.

A Political Economy Theory of Central Bank Independence

Gauti B. Eggertsson and Eric Le Borgne

August 2009

Abstract

We propose a simple theory to explain why, and under what circumstances, a politician delegates policy tasks to a technocrat in an independent institution and then analyze under what conditions delegation is optimal for society. Our theory builds on Holmström's (1982, 1999) "hidden effort" principal-agent model. The election pressures that politicians face, and the absence of such pressures for technocrats, give rise to a dynamic incentive structure that formalizes two rationales for delegation, one highlighted by Hamilton (1788) and the other by Blinder (1998). Delegation trades off the cost of having a possibly incompetent technocrat with a long-term job contract against the benefit of having a technocrat who 1) invests more effort into the specialized policy task and 2) is better insulated from the whims of public opinion. A natural application of our framework suggests a new theory of central bank independence.

KEY WORDS: delegation, elections, career concerns, learning-by-doing, insulation, central bank independence.

JEL CLASSIFICATION NUMBERS: E58, E61, H11, J45

Eggertsson: Federal Reserve Bank of New York, Gauti.Eggertsson@ny.frb.org; Le Borgne: IMF and World Bank, ELeBorgne@imf.org. The authors thank several colleagues and seminar participants, especially Allan Drazen, an anonymous referee and the editor for many useful comments. Patricia Egner gave excellent editorial advice. The views expressed here are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York, the Federal Reserve System, the International Monetary Fund, or the World Bank Group.

"It is a general principle of human nature, that a man will be interested in whatever he possesses, in proportion to the firmness or precariousness of the tenure by which he holds it; will be less attached to what he holds by a momentary or uncertain title, than to what he enjoys by a durable or certain title; and, of course, will be willing to risk more for the sake of the one, than for the sake of the other."

Alexander Hamilton (1788), Federalist Paper 71: The Duration in Office of the Executive

"Many governments wisely try to depoliticize monetary policy by, for example, putting it in the hands of unelected technocrats with *long terms of office* and *insulation* from the hurly-burly of politics" (emphasis added).

Alan S. Blinder (1998), former Vice Chairman of the Federal Reserve, pp. 56-57.

1 Introduction

In most countries, certain important and prestigious public policy tasks, such as interpreting the constitution or conducting monetary policy, are delegated to public officials who are insulated from both job insecurity and political interferences thanks to long-term employment contracts. The judges at the U.S. Supreme Court, for example, are appointed for life and are independent of the elected executive. Similarly, many governments delegate monetary policy to an independent institution that is ruled by career public officials whose terms of office are much longer than the average elected official's.

In some cases, delegation of political power is constitutionally mandated, as in the case of the Supreme Court, while in others politicians find it in their own interest to delegate power, even if they are not constitutionally required to do so. A voluntary transfer of power has been particularly noticeable in the case of monetary policy in recent decades. Granting the Bank of England independence, for example, was one of the first acts of the Labor government when it took power in the United Kingdom in 1997. That politicians give away power of their own accord may be considered somewhat of a puzzle. Most people, and many political economy models, assume that politicians are in the game of accumulating power rather than giving it away.

In this paper, we tackle the following questions. First, why do politicians voluntarily give away some power and delegate critical public policy tasks to technocrats in independent institutions? Second, do politicians delegate enough tasks—i.e., can delegation be optimal from the perspective of society while not being in the interest of the politician? And if so, under what conditions?

To analyze these questions, we model the incentives that technocrats and politicians face. Our model assumes a representative democracy where *all* agents have the same preferences. The difference between a politician and a technocrat is that the politician is subject to elections while the technocrat is not. This difference, captured by short-term election pressures, drives all of our

results. Election pressures create a rich, dynamic incentive structure in our model through the interaction of private and public signals and unobservable effort levels by the policymakers.

The framework we propose underlines a basic cost-benefit trade-off for delegation. The cost of delegation is that the technocrat cannot be dismissed from his job, even if he is incompetent. In the case of the Federal Reserve Board this cost is not trivial, since each governor is appointed by the president for a fourteen-year term (with the exception of the chairman).[1] In the case of the Supreme Court this cost is even starker because justices are appointed for life. Delegation, of course, is more costly if the ability of job candidates cannot be ascertained *perfectly* prior to hiring (a realistic feature of any hiring decision). On the benefit side, we identify two positives of delegation. They formalize Alexander Hamilton's argument for long-term contracts for the executive and Alan Blinder's case for having a technocrat running monetary policy.

Hamilton's rationale for delegation to a technocrat with a long-term employment contract is that he will be ready to "risk more" and/or work harder for a policy task that he holds, with certainty, for a considerable amount of time. In our model, a long-term employment contract enables the policymaker to have a long-term horizon that may improve his performance. In particular, we show that a longer employment contract gives the long-term appointee an incentive to invest more effort into his *job-specific* decision making, due to a learning effect, thereby increasing the quality of his decisions. We label this the *learning-by-doing* effect.

Meanwhile, in his argument for delegation, Blinder suggests that monetary policy should be put in the hands of unelected technocrats with *long terms of office* and *insulation* from the hurly-burly of politics. We model Blinder's idea by showing that a politician has an incentive to follow the whims of public opinion. In particular, we show that when there is a public signal about the policy problem, the politician has a perverse incentive to follow public opinion even if he has *superior* private information that contradicts it. The reason is that the politician is reluctant to go against public opinion, because if he is wrong (and the public is correct) this would signal low ability and lack of good judgment and prevent the politician from being reelected. In contrast, by following public opinion, the politician reveals nothing about his ability, so the public has no incentive to throw him out of office. A technocrat with a long-term job contract, however, does not face this dilemma because he is insulated from the electorate. We label this the *insulation effect*.

We draw several lessons from the model, besides identifying and formalizing the basic trade-off outlined above. One is that the more complex the policy task, the more desirable it is to delegate it to a technocrat. Another is that the larger the degree of private benefit from political office (which may include outright corruption), the lower the politician's incentive to delegate. Uncertainty about the competency of an officeholder works in the same direction. We are also able to show that delegation is always socially optimal when the policymaker bases his decisions on public opinion, even if the politician may not always find it privately optimal to delegate power in

[1] If a governor does not complete a term, his or her successor does not get a new fourteen-year term, but first has to serve the remaining time of the existing term. Hence, many governors have, in practice, a shorter term once appointed.

these circumstances. Furthermore, we show that if delegation occurs, it is always socially optimal in our model. The converse, however, is not true; i.e., some tasks are not delegated in equilibrium even though their delegation would improve social welfare.

Our theory of delegation is motivated by the observation that the job security of technocrats is typically greater than that of politicians. The longer tenure, and especially the dynamic incentive effect that arises from the absence of election pressures, is what mainly distinguishes technocrats from politicians. To the extent that short-term election pressures are important in differentiating politicians from technocrats, our model offers interesting insights.

