

Overconfidence, Monetary Policy Committees and Chairman Dominance

Carl Andreas Claussen*, Egil Matsen[†]

Øistein Røisland[‡] and Ragnar Torvik[§]

December 1, 2011

Abstract

Monetary policy decisions are typically characterized by three features: (i) decisions are made by a committee, (ii) the committee members often disagree, and (iii) the chairman is almost never on the losing side in the vote. We show that the combination of overconfident policymakers and a chairman with agenda-setting rights can explain all these features. The optimal agenda-setting power to the chairman is a strictly concave function of the degree of overconfidence. We also show that the quality of advice produced by the central bank staff is higher in a flat organization than in a hierarchical one.

Keywords: Central Bank Governance, Monetary Policy Committees, Overconfidence, Agenda-setting.

JEL Classification: D02, D71, E58.

*Sveriges Riksbank. carl-andreas.claussen@riksbank.se

[†]Norwegian University of Science and Technology. egilm@svt.ntnu.no

[‡]Norges Bank. oistein.roisland@norges-bank.no

[§]Norwegian University of Science and Technology. ragnar.torvik@svt.ntnu.no

1 Introduction

Overconfidence is arguably the best established cognitive bias in the psychology of judgment.¹ DellaVigna (2009, p.341) compactly summarize the bias, arguing that people tend to "...over-estimate their performance in tasks requiring ability, including the precision of their knowledge." Overconfidence has been documented among decision makers in many professions, including physicians, investment bankers, engineers, lawyers and managers.² In this paper, we investigate the consequences of possible overconfidence among decision makers involved in monetary policy decisions.³ We show that overconfidence yields predictions about monetary policymaking that is consistent with a set of stylized facts that cannot be easily explained with existing theories. These facts include (i) disagreement within monetary policy committees (MPCs) after deliberations, (ii) provision of decision power to MPC members, and (iii) that chairmen of MPCs are (almost) never on the losing side when the committees vote.

According to our model, the typical decision structures in contemporary central banks can be seen as an example of "Behavioral Institutional Design" (DellaVigna, 2009). The structures are designed to counteract the effects of cognitive biases and thereby improve welfare.

An important trend in practical monetary policy is the move from indi-

¹Researchers have documented many other biases in information processing (see, e.g., the surveys by Rabin, 1998 and DellaVigna, 2009), but according to DeBondt and Thaler (1995, p.389) overconfidence is perhaps the most robust finding in the psychology of judgment.

²See Odean (1998 p.1892) for references to studies of these and other professions.

³To our knowledge, there are no investigations of monetary policymakers targeted directly at testing overconfidence, but it would be hard to argue that they are exempted from such a common cognitive bias. Apel *et al.*'s (2010) questionnaire survey evidence from Swedish monetary policymakers contains information that is clearly consistent with overconfidence. More generally, the low-predictability, fluid environment in which monetary policymakers typically operate is exactly the type of situation where overconfidence can easily prevail (see, e.g., Odean, 1998). We discuss these and other studies in more detail in Section 2 below.

vidual decision making to committee decision making. The main explanation for this trend in the literature is simple: "two heads are better than one". Monetary policy committees (MPCs) improve decisions by pooling members' information and knowledge (see, e.g., Blinder 2007). Although information pooling within the committee is relevant to understand the transition from individual decision making, it cannot alone explain the use of MPCs. To see this, it is useful to distinguish between two types of information pooling, which we will denote 'pooling by talking' and 'pooling by voting'. 'Pooling by talking' refers to the sharing of views and information among MPC members during deliberations. 'Pooling by voting' refers to the implicit pooling that takes place after deliberations when the MPC votes, or use some other aggregation mechanism, to aggregate the different opinions into one decision. Following Condorcet's famous jury theorem, a huge literature on 'pooling by voting' ('Condorcet effects') has emerged. This literature describes under what conditions voting improves on decisions, see e.g. Koriyama and Szentes (2009) and references therein. Gerlach-Kristen (2006) uses a theoretical macroeconomic model to study Condorcet-effects in MPCs when there is uncertainty and disagreement about the size of the output gap.⁴

If there are no frictions in 'pooling by talking', each member should take the other members' information and arguments into account, and full agreement would result.⁵ As Blinder (2007) also points out, then you do not need a decision-making committee to achieve the pooling benefits. The pooling gains can be achieved by having independent board members serving as mere advisors to the chairman (as is the arrangement at the Reserve

⁴Blinder and Morgan (2005, 2008) and Lombardelli et al. (2005) provide experimental support for pooling by talking and pooling by voting in MPCs.

⁵We assume that the differences in preferred policy decisions before 'pooling by talking' reflects different judgments and information and not different preferences. This is a reasonable assumption, as most MPCs today consist of economic experts and not (former) politicians.

Bank of New Zealand). Alternatively the pooling benefits can be captured by the central bank staff on behalf of the central bank governor. If there are frictions in 'pooling by talking', the MPC members may end up disagreeing also after the deliberation round. We observe extensive disagreement among MPC members in practice, suggesting that 'pooling by talking' is not frictionless. This creates a potential role for 'pooling by voting'. MPC members are distinguished from central bank staff members in that they have *decision power*, whereas staff members have only *advisory power*. The staff can contribute to decisions through 'pooling by talking', while MPC members can contribute through both 'pooling by talking' and 'pooling by voting'. The common institutional setup in central banks is that there is an MPC where each member has decision power, but where the chairman (and other internal members) has access to a staff.⁶ An additional stylized fact is that the chairman is almost always in the majority coalition.⁷

How can overconfidence help explain the use of MPCs? Consider a central bank chairman who receives information and judgments from his staff, but who also has a private signal about the unknown "optimal" interest rate. If he is an unbiased information aggregator, he will optimally weigh the staff's advice and his own signal. To the extent that more people should be involved in the monetary policy decisions, these can be hired as advisors because the chairman will take their views properly into account. If, however, the chairman is overconfident, he will place a too high weight on his own signal and underweight the advice from his staff. Thus, an overconfident chairman does not extract all potential pooling gains inherent in

⁶Ordinary MPC members may also have some access to the staff, or have a small private staff, but staff resources are generally unequally distributed between the chairman and the ordinary members.

