period.

A.1.5 Fiscal policy

For most countries there is quarterly data on the ratio of government expenditures to GDP. Alas, some countries only report, to the best of our knowledge, yearly data on GY. These countries are: Bangladesh, Ghana, Kenya, Nigeria, Pakistan and Tunisia. In addition, at times there was not data available for the whole time series. This resulted in the loss of 100 observations, from 3881 in Table 1 to 3781 in Table 5.

A.2 Driscoll-Kraay standard errors

The results in the empirical analysis report the usual robust standard errors in every table. However, the features of the data are such that could be reason to worry about serial correlation or spatial dependence in the error term, the latter being a problem associated with country panels as the cross-sectional unit is nonrandom and countries are likely to be subject to common disturbances. There are different ways of addressing this issue within a panel. The most common approach in the microeconometric literature is to control for clustering within the cross section variable, countries in our case. A more popular approach when dealing with countries specifically is to use Driscoll and Kraay (1998) errors.

Table 15: Driscoll-Kraay errors: $i_{c,t}$ is the dependent variable

# Meetings	2	3
$FirstMeet$ (β_F)	0.075**	0.061**
	[0.014]	[0.041]
$LastMeet$ (β_L)	0.076*	0.088**
	[0.083]	[0.014]
Country FE	Y	Y
Year Dummy	Y	Y
# Obs	3881	3881

P-value between [ ], calculated with Driscoll-Kraay errors.

The advantage of Driscoll-Kraay errors is that they are robust to "very general forms of spatial and temporal dependence as the time dimension becomes large". In other words, its asymptotic properties rely on large T holding N fixed, which is a closer description of our panel data. In fact, our data comprise a small number of countries but long time periods for a given country. In contrast, clustered errors are consistent as the number of clusters goes to infinite, which is hardly the case here. Moreover, one needs an even greater number of clusters when some countries span decades while others span only a few years.

We report in Table 15 results analogous to Table 1 but with Driscoll-Kraay standard errors instead. Results are still significant. Although p-values increase a bit, they are fairly small considering the large set of errors dependence Driscoll-Kray method corrects for.

A.3 Alternative Taylor rules

A.3.1 Two lags of the interest rate

In this paper, the Taylor rule with lagged interest rate follows the standard specification in the literature. This lag captures the fact that interest rates are very persistent due to many factors that lead central banks to avoid abrupt policy changes.[18] However, Coibion and Gorodnichenko (2012) point out the empirical need to include two lags in the Taylor rule. Table 16 reports results analogous to Table 1 but considering two lags of the interest rate.

Table 16: Taylor rule has two lags of $i_{c,m}$ as regressors

# Meetings	2	3
$FirstMeet$ (β_F)	0.060**	0.048
	[0.019]	[0.133]
$LastMeet$ (β_L)	0.056*	0.060**
	[0.063]	[0.014]
Country FE	Y	Y
Year Dummy	Y	Y
# Obs	3881	3881

P-value between [ ], calculated with robust standard errors.

Results are weakened but they survive. As we add an extra lag to absorb part of the variability in the data, this weakening is expected.

A.3.2 Lagged values of inflation and activity

In this section we assess how the baseline results would change if we consider lagged values of inflation and activity, π_{m-1} and y_{m-1}, in the Taylor rule instead of current values, π_m and y_m. This exercise captures the idea that current inflation and activity might not be available information for the central bank when it chooses interest rate i_m. Of course this an extreme consideration as

[18]See Cukierman (1991), Woodford (2003) and Riboni and Ruge-Murcia (2010).

central banks have a fairly good idea of the current state of the economy. As expected, results reported in Table 17 are indistinguishable from the baseline results displayed in Table 1.

Table 17: π_{m-1} and y_{m-1} in the Taylor rule instead of π_m and y_m

# Meetings	2	3
$FirstMeet$ (β_F)	0.069***	0.064**
	[0.007]	[0.038]
$LastMeet$ (β_L)	0.086**	0.096***
	[0.017]	[0.001]
Country FE	Y	Y
Year Dummy	Y	Y
# Obs	3881	3881

P-value between [ ], calculated with robust standard errors.

B Model Appendix

B.1 Full information

In this section, we assume that θ is publicly known. Although we are mainly interested in the incomplete information scenario, we compute the equilibrium under full information, which serves as benchmark for the rest of the analysis. By defining $\kappa \equiv a\lambda$, which is the inflationary bias in the basic setup, and substituting (3) in (7), the loss functions of the NCB at $t = 1$ and $t = 2$ are given by

$$L_1 = \frac{\theta(\pi_1^c)^2}{2} + \frac{\pi_1^2}{2} - \kappa(\pi_1 - E[\pi_1]) + \beta\left(\frac{\theta(\pi_2^c)^2}{2} + \frac{\pi_2^2}{2} - \kappa(\pi_2 - E[\pi_2])\right)$$

and

$$L_2 = \frac{\theta(\pi_2^c)^2}{2} + \frac{\pi_2^2}{2} - \kappa(\pi_2 - E[\pi_2]),$$

respectively.

We use backward induction to solve the model. First we solve NCB's problem at $t = 2$: given the state variable π_1, the NCB chooses π_2^c in order to maximize L_2 subject to (5). Then, bearing in mind that π_2^c is a function of π_1, we solve NCB's problem at $t = 1$: given the state variable π_0, the NCB chooses π_1^c in order to maximize L_1 subject to (5). Hence, we can compute equilibrium inflation rates π_1 and π_2 as functions of π_0.

B.1.1 Problem at $t = 2$

At $t = 2$, given expected inflation under control $E[\pi_2^c]$ and past inflation π_1, the NCB solves:

$$\min_{\pi_2^c} \frac{\theta(\pi_2^c)^2}{2} + \frac{\pi_2^2}{2} - \kappa(\pi_2 - E[\pi_2])$$

$$\text{s.t.} \quad \pi_2 = \gamma\pi_1 + (1-\gamma)\pi_2^c \quad \text{and} \quad E[\pi_2] = \gamma\pi_1 + (1-\gamma)E[\pi_2^c].$$

After inserting the restrictions into the objective function, the first order condition (FOC) with respect to π_c^2 yields:

$$\pi_2^c = \frac{(1-\gamma)\{\kappa - \gamma\pi_1\}}{\theta + (1-\gamma)^2}. \tag{8}$$

By plugging (8) at (5) with $t = 2$, one obtains:

$$\pi_2 = \gamma\pi_1 + (1-\gamma)\pi_2^c = \frac{\kappa(1-\gamma)^2 + \theta\gamma\pi_1}{\theta + (1-\gamma)^2}. \tag{9}$$

Notice that $\frac{\partial\pi_2^c}{\partial\pi_1} < 0$ and $\frac{\partial\pi_2}{\partial\pi_1} > 0$. The intuition is straightforward. An increase in π_1 raises the marginal cost of inflating the economy and, thus, entailing a lower π_2^c. However, this decrease in π_2^c is not large enough to compensate the direct increase in π_2 due to the inertial effect. Algebraically,

$$\frac{\partial \pi_2}{\partial \pi_1} = \underbrace{\gamma}_{\text{inertial effect} > 0} + \underbrace{(1-\gamma)\frac{\partial \pi_2^c}{\partial \pi_1}}_{\text{marginal cost effect} < \gamma} > 0.$$

B.1.2 Problem at $t = 1$

At $t = 1$, given $E[\pi_1]$, the NBC takes into account its future choice π_2^c as a function of π_1, computed above, and solves its problem. That is,

$$\min_{\pi_1^c} \frac{\theta(\pi_1^c)^2}{2} + \frac{\pi_1^2}{2} - \kappa(\pi_1 - E[\pi_1]) + \beta\left(\frac{\theta(\pi_2^c)^2}{2} + \frac{\pi_2^2}{2}\right)$$

$$\text{s.t.} \quad \pi_2 = \gamma\pi_1 + (1-\gamma)\pi_2^c, \quad \pi_1 = \gamma\pi_0 + (1-\gamma)\pi_1^c \quad \text{and} \quad \pi_2^c = \frac{\kappa(1-\gamma) - (1-\gamma)\gamma\pi_1}{\theta + (1-\gamma)^2}.$$

Recall that we assume that $E[\pi_{t+1}]$ is set after NCB's choice at t. Hence, rational expectations require that $E[\pi_2] = \pi_2$ be incorporated in the NBC's problem at $t = 1$.