It is straightforward to embed the basic channel into other models. We illustrate this by embedding our political economy channel into a standard monetary model. This example provides a new theory of central bank independence (CBI), which is the original motivation of this study. While our theory gives an alternative to the standard time-inconsistency rationale for CBI, it does not contradict it. Our theory, therefore, complements existing work on CBI.

2 Related literature

Our paper provides a formalization of the literature on constitutionalism, whereby certain issues are taken "off the table" in everyday policymaking in order to insulate policymakers from the short-term pressures of public opinion. One criteria for insulation in this literature is that there should be a broad agreement on the long-term goals of policies that are taken off the table (see, e.g., Elster (1995) and Drazen (2002)), a criteria that our model satisfies, since all agents have the same preferences.

Our model of dynamic incentives in public policy builds on the seminal "career concerns" model of Holmström (1982, 1999). Holmström shows that, within a firm, a manager of unknown ability can be induced to supply effort that is not directly observed by relying on the manager's concerns about his future career. We extend Holmström's model in three ways. First, the manager/policymaker is disciplined by elections, whereas in Holmström's model incentives are shaped by the expectation of higher future wages. Second, we introduce learning. Third, we introduce a public signal about the desirability of the policy decisions. These three extensions relate to several papers. The first extension is done in Persson and Tabellini (2000).[2] They do not, however, address the issue of delegation of policy tasks. The second extension is related to the large literature on learning-by-doing on the job that has been applied in several areas of economics (starting with Arrow (1962)), although we are not aware of applications in our political economy context. The third extension should be interpreted as an application of the seminal paper on investors' herding behavior by Scharfstein and Stein (1990). They show that investors may choose to ignore superior information in favor of following the "herd," in the same way the politician in our model may

[2] See also Le Borgne and Lockwood (2006), who study questions that are more similar to ours, although they do not model the learning or insulation effects. Instead, they assume that effort and ability are multiplicative, which also gives a rationale for delegation (through an "experimentation" channel).

choose to ignore superior private information in favor of public opinion.[3]

Maskin and Tirole (2004) address questions similar to ours. They investigate which policy decisions should be made by politicians and which ones should be decided by judges.[4] A similarity between our paper and theirs is that both illustrate that elections give a politician perverse incentives to follow public opinion instead of taking advantage of superior private information. In Maskin and Tirole's model a politician may have a pandering motive because he has a hidden agenda. Our insulation motive is unrelated to hidden agendas, since everyone has the same preferences and therefore relates to a different set of policy tasks. It arises because elections give rise to perverse *dynamic incentives* even if all agents have the same political agenda.

Alesina and Tabellini (2007) also address questions similar to ours, i.e., when should a policy task be delegated to a technocrat? A key difference is that Alesina and Tabellini assume that the trade-off between technocrats and politicians stems from the different intrinsic preferences of these agents: Some agents are born politicians while others are born technocrats. In contrast, we assume a representative democracy where everyone has, ex ante, the same preferences. Another difference is that while we illustrate the role of insulation and learning, Alesina and Tabellini focus on other reasons for delegation, such as issues related to dynamic inconsistency of public policy as a rationale for delegation, which is an important consideration that we abstract from.

3 The framework

The economy is populated by a large number of citizens and evolves over two time periods, $t = 1, 2$. There is a political office that can be occupied by only one citizen, the "officeholder." In this representative democracy, the officeholder is entrusted with (and held accountable for) policy decisions. The game begins with a random selection/election of a citizen from the population to become the elected officeholder (called "politician" thereafter). The politician can decide whether to make the policy decision himself or to delegate it to an appointed agent, whom we call a "technocrat," who is not subject to reelection (i.e., he has a longer term in office).

Citizens' utilities are a function of a policy decision. The quality of this policy decision depends both on ability and effort of the officeholder. Utility is given by $U_t^j = E\sum_{t=1}^{2} v_t^j$, where the index j refers either to the officeholder ("o") or a citizen ("c"). The period payoff is

$$v_t^j = v_t^c\left(i_t, \bar{r}_t\right) + \Omega^j - \frac{\alpha}{2}(e_t^j)^2, \tag{1}$$

where $v_t^c\left(.\right)$ is a payoff due to a public policy that accrues to all citizens at time t; this payoff depends on the interaction of the policy decision i_t and a shock $\bar{r}_t$. The term $\frac{\alpha}{2}(e_t^j)^2$ represents the cost of effort, and the effort level of each citizen is *private information*. The quality of the policy decision (i.e., the choice of i_t) is a function of how well the policymaker can predict $\bar{r}_t$.

[3] In the political economy literature, the herding incentive of politicians has also been recently modeled by Canes-Wrone, Herron, and Shotts (2001) in a somewhat different context.

[4] Besley and Coate (2003) contrast direct election with political appointment of regulators in a model where electing regulators produces more pro-consumer regulators.

This prediction depends on the effort and ability of the policymaker in a way described in the next subsection. Since effort is useful only to the officeholder, $e^j = 0$ when $j = c$. Ω^j is a binary variable that is equal to Ω if citizen j is the officeholder ($j = o$) and is zero otherwise. Ω is an "ego rent" from being in office and managing public policy (as in Rogoff and Sibert (1988)), deriving from the prestige associated with managing public affairs. We abstract from discounting.

To clarify the game, we now put more structure on (1), although we take advantage of the more general formulation in Section 6. In each period $t = 1, 2$, whoever is the officeholder has to make a binary policy decision by choosing i_t^H or i_t^L. This decision can refer to any public policy; one may, for example, think of it as choosing between whether or not to fight a war, to invest in a public project, or to raise or lower interest rates. $\bar{r}_t$ can take on only two values, either $\bar{r}^H$ or $\bar{r}^L$ with equal probability, i.e., $Prob(\bar{r}_t = \bar{r}^H) = Prob(\bar{r}_t = \bar{r}^L) = \frac{1}{2}$. The policymaker (and the public) does not know this variable before choosing i_t. At the end of each period, the realization of $\bar{r}_t$ becomes publicly available and hence it becomes known if the policy decision was good or bad. The policy problem is thus to select i_t optimally, conditional on a forecast about $\bar{r}_t$. In the baseline illustration we normalize the payoffs so that if the policymaker correctly predicts $\bar{r}_t$ then $v^c\left(i^H, \bar{r}^H\right) = v^c\left(i^L, \bar{r}^L\right) = 0$, and if he makes a mistake then $v^c\left(i^H, \bar{r}^L\right) = v^c\left(i^L, r^H\right) = -1$.