⁷The only known example of the chairman being outvoted is the MPC at the Bank of England where the Governor has been in minority twice (out of 157 meetings, see Section 2).

his staff's advice. This increases the risk of bad policy decisions if he alone decides. An MPC with decision power can reduce the risk induced by overconfidence partly because it can intervene against extreme policy proposals, but also because a chairman who has to bring his views to a committee will moderate his proposals. Giving decision power to the MPC is a necessary condition for such moderation to take place. These results hold even though all committee members are subject to the same overconfidence bias. Our approach suggests a different understanding of the role of MPC members: Rather than thinking of MPCs primarily as tools for information pooling, we interpret them primarily as an insurance mechanism against extreme actions from a single policymaker.⁸

Overconfidence precludes agreement about policy in a committee, and it has consequences for the optimal allocation of decision power in the MPC. Through the chairman's unique access to the central bank staff (and perhaps superior competence), the chairman's policy view should on average carry a higher weight than rank-and-file members'. However, overconfidence gives him a suboptimal influence on policy if it is set through simple majority voting. Giving the chairman the agenda-setting right (i.e., the right to propose a policy action that other members must vote for or against) yields an extra layer of decision power, and is a mechanism for restoring (or approaching) his optimal influence.

In addition to the papers mentioned above, our model is related to work by Lohmann (1992), Riboni and Ruge-Murcia (2008), and Gerlach-Kristen (2008), but as we discuss below it differs in important respects. The most

⁸In a rather provocative paper, Romer and Romer (2008) empirically show that the FOMC has not added any value to the forecasts made by the Fed's staff, by adding their own judgments. This result clearly calls into doubt the importance of the MPC as an information pooling device, and thus points to other reasons for having committees (formally) in charge of monetary policy (see also Ellison and Sargent, 2009).

closely related contribution is Gerlach-Kristen (2008), who studies a model with communication errors between MPC members which also yield disagreement among MPC members after deliberations. Although we have another microfoundation and our model of voting and agenda setting is less reduced form, it shares the property that the chairman adjusts his proposal so as to achieve a majority in the MPC. Also in contrast to Gerlach-Kristen (2008) we study normative implications with regard to agenda setting power and the organization of the central bank staff.

The remainder of the paper is organized as follows: In Section 2 we review the evidence on leader dominance and dissent in MPCs. We also briefly discuss the evidence of overconfidence among decision makers, and make the case for its relevance in monetary policy making. In Section 3 we develop a simple model of policy opinions. We show how overconfidence leads to suboptimal use of other people's views and how it precludes agreement among policymakers. With disagreement about policy also after deliberations, there is need for a mechanism to aggregate individual judgments into a policy decision. In Section 4, we explore such a mechanism by developing an agenda-setting model for monetary policy. In Section 5, we turn to normative implications of our model. We discuss the optimal power of the chairman in MPCs, the merits of having central bank insiders on the MPC, and, in particular, the implications of overconfidence for the organization of advice transmission within the central bank staff. Section 6 concludes the paper.

2 Motivating evidence

Our theory is motivated by two strands of evidence. The first set of facts shows that disagreement about policy is common in MPCs where voting

records are available, yet chairmen's views have a strong tendency to prevail. The second line of evidence is the prevalence of overconfidence among decision makers. We now briefly review both sets of evidence, and also discuss why overconfidence is relevant for monetary policymaking.

2.1 Leader dominance and dissent in monetary policy committees

The best known case of leader dominance in MPCs is probably the FOMC under Alan Greenspan's leadership. According to Blinder (2007, p.111), FOMC members under Greenspan's tenure had only one real choice: "to go on record as supporting or opposing the chairman's recommendation, which was certain to prevail." Greenspan chaired the FOMC for over 18 years and was never on the losing side of a vote. The Greenspan period is not unique in the history of the Federal Reserve System. Chappell et al. (2004; 2005 ch. 7) empirically analyze the power of Arthur Burns in his period as chairman of the FOMC. They conclude that Burns' opinion counted about as much as the 18 other committee members put together. An important source of this policymaking weight is reluctance among FOMC members to challenge the proposal offered by an agenda-setting chairman (Chappell et al., 2005, p.101). Like Greenspan, Burns was never on the losing side of a vote in the FOMC. In general, the historical records of the FOMC, as documented by Chappell et al. (2005), indicate a tradition of a strong chairman in the FOMC.

One may argue that the phenomenon of a strong chairman is special for the FOMC; after all Blinder (2004) classified the (Greenspan) FOMC as an *Autocratically-collegial committee*, where "the chairman came close to dictating the committee's decision". At the other side of the central bank

spectrum in terms of chairman influence is the Bank of England's MPC, labeled by Blinder as an *Individualistic committee*. And indeed, the minutes from this MPC reveal a great deal of dissent about monetary policy actions. Between June 6, 1997 and May 10, 2010, the Bank of England had 157 MPC meetings, and there was dissent on the interest rate decision at 93 (59 percent) occasions. But even so, the Governor (Mervyn King) lost the vote at two meetings only. A reasonable interpretation is that the Governor carries a big policyweight also at England's MPC. Minutes from other central banks' MPC meetings strengthen the impression of strong chairmen.⁹ The Bank of Japan's MPC, for instance, held 201 meetings from March 3, 1998 to April 30, 2010. There was dissent on policy on 97 occasions (48 percent), but the chairman was never on the losing side of the vote. Sweden's central bank (The Riksbank) has available minutes from 98 MPC meetings covering January 4, 1999 to April 19, 2010. It was dissenting votes about policy at 33 meetings (34 percent of the time), but again the chairman's proposal always prevailed.¹⁰

This mixture of anecdotal and more careful empirical evidence (as in Chappell et al. 2005) points to the chairman's agenda-setting power as a key source of his heavy policy influence. In the MPCs discussed above, as in many others, the chairman typically proposes a policy decision that the other members must accept or reject. The other members are often reluctant to challenge the chairman's proposal, and this gives him an extra layer of

⁹Most MPCs suppress internal dissent from public view. The Governing Council at the European Central Bank, for example, claims to make decisions by consensus, but offers no voting records against to which assess this claim (Crowe and Meade, 2007).

¹⁰On the face of it, the degree of dissent appears smaller at the FOMC than at the other three MPCs discussed here. Chappell et al. (2005) report that dissents represent 7.8 percent of voting observations over the 1966-96 period. According to Meade (2005), however, the FOMC's internal rates of disagreement are quite similar to dissent rates at the Bank of England, if one looks at opinions expressed during the discussion of policy proposals.

decision power.

2.2 The case for overconfidence in monetary policymaking

A substantial literature in cognitive psychology establishes that individuals tend to be overconfident about the accuracy of their information (Lichtenstein et al., 1982 reviews this calibration literature).¹¹ As mentioned in the Introduction, such overconfidence has been observed in many professional fields, and, as argued below, monetary policymakers operate in an environment where overconfidence may well prevail.¹² One implication of overconfidence is that people listen too little to other people's views and information when forming judgments. Apel *et al.* (2010) present evidence that are consistent with this implication of overconfidence among monetary policymakers. Through a questionnaire survey, these authors asked members of the Swedish Riksbank's Executive Board about how important their peer Board members were to them when forming their views about the appropriate decision. The responses showed that the Board members had little influence on each others' judgments.