In the previous section, we show that $\frac{\partial \pi_2^c}{\partial \pi_1} < 0$ and $\frac{\partial \pi_2}{\partial \pi_1} > 0$. In addition to the inconsistency that generates the inflationary bias and the effect of inherited inflation, the choice of inflation under control π_1^c also balances the trade-off between its opposing effects on π_2^c and π_2. The following equation, derived after some manipulation of the FOC with respect to π_1^c, highlights these forces.

$$((1-\gamma)^2+\theta)\pi_1^c = \underbrace{(1-\gamma)}_{\frac{\partial \pi_1}{\partial \pi_1^c}}\left[\underbrace{\kappa}_{\text{infl. bias}} - \underbrace{\gamma\pi_0}_{\text{inherited infl.}} - \beta\overbrace{\left(\frac{\partial \pi_2}{\partial \pi_1}\underbrace{\pi_2}_{\text{mg cost of } \pi_2} + \underbrace{\theta\pi_2^c}_{\text{mg cost of } \pi_2^c}\frac{\partial \pi_2^c}{\partial \pi_1}\right)}^{\text{cost of an increase of } \pi_1 \text{ in } t=2}\right]. \tag{10}$$

After some algebra, equation (10) yields:

$$\pi_1^c = \frac{(1-\gamma)\{\kappa - \gamma\pi_0(1+\beta\varphi)\}}{\theta + (1-\gamma)^2(1+\beta\varphi)}, \tag{11}$$

where $\varphi \equiv \frac{\theta\gamma^2}{\theta+(1-\gamma)^2}$. By plugging (11) at (5) with $t = 1$, one obtains:

$$\pi_1 = \gamma\pi_0 + (1-\gamma)\pi_1^c = \frac{\kappa(1-\gamma)^2 + \theta\gamma\pi_0}{\theta + (1-\gamma)^2(1+\beta\varphi)}. \tag{12}$$

Notice that $\frac{\partial \pi_1^c}{\partial \pi_0} < 0$ and $\frac{\partial \pi_1}{\partial \pi_0} > 0$. A similar intuition applies. The decrease in π_1^c due to a higher marginal cost to inflate the economy is not large enough to compensate for the direct inertial effect in π_1.

Equations (12) and (9) characterize the equilibrium levels of inflation at $t = 1$ and $t = 2$, respectively.

B.2 Incomplete information

To solve the model under incomplete information, we use backward induction. First, in Section B.2.1, we characterize equilibrium inflation at $t = 2$ for both types of NCB, Dove and Hawk, under two cases, whether the actions at $t = 1$ are pooled or separated. Second, in Section B.2.2, we characterize the inflation levels at $t = 1$ for both types, Dove and Hawk, in both pooling and separating equilibria.

Recall that agents start with a prior $\mu \in (0, 1)$ that the NCB is Hawk. After observing the outcomes at $t = 1$, agents update their prior using Bayes' rule. In pure strategy equilibria, only two forms of update are possible. If actions are separated, agents observe either π_{1S}^{cH} or π_{1S}^{cD}, and thus, the probability that the NBC is Hawk is updated to one if π_{1S}^{cH} is oberved or zero otherwise. If actions are pooled instead, the prior is not updated and, thus, μ is still the probability attached for a NBC being Hawk.

B.2.1 Problem at $t = 2$

A - Actions are separated at $t = 1$. This is the simplest case. As agents know the type $i \in \{H, D\}$ of the NCB, the problem is equivalent to the full information case at $t = 2$, except that inherited inflation is π_{1S}^i.[19] See equations (8) and (9).

$$\pi_{2S}^{ci} = \frac{(1-\gamma)\{\kappa - \gamma \pi_{1S}^i\}}{\theta^i + (1-\gamma)^2} \tag{13}$$

$$\pi_{2S}^{i} = \gamma \pi_{1S}^i + (1-\gamma)\pi_{2S}^{ci} = \frac{\kappa(1-\gamma)^2 + \theta^i \gamma \pi_{1S}^i}{\theta^i + (1-\gamma)^2} \tag{14}$$

B - Action are pooled at $t = 1$. In this case, agents' prior μ was not updated as both types chose the same level of inflation at $t = 1$. Hence, $E[\pi_{2P}^c] = \mu \pi_{2P}^{cH} + (1-\mu)\pi_{2P}^{cD}$.[20] Note, however, that the NCB takes expectations as given and, thus, the FOC of its problem does not rely on them. Therefore, π_{2P}^{ci} is also equivalent to the full information case at $t = 2$, except that inherited inflation

[19] Note that π_{tS}^i is the equilibrium inflation at t if the NCB is of type i, given that previous period actions were separated. If we add the superscript c, π_{tS}^{ci} refers to controllable part of inflation.

[20] Notice that the subscript P refers to the case that actions are pooled at $t = 1$.

is π_{1P}.

$$\pi_{2P}^{ci} = \frac{(1-\gamma)\{\kappa - \gamma\pi_{1P}\}}{\theta^i + (1-\gamma)^2} \tag{15}$$

$$\pi_{2P}^{i} = \gamma\pi_{1P} + (1-\gamma)\pi_{2P}^{ci} = \frac{\kappa(1-\gamma)^2 + \theta^i\gamma\pi_{1P}}{\theta^i + (1-\gamma)^2} \tag{16}$$

Since $\theta^H > \theta^D$, it is easy to verify that both $\pi_{2P}^{cD} > \pi_{2P}^{cH}$ and $\pi_{2P}^{D} > \pi_{2P}^{H}$.[21] Facing the same expected and previous inflation rates, the Dove NBC chooses a higher inflation. As a consequence of $E[\pi_{2P}^c]$ being a weighted average of both equilibrium inflation levels under a Dove and Hawk NBC, the Dove NBC manages to stimulate output above its natural level, while the Hawk NBC brings output below its natural level.

B.2.2 Problem at $t = 1$

A - Separating equilibrium. Recall that we call an equilibrium separating (pooling) if actions are separated (pooled) in the firs period. In a separating equilibrium, each type $i \in \{D, H\}$ solves the following problem.

$$\min_{\pi_{1S}^{ci}} \frac{\theta^i(\pi_{1S}^{ci})^2}{2} + \frac{(\pi_{1S}^{i})^2}{2} - \kappa(\pi_{1S}^i - E[\pi_{1S}]) + \beta\left(\frac{\theta^i(\pi_{2S}^{ci})^2}{2} + \frac{(\pi_{2S}^{i})^2}{2}\right)$$
$$\text{s.t.} \quad (13), \quad (14) \quad \text{and} \quad \pi_{1S}^i = \gamma\pi_0 + (1-\gamma)\pi_{1S}^{ci}.$$

As usual, the NCB takes expectations and π_0 as given. Moreover, the NCB understands that a action being separated at $t = 1$ implies that output in the second period cannot be stimulated. Hence, we impose $\pi_{2S}^i = E[\pi_{2S}]$ in the loss function above. This argument relies on the timing of the model, which states that current public expectations are set before the NCB's current choice of inflation, but after then previous NBC's choice. Except for expected inflation levels, e.g. $E[\pi_{1S}^c] = \mu\pi_{1S}^{cH} + (1-\mu)\pi_{1S}^{cD}$, this problem is analogous to the full information case. However, since expectations are taken as given, equilibrium inflation levels for each type are equivalent to the full information case at $t = 1$. See equations (11) and (12).

$$\pi_{1S}^{ci} = \frac{(1-\gamma)\{\kappa - \gamma\pi_0(1+\beta\varphi^i)\}}{\theta^i + \beta(1+(1-\gamma)^2\varphi^i)}, \tag{17}$$

where $\varphi^i \equiv \frac{\theta^i\gamma^2}{\theta^i+(1-\gamma)^2}$.

$$\pi_{1S}^{i} = \gamma\pi_0 + (1-\gamma)\pi_{1S}^{ci} = \frac{\kappa(1-\gamma)^2 + \theta^i\gamma\pi_0}{\theta^i + (1-\gamma)^2(1+\beta\varphi^i)}. \tag{18}$$

[21] This holds as long as π_{2P}^{ci} is positive, which is true in equilibrium.

B - Pooling equilibrium. In a pooling equilibrium, both types choose the same action at $t = 1$, π^c_{1P}, and thus, agents must expect this level of inflation. Given expectations, both Dove and Hawk NCBs prefer different inflation levels. However, if they did act on these preferences, their types would be revealed, undermining the pooling equilibrium. Hence, at least one type of the NCBs must choose inflation different from its preferred level. The only type that has incentives to do so is the Dove NCB so it can face lower expected inflation and, thus, a better tradeoff at $t = 2$.

Consequently, in order to find the pooling equilibrium, we focus on the problem faced by the Hawk NCB, whose behavior is mimicked by the Dove NCB. When the Hawk NCB solves its problem, it takes into account that π^c_{1P} will be a state variable at $t = 2$ for the subcase detailed in B.2.1.B. Its problem reads:

$$\min_{\pi^c_{1P}} \frac{\pi^2_{1P}}{2} + \frac{\theta^H(\pi^c_{1P})^2}{2} - \kappa(\pi_{1P} - E[\pi_{1P}]) + \beta\left(\frac{(\pi^H_{2P})^2}{2} + \frac{\theta^H(\pi^{cH}_{2P})^2}{2} - \kappa(\pi^H_{2P} - E[\pi_{2P}])\right)$$

$$\text{s.t.} \quad \pi^H_{2P} = \gamma\pi_{1P} + (1-\gamma)\pi^{cH}_{2P}, \quad (16) \quad \text{and} \quad \pi_{1P} = \gamma\pi_0 + (1-\gamma)\pi^c_{1P}$$

In contrast with the separating equilibrium, there will be an inflation surprise at $t = 2$, $\pi^H_{2P} - E[\pi_{2P}]$. The inflation surprise stems from the fact that the belief μ is not updated when both types pool their actions. Hence, at $t = 2$, expected inflation averages both equilibrium inflation levels under a Dove and Hawk NCB. In this case, the Hawk NCB can affect this inflation surprise through the choice of π^c_{1P}.