Private signal The officeholder receives a private signal σ_t about the realization of the shock $\bar{r}_t$. The quality of the private signal depends both on effort and ability of the officeholder two signals can occur that correspond to the two shocks, i.e., $\sigma_t \in \{\sigma^H, \sigma^L\}$. If the officeholder bases his decision on his private signal, his decision is trivial: He will set i^H if he observes σ^H and i^L if he observes σ^L. The combination of the officeholder's forecast and the realization of the shock produces a state s. Four possible states can therefore arise in a given period, depending on whether the shock (labeled as H or L) has been rightly (R) or wrongly (W) predicted: i.e., $s = r\sigma \in \{HR, HW, LR, LW\}$. Since payoffs are symmetric in the baseline example for HR, LR on one hand and HW, LW on the other, the only payoff-relevant events are whether the officeholder is right or wrong. The probability that a new officeholder receives a right signal σ conditional on a shock $\bar{r}$, his effort level e, and expected ability θ is:

$$\begin{aligned} Prob(\text{Right signal for the officeholder}) &= Prob(\sigma^H \mid \bar{r}_1 = \bar{r}^H, e_1, \theta) = Prob(\sigma^L \mid \bar{r}_1 = \bar{r}^L, e_1, \theta) \\ &= \frac{1}{2} + \theta + e_1 \end{aligned} \tag{2}$$

and the probability of the wrong signal is:

$$\begin{aligned} Prob(\text{Wrong signal for the officeholder}) &= Prob(\sigma^H \mid \bar{r}_1 = \bar{r}^L, e_1, \theta) = Prob(\sigma^L \mid \bar{r}_1 = r^H, e_1, \theta) \\ &= \frac{1}{2} - \theta - e_1. \end{aligned} \tag{3}$$

Here, θ is the forecasting ability of the officeholder. It is a random draw from a distribution that can take two values, "good" or "bad," i.e., $\theta_G > \theta_B = 0$ with equal probabilities $\frac{1}{2}$. We refer to $a \in \{G, B\}$ as the (ability) types of the citizens. We assume that citizens do not know θ but they all know the joint distribution of θ, i.e., there is *symmetric incomplete information*, following

Holmström (1982, 1999). This is a key simplification, because in equilibrium this means that expected utility of the officeholder will then be the same for the policymaker and the public. This is one reason why we will obtain closed-form solutions. The alternative assumption of asymmetric information about the type of officeholder leads to significantly more complicated analysis with complications unrelated to the key mechanism we want to emphasize. This uncertainty about types can also apply to a subset of the pool of candidates for public office (e.g., citizens with a Ph.D. in economics). All that is needed for our results is the fact of some remaining uncertainty among candidates. Of course, the probability needs to be between zero and one, which implies bounds on the feasible values of e and θ. When the first-period officeholder is in office in period two, we assume there is a *learning-by-doing* effect, so that

$$\begin{aligned} Prob(\text{Right signal for officeholder}) &= Prob(\sigma^H \mid \bar{r}_2 = \bar{r}^H, e_2, \theta) = Prob(\sigma^L \mid \bar{r}_1 = \bar{r}^L, e_1, \theta) \\ &= \frac{1}{2} + \theta + \beta e_1 + e_2 \end{aligned} \tag{4}$$

and then the probability of the wrong signal is

$$\begin{aligned} Prob(\text{Wrong signal for officeholder}) &= Prob(\sigma^L \mid \bar{r}_1 = r^L, \bar{r}_1, \theta) = Prob(\sigma^H \mid \bar{r}_1 = \bar{r}^H, e_1, \theta) \\ &= \frac{1}{2} - \theta - \beta e_1 - e_2, \end{aligned} \tag{5}$$

where the coefficient $\beta > 0$. We interpret this coefficient as corresponding to the complexity of the task at hand since it implies that effort put in the first period translates into better learning for period two forecasting. One would expect job experience and learning on the job to be extremely important for a complex task and less relevant for a simple or a manual task. Adding the term βe_1 in the second-period production function is a simple and transparent way of introducing a learning-by-doing effect, i.e., experience in office increases the ability of the officeholder. Thus, effort in the first period is akin to a (sunk cost) investment; if the first-period officeholder is also in office in the second period, his effort invested in learning (say, about the functioning of an economy, the structure of the central bank, etc.) means his accumulated (task specific) knowledge gives him an incumbency advantage compared to a new officeholder.

Meanwhile, a new officeholder in period two uses the forecasting technology in (3); a period-one officeholder who is thrown out of office cannot pass on his *individual-and-job-specific* knowledge to his successor.

Again the probabilities (4) and (5) have to be between zero and one. Because e_1 and e_2 are endogenously determined (as we will show below) and depend on the structural parameters, we will assume that the values of θ,β, α, and Ω are such that the probabilities are always between zero and one. We do not derive the bounds that this restriction implies for the structural parameters, but simply assume that their values are such that there is an interior solution for e_1 and e_2. Deriving these bounds is relatively straightforward, but does not yield any extra insights.

Public signal We also assume that there is a public signal ρ_t about the realization of the shock $\bar{r}_t$ that can take the values ρ^H and ρ^L, and that this signal is revealed at the beginning of each

period (before the decision maker decides whether or not to exert any effort to receive a private signal). We assume that this signal is created through a process similar to the private signal, i.e., $Prob$(Right public signal)=$Prob(\rho^H \mid \bar{r}_1 = \bar{r}^H) = Prob(\rho^L \mid \bar{r}_1 = \bar{r}^L) = \frac{1}{2} + E_p\theta = \frac{1}{2} + \frac{1}{2}\theta_G$ and the probability of the wrong signal is $Prob$(Wrong public signal) $= Prob(\rho^H \mid \bar{r}_1 = \bar{r}^L) = Prob(\rho^L \mid \bar{r}_1 = r^H) = \frac{1}{2} - E_p\theta = \frac{1}{2} - \frac{1}{2}\theta_G$, where $E_p\theta$ is the expected ability of the whole population. This captures the idea that public opinion is formed by the average cognitive abilities of the citizens. Not expecting to be pivotal in the policy decision, each citizen does not exert any effort to improve the quality of the public signal. Hence the signal, on average, is of worse quality than the private signal of the officeholder as long as he exerts any effort.