Monetary policymakers try to assess the appropriate interest rate in a complex and often fluid environment. It is precisely in such difficult tasks that people exhibit the greatest overconfidence (Odean, 1998). Griffin and Tversky (1992) report that when predictability is low, as is often the case in monetary policy, experts may even be more prone to overconfidence than

¹¹Miscalibration of probabilities is only one manifestation of overconfidence. Others include overestimation of own ability to do well on a task, unrealistic optimism about pure chance events, and overestimation of own contributions to past positive outcomes. See Odean (1998, Section II) for an overview and discussion. Malmendier and Tate (2005) is an excellent recent example of how overconfidence can shed light on economic phenomena.

¹²Angner (2006) argues that economists in general are likely to be victims of significant overconfidence, when acting as experts in matters of public policy (e.g. monetary policy). He bases his case on the nature of the task facing economists and on the institutional constraints under which they operate.

novices, since monetary experts have theories and models of how the economy works which they tend to overweight.

While an extensive experimental literature documents the tendency of overconfidence, there is less research on *why* individuals might be overconfident. Bénabou and Tirole (2002) apply elements from psychology within an economic analysis, and show that various seemingly 'irrational' features of human beings, including overconfidence, can be explained by various 'rational' factors. They focus in particular on the *motivation value* of self-confidence. Being self-confident enhances the ability to undertake difficult tasks. For example, the decision to do a Ph.D. degree implies high costs in terms of time and effort during the process, but with the potential of high return when the degree has been awarded. The student is relatively certain about the costs, while the return depends on the student's ability, on which the student is uncertain. The more self-confident the individual is, the higher is the expected return, and the more motivated is the student for finishing the degree. An implication of this is that experts who have invested much effort in accumulating human capital, are likely to be more self-confident than others. The complementarity between confidence and ability has long been recognized in pedagogics. Moreover, from a demand-side perspective, experts that are self-confident are often more highly valued than experts that appear uncertain. (Politicians want "one-handed economists".) Mechanisms as described above may lead to an equilibrium selection where experts, including monetary policy experts, are overconfident.

In summary, overconfidence may be relevant for central bankers because they are likely to be subject to the same cognitive biases as other decision makers, and because they operate in an environment where such traits can easily prevail. We explore what happens to monetary policy decisions if

policymakers are overconfident, and show that overconfidence can in fact explain many stylized facts about these decisions that existing theories cannot explain.

3 A simple model of policy opinions

3.1 The loss function

The aim of monetary policy is to set the key interest rate r_t to minimize the loss

$$L_t = L(W_t),$$

where W_t is a vector of target variables dependent on r_t . For example, we could have that $W_t = (\pi_t, y_t)$ where π_t and y_t are the inflation gap and the output gap respectively, and $L(\pi_t, y_t) = (\pi_t^2 + \lambda y_t^2)$ as is usual in many models of monetary policy. In order to keep the analysis simple we assume that the decision problem is static so that we can focus on the period loss function and disregard expected future losses. This would, for example, be the case within a standard New Keynesian model without persistence, and where the central bank follows a time-consistent (discretionary) policy.

In practice, monetary policy decisions have dynamic properties in several respects. First, there are "long and variable lags" in the monetary transmission process. Second, monetary policy decisions can be seen as a repetitive task, as the MPC meets regularly. And third, the MPC members receive (imperfect) feedback over time on the quality of previous decisions, allowing policymakers to update their priors and recalibrate their assessment of own and others' competence. For example, the financial crisis starting in 2008 have probably made many policymakers adjust some of their pre-crisis priors on the workings of the economy and on the uncertainty related to future

economic developments.

We note that these dynamic characteristics of monetary policy decisions are inconsequential for our qualitative results, as long as the dynamic learning process does not converge to full certainty about the optimal policy decision. It seems reasonable to assume that full convergence of the learning process do not occur, since the workings of the economy is constantly changing so that the knowledge accumulation will never fully catch up with the economic developments. Thus, even if the confidence in estimates of the optimal interest rate may vary over time depending on the information and feedback to MPC members, the following fact holds: At a given monetary policy meeting, committee members come with a pre-meeting view which they share with the other members, and during the deliberations they form a revised view. We will restrict our analysis to such a "typical" monetary policy meeting and thus drop the time subscripts henceforth.

Monetary policy is conducted in an environment of uncertainty where the interest rate that minimizes L is unknown. Denote this (unobservable) interest rate r^* . Using a second order Taylor approximation of the loss we have that the excess loss by setting a sub-optimal interest rate can be written as $L - L^* = (r - r^*)^2$ where r is the (sub-optimal) interest rate.¹³ In the following we let the expected excess loss $E(L - L^*)$ as given by

$$E(\widehat{L}) = E(r - r^*)^2 \tag{1}$$

¹³Let $L^* = L(W_t(r^*)) \equiv M(r^*)$. A second-order Taylor approximation of $M(r)$ gives

$$L = L^* + M'(r^*)(r - r^*) + \frac{1}{2}M''(r^*)(r - r^*)^2 = L^* + \frac{1}{2}M''(r^*)(r - r^*)^2,$$

where the second equality follows from the first-order condition for minimizing the loss. In linear-quadratic models, $M''(r^*)$ will be constant, and depend of the parameters of the model. For the purpose of this paper, we may, without loss of generality, normalize the second derivative by setting $\frac{1}{2}M''(r^*) = 1$.

be the normative criterion and call (1) the loss function. We assume that those involved in monetary policy decisions know and share this loss function so that there is no disagreement about the goal of policy.

3.2 Timing of events and equilibrium concept

Our set-up thus assumes the following timing of events:

1. Those involved in the monetary policy decision receive an individual noisy signal on the optimal interest rate.
2. Those involved in the monetary policy decision simultaneously reveal their signal (i.e., exchange information) and form a revised individual opinion about the optimal interest rate.¹⁴
3. The interest rate is decided according to the institutional setting in place (e.g., majority voting).

To proceed we thus need to specify who are involved in monetary policy, how they receive their individual signal on the optimal interest rate, and how they revise their signal when they interact with others. In turn, the mapping from this information and communication process to the actual interest rate depends on the decision rule, which is where institutional design enters the analysis.