After taking the FOC and rearranging the terms, one obtains:

$$\pi^c_{1P} = \frac{(1-\gamma)\{\kappa\left(1 + \frac{\beta(1-\mu)\gamma(1-\gamma)^2(\theta^H-\theta^D)}{(\theta^H+(1-\gamma)^2)(\theta^D+(1-\gamma)^2)}\right) - \gamma\pi_0\left(1+\beta\varphi^H\right)\}}{\theta^H + (1-\gamma)^2\left(1+\beta\varphi^H\right)} \tag{19}$$

In a pooling equilibrium, both Dove and Hawk NCBs choose inflation equal to π^c_{1P}. At $t = 1$, as expectations are proven correct, output remains at its natural level. At $t = 2$, a negative (positive) inflation surprise arises if a Hawk (Dove) NCB is in office, bringing output below (above) its natural level. Notice that if $\mu = 1$, there would be no inflation surprise at $t = 2$ and, thus, $\pi^c_{1P} = \pi^{cH}_{1S}$.

Finally, it is noteworthy that $\pi^c_{1P} \geq \pi^{cH}_{1S}$. This holds because an increase in inflation at $t = 1$ makes, under a Hawk NCB, the inflation surprise less negative at $t = 1$.

B.2.3 Mixed equilibrium

In a mixed strategies equilibrium, the Dove NCB randomizes between pooling and separating their actions. Let $\alpha \in (0, 1)$ be the probability the separating action is played. If the actions are separated at $t = 1$, beliefs that the NCB is Hawk are updated to zero. If actions are pooled instead, Bayes' rule implies that beliefs at $t = 2$ is given by $\mu_{post} = \frac{\mu}{\mu+(1-\mu)(1-\alpha)}$.

By minimizing its loss function properly accounting for the evolution of beliefs, the Hawk NCB's FOC, after some algebra, yields:

$$\pi_{1M}^{c} = \frac{(1-\gamma)\kappa\left(1+\frac{\beta(1-\alpha)(1-\mu_{post})\gamma(1-\gamma)^2(\theta^H-\theta^D)}{(\theta^H+(1-\gamma)^2)(\theta^D+(1-\gamma)^2)}\right)-\gamma(1-\gamma)\pi_0\left(1+\beta\varphi_H\right)}{\theta^H+(1-\gamma)^2\left(1+\beta\varphi_H\right)}. \tag{20}$$

Finally, α is pinned down by the indifference of the Dove NCB between playing either a pooling action, π_{1M}^{c}, or a separating action, π_{1S}^{cD} given by equation (17). In other words, equilibrium α equalizes the loss functions associated with both pooling and separating equilibria at $t=1$.

As actions are always separated in the last period and the NCB takes expectations as given, the π_{2M}^{cH} is also equivalent to the full information case at $t=2$, except that inherited inflation is π_{1M}.

$$\pi_{2M}^{cH} = \frac{(1-\gamma)\{\kappa-\gamma\pi_{1M}\}}{\theta^H+(1-\gamma)^2} \tag{21}$$

The Dove NCB has a similar expression, with θ^D rather than θ^H, and inherited inflation being either π_{1M}^{c} or π_{1S}^{cD} depending on actions played in the first period.

B.3 Multiple equilibria

Recall that our refinement criterion guarantees uniqueness of the equilibrium. In this section we discuss to what extent different criteria alter the model results. While some implications obviously change, we assess the robustness of the key mechanism behind the reputation transfer's concept: a reduction in π_0 makes it harder to sustain pooling equilibria and easier to sustain separating ones.

First, consider the separating equilibria. In every separating equilibrium, the Dove NCB chooses his favorite choice π_{1S}^{cD}. Given that his type is revealed, there is no benefit of choosing something different. In contrast, other levels for the Hawk NCB's choice of inflation can be sustained. If agents believe that the Hawk NCB plays a given level π_1^*, the Hawk NCB will be considered a Dove at $t = 2$ if it does not play π_1^*. As long as π_1^* entails a smaller loss than playing π_{1S}^{cH} at $t = 1$ and be treated as a Dove at $t = 2$, π_1^* can be sustained as a separating equilibrium.

The discussion above sheds light on the refinement criterion used in the paper. If one must choose what the agents expect a Hawk NCB to do, it is intuitive to assume that the Hawk NCB chooses its favorite action without worrying about being mistake for a Dove. This is the logic of the criterion suggested by Cukierman and Liviatan (1991) and Walsh (2000) which is espoused in this paper.

In Figure 10, for a given parametrization of the model,[22] we plot the set of separating equilibria (the interval between the blue lines) against initial inflation π_0. The upper blue line plots, as a function of π_0, the maximum level of inflation a Hawk NCB can choose without making the Dove NCB deviate and pretend to be Hawk. The bottom blue line plots the minimum level of inflation a Hawk NCB is willing to choose in order not to be considered a Dove at $t = 2$. Any level of inflation below the bottom blue line is dominated by its preferred level of inflation even at the cost of being considered a Dove at $t = 2$. Notice that both blue lines decrease in π_0. In particular, as π_0 falls, the maximum level of inflation that can be sustained without the Dove NCB deviating increases. Hence, loosely speaking, monetary policy contraction makes it easier to sustain a separating equilibrium, which is the key mechanism behind the reputation transfer's concept. This idea also echoes on the fact that reduction in π_0 increases the set of separating equilibria (i.e, the interval between the blue lines).

This kind of analysis only makes sense if π_0 is a parameter rather than a choice. If one is willing to properly model the OCB's behavior, the model must specify a unique equilibrium so that the OCB can choose optimally. Nevertheless, Figure 10 reveals that type revelation through a monetary policy contraction must be present even when different refinement criteria are considered.

Consider the set of pooling equilibria instead. For each π_0, this set is the interval between the green lines in Figure 10. The lower green line plots the minimum level of inflation for which a Dove

[22] The specific parametrization is unimportant for the ideas conveyed here. Codes are available upon request.

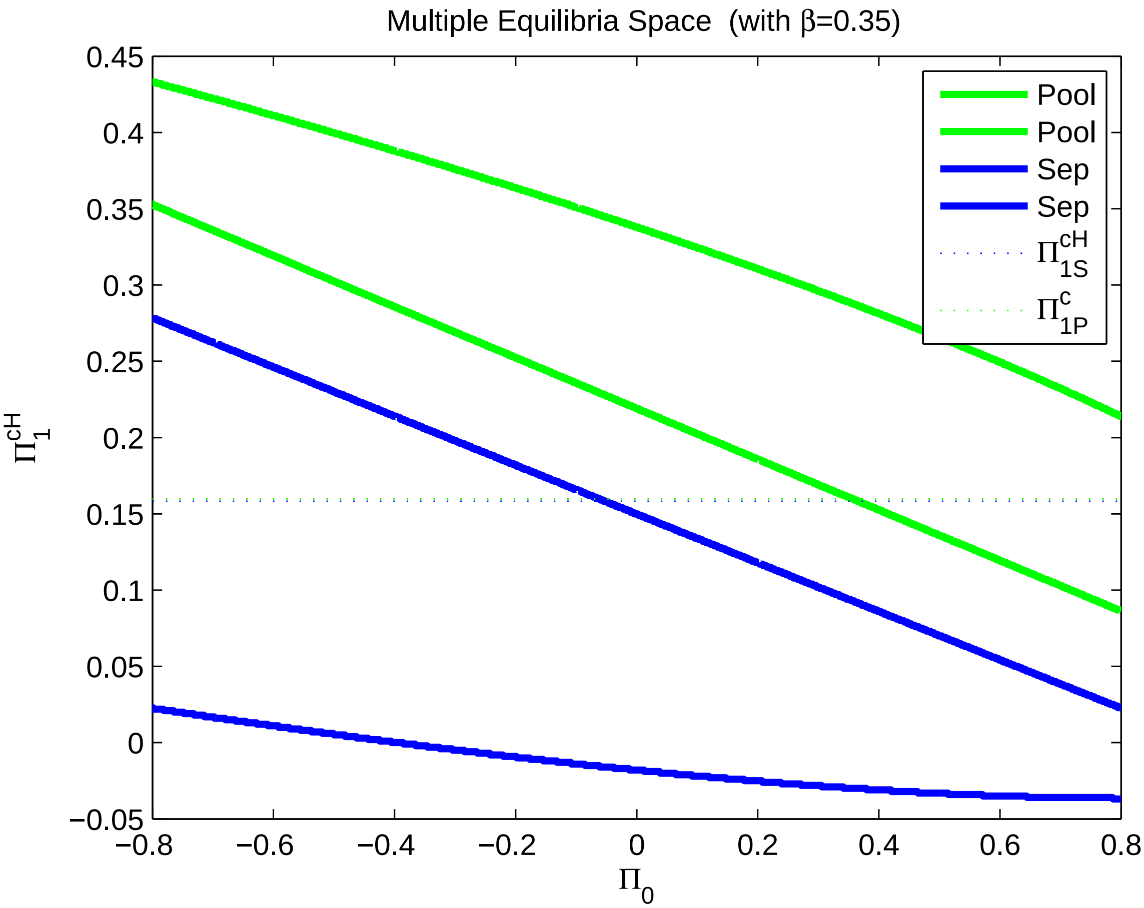

Figure 10: Multiple Equilibria

NCB pools its action with the Hawk one. For any level of inflation bellow the bottom green line, the Dove NCB prefers to play π_{1S}^{cD} and have its type revealed.[23] The upper green line plots the maximum level of inflation a Hawk NCB is willing to choose in order not to be considered a Dove for sure at $t = 2$. Notice that a contraction in π_0 reduces the set of pooling equilibrium. Therefore, the idea behind reputation transfer survives: a contraction in monetary policy makes it harder to sustain a pooling equilibrium.