Institutions The politician (randomly selected from the set of citizens), once "elected" at the beginning of period $t = 1$, decides between two possible institutions/regimes:
Delegation: At the beginning of period $t = 1$, the politician appoints an agent (a technocrat) to be the officeholder and to make the policy decision. Since all citizens are, ex ante, identical, this agent is randomly selected by the politician from the set of citizens and is (contractually) in office for *both* periods. When the politician delegates the policy decision, his utility then becomes the same as that of a representative citizen (he is not in charge of policy, so $j \neq o$ in (1)) except that he receives a small political rent from office Ω^D. This captures the idea that elected office confers some prestige to the politician and the fact that voters, who in equilibrium understand the motives for delegation, view favorably a politician who delegates a task (say, monetary policy), because citizens know that this is welfare-increasing for them (shown later in the paper), or the fact that the politician is the principal of the technocrat, a position that confers some "ego rent." We assume that delegation is credible, i.e., that the politician can make a binding promise not to fire the person he delegates powers to for the duration of the game. This is an important—and entirely reasonable—assumption. To remove a Federal Reserve Board governor or a Supreme Court justice, for example, is exceedingly hard, and while not impossible, it would be prohibitively costly for most elected politicians.
Political control: At the beginning of period $t = 1$, the elected politician decides to make the policy decision himself. The politician is in office during period $t = 1$, but faces an election at the beginning of period two. At this stage, an opponent is randomly selected from the set of remaining citizens. Citizens then vote on the opponent versus the incumbent, and the winner is the officeholder in period $t = 2$. Under both delegation and political control, we impose the individual rationality condition that the officeholder prefers to be in office than not, which boils down to assuming high enough Ω.

Timing of events At the beginning of the game, the politician decides whether or not to delegate the policy task. The probability that people put on the decision maker being of high ability is denoted by q_1 (it is 1/2 under our assumptions). We now enter period one. The public signal, ρ_1, about the shock is revealed at the beginning of the period. The decision maker then has to choose effort level, e_1, to obtain his private information. The private signal, σ_1, is then

revealed. Observe that the private signal already incorporates whatever information is revealed by the public signal. After receiving the private signal, the policy maker has to make the policy decision i_1. Finally, the state of the world, S_1, is revealed and the public observes if the policy decision was correct. After the public observes the state of the world, it updates its beliefs, q_1, about the ability of the officeholder by Bayes' rule. We denote the new belief by q_2. If the decision maker faces an election, then it takes place between periods one and two, taking into account the updated beliefs, q_2. The sequence of events for period two is the same as in period one.

4 Hamilton's hypothesis: Learning

This section illustrates the basic trade-off between delegation and political control, abstracting from the public signal. We show this by illustrating the learning effect that arises from delegation, which formalizes Hamilton's hypothesis. In Section 4, we introduce a public signal into the model which generates an insulation incentive.

In order for the politician to decide whether to make the policy decision himself or to delegate it, he compares his utility under both regimes. This, in turn, is a function of the equilibrium level of effort and ability. We first derive the effort level chosen in period two. This effort level is the same whether or not there is delegation, since the officeholder faces the same maximization problem in either case. We then calculate the effort choice in period one, which depends on the policy regime.

Period-two effort We solve the officeholder's decision problem by backward induction. In the last period ($t = 2$), we have a static game. The maximization problem of the officeholder (politician or technocrat) is to select e_2 to maximize utility in period two that is given by $U_2 = \left\{\Omega - \frac{1}{2}\alpha e_2^2\right\} - \left\{\frac{1}{2} - E_a\theta - e_2\right\}$. The first bracket is the net rent from office. The second bracket is the probability of the officeholder making the wrong policy decision. This probability is decreasing in effort and expected ability of the officeholder, denoted $E_a\theta$. If the officeholder was not in office in the previous period, then $E_a\theta = \frac{1}{2}\theta_G$. If the officeholder was in office in the previous period, then $E_a\theta = q_2\theta_G + \beta e_1$ because of the learning-by-doing effect explained above. Observe that q_2 is the updated probability of the person in office being of high ability (see Section 4, where this number is computed explicitly using Baysian updating).[5] Recall that we assume, for simplicity, that $v^c\left(i^H, \bar{r}^H\right) = v_t^c\left(i^L, \bar{r}^L\right) = 0$, so only the probability of making the wrong decision appears explicitly in this utility function. Also recall that we normalized $v^c\left(i^H, \bar{r}^L\right) = v^c\left(i^L, \bar{r}^H\right) = -1$ so that there is nothing that scales the probability in the utility function. The first-order condition

[5] In equilibrium, the public's and the office holder's expectations about his utility coincide because of the assumption of symmetric incomplete information.

of this maximization problem is $\alpha e_2 - 1 = 0$, so the equilibrium effort, denoted e^*, is:[6]

$$e^* = \alpha^{-1} > 0. \tag{6}$$

The effort choice does not depend on $E_a\theta$; thus it is independent of the officeholder's ability and his tenure.

Period-one effort under delegation When the politician delegates the policy decision, the first-period maximization problem of the technocrat (superscript "T") is to select e_1 (taking the optimal choice of e_2 as given) to maximize expected utility in period one, given by

$$U_1^T = \left\{2\Omega - \frac{1}{2}\alpha e_1^2 - \frac{1}{2}\alpha e_2^2\right\} - \{\frac{1}{2} - \frac{1}{2}\theta_G - e_1\} - \{\frac{1}{2} - \frac{1}{2}\theta_G - \beta e_1 - e_2\}. \tag{7}$$

The first bracket captures, again, the technocrat's net rent from office. The second (third) bracket is the probability of making the wrong policy decision in period one (two). The probability of making a mistake in period two is also a function of effort in period one because of the learning-by-doing effect. This leads to first-order condition $\alpha e_1^T - (1+\beta) = 0$, so a technocrat's first-period equilibrium level of effort is

$$e_1^T = \frac{1}{\alpha}(1+\beta) > 0. \tag{8}$$

The first-period effort level is composed of two terms: The first (1) is the first-period gain from increasing effort, while the second (β) represents the marginal gain due to the value of effort in period one on the quality of the policy decision in period two; we call this term the *learning-by-doing effect*, which arises because effort in period one has a long-lasting effect on the officeholder's forecasting ability. This implies that effort is greater in period one than two, i.e.,

$$e_1^T = \frac{1+\beta}{\alpha} > e^* = \frac{1}{\alpha} > 0 \tag{9}$$

Period-one effort without delegation The analysis is more complex without delegation (i.e., political control) because we have a game of incomplete information. We characterize the perfect Bayesian equilibria (PBE) of this game. The effort choice of the officeholder in period two, whether new in office or not, is $e^* = \alpha^{-1}$. So, the expected utility from electing a new politician in office is $U_2^{NP} = -\{\frac{1}{2} - \frac{1}{2}\theta_G - e^*\}$, where NP stands for electing a "new politician." In deciding between a new officeholder and the incumbent, voters consider the expected utility of retaining the incumbent politician, which is given by $U_2^{IP} = -\{\frac{1}{2} - q(s_1)\theta_G - \beta e_1 - e^*\}$, where IP stands for reelecting the "incumbent politician". The variable $q(s_1)$ is indexed by the realization of the state in period one. It is the probability voters put on whether the policymaker is a good type.