In the voting game the set of subgame perfect equilibria includes many equilibria in which those involved use weakly dominated strategies, for instance voting against a proposal they support because the others are doing

¹⁴In our model MPC members have nothing to gain by reporting strategically as long as there is simultaneous announcements of signals. The reason is simply that for any individual member, the expected signal from her peers is equal to her own signal. A combination of strategic reporting and sequential announcements would implicitly give higher influence to MPC members that announce late. We focus on committees where all members, except the chairman, have equal opportunities to influence the policy decision. It is well known that information aggregation might fail when individuals act sequentially and observe the actions of all previous agents (see, e.g., Banerjee, 1992).

so (and therefore the proposal will be defeated in any case). It seems unreasonable to assume that monetary policymakers employ (weakly) dominated strategies, and we therefore focus on equilibria in undominated strategies. This implies that the equilibrium of our voting game will be unique.

3.3 Policy opinions

MPC members' task is to form a judgment on the optimal interest rate given by (1). Each member j receives a (noisy) independent signal of optimal interest rate:

$$r_j = r^* + \varepsilon_j, \quad (2)$$

where ε_j is the judgment error, which is characterized by

$$\varepsilon_j \sim N(0, 1/\alpha), \quad \text{all } j = 1, \dots, n + 1.$$

There are $n + 1$ members of the committee and α measures the precision of the members' signals, which may also be interpreted as the competence of the MPC members. (Throughout, we assume that all parameters of the model are constant.) For simplicity, but without changing the qualitative results, we assume that the members do not have any informative priors on the appropriate interest rate, so that their best individual estimate is equal to their signal (2).

3.3.1 Bayesian updating

We will first see that the benchmark case of perfect information updating makes the interest rate decision particularly simple, and that it has a straightforward implication for institutional design.

Symmetrical case. If all MPC members are equally competent and

they have no prior information about the distribution of r^* , their best linear unbiased estimate of the optimal interest rate is:

$$r = \frac{1}{n+1} \sum_{i=1}^{n+1} r_i. \quad (3)$$

It follows from (3) that if all members have the correct perception of their own and others competence, and if they share their individual signals, all members will combine the signals equally and thus end up with the same judgment on r^* . In other words, they will always agree. The institutional aggregation rule from individual opinions to the actual policy decision is irrelevant. Delegating decision power to more than one person in a group will not affect policy.

The precision (inverse of the variance) of the estimate (3) is

$$(n+1)\alpha.$$

The more members of the committee, the better the quality of policy. This is the pooling (Condorcet) argument for committees discussed in the Introduction. Note that in this benchmark case, pooling of policy judgments does not imply that decisions can be improved by delegating decision power to more than one person - advisory power is sufficient.

A chairman with staff. Assume now that one MPC member (the chairman) has better access to a group of m advisors (the staff) than the rest of the committee. For simplicity, we assume that the chairman's access to the staff is unique,¹⁵ and that individual staff members have the same

¹⁵In practice, other MPC members also have some access to staff resources. For example, the staff forecasts and analyses in the Federal Reserve's Greenbook are made available to all FOMC members prior to their meetings. However, our model can be interpreted as a situation where the chairman has access to *more* staff resources than the other members, which is a realistic case for most central banks. The parameter m may then be interpreted

competence as MPC members.

The chairman's optimal combination of his individual signal r_c and his staff's signals is

$$\tilde{r}_c = \frac{1}{m+1} \left(r_c + \sum_{i=1}^m r_i \right), \quad (4)$$

while his optimal posterior (i.e., after MPC deliberations) becomes

$$\hat{r}_c = \frac{1}{n+m+1} \left(r_c + \sum_{i=1}^{n+m} r_i \right).$$

Other members of the MPC can not observe the chairman's individual signal r_c , but only \tilde{r}_c . These members optimal estimate then becomes:

$$\hat{r}_j = \frac{1}{n+m+1} \left((m+1)\tilde{r}_c + \sum_{i=1}^n r_i \right). \quad (5)$$

By substituting from (4), we can immediately see that $\hat{r}_j = \hat{r}_c$. Optimal information aggregation implies that "ordinary" MPC members will take into account that the chairman has (better) access to information from the staff, and end up with the same opinion about optimal policy as the chairman. It is thus still the case that the allocation of decision power is inconsequential for policy. One person with advisors will make the same decision as a committee.

3.3.2 Overconfidence

Consider then the case where policymakers are overconfident. Let $\tilde{\alpha}_j$ be MPC member j 's perception of the precision of his own signal. Following

 as a measure of the chairman's excess staff resources relative to the other members.

Odean (1998), we specify overconfidence as follows:

$$\tilde{\alpha}_j = \alpha k, \quad k \geq 1.$$

The parameter k characterizes the degree of overconfidence. When $k = 1$ policymaker j is an error-free Bayesian, while $k > 1$ implies that he uses the wrong weights when updating his interest rate judgment after receiving new information.

Symmetrical case. Start again with the symmetrical case where all MPC members are truly equally competent. Suppose that MPC members reveal their true signal. Given member j 's perception, the subjectively optimal combination of his own and the other members' signals is:

$$\hat{r}_j = \frac{1}{(n+k)} \left(kr_j + \sum_{i=1}^n r_i \right), \quad i \neq j. \quad (6)$$

Compared to the case of perfect updating, all members overweight their own signal, $k/(n+k) \geq 1/(1+n)$, and underweight the signals of their peers. Even if the members have the correct perception of the their peers' competence, overconfidence implies underweighting the peers' signals. Member j 's perceived precision of his own posterior estimate is

$$(n+k)\alpha \geq (n+1)\alpha,$$

while the true precision of estimate (6) is

$$\frac{(n+k)^2\alpha}{n+k^2} \leq (n+1)\alpha.$$

Overconfidence deteriorates the quality of policy decisions, in the sense that it lowers the true precision in MPC members' judgments.

Equation (6) implies that individual MPC members generally have different posterior judgments on the optimal interest rate; they end up disagreeing even if they share all information.

A chairman with staff. Let us finally look at policy opinions with overconfidence and a staff. The chairman now combines his individual signal and his staff's signals according to

$$\tilde{r}_c = \frac{1}{k+m} \left(kr_c + \sum_{i=1}^m r_i \right). \quad (7)$$

Compared to the case of perfect information updating above, the chairman overweights his own signal and underweights the signals (i.e., advice) of the staff. The chairman treats the signals from his staff and from his MPC colleagues symmetrically, implying that his subjectively optimal posterior estimate becomes:

$$\hat{r}_c = \frac{1}{n+m+k} \left(kr_c + \sum_{i=1}^{n+m} r_i \right). \quad (8)$$

As before, the other members of the MPC can only observe the combination of the chairman's individual signal and his advisors' signal, as given in (7). The ordinary members' subjectively optimal estimate then becomes:

$$\hat{r}_j = \frac{1}{n-1+k+\gamma} \left(kr_j + \gamma \tilde{r}_c + \sum_{i=1}^{n-1} r_i \right), \quad i \neq j, \quad (9)$$

where

$$\gamma = \frac{(m+k)^2}{m+k^2} > 1.$$

Comparing (9) to (5), we see that members overweight their own priors and underweight the judgments of their peers.¹⁶ Again, we see how over-

¹⁶The condition for members' overweighting their own opinion is given by $\frac{k}{n-1+k+\gamma} >$

confidence leads to disagreement about policy even if decision makers share information. MPC members do take into account that the chairman is better informed through his better access to the staff, but put too little weight on this. In addition, the chairman puts too little weight on his staff. For both reasons, the competence of the staff becomes underutilized.