Finally, notice that this discussion focuses only on pure strategy equilibria. However, given our refinement criterium, there are values of β and π_0 for which a pure strategy equilibrium does not

[23]Recall that the upper blue line is the limit of what a Dove NCB can put up to pass itself as a Hawk in a separating equilibrium. Notice that the upper blue line is a parallel downward shift of the lower green line. After all, deviating from a separating equilibrium is more attractive than a pooling equilibrium, since in the first agents believe the NCB is a Hawk with probability one, whereas in the second agents expect a Hawk with probability μ.

exist. For the parametrization used to generate Figure 10, for example, there are certain values of π_0 such that only a mixed strategy equilibrium exists. This can be seen in the plot as the horizontal lines, i.e. the levels of equilibrium inflation for both separating and pooling equilibria given our refinement criterium, lie outside the space between blue or green lines for certain values of π_0. In these cases, we showed in the main text that reputation transfer has a straightforward intuition: it increases the probability α of the Dove NCB choosing the separating action and, consequently, increases the posterior probability of the NCB being seen as a Hawk.

C Proofs

Proposition 1. *For γ small enough, there exists $\beta^S \geq 0$ such that $L_S^D \leq L_{SD}^D$ for all $\beta \in [0, \beta^S]$, and that $L_S^D > L_{SD}^D$ for all $\beta \in (\beta^S, \infty)$.*

Proof. First note that

$$\lim_{\beta \to 0} \left[L_S^D - L_{SD}^D\right] = -\frac{(1-\gamma)^2(\theta^D - \theta^H)^2(\kappa - \gamma\pi_0)^2}{2\left((1-\gamma)^2 + \theta^D\right)\left((1-\gamma)^2 + \theta^H\right)^2} < 0$$

$$\Rightarrow L_S^D < L_{SD}^D$$

Hence, for β small, there is a separating equilibrium.

Also note that

$$\lim_{\beta \to +\infty} \left[L_S^D - L_{SD}^D\right] = +\infty$$

For β large, the separating equilibrium does not exist.

Due to the Intermediate Value theorem, there is β_S such that $L_S^D - L_{SD}^D = 0$. To conclude the proof, we need to show that for γ small, $\left[L_S^D - L_{SD}^D\right]$ is monotonous in β.

The derivatives are:

$$\frac{\partial L_S^D}{\partial \beta} = \frac{1}{2}\Bigg[\frac{\left((1-\gamma)^2\kappa\left((1-\gamma)^2(\beta\varphi_D+1)+(\gamma+1)\theta^D\right)+\gamma^2\pi_0\theta^{D^2}\right)^2}{\left((1-\gamma)^2+\theta^D\right)^2\left((1-\gamma)^2(\beta\varphi_D+1)+\theta^D\right)^2} -$$

$$\frac{2\beta\gamma^2(1-\gamma)^2\theta^D\varphi_D\left((1-\gamma)^2\kappa+\gamma\pi_0\theta^D\right)^2}{\left((1-\gamma)^2+\theta^D\right)\left((1-\gamma)^2(\beta\varphi_D+1)+\theta^D\right)^3} - \frac{2(\gamma-1)^4\theta^D\varphi_D(\kappa-\gamma\pi_0(\beta\varphi_D+1))^2}{\left((1-\gamma)^2(\beta\varphi_D+1)+\theta^D\right)^3} +$$

$$\frac{(1-\gamma)^2\theta^D\left(\kappa-\frac{\gamma\left((1-\gamma)^2\kappa+\gamma\pi_0\theta^D\right)}{(1-\gamma)^2(\beta\varphi_D+1)+\theta^D}\right)^2}{\left((1-\gamma)^2+\theta^D\right)^2} + \frac{2\gamma(1-\gamma)^2\pi_0\theta^D\varphi_D(\gamma\pi_0(\beta\varphi_D+1)-\kappa)}{\left((1-\gamma)^2(\beta\varphi_D+1)+\theta^D\right)^2} -$$

$$\frac{2(1-\gamma)^2\varphi_D\left((1-\gamma)^2\kappa+\gamma\pi_0\theta^D\right)^2}{\left((1-\gamma)^2(\beta\varphi_D+1)+\theta^D\right)^3}$$

$$-2(1-\gamma)^2\kappa\mu\left(\frac{\varphi\mathrm{h}\left((1-\gamma)^2\kappa+\gamma\pi_0\theta^H\right)}{\left((1-\gamma)^2(\beta\varphi\mathrm{h}+1)+\theta^H\right)^2} - \frac{\varphi_D\left((1-\gamma)^2\kappa+\gamma\pi_0\theta^D\right)}{\left((1-\gamma)^2(\beta\varphi_D+1)+\theta^D\right)^2}\right)\Bigg]$$

$$\frac{\partial L_{DS}^D}{\partial \beta} = \frac{1}{2}\Bigg[\frac{2(1-\gamma)^2\kappa(\theta^D-\theta^H)\left(\kappa\left(\theta^H-(1-\gamma)^2(-\beta\varphi_H+\gamma-1)\right)-\gamma^2\pi_0\theta^H\right)}{\left((1-\gamma)^2+\theta^D\right)\left((1-\gamma)^2+\theta^H\right)\left((1-\gamma)^2(\beta\varphi_H+1)+\theta^H\right)}$$

$$-\frac{2\beta\gamma(1-\gamma)^2\varphi_H\left((1-\gamma)^2\kappa+\gamma\pi_0\theta^H\right)}{\left((1-\gamma)^2+\theta^D\right)\left((1-\gamma)^2+\theta^H\right)\left((1-\gamma)^2(\beta\varphi_H+1)+\theta^H\right)^3}\Bigg[$$

$$(1-\gamma)^2\kappa\left((\gamma-1)\theta^D((\gamma-1)(-\beta\varphi_H+\gamma-1)+\theta^H)+\theta^H\left((1-\gamma)^2(\beta\varphi_H+1)+\theta^H\right)\right)+$$

$$+\gamma^2\pi_0\theta^D\theta^H\left((1-\gamma)^2+\theta^H\right)\Bigg]$$

$$+\frac{\left((1-\gamma)^2\kappa\left(\gamma^2(\beta\varphi_H+1)+\beta\varphi_H+\gamma(-2\beta\varphi_H+\theta^D-2)+\theta^H+1\right)+\gamma^2\pi_0\theta^D\theta^H\right)^2}{\left((1-\gamma)^2+\theta^D\right)^2\left((1-\gamma)^2(\beta\varphi_H+1)+\theta^H\right)^2}$$

$$-2(1-\gamma)^2\kappa(1-\mu)\left(\frac{\varphi_D\left((1-\gamma)^2\kappa+\gamma\pi_0\theta^D\right)}{\left((1-\gamma)^2(\beta\varphi_D+1)+\theta^D\right)^2} - \frac{\varphi_H\left((1-\gamma)^2\kappa+\gamma\pi_0\theta^H\right)}{\left((1-\gamma)^2(\beta\varphi_H+1)+\theta^H\right)^2}\right)$$

$$-\frac{2(\gamma-1)^4\theta^D\varphi_H(\kappa-\gamma\pi_0(\beta\varphi_H+1))^2}{\left((1-\gamma)^2(\beta\varphi_H+1)+\theta^H\right)^3} + \frac{(1-\gamma)^2\theta^D\left(\kappa-\frac{\gamma\left((1-\gamma)^2\kappa+\gamma\pi_0\theta^H\right)}{(1-\gamma)^2(\beta\varphi_H+1)+\theta^H}\right)^2}{\left((1-\gamma)^2+\theta^D\right)^2}$$

$$+\frac{2\gamma(1-\gamma)^2\pi_0\theta^D\varphi_H(\gamma\pi_0(\beta\varphi_H+1)-\kappa)}{\left((1-\gamma)^2(\beta\varphi_H+1)+\theta^H\right)^2} - \frac{2(1-\gamma)^2\varphi_H\left((1-\gamma)^2\kappa+\gamma\pi_0\theta^H\right)^2}{\left((1-\gamma)^2(\beta\varphi_H+1)+\theta^H\right)^3}\Bigg]$$

As we can see from above, $\frac{\partial L_S^D}{\partial \beta}$ and $\frac{\partial L_{DS}^D}{\partial \beta}$ are continuous in γ. Now note that:

$$\lim_{\gamma\to 0}\frac{\partial(L_S^D - L_{DS}^D)}{\partial \beta} = \frac{\kappa^2(\theta^H-\theta^D)}{\left(1+\theta^D\right)\left(1+\theta^H\right)} > 0$$

Therefore, for γ small enough, $L_S^D - L_{DS}^D$ will be strictly increasing in β, which concludes the proof.