[6] We imposed the restriction on the structural parameters that the probabilities in (2)-(5) have to be between zero and one and that the solution for effort is interior. Accordingly, the condition above implies some restrictions on the feasible range for α^{-1}.

Note that this probability depends on the realization of the policy decision and the shock $\bar{r}_1$, both of which are observed before voting. The politician is thrown out of office if $U_2^{IP} < U_2^{NP}$, i.e.,

$$q(s_1)\theta_G + \beta e_1 < \frac{1}{2}\theta_G. \tag{10}$$

Note that the learning-by-doing effect, i.e., βe_1, gives the politician an incumbency advantage. We compute the probability $q(s_1)$ by using Bayes' rule. In the case where the politician makes a wrong forecast, for example, it can be calculated by computing the probability that the shock is low when the signal is high (this probability is the same, under our assumption, as the one where the shock is high but the signal is low), i.e.,

$$\begin{aligned} q(s_1 \in \text{Wrong signal}) &= q_W(\theta_G, e_1) = \Pr ob(\theta_G \mid \bar{r}_1 = \bar{r}^L, \sigma^H, e_1) \\ &= \frac{\frac{1}{2}Prob(\sigma^H \mid \bar{r}_1 = \bar{r}^L, e_1, \theta_G)}{\frac{1}{2}Prob(\sigma^H \mid \bar{r}_1 = \bar{r}^L, e_1, \theta_G) + \frac{1}{2}Prob(\sigma^H \mid \bar{r}_1 = \bar{r}^L, e_1, \theta_B)} = \frac{\frac{1}{2} - \theta_G - e_1}{2(\frac{1}{2} - \frac{1}{2}\theta_G - e_1)}. \end{aligned}$$

If we substitute this into (10), the officeholder is thrown out of office if, and only if, the incumbency effect is small enough, i.e.,

$$\beta e_1 < \frac{\frac{1}{2}\theta_G^2}{1 - \theta_G - 2e_1}. \tag{11}$$

In order to focus on election versus appointment of officeholder, we assume that this condition is satisfied, *in equilibrium*, since if it were violated, the incumbent would always be elected.[7] We can always choose β small enough such that this condition is satisfied (since in the limit, as $\beta = 0$, there is no incumbency advantage).

Using Bayes' rule, the probability of the policymaker being a good type conditional on receiving the right signal is given by

$$q(s_1 \in \text{Right signal}) = q_R(e_1, \theta_G) = \frac{\frac{1}{2} + \theta_G + e_1}{2(\frac{1}{2} + \frac{1}{2}\theta_G + e_1)}. \tag{12}$$

The choice of first-period policies and effort of the incumbent politician is e_1 (taking e_2 as given) that maximizes expected utility in period one, given by

$$\begin{aligned} U_1^P = &\left\{\Omega - \frac{1}{2}\alpha e_1^2\right\} + \{\frac{1}{2} + \frac{1}{2}\theta_G + e_1\}(\Omega - \frac{1}{2}\alpha e_2^2) - \{\frac{1}{2} - \frac{1}{2}\theta_G - e_1\} \\ &- \{\frac{1}{2} + \frac{1}{2}\theta_G + e_1\}\{\frac{1}{2} - q_R(e_1, \theta_G)\theta_G - \beta e_1 - e_2\} - \{\frac{1}{2} - \frac{1}{2}\theta_G - e_1\}\{\frac{1}{2} - \frac{1}{2}\theta_G - e^*\}. \end{aligned} \tag{13}$$

The difference between this expression and the utility of the technocrat (U_1^T) is that the politician is uncertain whether or not he will stay in office in period two. This is the reason why the net rent from office in period two, i.e., $(\Omega - \frac{\alpha}{2}e_2^2)$, is multiplied by $\{\frac{1}{2} + \frac{1}{2}\theta_G + e_1\}$, which is the probability of the officeholder being right in period one. If wrong, he loses the election in the beginning of

[7] A simple way of restating this condition in terms of variables that are not endogenously determined is to observe that e_1 needs to be such that the probabilities in (2)-(5) are between zero and one. We also need to ensure that it is individually rational for both the incumbent and the opponent to stand for election, given the voters' cutoff rule. This requires that this gain be positive, which can be ensured by selecting high enough Ω.

period two. The third curly bracket corresponds, as in the delegation case, to the probability of making a wrong decision in period one. The first term of the second line corresponds to the loss of utility associated with making a wrong forecast in period two, and this term is weighted by the probability of the politician still being in office. Note that the probability of the politician still being in office needs to be weighted by the probability of being a good type conditional on the politician being reelected (i.e., he made the right decision in period one) using the Bayes' rule. Finally, the last term is the utility of the politician in period two if he gets the policy wrong in period one, weighted by the probability of that happening. In this case, he will receive the same utility as the average citizen and the outcome does not depend on his effort in period one or two.

The first-order condition with respect to effort is:

$$\alpha e_1 = (\Omega - \frac{1}{2\alpha}) + 1 + \beta(\frac{1}{2} + \frac{1}{2}\theta_G + 2e_1) + \theta_G\{q_R - \frac{1}{2}\} + \frac{\partial q_R}{\partial e_1}\theta_G\{\frac{1}{2} + \frac{1}{2}\theta_G + e_1\}. \tag{14}$$

Using (12), we observe that both last two terms equal $\frac{1}{4}\frac{\theta_G^2}{(\frac{1}{2}+\frac{1}{2}\theta_G+e_1)}$ but with the opposite sign and thus cancel out. Assuming that $\alpha \neq 2\beta$ to ensure that a solution exists, the effort level that solves (14), denoted e_1^P, is

$$e_1^P = \frac{\{1 + \beta(\frac{1}{2} + \frac{1}{2}\theta_G)\} + \{\Omega - \frac{1}{2\alpha}\}}{\alpha - 2\beta} > 0. \tag{15}$$

We show in the appendix that the denominator of this expression has to be positive. Recall that the structural parameters assumed are such that this solution is interior (i.e., does not violate the restriction that the probabilities (2)-(5) are between zero and one). There are two effects that influence the politician's effort: the *learning-by-doing* effect and the *career-concerns* effect. As already discussed, the former is related to both the complexity of the task and the expected first-period ability of the officeholder, while the latter is an increasing function of the net rent from office. Under the delegation regime, the career-concerns effect was not present since the officeholder is in office in period two with probability one (so that his net rent is guaranteed).