4 Monetary policy decisions

How can an MPC with members who disagree after deliberations reach a decision? Earlier literature has commonly assumed that the MPC aggregates by a simple majority vote (e.g. Blinder and Morgan 2005; Gerlach-Kristen 2006). The policy decision then corresponds to the interest rate preferred by the median MPC member. Although this median-voter perspective on monetary policy decisions is consistent with disagreement in the committee, and thus the need for voting, it is not consistent with the pattern we observe in the outcomes of these votes. In particular, as we discussed in Section 2, MPC chairmen is almost never on the losing side of the vote. To account for this we need to specify an institutional structure where the chairman has the agenda setting power and to specify what happens if he is voted down.

With an agenda-setting chairman, the role of the other MPC members is somewhat different compared to the standard majority voting model. Instead of proposing their own preferred interest rate decisions, as is implicitly assumed in models with majority voting, their role is to assess the chairman's proposal and vote in favor or against it.

Denote the median of the post deliberation opinions $\hat{r}_1, \dots, \hat{r}_{n+1}$ by \hat{r}_{med} . In the beginning of the aggregation stage the chairman proposes a final de-

$\frac{1}{n+m+1}$. After straightforward calculations and inserting for γ this reduces to $(k-1)n(m+k^2) + m(k(m+k^2-2) - m+1) > 0$, which is always fulfilled for $k > 1$.

cision $r_{proposal}$. If the proposal is not adopted by a majority of the members there will be voting resulting in the interest rate \hat{r}_{med} . We assume that for each member there is a cost θ of being in a majority that votes down the chairman's proposal.¹⁷ The cost θ could capture several aspects of collective decision-making. For example, for individualistic committees with voting records, voting down the chairman could be interpreted by the public as lack of confidence in the chairman among the MPC members, and this may affect the credibility of the central bank. This effect is likely to be internalized by individualistic MPC members. For collegial committees without voting records, voting down the chairman's proposal may hurt the collegial spirit, and removing the chairman as a facilitator for unanimous decisions may make it more difficult to agree on a decision.

As usual we solve the game with backwards induction, starting at the final stage with the voting in the MPC. A weakly undominated strategy (recall our definition of equilibrium in Section 3.2) for member j is then to vote against the chairman's proposal if $(r_{proposal} - \hat{r}_j)^2 > (\hat{r}_{med} - \hat{r}_j)^2 + \theta$. Similarly, supporting the chairman's proposal is a weakly undominated strategy as long as $(r_{proposal} - \hat{r}_j)^2 \leq (\hat{r}_{med} - \hat{r}_j)^2 + \theta$.¹⁸ In particular, note that this implies that the median member of the MPC will vote against the chairman's proposal when $(r_{proposal} - \hat{r}_{med})^2 > \theta$, while the median member will support the proposal of the chairman when $(r_{proposal} - \hat{r}_{med})^2 \leq \theta$.

Since for each MPC member j , the expected loss is single peaked around

¹⁷Gerlach-Kristen and Meade (2010) find that dissenting votes on the FOMC is partly explained by the alphabetical voting order. Members who vote early are more likely to dissent, suggesting that it is not voting against the chairman *per se* that carries a cost, but rather being a part of a *majority* that votes against him.

¹⁸This strategy is consistent with Meade's (2005) finding that there is more voiced dissent in the FOMC's deliberation stage than formal dissent in the voting stage. In our model, a member will not cast a dissenting vote if the policy proposal is sufficiently close to his own preferred rate, but the member may still argue that the proposal is wrong in the deliberation stage.

\hat{r}_j , and since the interest rate decided will be \hat{r}_{med} if the proposal of the chairman is voted down, a sufficient and necessary condition for the proposal of the chairman to receive a majority is that it gets the support of the median member of the MPC. Thus, given the voting behavior by the MPC members, the chairman's strategy is as follows: If $(\hat{r}_c - \hat{r}_{med})^2 \leq \theta$ he proposes his post deliberation opinion \hat{r}_c , as given in (8). If $(\hat{r}_c - \hat{r}_{med})^2 > \theta$, he puts forward a *modified proposal* $\hat{r}_{proposal}$, such that $(\hat{r}_{proposal} - \hat{r}_{med})^2 = \theta$, so as to form a minimum winning coalition. It follows that the chairman's proposal will always be supported by a majority in the MPC. Note that we do not specify a cost perceived by the chairman of being voted down beyond the cost of accepting a decision further from his own preferred decision. Even if such extra perceived costs might be high, this will not affect the behavior of the chairman, as he in any case has incentives to adjust his proposal to avoid being voted down.

The interest rate r_D actually set by the MPC is thus the following:

$$r_D = \begin{cases} \hat{r}_c & \text{if } (\hat{r}_c - \hat{r}_{med})^2 \leq \theta \\ \hat{r}_{med} + \sqrt{\theta} & \text{if } (\hat{r}_c - \hat{r}_{med})^2 > \theta \text{ and } \hat{r}_{med} < \hat{r}_c \\ \hat{r}_{med} - \sqrt{\theta} & \text{if } (\hat{r}_c - \hat{r}_{med})^2 > \theta \text{ and } \hat{r}_{med} > \hat{r}_c \end{cases} \quad (10)$$

The fact that the chairman's proposal will prevail does not mean that the ordinary MPC members are without power in the committee. The chairman will modify his proposal if his individually preferred interest rate is sufficiently far from the median view in the MPC, and this clearly gives the other members influence on the decision. The power of the chairman will be higher the higher is θ ; when $\theta \rightarrow \infty$ the chairman always gets his individually preferred rate through, and the rest of the MPC has advisory power only. At the opposite extreme with $\theta = 0$ we are back in the standard

median voter case.