It is interesting to note that:

$$\lim_{\beta \to +\infty} \frac{\partial (L_S^D - L_{DS}^D)}{\partial \beta} = \frac{\kappa^2(\theta^H - \theta^D)}{(\theta^D + (1-\gamma)^2)\,(\theta^H + (1-\gamma)^2)} > 0$$

So the restriction on γ is need for intermediate values of β. With γ high, it is possible some non monotonicity to be present.

$\square$

Proposition 2. *For γ small enough, there exists $\beta_P \geq 0$ such that $L_P^D \leq L_{PD}^D$ for all $\beta \in (\beta^P, \infty)$, and that $L_P^D > L_{PD}^D$ for all $\beta \in [0, \beta^P]$.*

Proof. This proof borrows a lot from the proof of Proposition 1.

First note that

$$\lim_{\beta \to 0} \left[L_P^D - L_{PD}^D\right] = \frac{(1-\gamma)^2(\theta^D - \theta^H)^2(\kappa\left((1-\gamma)^2(\gamma(\mu-1)+1) + \theta^H\right) - \gamma\left((1-\gamma)^2 + \theta^H\right)\pi_0)^2}{2\left((1-\gamma)^2 + \theta^D\right)\left((1-\gamma)^2 + \theta^H\right)^4} > 0$$

$$\Rightarrow L_P^D > L_{PD}^D$$

Hence, for β small, there is no pooling equilibrium.

Also note that

$$\lim_{\beta \to +\infty} \left[L_P^D - L_{PD}^D\right] = -\infty$$

For β large, the pooling equilibrium shall exist.

As $\left[L_P^D - L_{PD}^D\right]$ is a continuous function of β, the Intermediate Value theorem implies that there is a β_P such that $L_P^D - L_{PD}^D = 0$. To conclude the proof, all we need to show is that, for γ small, $\left[L_P^D - L_{PD}^D\right]$ is monotonous in β. This means we can divide the parameter space for β in areas where the pooling equilibrium is sustained and areas where the Dove NCB will deviate.

The derivatives $\frac{\partial(L_P^D)}{\partial\beta}$ and $\frac{\partial(L_{DP}^D)}{\partial\beta}$ are omitted due to space convenience[24]. We are, however, interested in the limit of their difference:

$$\lim_{\gamma \to 0} \frac{\partial (L_P^D - L_{DP}^D)}{\partial \beta} = -\frac{\kappa^2(\theta^H - \theta^D)\mu}{(1+\theta^D)\,(1+\theta^H)} < 0$$

Therefore, for γ small enough, $L_P^D - L_{PD}^D$ will be strictly decreasing in β, which concludes the proof.

[24]Their expressions are available under request in a Wolfram Mathematica file.

$\square$

Proposition 3. *For $\pi_0 < \bar{\pi}_0$, $\Delta^S(\pi_0)$ decreases in π_0 and $\Delta^P(\pi_0)$ increases in π_0.*

Proof. .

Part I: Separating

First, note:

$$\frac{\partial L_S^D}{\partial \pi_0} = \left(\frac{\gamma^2\theta^D(1+\beta\varphi_D)}{\theta^D + (1-\gamma)^2(1+\beta\varphi_D)}\right)\pi_0 - \kappa(1-\gamma)\mu\left(\frac{\partial\left(\pi_{2S}^{cD} - \pi_{2S}^{cH}\right)}{\partial \pi_0}\right)$$

and

$$\frac{\partial L_{DS}^D}{\partial \pi_0} = \frac{\kappa(1-\gamma)^2\gamma\left[\theta^H(1+\beta\varphi_D) - \theta^D(1+\beta\varphi_H)\right] + \left[(\theta^H)^2(1+\beta\varphi_D) + \theta^D(1-\gamma)^2(1+\beta\varphi_H)^2\right]\gamma^2\pi_0}{\left[\theta^H + (1-\gamma)^2(1+\beta\varphi_H)\right]^2} - \kappa(1-\gamma)(1-\mu)\left(\frac{\partial\left(\pi_{1S}^{cH} - \pi_{1S}^{cD}\right)}{\partial \pi_0}\right) - \beta\kappa(1-\gamma)\left(\frac{\partial\left(\pi_{2SD}^{cD} - \pi_{2S}^{cH}\right)}{\partial \pi_0}\right)$$

Tedious algebra shows that

$$\frac{\left[(\theta^H)^2(1+\beta\varphi_D) + \theta^D(1-\gamma)^2(1+\beta\varphi_H)^2\right]}{\left[\theta^H + (1-\gamma)^2(1+\beta\varphi_H)\right]^2} \equiv h(\theta^H)$$

increases in θ^H.

Therefore,

$$h(\theta^H) > h(\theta^D) = \frac{\theta^D(1+\beta\varphi_D)}{\theta^D + (1-\gamma)^2(1+\beta\varphi_D)}$$

Define:

$$\frac{\Delta^S(\pi_0)}{\partial \pi_0} \equiv \frac{\partial L_{DS}^D}{\partial \pi_0} - \frac{\partial L_S^D}{\partial \pi_0}$$

$$\frac{\Delta^S(\pi_0)}{\partial \pi_0} = A\pi_0 - \kappa(1-\gamma)\left(\frac{\partial\left(\pi_{1S}^{cH} - \pi_{1S}^{cD}\right)}{\partial \pi_0} + \beta\frac{\partial\left(\pi_{2SD}^{cD} - \pi_{2S}^{cH}\right)}{\partial \pi_0} - \frac{(1-\gamma)\gamma\left[\theta^H(1+\beta\varphi_D) - \theta^D(1+\beta\varphi_H)\right]}{\left[\theta^H + (1-\gamma)^2(1+\beta\varphi_H)\right]^2}\right)$$

where

$$A \equiv \gamma^2\left[\frac{\left[(\theta^H)^2(1+\beta\varphi_D) + \theta^D(1-\gamma)^2(1+\beta\varphi_H)^2\right]}{\left[\theta^H + (1-\gamma)^2(1+\beta\varphi_H)\right]^2} - \frac{\theta^D(1+\beta\varphi_D)}{\theta^D + (1-\gamma)^2(1+\beta\varphi_D)}\right] > 0$$

Also:

$$\frac{\partial\left(\pi_{1S}^{cH}-\pi_{1S}^{cD}\right)}{\partial\pi_0}=\gamma(1-\gamma)\frac{\left[\theta^H(1+\beta\varphi_D)-\theta^D(1+\beta\varphi_H)\right]}{\left[\theta^H+(1-\gamma)^2(1+\beta\varphi_H)\right]\left[\theta^D+(1-\gamma)^2(1+\beta\varphi_D)\right]}$$

$$\frac{\partial\left(\pi_{2SD}^{cD}-\pi_{2S}^{cH}\right)}{\partial\pi_0}=\gamma^2(1-\gamma)\frac{(\theta^D-\theta^H)\theta^H}{\left[\theta^H+(1-\gamma)^2(1+\beta\varphi_H)\right]\left[\theta^H+(1-\gamma)^2\right]\left[\theta^D+(1-\gamma)^2\right]}$$

Hence we can rewrite as

$$\frac{\Delta^S(\pi_0)}{\partial\pi_0}=A\pi_0-B+\beta C$$

where

$$B=\frac{\kappa(1-\gamma)^2\gamma\left[\theta^H(1+\beta\varphi_D)-\theta^D(1+\beta\varphi_H)\right]}{\left[\theta^H+(1-\gamma)^2(1+\beta\varphi_H)\right]}\left[\frac{1}{\left[\theta^D+(1-\gamma)^2(1+\beta\varphi_D)\right]}-\frac{1}{\left[\theta^H+(1-\gamma)^2(1+\beta\varphi_H)\right]}\right]$$

$$C=\kappa\gamma^2(1-\gamma)^2\frac{(\theta^H-\theta^D)\theta^H}{\left[\theta^H+(1-\gamma)^2(1+\beta\varphi_H)\right]\left[\theta^H+(1-\gamma)^2\right]\left[\theta^D+(1-\gamma)^2\right]}$$

Clearly,

$$\frac{\Delta^S(\pi_0)}{\partial\pi_0}<0\iff\pi_0<\frac{B-\beta C}{A}\equiv\bar{\pi_0}^S$$

This concludes the part of the proof referring to Separating equilibrium.