Welfare and the existence of trade-offs between regimes We can now analyze whether a newly elected politician has any incentive to delegate the policy decision. Recall that if a politician delegates the policy decision but still remains in office, he gets a rent $\Omega^D > 0$ (as the principal of the technocrat) that is strictly less than the *net* rent from office that the technocrat obtains. Let us call the politician's utility when he delegates U^D and let it be called U^P if he does not. Thus the politician delegates if, and only if,

$$U^D > U^P, \tag{16}$$

where, after substituting the equilibrium effort values into the above equation, we get

$$U^D = 2\Omega^D - 1 + \theta_G + \frac{(1+\beta)^2 + 1}{\alpha}, \tag{17}$$

which is only a function of the structural parameters Ω^D, θ_G, β, and α. Note that this utility is an increasing function of θ_G with a slope of 1. The utility of the politician if he does not delegate

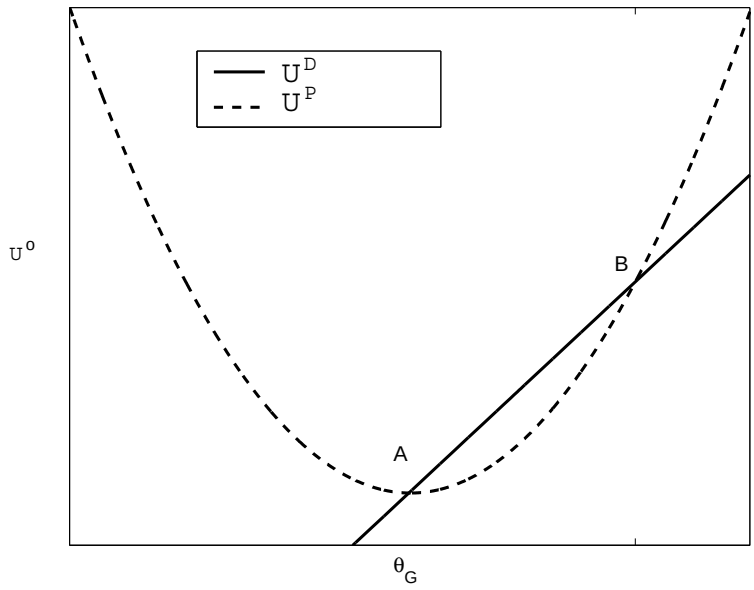

Figure 1: Possible cutoff for welfare

can also be written in terms of the structural parameters by substituting (6) and (15) into (13). We can express U^P in terms of a quadratic function of θ_G so that

$$U^P = \gamma_1 + \gamma_2 \theta_G + \gamma_3 \theta_G^2, \tag{18}$$

where

$$\gamma_3 = \left(\frac{1}{4} \frac{\beta^2}{\alpha - 2\beta} + \frac{1}{4} \frac{\beta^3}{(\alpha - 2\beta)^2} - \frac{1}{8} \alpha \frac{\beta^2}{(\alpha - 2\beta)^2} + \frac{1}{4} \right) > 0$$

and we show in the appendix that γ_3 must be positive for all permissible parameters in the model.[8]

The generic form of U^P and U^D as a function of θ_G is shown in Figure 1. The curve when the politician chooses not to delegate is denoted U^P for "politician" while the curve when he chooses to delegate is denoted U^D for "delegation" to a technocrat.

As can be seen, there are, in general, two possible intersections for these curves. However, only one of them corresponds to a region in the parameter space for an admissible equilibrium of our model, namely point B. To see this, suppose that point A is an equilibrium. Note that the slope of the U^P curve is smaller than that of U^D. Consider now the utility under delegation. Using (17), we have $\frac{\partial U^D}{\partial \theta_G} = 1$ and, using the envelope theorem, we can calculate the slope of U^P by taking a partial derivative of (13) to obtain[9]

$$\frac{\partial U^P}{\partial \theta_G} = 1 + \frac{1}{2}(\Omega - \frac{\alpha}{2}(e_2)^2) + \frac{1}{2}\theta_G \beta e_1 + \frac{1}{2}(q_R - \frac{1}{2}) + \frac{\partial q_R}{\partial \theta_G}\{\frac{1}{2} + \frac{1}{2}\theta_G + e_1\} > \frac{\partial U^D}{\partial \theta_G}. \tag{19}$$

[8] Note that the value of γ_1 and γ_2 are not necessary for the argument that follows.

[9] One can confirm by (12) that $\frac{\partial q_R}{\partial \theta_G} > 0$ and we know from Bayes' rule that $q_R > \frac{1}{2}$.

We know, however, that this cannot be the case in point A (because the slope of the U^P line in Figure 1 is smaller than the technocrat line). So this cannot be an equilibrium for the parameter values we assume. What does this mean? It means that one of the assumptions we used to derive condition (19) has to be violated at or to the left of point A. The assumptions we used were that e is positive, θ_G is between zero and one and the probabilities of the model have to be between zero and one.[10] This suggest that parameter values that generate the curves corresponding to the region equal to or to the left of A in figure (1) are not of interest for our analysis (that region, for example, would correspond to negative values of θ_G or e).

The discussion above indicates that, for a given set of parameters, equilibrium will be on the right-hand side of point A in Figure 1. If the ability spread is in the range between A and B, delegation dominates political control, while political control dominates delegation to the right-hand side of B. To have an interesting theory of endogenous delegation, we must show that an equilibrium exists on both sides of point B, depending on the parameters. If this can be established, we can discuss how the equilibrium depends on parameter values and whether or not delegation occurs in equilibrium. Thus the question of interest is whether there *exists* a trade-off between delegation and political control, i.e., are there configurations of the parameters on both sides of B that satisfy all the restrictions of the model?[11] The proof of existence of this policy trade-off is trivial. We only need to establish a numerical example that shows that for one set of parameters delegation dominates and for another set it does not. This example is given below. With existence of a solution established, we can discuss how the welfare trade-off depends on the different parameters of the model. That discussion does not require any of the parameter values in Example 1.

Example 1. Suppose that the coefficients of the model are $(\theta_G, \Omega, \Omega^D, \alpha, \beta) = (0.2, 0.2, 0.15, 100, 1)$. This implies that $(e_1^T, e_1^P, e^*, U^D, U^P) = (0.02, 0.018, 0.01, -0.45, -0.46)$, i.e., $U^D > U^P$, the necessary conditions are satisfied, and delegation dominates. If we assume instead that $\theta_G = 0.4$ (and the same values for the other parameters), then $(e_1^T, e_1^P, e^*, U^D, U^P) = (0.02, 0.0193, 0.01, -0.25, -0.2402)$, i.e., $U^D < U^P$, the necessary conditions are satisfied, and delegation is dominated so that the politician retains control.