Like us, Riboni and Ruge-Murcia (2008) apply an agenda-setting approach to interest rate decisions. However, we depart from Riboni and Ruge-Murcia by assuming that the *reversion point* is not the *status quo*, but the value preferred by the majority of the MPC (the median judgment). We argue that the *status quo* is not a realistic reversion point for monetary policy decisions, although it can be so for some political decisions, such as voting on economic reforms. To illustrate our point, consider the FOMC meeting in February 1994, where, according to Blinder (2007, p.111), the transcripts clearly indicated that a majority of the FOMC members wanted to raise the funds rate by 50 basis points, while Greenspan proposed a 25 basis point increase. Since Greenspan used his power to get his will through, we will never know what would have happened if his proposal was rejected. Nevertheless, if the chairman proposes a 25 basis point interest rate hike which is voted down because the MPC members see this as too little, it is unlikely that the effect of voting down the chairman would be to leave the interest rate unaltered. Since the FOMC formally reaches decisions by majority voting, it is more reasonable to believe that the FOMC, if rejecting Greenspan's proposal, would have voted for a 50 basis point rise.

The argument against *status quo* as the reversion point in monetary policy decisions is also clear if we assume that the majority of the committee wants an increase in the interest rate, while the chairman proposes an unchanged interest rate. If the reversion point is the *status quo*, it is impossible for the majority of the committee to achieve their preferred decision, which implies that the chairman has unlimited voting power in all situations where he prefers an unchanged interest rate. We will thus argue that the common assumption of *status quo* as the reversion point in traditional agenda-setting

models such as Romer and Rosenthal (1978) and Baron and Ferejohn (1989) reflects the type of decisions these models were applied to, while interest rate decisions are of a somewhat different character.

In this sense our model is more closely related to Gerlach-Kristen (2008), who assumes that MPC members oppose the chairman if they disagree sufficiently with his proposal. In her model the power of the chairman arises from two sources: First, the chairman chairs the discussions to facilitate the communication between the other members, which may limit the disagreement within the MPC. Second, the chairman is more skilled, so it is optimal for the members to place a higher weight on the chairman's judgment when updating the priors. The chairman thus improves information pooling in the deliberation process. Although we have a different reason for disagreement between MPC members and a specified agenda setting procedure, we find an interest rate decision that closely resembles monetary policy in what Gerlach-Kristen (2008) (due to Blinder, 2004) labels Automatically Collegial Committee. The different sources of disagreement in our model and in Gerlach-Kristen's model - overconfidence and communication errors respectively - has, however, implications for the interpretation of the agenda-setting mechanism. In both cases the chairman adjusts his proposal in order to get the median voter indifferent between accepting the chairman's proposal and voting against the chairman. This is unproblematic when the reason for disagreement is overconfidence, since the chairman then can observe the median voter's preferred interest rate. When the reason for disagreement is communication errors, as assumed by Gerlach-Kristen, it is not clear how the chairman can observe the median voter's preferred interest rate and thereby adjust his proposal optimally towards the median voter.

Our agenda-setting model also has similarities with the model of the

central bank and the government in Lohmann (1992). She assumes that the government can override the central bank's decision, but has to pay a fixed cost. The cost of overriding the central bank could be interpreted as the degree of central bank independence. The focus in Lohmann's paper is to show that Rogoff's (1985) solution to the time-inconsistency problem by delegating monetary policy to an independent but 'conservative' central bank could be improved upon by limiting the degree of independence. Her point is that a 'conservative' central bank works well for moderate supply shocks, but when sufficiently large shocks occur, the cost of having a 'conservative' central bank dictating monetary policy becomes larger than the gain, because a 'conservative' central bank stabilizes output too little relative to what society prefers. By having the opportunity to override the central bank when large shocks occur, the game between the central bank and the government acts as an insurance against bad monetary policy when extreme shocks occur.

In our model, there is a judgment aggregation problem that calls for an insurance against extreme decisions by the chairman. Similarly to Lohmann's model, the MPC members are expected to override the chairman when they think the chairman tries to force through a bad decision. There is, however, a difference between the two mechanisms. In Lohmann, it always leads to a better policy when the government forces the central bank to adjust policy. In our model, the MPC's influence on the decision can deteriorate the quality of monetary policy, since the chairman, through his staff access, on average is better informed than the other committee members.

5 Normative implications

The predictions of the simple approach developed above is consistent with the actual operation of MPCs that we discussed in Section 2. Moreover, the approach highlights that a main role for MPCs is to step in if the chairman is astray, i.e., to provide insurance against extreme policy errors. We now proceed to show that our framework has additional important implications for the institutional design of central banks. We discuss three such design issues: the optimal degree of chairman agenda setting power, the merits of having central bank insiders on the MPC, and, in particular, the organization of advice transmission in the central bank staff.

5.1 The optimal power of the chairman

In designing monetary policy institutions, a key question that has been little studied is what is the *optimal* degree of agenda setting power. A straightforward implication from the analysis above is that the optimal degree of agenda setting power is decreasing in committee size n and increasing in staff size m . The reason for this is simply that the quality of the signal of the chairman relative to that of the median member of the MPC is decreasing in n and increasing in m .

The effect of the degree of overconfidence on the optimal agenda setting power is less obvious. Naturally, if there is no overconfidence there is no need for agenda setting power. The more overconfidence there is the less the MPC improves the quality of the decision from the chairman, which viewed in isolation pulls in the direction of allocating stronger agenda setting power to the chairman. On the other hand, the more overconfidence there is, the poorer the chairman utilizes the signals from his staff. Therefore, with much overconfidence, the pre-deliberation interest rate the chairman prefers has

(on average) a poorer quality. Viewed in isolation, this pulls in the direction of allocating less agenda setting power to the chairman. The question is therefore not how the degree of overconfidence affects the quality of the preferred interest rate of the chairman or the MPC, but how the *relative* quality of the preferred policy by the chairman and the MPC is affected.

To investigate this question we have to rely on numerical methods because, under agenda setting, the interest rate decision involves taking the median of random variables from distributions with different second order moments. There is no explicit mathematical expression for the median in such cases. In the simulations we fix the true precision α to one and impose normally distributed judgment errors ε . We calculate the optimal agenda-setting power of the chairman, measured by θ , as a function of overconfidence k for various combinations of committee and staff size, all with $m > n$. Each simulation is based on 10,000 draws.

Figure 1 here

For all combinations of m and n , the pattern that emerges is as depicted in Figure 1: a hump-shaped relationship between the degree of overconfidence and the optimal cost of going against the chairman's proposal. This pattern occurs because of the channels described above: Overconfidence leads to poor use of staff advice by the chairman, but also to less precise policy opinions among MPC members after deliberations. When overconfidence is mild the latter effect dominates, the optimal agenda setting power is increasing in the degree of overconfidence. A marginal increase in k from a low level means that the chairman lowers the staff influence, but he still gives it considerable weight; he is significantly better informed than the

other MPC members. Meanwhile, these members give less weight to the chairman's opinion as the degree of overconfidence increases. When the distortions due to overconfidence increases from a low level it is optimal to increase the power of the chairman.