To provide a bit of intuition for this upper bound, we analyze the case where $\beta=0$:

$$\beta=0\Rightarrow\bar{\pi_0}^S=\frac{\kappa}{\gamma}$$

This is the upper bound to keep π_1^c positive. The reason why this is so stems from the intuition given in Figure 6. If π_0 is too large, the MgB line becomes negative and the gray triangle goes to the third quadrant. There, an increase in π_0 increases the triangle's area further and thus makes pooling harder.

Note, however, that in equilibrium OCB will never choose $\pi_0>\frac{\kappa}{\gamma}$.

Part II: Pooling

Define:

$$\kappa^*\equiv\kappa\left(1+\frac{\beta(1-\mu)\gamma(1-\gamma)^2(\theta^H-\theta^D)}{\left(\theta^H+(1-\gamma)^2\right)\left(\theta^D+(1-\gamma)^2\right)}\right)$$

Then:

$$\pi_{1P}^{c} = \frac{(1-\gamma)\kappa^{*} - \gamma(1-\gamma)\pi_0\left(1+\beta\varphi_H\right)}{\theta^H + (1-\gamma)^2\left(1+\beta\varphi_H\right)}$$

$$\pi_{1P} = \frac{(1-\gamma)^2\kappa^{*} + \gamma\pi_0\theta^H}{\theta^H + (1-\gamma)^2\left(1+\beta\varphi_H\right)}$$

It is easy to see that the difference between π_{1P}^{c} and π_{1S}^{cH} stems from the difference between κ and κ^{*}. Hence, we can write:

$$\frac{\partial L_P^D}{\partial \pi_0} = \frac{\kappa^{*}(1-\gamma)^2\gamma\left[\theta^H(1+\beta\varphi_D) - \theta^D(1+\beta\varphi_H)\right] + \left[(\theta^H)^2(1+\beta\varphi_D) + \theta^D(1-\gamma)^2(1+\beta\varphi_H)^2\right]\gamma^2\pi_0}{\left[\theta^H + (1-\gamma)^2(1+\beta\varphi_H)\right]^2} - \beta\kappa\mu(1-\gamma)\left(\frac{\partial\left(\pi_{2P}^{cD} - \pi_{2P}^{cH}\right)}{\partial\pi_0}\right)$$

If a Dove NCB deviates from Pooling Equilibrium, it will choose the same inflation as its separating inflation. Note, however, that inflations expectations are different.

$$\frac{\partial L_{PD}^D}{\partial \pi_0} = \left(\frac{\gamma^2\theta^D(1+\beta\varphi_D)}{\theta^D + (1-\gamma)^2(1+\beta\varphi_D)}\right)\pi_0 - \kappa(1-\gamma)\left(\frac{\partial\left(\pi_{1S}^{cD} - \pi_{1P}^{c}\right)}{\partial\pi_0}\right)$$

Also:

$$\frac{\partial\left(\pi_{1S}^{cD} - \pi_{1P}^{c}\right)}{\partial\pi_0} = -\gamma(1-\gamma)\frac{\left[\theta^H(1+\beta\varphi_D) - \theta^D(1+\beta\varphi_H)\right]}{\left[\theta^H + (1-\gamma)^2(1+\beta\varphi_H)\right]\left[\theta^D + (1-\gamma)^2(1+\beta\varphi_D)\right]}$$

$$\frac{\partial\left(\pi_{2P}^{cD} - \pi_{2P}^{cH}\right)}{\partial\pi_0} = \gamma^2(1-\gamma)\frac{(\theta^D - \theta^H)\theta^H}{\left[\theta^H + (1-\gamma)^2(1+\beta\varphi_H)\right]\left[\theta^H + (1-\gamma)^2\right]\left[\theta^D + (1-\gamma)^2\right]}$$

Define:

$$\frac{\Delta^P(\pi_0)}{\partial\pi_0} \equiv \frac{\partial L_{DP}^D}{\partial\pi_0} - \frac{\partial L_P^D}{\partial\pi_0}$$

Hence we can rewrite as

$$\frac{\Delta^P(\pi_0)}{\partial\pi_0} = -A\pi_0 + B + \beta D$$

where A, B were defined above and:

$$D = \kappa\gamma^2(1-\gamma)^2 \frac{(\theta^H - \theta^D)}{[\theta^H + (1-\gamma)^2(1+\beta\varphi_H)]\,[\theta^H + (1-\gamma)^2]\,[\theta^D + (1-\gamma)^2]} \times$$
$$\times \left(\mu\theta^H - (1-\mu)(1-\gamma)^2 \frac{\left[\theta^H(1+\beta\varphi_D) - \theta^D(1+\beta\varphi_H)\right]}{[\theta^H + (1-\gamma)^2(1+\beta\varphi_H)]} \right)$$

Clearly,

$$\frac{\Delta^P(\pi_0)}{\partial\pi_0} > 0 \iff \pi_0 < \frac{B + \beta D}{A} \equiv \bar{\pi_0}^P$$

Defining

$$\bar{\pi_0} \equiv \min(\bar{\pi_0}^P, \bar{\pi_0}^S)$$

concludes the proof.

$\square$

Proposition 4. *For $\pi_0 < \bar{\bar{\pi}}_0$ and $\mu > \underline{\mu}$, α decreases in π_0.*

Proof. α is the point where Dove NCB is indifferent, in a mixed equilibrium, between playing the same action as Hawk NCB , π_{1M}^c or choosing its favorite inflation level π_{1S}^{cD}. After all, if Dove NCB was not indifferent, α would be either 0 or 1. Consequently the loss functions must equal:

$$\frac{\left(\pi_{1M}^2 + \theta^D(\pi_{1M}^c)^2\right)}{2} - \kappa(1-\gamma)(\pi_{1M}^c - \cancel{E[\pi_1^c]}) + \beta\left(\frac{\left((\pi_{2M}^D)^2 + \theta^D(\pi_{2M}^{cD})^2\right)}{2} - \kappa\left(\pi_{2M}^D - E[\pi_{2M}]\right)\right) =$$
$$= \frac{\left((\pi_{1S}^D)^2 + \theta^D(\pi_{1S}^{Dc})^2\right)}{2} - \kappa(1-\gamma)(\pi_{1S}^{cD} - \cancel{E[\pi_1^c]}) + \beta\left(\frac{\left((\pi_{2S}^D)^2 + \theta^D(\pi_{2S}^{cD})^2\right)}{2} - \kappa\cancelto{0}{\left(\pi_{2S}^D - E[\pi_{2S}]\right)}\right)$$

Taking the derivative of π_0, we have:

$$\pi_{1M}\frac{d\pi_{1M}}{d\pi_0} + (\theta^D\pi_{1M}^c - \kappa(1-\gamma))\frac{d\pi_{1M}^c}{d\pi_0} + \beta\left(\frac{d\pi_{1M}}{d\pi_0}\left(\pi_{2M}^D\frac{d\pi_{2M}^D}{d\pi_{1M}} + \theta^D\pi_{2M}^{cD}\frac{d\pi_{2M}^{cD}}{d\pi_{1M}} - \kappa\frac{d\left(\pi_{2M}^D - E[\pi_{2M}]\right)}{d\pi_{1M}}\right)\right)$$
$$= \pi_{1S}^D\frac{d\pi_{1S}^D}{d\pi_0} + (\theta^D\pi_{1S}^{Dc} - \kappa(1-\gamma))\frac{d\pi_{1S}^{cD}}{d\pi_0} + \beta\left(\frac{d\pi_{1S}^D}{d\pi_0}\left(\pi_{2S}^D\frac{d\pi_{2S}^D}{d\pi_{1S}^D} + \theta^D\pi_{2S}^{cD}\frac{d\pi_{2S}^{cD}}{d\pi_{1S}^D}\right)\right) \tag{22}$$

Note that:

$$\left(\pi_{2M}^D - E[\pi_{2M}]\right) = \mu_{post}\left(\pi_{2M}^D - \pi_{2M}^H\right) = (1-\gamma)\mu_{post}\left(\pi_{2M}^{cD} - \pi_{2M}^{cH}\right)$$

$$\left(\pi_{2M}^{D}\frac{d\pi_{2M}^{D}}{d\pi_{1M}}+\theta^{D}\pi_{2M}^{cD}\frac{d\pi_{2M}^{cD}}{d\pi_{1M}}\right)=\varphi^{D}\pi_{1M}$$

$$\left(\pi_{2S}^{D}\frac{d\pi_{2S}^{D}}{d\pi_{1S}^{D}}+\theta^{D}\pi_{2S}^{cD}\frac{d\pi_{2S}^{cD}}{d\pi_{1S}^{D}}\right)=\varphi^{D}\pi_{1S}^{D}$$