The welfare trade-off The previous section established that endogenous delegation can occur and that whether or not it happens depends on the parameters of the model. In other words, the nature of the policy task has an effect on whether or not delegation takes place. We established that it can be *individually rational* for a self-interested politician to delegate policy decisions. In this case, the officeholder becomes the agent of the elected politician, who himself is the representative of citizens at large. It is also clear that the individual-rationality constraint of an

[10] This latter restriction implies that, for a given θ_G, α is constrained so that the equilibrium e is small enough to keep the probabilities between zero and one. We can always choose α high enough to satisfy this.

[11] In addition to the welfare functions derived in equations (17) and (13), a candidate solution has to satisfy the conditions that 1) the implied probabilities of every event are between zero and one; 2) effort is positive; and 3) the individual rationality constraints of the politician and technocrat are satisfied.

appointed officeholder is satisfied because he also gets a strictly positive private rent from office such that he is better off than the rest of the citizenry. We also showed that the utility of the officeholder is increasing in θ_G whether or not he delegates power, but it increases *more* if he retains power, i.e., the slope of U^P is greater than U^D.

The relationship between the utility of the politician with and without delegation is shown in Figure 2. This figure corresponds to the right-hand side of point A in Figure 1. For small values of the ability parameter, delegation to a technocrat dominates. As θ_G increases (or, more generally, the spread between θ_G and θ_B), the utility of keeping power increases until the two lines cross and the politician does not delegate any more. The intuition is straightforward and captures the main trade-off between delegation and democratically accountable power: If the politician appoints an independent technocrat, he cannot fire him (even if he turns out not to be as competent as expected)! If there is significant uncertainty about the ability of the technocrat, the politician will be increasingly reluctant to give away power that extends beyond his election term because society can be stuck with an incompetent technocrat. On the other hand, if there is a small difference between θ_G and θ_B, then the cost of being stuck with a "bad" technocrat is smaller and the politician is more likely to delegate.

The slope of the two curves therefore indicates that the politician will be more willing to delegate power to technocrats if he can easily identify whether or not the candidates are qualified to handle it, since this implies a lower spread between θ_G and θ_B. What is relevant here is not necessarily the ability of all citizens, but the relevant pool of candidates among whom the politician has to choose.[12] While one may argue that it is easy to identify the qualifications needed for a central bank, the ability of technocrats to perform some other tasks may be subject to greater uncertainty. Some public offices – fiscal policy, for example – depend on interpreting the wishes and needs of the electorate, which may frequently shift over time. In this case, technical qualification may not be sufficient for determining whether a candidate is suitable for appointment or not, which is an argument for the politician to maintain power in his own hands.

Other parameters, however, can lead to endogenous delegation: namely, the complexity of the task and the rents from office. Either of these can explain why some tasks are delegated and others not, or even why the same tasks are delegated in some countries and not in others.

Having established the existence of a trade-off between delegating power and retaining it, and having shown how it depends on the expected ability of the officeholder, we now turn to analyzing the effect of the learning-by-doing effect, the cost of effort, and rents.

Learning-by-doing (β) effect We are interested in knowing the effect at the margin (i.e., around the cutoff point of the two curves in Figure 2) of increasing the learning-by-doing effect

[12] Hence a sensible interpretation is that, in this case, the discrepancy between θ_G and θ_B is low. This does not imply, however, that the unconditional variance in the whole population is small; e.g., most people would probably be very badly suited for the role of a Fed governor. This just means that the degree of uncertainty among the people who qualify for the job is relatively small.

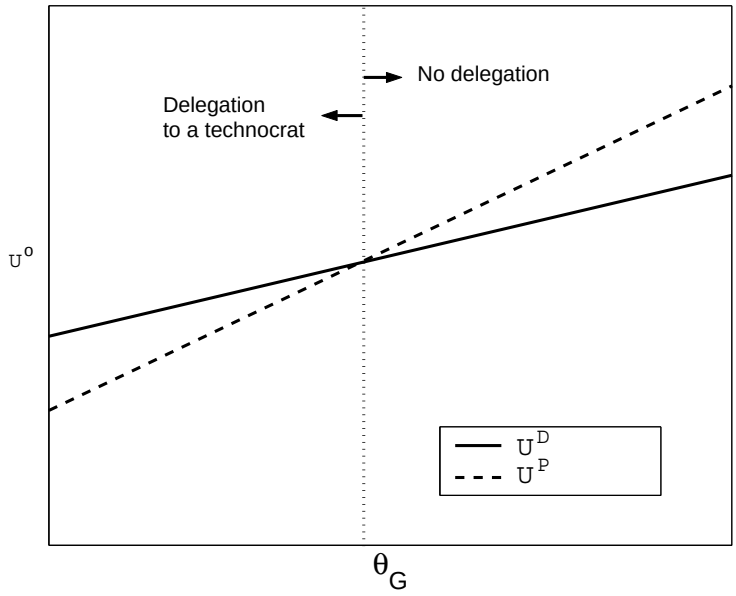

Figure 2: Basic trade-off in delegation

in the forecasting task. To do this, we take the partial derivative of U^D, i.e.,

$$\frac{\partial U^D}{\partial \beta} = e_1 = 2\frac{1+\beta}{\alpha}.$$

We can use the envelope theorem to calculate the derivative of U^P with respect to β, i.e.,

$$\frac{\partial U^P}{\partial \beta} = e_1(\frac{1}{2} + \frac{1}{2}\theta_G + e_1).$$

Consider the point at which the two curves intersect. At this point, we know that the effort level expanded by the technocrat must be higher than that of the politician, so that $e_1^P < e_1^T$, since the technocrat needs to compensate the politician for the loss of rent. This implies that

$$\frac{\partial U^P}{\partial \beta} = e_1^P(\frac{1}{2} + \frac{1}{2}\theta_G + e_1^P) < \frac{1+\beta}{\alpha}(\frac{1}{2} + \frac{1}{2}\theta_G + e_1^P) < \frac{\partial U^D}{\partial \beta}$$

and the inequality follows because we know that $(\frac{1}{2} + \frac{1}{2}\theta_G + e_1) < 1$.