When overconfidence is severe, on the other hand, more overconfidence pulls in the direction of less agenda setting power to the chairman. To understand this result, note that the optimal agenda setting power of the chairman goes to zero as k approaches infinity. In the limit the chairman is so overconfident that he has no better signal than the other MPC members because he completely ignores the inputs from his staff. Allocating him agenda setting power in such a case reduces the quality of monetary policy, as the policy view of the median MPC member is on average better than that of the chairman. The gradually less influence of staff advice as k increases is the dominating factor along the negatively sloped parts of the lines in Figure 1.

The earlier literature on monetary policy decision making has mainly compared simple majority voting to decisions taken by the chairman alone. Our analysis above shows that as long as there is positive but not an infinite degree of overconfidence, neither of these corner solutions are optimal. This result stands in clear contrast to the conclusion in Gerlach-Kristen (2008). She finds that interest rate setting is worse in committees with heavy chairman influence (autocratically collegial committee) than in individualistic committees. According to our analysis, a committee with a strong chairman is optimal as long as decision makers have bounded overconfidence and the chairman has better access to staff advice.

The optimal agenda setting power balances the better access to information by the chairman and the insurance involved in having another look

at the chairman's preferred interest rate. The chairman on average makes a better projection of the optimal interest rate than ordinary MPC members due to his closer interaction with the staff. However, an overconfident chairman may sometimes be terribly wrong even after consulting with the staff. Agenda setting trades off these conflicting arguments because it gives a higher weight to the person with the expected best policy signal at the same time as it works as an insurance against letting the possible mistakes of one individual have a too strong impact on policy decisions.

5.2 Central bank insiders as MPC members

Our model also sheds light on whether all MPC members should be full-time central bank insiders. Viewed in isolation, a normative implication of the model is that all MPC members should have the same opportunity as the chairman to get information from and interact with the central bank staff. An 'indoor' MPC will improve the average quality of policy judgments among MPC members. A possible paradox, however, is that in such a case the model suggests that the optimal agenda setting power of the chairman should be lower, while in practice such an arrangement may make it more costly for MPC members to vote against the chairman. This is especially relevant if career concerns for insiders become dependent on how their competence is viewed by the chairman. (Obviously there may also be other counterarguments against such a proposal that is not captured by the model, such as the danger of conformity and group thinking.) Thus if one chooses to have an 'indoor' MPC, it is important with arrangements that makes the governor 'weak' in the sense that MPC members will know that there is a low cost of opposing him.

5.3 Central bank hierarchy

A related issue is the optimal organization of the chain of policy advice within the central bank. In the simple model of policy opinions laid out in Section 3, we implicitly assumed that the central bank governor receives opinions and advice directly from all staff members. In practice, however, central banks are organized as hierarchies. Given the scale of modern central banks, hierarchical structures are probably effective in decentralizing the production (of analyses, speeches, liquidity management, etc.) tasks of the bank's staff (e.g., Radner, 1992). It is far from obvious, however, that a hierarchical chain of advice transmission maximizes the quality of the advice that is ultimately given to the chairman.

To analyze this issue, we discuss two alternative organizations of the transmission of staff advice. Recall that there are m staff members in total. In Section 3, we assumed a *flat* organization of advice transmission, where the chairman receives m independent signals (pieces of advice).

Consider instead a *hierarchical* organization, where we introduce an additional layer in the structure of the staff. Let $m = m_d + m_e$, where m_d is the number of "directors", who give advice directly to the chairman, and m_e is the number of "economists", who give advice to the director of their respective unit. Each director is in charge of a "unit" consisting of $\frac{m_e}{m_d}$ economists who give their individual advice directly to their director. We assume that the directors are prone to the same type of overconfidence as the MPC members. (Overconfidence among economists is irrelevant matter for the outcome here, since their task is not to form posterior advice based on signals from others).

Director h 's posterior advice, based on her prior signal and the advice from her subordinates, is

$$\hat{r}_h = \frac{1}{k + \frac{m_e}{m_d}} \left(kr_h + \sum_{i=1}^{\frac{m_e}{m_d}} r_i \right).$$

The true precision of director h 's estimate is

$$\frac{\left(\frac{m_e}{m_d} + k\right)^2 \alpha}{\frac{m_e}{m_d} + k^2}.$$

As above, we assume that the chairman knows the true precision of the directors' advice (but is overconfident about his own precision). He then combines the advice from the m_d directors with his own signal according to

$$\check{r}_c = \frac{1}{k + \hat{\gamma} m_d} \left(kr_c + \hat{\gamma} \sum_{i=1}^{m_d} r_i \right),$$

where

$$\hat{\gamma} = \frac{\left(\frac{m_e}{m_d} + k\right)^2}{\frac{m_e}{m_d} + k^2}.$$

Let P^s denote the precision of the chairman's policy view after receiving the staff's advice in an organization of type s , where $s = H$ denotes an hierarchical organization and $s = F$ denotes a flat organization.

In a hierarchical organization, upon getting advice from the m_d directors, the precision is given by

$$P^H = m_d \hat{\gamma} \alpha = \frac{(m + m_d(k - 1))^2 \alpha}{m + m_d(k^2 - 1)}. \quad (11)$$

With a flat organization the precision of the chairman's policy view is

$$P^F = m \alpha.$$

The difference between a flat and a hierarchical organization's precision

becomes (after some algebra)

$$P^F - P^H = \frac{m_d(m - m_d)(k - 1)^2\alpha}{m + m_d(k^2 - 1)}. \quad (12)$$

Recall that $m_d = 0$ or $m_d = m$ would imply a flat organization, so a hierarchical organization in our model is characterized by $0 < m_d < m$.

Equation (12) contains two unambiguous normative implications: First, a flat organization gives better aggregate advice than a hierarchical organization when directors are overconfident ($k > 1$).¹⁹ When the directors have a correct perception of own competence ($k = 1$) the two types of organizations are equally good in terms of quality of advice. Second, it is straightforward to show that the difference in advice quality increases in k . Thus, the more severe are frictions in information aggregation, such as overconfidence among professionals, the more does a hierarchical structure of advice transmission pollute the quality of the views that reach the top of the organization.