Then (22) becomes:

$$\begin{aligned}\frac{d\pi_{1M}}{d\pi_0}\left(\pi_{1M}(1+\beta\varphi^{D})-\beta\kappa(1-\gamma)\mu_{post}\frac{d\left(\pi_{2M}^{cD}-\pi_{2M}^{cH}\right)}{d\pi_{1M}}\right)-\beta\kappa(1-\gamma)\frac{d\mu_{post}}{d\pi_0}\left(\pi_{2M}^{cD}-\pi_{2M}^{cH}\right)+\\+\frac{d\pi_{1M}^{c}}{d\pi_0}(\theta^{D}\pi_{1M}^{c}-\kappa(1-\gamma))=\frac{d\pi_{1S}^{D}}{d\pi_0}\left(\pi_{1S}^{D}(1+\beta\varphi^{D})\right)+\frac{d\pi_{1S}^{cD}}{d\pi_0}(\theta^{D}\pi_{1S}^{cD}-\kappa(1-\gamma))\end{aligned}\tag{23}$$

Note that:

$$\frac{d\mu_{post}}{d\pi_0}=\frac{\partial\mu_{post}}{\partial\alpha}\frac{d\alpha}{d\pi_0}$$

$$\frac{d\pi_{1M}}{d\pi_0}=\gamma+(1-\gamma)\frac{d\pi_{1M}^{c}}{d\pi_0}$$

$$\frac{d\pi_{1M}^{c}}{d\pi_0}=\frac{\partial\pi_{1M}^{c}}{\partial\pi_0}+\frac{\partial\pi_{1M}^{c}}{\partial\alpha}\frac{d\alpha}{d\pi_0}$$

Collecting $\frac{d\alpha}{d\pi_0}$, (23) becomes:

$$DEN\frac{d\alpha}{d\pi_0}=NUM$$

where:

$$\begin{aligned}DEN\equiv\frac{\partial\pi_{1M}^{c}}{\partial\alpha}\left[(1-\gamma)\left(\pi_{1M}(1+\beta\varphi^{D})-\beta\kappa(1-\gamma)\mu_{post}\frac{d\left(\pi_{2M}^{cD}-\pi_{2M}^{cH}\right)}{d\pi_{1M}}\right)+(\theta^{D}\pi_{1M}^{c}-\kappa(1-\gamma))\right]\\-\frac{\partial\mu_{post}}{\partial\alpha}\left[\beta\kappa(1-\gamma)\left(\pi_{2M}^{cD}-\pi_{2M}^{cH}\right)\right]\end{aligned}\tag{24}$$

$$NUM \equiv \left[\frac{d\pi_{1S}^{D}}{d\pi_0}\left(\pi_{1S}^{D}(1+\beta\varphi^{D})\right) + \frac{d\pi_{1S}^{cD}}{d\pi_0}(\theta^{D}\pi_{1S}^{cD} - \kappa(1-\gamma))\right] -$$

$$\left[\left(\gamma + (1-\gamma)\frac{\partial\pi_{1M}^{c}}{\partial\pi_0}\right)\left(\pi_{1M}(1+\beta\varphi^{D}) - \beta\kappa(1-\gamma)\mu_{post}\frac{d\left(\pi_{2M}^{cD} - \pi_{2M}^{cH}\right)}{d\pi_{1M}}\right) + \frac{\partial\pi_{1M}^{c}}{\partial\pi_0}(\theta^{D}\pi_{1M}^{c} - \kappa(1-\gamma))\right] \tag{25}$$

and:

$$\frac{\partial\mu_{post}}{\partial\alpha} = \frac{\mu(1-\mu)}{(\mu + (1-\mu)(1-\alpha))^2}$$

$$\frac{d\left(\pi_{2M}^{cD} - \pi_{2M}^{cH}\right)}{d\pi_{1M}} = \frac{-\gamma(1-\gamma)(\theta^{H} - \theta^{D})}{(\theta^{H} + (1-\gamma)^2)(\theta^{D} + (1-\gamma)^2)}$$

$$\left(\pi_{2M}^{cD} - \pi_{2M}^{cH}\right) = \frac{(1-\gamma)(\kappa - \gamma\pi_{1M})(\theta^{H} - \theta^{D})}{(\theta^{H} + (1-\gamma)^2)(\theta^{D} + (1-\gamma)^2)}$$

We want to find the conditions in which $\frac{NUM}{DEN} < 0$ is true. In order to do so, we rewrite NUM and DEN by collecting π_0 such that:

$$NUM = A_N + B_N\pi_0$$

$$DEN = A_D + B_D\pi_0$$

Rearranging and simplifying, we have:

$$
\begin{aligned}
A_N \equiv \kappa(1-\gamma)\gamma &\left(-\frac{(1-\gamma)\theta^H(\beta\varphi^D+1)\left(\frac{(\alpha-1)^2\beta\gamma(\gamma-1)^2(\mu-1)(\theta^D-\theta^H)}{(\alpha(\mu-1)+1)((1-\gamma)^2+\theta^D)((1-\gamma)^2+\theta^H)}+1\right)}{\left((\gamma-1)^2(\beta\varphi^H+1)+\theta^H\right)^2}\right) \\
+\kappa(1-\gamma)\gamma &\left(\frac{(1-\gamma)\theta^D(\beta\varphi^H+1)\left(\frac{(\alpha-1)^2\beta\gamma(\gamma-1)^2(\mu-1)(\theta^D-\theta^H)}{(\alpha(\mu-1)+1)((1-\gamma)^2+\theta^D)((1-\gamma)^2+\theta^H)}+1\right)}{\left((\gamma-1)^2(\beta\varphi^H+1)+\theta^H\right)^2}\right)+ \\
\kappa(1-\gamma)\gamma &\left(\frac{(\alpha-1)\beta(\gamma-1)\gamma\mu\theta^H(\theta^D-\theta^H)}{(\alpha(\mu-1)+1)\left(\theta^D+(1-\gamma)^2\right)\left(\theta^H+(1-\gamma)^2\right)\left((1-\gamma)^2(\beta\varphi^H+1)+\theta^H\right)}\right)- \\
\kappa(1-\gamma)\gamma &\left(\frac{(\gamma-1)\theta^D(\beta\varphi^D+1)}{\left(\gamma^2(\beta\varphi^D+1)-2\gamma(\beta\varphi^D+1)+\beta\varphi^D+\theta^D+1\right)^2}-\frac{(\gamma-1)(\beta\varphi^D+1)}{(\gamma-1)^2(\beta\varphi^D+1)+\theta^D}\right)+ \\
\kappa(1-\gamma)\gamma &\left(\frac{(\gamma-1)\theta^D(\beta\varphi^D+1)}{\left((\gamma-1)^2(\beta\varphi^D+1)+\theta^D\right)^2}-\frac{(1-\gamma)(\beta\varphi^H+1)}{(\gamma-1)^2(\beta\varphi^H+1)+\theta^H}\right)
\end{aligned}
\tag{26}
$$

$$
B_N \equiv -\frac{(\gamma-1)^2\gamma^2(-\theta^H(\beta\varphi^D+1)+\beta\varphi^h\theta^D+\theta^D)^2}{\left((\gamma-1)^2(\beta\varphi^D+1)+\theta^D\right)\left((\gamma-1)^2(\beta\varphi^h+1)+\theta^H\right)^2}<0 \tag{27}
$$

$$
\begin{aligned}
A_D \equiv &-\frac{(1-\alpha)\beta^2\gamma^2(1-\gamma)^6\kappa^2(1-\mu)\mu(-\alpha(1-\mu)+\mu+1)(\theta^H-\theta^D)^2}{(1-\alpha(1-\mu))^2((1-\alpha)(1-\mu)+\mu)\left((1-\gamma)^2+\theta^D\right)^2\left((1-\gamma)^2+\theta^H\right)^2\left((1-\gamma)^2(\beta\varphi^h+1)+\theta^H\right)} \\
&-\frac{(1-\alpha)\beta\gamma(1-\gamma)^6\kappa^2(1-\mu)(-\alpha(1-\mu)+\mu+1)(\beta\varphi^D+1)(\theta^H-\theta^D)\left(\frac{(1-\alpha)\beta\gamma(1-\gamma)^2\left(1-\frac{\mu}{(1-\alpha)(1-\mu)+\mu}\right)(\theta^H-\theta^D)}{\left((1-\gamma)^2+\theta^D\right)\left((1-\gamma)^2+\theta^H\right)}+1\right)}{(1-\alpha(1-\mu))^2((1-\gamma)^2+\theta^D)((1-\gamma)^2+\theta^H)\left((1-\gamma)^2(\beta\varphi^h+1)+\theta^H\right)^2} \\
&+\frac{(1-\alpha)\beta\gamma(1-\gamma)^4\kappa^2(1-\mu)(-\alpha(1-\mu)+\mu+1)(\theta^H-\theta^D)}{(1-\alpha(1-\mu))^2\left((1-\gamma)^2+\theta^D\right)\left((1-\gamma)^2+\theta^H\right)\left((1-\gamma)^2(\beta\varphi^h+1)+\theta^H\right)} \\
&+\frac{\beta\gamma(1-\gamma)^4\kappa^2(1-\mu)\mu\left(\frac{(1-\alpha)\beta\gamma(1-\gamma)^2\left(1-\frac{\mu}{(1-\alpha)(1-\mu)+\mu}\right)(\theta^H-\theta^D)}{((1-\gamma)^2+\theta^D)((1-\gamma)^2+\theta^H)}+1\right)}{((1-\alpha)(1-\mu)+\mu)^2\left((1-\gamma)^2+\theta^D\right)\left((1-\gamma)^2+\theta^H\right)\left((1-\gamma)^2(\beta\varphi^h+1)+\theta^H\right)} \\
&-\frac{(1-\alpha)\beta\gamma(1-\gamma)^4\kappa^2(1-\mu)\theta^D(-\alpha(1-\mu)+\mu+1)(\theta^H-\theta^D)\left(\frac{(1-\alpha)\beta\gamma(1-\gamma)^2\left(1-\frac{\mu}{(1-\alpha)(1-\mu)+\mu}\right)(\theta^H-\theta^D)}{((1-\gamma)^2+\theta^D)((1-\gamma)^2+\theta^H)}+1\right)}{(1-\alpha(1-\mu))^2\left((1-\gamma)^2+\theta^D\right)\left((1-\gamma)^2+\theta^H\right)\left((1-\gamma)^2(\beta\varphi^h+1)+\theta^H\right)^2} \\
&-\frac{\beta(1-\gamma)^2\kappa^2(1-\mu)\mu}{((1-\alpha)(1-\mu)+\mu)^2\left((1-\gamma)^2+\theta^D\right)\left((1-\gamma)^2+\theta^H\right)}
\end{aligned}
\tag{28}
$$

$$
\begin{aligned}
B_D \equiv & \frac{\beta(\gamma-1)^2\gamma^2\kappa(1-\mu)}{(\alpha(\mu-1)+1)^2\left(\gamma^2-2\gamma+\theta^D+1\right)\left(\gamma^2-2\gamma+\theta^H+1\right)} \times ... \\
... & \left[\frac{(\alpha-1)(\gamma-1)^4(\alpha(\mu-1)+\mu+1)(\beta\varphi^D+1)(\beta\varphi^H+1)(\theta^D-\theta^H)}{\left(\gamma^2(\beta\varphi^H+1)-2\gamma(\beta\varphi^H+1)+\beta\varphi^H+\theta^H+1\right)^2}\right. \\
& +\frac{(\alpha-1)(\gamma-1)^2(\alpha(\mu-1)+\mu+1)(\beta\varphi^D+1)(\theta^H-\theta^D)}{\gamma^2(\beta\varphi^H+1)-2\gamma(\beta\varphi^H+1)+\beta\varphi^H+\theta^H+1} \\
& +\frac{(\alpha-1)(\gamma-1)^2\theta^D(\alpha(\mu-1)+\mu+1)(\beta\varphi^H+1)(\theta^D-\theta^H)}{\left(\gamma^2(\beta\varphi^H+1)-2\gamma(\beta\varphi^H+1)+\beta\varphi^H+\theta^H+1\right)^2} \\
& \left.-\frac{(\gamma-1)^2\mu(\beta\varphi^H+1)}{\gamma^2(\beta\varphi^H+1)-2\gamma(\beta\varphi^H+1)+\beta\varphi^H+\theta^H+1}+\mu\right]
\end{aligned}
\tag{29}
$$

Numerically, it is easy to see that B_D is positive unless both μ and α are very small. It is also possible to show numerically than when μ approaches 0, α approaches 1, which suggests that B_D might be always positive in equilibrium. However, we cannot be certain as we do not have an analytical expression for α. Therefore, we can only prove that B_D is positive assuming that μ is greater than a given threshold. This is what we shall do now.

Collecting μ and rearranging (29), we can show that $B_D > 0$ as long as:

$$
\begin{aligned}
\mu > & \left[(\alpha-1)^2(\gamma-1)^2\left((1-\gamma)^2+\theta^H\right)(\theta^D-\theta^H)^2\left(\theta^H(\gamma(\beta\gamma\theta^D+\gamma-2)+\theta^D+1)+(\gamma-1)^2\theta^D+(\gamma-1)^4\right)\right] \Big/ ... \\
& ... \left[\left((1-\gamma)^2+\theta^D\right)\left((\gamma-1)^2\theta^H\left(\beta\gamma^2+2\right)+(\gamma-1)^4+{\theta^H}^2\right)^2\right. \\
& \times\left(1+\frac{(\alpha-1)(\alpha+1)(\gamma-1)^2\left((1-\gamma)^2+\theta^H\right)(\theta^D-\theta^H)^2\left(\theta^H(\gamma(\beta\gamma\theta^D+\gamma-2)+\theta^D+1)+(\gamma-1)^2\theta^D+(\gamma-1)^4\right)}{\left((1-\gamma)^2+\theta^D\right)\left((\gamma-1)^2\theta^H(\beta\gamma^2+2)+(\gamma-1)^4+{\theta^H}^2\right)^2}\right. \\
& \left.\left.-\frac{(\gamma-1)^2\left(\gamma^2(\beta\theta^H+1)-2\gamma+\theta^H+1\right)}{\gamma^4(\beta\theta^H+1)-2\gamma^3(\beta\theta^H+2)+\gamma^2((\beta+2)\theta^H+6)-4\gamma(\theta^H+1)+(\theta^H+1)^2}\right)\right]
\end{aligned}
\tag{30}
$$

The case we are interested in is when the Right Hand Side (RHS) of (30) is positive. After all, if it is negative, every μ satisfies the inequality by construction and we have $B_D > 0$.

As α is endogenous to our model, we want a lower bound of μ which makes $B_D > 0$ for all α. In addition, it is easy to see that if RHS of (30) is positive, it decreases in α. Hence a lower bound of μ will be the RHS of (30) imposing $\alpha = 0$. In this case, RHS becomes $\underline{\mu}$:

$$\underline{\mu} \equiv \left[(\gamma-1)^2\left((1-\gamma)^2+\theta^H\right)\left(\theta^D-\theta^H\right)^2\left(\theta^H(\gamma(\beta\gamma\theta^D+\gamma-2)+\theta^D+1)+(\gamma-1)^2\theta^D+(\gamma-1)^4\right)\right] \Big/ \ldots$$

$$\ldots \left[\left((1-\gamma)^2+\theta^D\right)\left((\gamma-1)^2\theta^H\left(\beta\gamma^2+2\right)+(\gamma-1)^4+{\theta^H}^2\right)^2\right.$$

$$\times\left(1-\frac{(\gamma-1)^2\left((1-\gamma)^2+\theta^H\right)(\theta^D-\theta^H)^2\left(\theta^H(\gamma(\beta\gamma\theta^D+\gamma-2)+\theta^D+1)+(\gamma-1)^2\theta^D+(\gamma-1)^4\right)}{((1-\gamma)^2+\theta^D)\left((\gamma-1)^2\theta^H(\beta\gamma^2+2)+(\gamma-1)^4+{\theta^H}^2\right)^2}\right.$$

$$\left.\left.-\frac{(\gamma-1)^2\left(\gamma^2(\beta\theta^H+1)-2\gamma+\theta^H+1\right)}{\gamma^4(\beta\theta^H+1)-2\gamma^3(\beta\theta^H+2)+\gamma^2((\beta+2)\theta^H+6)-4\gamma(\theta^H+1)+(\theta^H+1)^2}\right)\right] \tag{31}$$

Naturally, we also want that $\underline{\mu} < 1$, but this restricts the parameter space very little. For instance, a sufficient condition is $(\theta^H - \theta^D) < \frac{1}{4}\left(1+\sqrt{1+8\theta^D}\right)$. Alternatively, one could prevent γ and β from being too high. It is worth noting that the sufficient condition we have just mentioned does not affect any other aspect of this paper and is not even necessary for reasonable values of the remaining parameters. Therefore we can consider $\underline{\mu} < 1$ without concerns about the other theoretical results presented in this paper.

Hence, for $\mu > \underline{\mu}$, it holds true that $B_D > 0$. Therefore:

As $B_N < 0$, we have $NUM > 0 \iff \pi_0 < \frac{-A_N}{B_N} \equiv \bar{\pi}_0{}^N$.

As $B_D > 0$, we have $DEN < 0 \iff \pi_0 < \frac{-A_D}{B_D} \equiv \bar{\pi}_0{}^D$. Define:

$$\bar{\bar{\pi}}_0 \equiv \min\{\bar{\pi}_0{}^N, \bar{\pi}_0{}^D\}$$

Hence, for $\pi_0 < \bar{\bar{\pi}}_0$, we have that:

$$\frac{NUM}{DEN} = \frac{d\alpha}{d\pi_0} < 0$$

$\square$