6 Conclusion

In contemporary central banking, the formal decision power over monetary policy is delegated to an MPC rather than a single individual. There is considerable disagreement about policy within MPCs, leading to a great deal of dissent in actual policy decisions. Yet, MPC chairmen almost never lose a vote about monetary policy.

In this paper, we have provided a theory for these stylized facts about the decision structure in modern central banks. Our theory rests on the notion that people are not perfect information aggregators, and in particular

¹⁹This results also holds if directors are *under*confident, i.e., $k < 1$.

that they may be subject to overconfidence. An MPC with decision power reduces the policy risk occurring when an overconfident chairman gives a suboptimal weight to staff judgments. Overconfidence also yields disagreement and dissent among decision makers, and this gives the chairman too little influence if policy is set through simple majority voting. Giving the chairman extra decision power through agenda-setting rights restores his influence, but also means that he generally will not lose when there is a vote in the MPC. We emphasize that the MPC still has important, but largely unobservable policy influence by inducing moderation from the chairman (and his staff).

We have seen that even though overconfidence provides a reason for an institutional setting where the chairman has agenda setting power, the extent of such power should be limited if overconfidence is perceived to be a severe problem. Neither a chairman deciding alone or an MPC with simple majority voting are optimal as long as there is positive but a bounded degree of overconfidence. Finally, we have shown that overconfidence implies that flat structures of information transmission within the central bank staff are superior to hierarchical structures in terms of the quality of advice reaching the bank's chairman.

References

- Angner, E. (2006). "Economists as experts: Overconfidence in theory and practice", *Journal of Economic Methodology* 13, 1-24.
- Apel, M., Claussen, C. A. and P. Lennartsdotter (2010). "Picking the Brains of MPC Members". Sveriges Riksbank Working Paper No. 237.
- Banerjee A. (1992). "A Simple Model of Herd Behavior", *Quarterly Journal of Economics* 107, 797-817.
- Baron, D. and J. Ferejohn (1989). "Bargaining in Legislatures", *American Political Science Review* 83, 1181-1206.
- Benabou, R.J-M. and J. Tirole (2002). "Self-Confidence and Personal Motivation", *Quarterly Journal of Economics* 117, 871-915.
- Blinder, A.S. (2004). *The Quiet Revolution: Central Banking Goes Modern*, Yale University Press, New Haven.
- Blinder, A.S. (2007). "Monetary Policy by Committee: Why and How?", *European Journal of Political Economy* 23 106-123.
- Blinder, A.S. and J. Morgan (2005). "Are two heads better than one? Monetary policy by committee", *Journal of Money, Credit and Banking* 37, 789-812.
- Blinder, A.S. and J. Morgan (2008). "Leadership in Groups: A Monetary Policy Experiment," *International Journal of Central Banking* 4, 117-150.
- Chappell, H.W., R.R. McGregor and T. Vermilyea (2005). "Majority Rule, Consensus Building, and the Power of the Chairman: Arthur Burns and the FOMC", *Journal of Money, Credit and Banking* 36, 407-422
- Chappell, H.W., R.R. McGregor and T. Vermilyea (2005). *Committee Decisions on Monetary Policy*, MIT Press, Cambridge.
- Crowe, C. and E.E. Meade (2007). "The Evolution of Central Bank Governance around the World", *Journal of Economic Perspectives* 21, 69-

90.

DeBondt, W.F.M and R.H. Thaler (1995). "Financial Decision-Making in Markets and Firms: A Behavioral Perspective", in: R Jarrow et al. (eds.), *Handbooks in Operations Research and Management* vol. 9, Elsevier Science, North-Holland, Amsterdam.

DellaVigna, S. (2009). "Psychology and Economics: Evidence from the Field", *Journal of Economic Literature* 47, 315–72.

Ellison, M. and T.J. Sargent (2009). "A defence of the FOMC," Economics Series Working Papers 457, University of Oxford, Department of Economics

Gerlach-Kristen, P. (2006). "Monetary policy committees and interest-rate setting", *European Economic Review* 50, 487–507.

Gerlach-Kristen, P. (2008). "The role of the Chairman in setting monetary policy: Individualistic vs autocratically collegial MPCs", *International Journal of Central Banking* 4, 119-143.

Gerlach-Kristen, P. and E.E. Meade. "Is There a Limit on FOMC Dissents? Evidence from the Greenspan Era", Working Paper, <http://home.intergga.ch/gerlach/ellen.pdf>

Griffin, D.W. and A. Tversky (1992). "The Weighing of Evidence and the Determinants of Confidence", *Cognitive Psychology* 24, 411-435.

Koriyama, Y. and B. Szentes (2009). "A resurrection of the Condorcet Jury Theorem", *Theoretical Economics* 4, 227–252.

Lichtenstein, S., B. Fischhoff and L. Phillips (1982). "Calibration of probabilities: The state of the art to 1980", in: D. Kahneman, P. Slovic, and A. Tversky (eds.) *Judgment Under Uncertainty: Heuristics and Biases*, Cambridge University Press, Cambridge.

Lohmann, S. (1992). "Optimal Commitment in Monetary Policy: Credibility versus Flexibility", *American Economic Review* 82, 273-286.

Lombardelli, C., J. Proudman and J. Talbot (2005). "Committees versus individuals: an experimental analysis of monetary policy decision-making", *International Journal of Central Banking* 1, 181–205.

Malmendier, U. and G. Tate (2005). "CEO Overconfidence and Corporate Investment", *Journal of Finance* 60, 2661-2700.

Meade, E.E. (2005). "The FOMC: Preferences, Voting, and Consensus", *Federal Reserve Bank of St. Louis Review* 87, 93–101.

Odean, T. (1998). "Volume, Volatility, Price, and Profit When All Traders Are above Average", *Journal of Finance* 53, 1887-1934.

Rabin, M. (1998). "Psychology and Economics", *Journal of Economic Literature* 36, 11-46.

Radner, R. (1992). "Hierarchy: The Economics of Managing", *Journal of Economic Literature* 30, 1382-1415.

Riboni, A. and F.J. Ruge-Murcia (2008). "The Dynamic (In)efficiency of Monetary Policy by Committee", *Journal of Money, Credit and Banking* 40, 1001-1032.

Rogoff, K.S. (1985). "The Optimal Degree of Commitment to an Intermediate Monetary Target", *Quarterly Journal of Economics* 100, 1169–1189.

Romer, T. and H. Rosenthal (1978). "Political Resource Allocation, Controlled Agendas and The Status Quo", *Public Choice* 33(4), 27–43.

Romer, C.D. and D.H. Romer, (2008). "The FOMC versus the Staff: Where Can Monetary Policymakers Add Value?", *American Economic Review* 98, 230-235.

Figure 1