So, the more learning by doing on the job, the more desirable it is to delegate power, as shown in Figure 2. First, observe that increasing β shifts both curves up such that it increases welfare in both regimes and the region where the politician delegates increases. The reason is that, in both cases, the officeholder will inherit his effort in period one (in the case of the politician, this increases welfare only if he is reelected), and the higher β, the more useful this will be for policy. The result above illustrates, however, that utility under the delegation regime (the solid curve) increases by more than utility if the politician retains power. The reason is that the technocrat has a higher stake in being in office (because he cannot be fired), so his effort in period one is

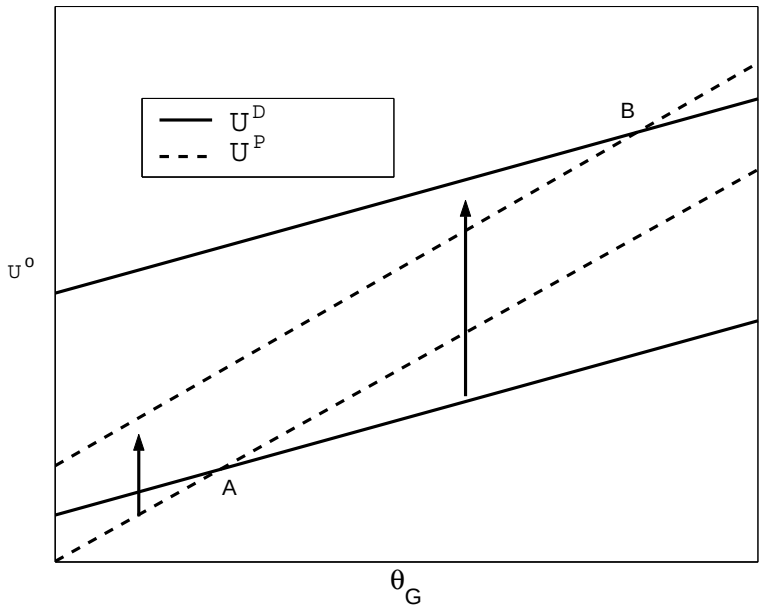

Figure 3: An increase in complexity of the task increases delegation.

more sensitive to the learning parameter. In contrast, the envelope theorem implies that if the politician conducts the policy task, then $\partial e_1^P/\partial\beta = 0$ (when this derivative is evaluated at the optimal e_1^P). Since the utility of the politician (and society in general) is increasing in the effort of the technocrat in period one, this increase in the technocrat's effort raises the politician's utility.

Thus, the more learning by doing, the more important it is to give the task to a technocrat with a long-term contract, because the long-term benefit of being in office gives him the *incentive* to invest in learning (supplying effort). This effect is reminiscent of Hamilton's claim in the Federalist Paper 71 about the "tenure of the executive" cited earlier. Consider, as an example, the setting of monetary policy, a task that requires highly specialized knowledge. The Federal Reserve is the single largest employer of Ph.D. economists in the world. The discussion at FOMC meetings is highly technical; district bank presidents and staff economists present a wealth of statistical and economic concepts before a policy decisions are made. It is not uncommon for policymakers to analyze results from sophisticated statistical models or to discuss different economic theories when interpreting different data series. Few if any public policy tasks involve technical discussions of this kind. In the case of fiscal policy, for example, it is rare that economic or statistical theory is discussed in the context of everyday decision making at the U.S. Treasury. Indeed, a Secretary of the Treasury may not even be familiar with the technical language.

One possible interpretation of β is that it reflects "complexity" of the task. The variable β measures the degree to which tomorrow's tasks are similar to today's tasks, meaning that expertise gathered today is useful tomorrow. It seems reasonable to assume that learning on the job is more important in tasks that are complex than in those that are relatively simple.

Private information $(1/\alpha)$ **effect** We next consider comparative static with respect to the parameter α^{-1}, which indexes the cost of effort. The higher α^{-1}, the more effort both the technocrat and the politician will exert and thus the better informed they will be when called on to make policy. As α^{-1} reaches the limit of feasible range for the model parameters, the policymaker will exert enough effort to be perfectly informed and to reach the best possible decision. We are interested in knowing the effect of increasing $1/\alpha$ at the margin. The partial derivative of U^D is $\frac{\partial U^D}{\partial \alpha^{-1}} = (1+\beta)^2+1$.We can use the envelope theorem to calculate the derivative of U^P with respect to α^{-1} :

$$\frac{\partial U^P}{\partial \alpha^{-1}} = \frac{1}{2}\alpha^2(e_1^P)^2 + \{\frac{1}{2} + \frac{1}{2}\theta_G + e_1^P\}\frac{1}{2}.$$

Consider now the point at which the two curves U^P and U^D intersect. We know that $e_1^P < e_1^T$ by the same argument as in the last section, so

$$\frac{\partial U^P}{\partial \alpha^{-1}} = \frac{1}{2}\alpha^2(e_1^P)^2 + \{\frac{1}{2} + \frac{1}{2}\theta_G + e_1^P\}\frac{1}{2} < \frac{1}{2}(1+\beta)^2 + \{\frac{1}{2} + \frac{1}{2}\theta_G + e_1^P\}\frac{1}{2} < \frac{\partial U^D}{\partial \alpha^{-1}},$$

which suggests that the comparative static of α^{-1}, is the same as that of β. Thus our analysis suggests that the higher α^{-1} the higher the incentive for delegating policy to a technocrat, and the graphical analysis in Figure 2 applies once again. We suggested in the last subsection that the parameter β was a measure of the "complexity" of the policy task. We also wish to suggest a similar interpretation of α^{-1} because it measures the extent of "private information" the policymaker has relative to the public. We defer a further discussion of this interpretation to the next section.

Rents and corruption effect Finally, we can also see that increasing Ω will shift the "politician curve" up, which makes delegation less likely. The intuition is immediate: Delegation is a way for the politician to increase his utility (and that of a representative citizen) because it depends on the quality of the policy decision. However, delegation comes at the expense of forsaking the net private rent associated with managing public policy (this private benefit can range from an ego rent to outright corruption). To the extent that this private rent is large, the concern of the politician for the common good of society will be smaller and he will be less likely to delegate.

Proposition 1 (endogenous delegation of political power): *In the equilibrium of the model, the set of variables* $(e_1^T, e_1^P, e^*, U^D, U^P)$ *depends on the set of parameters* $(\theta_G, \Omega, \Omega^D, \alpha, \beta)$. *Delegation occurs if* $U^D > U^P$. *Both delegation and nondelegation equilibria exist. Which arises depends on the value of the parameters* $(\theta_G, \Omega, \Omega^D, \alpha, \beta)$. *For the policy task, the following holds at the margin:*

i) An increase in the complexity of the task (measured by α^{-1}*or* β*) increases the incentive to delegate.*

ii) An increase in the rent from office (Ω) *reduces the incentive to delegate.*

iii) An increase in the uncertainty of the officeholder's ability $(\theta_G - \theta_B)$ *reduces the incentive to delegate.*

The proof of existence of both equilibria follows from Example 1 and the rest from the discussion in the sections above. Proposition 1 establishes a theory of delegation. It shows that the decision to delegate depends on the properties of the policy task and how it affects political rents. The question we now turn to is whether or not delegation is optimal for society as a whole rather than from the perspective of the politician. Thus, the question is under what circumstances $U_c^D > U_c^P$ where D denotes delegation, P that the politician retains power, and c that utility refers to that of the representative citizen. Observe that the politician's utility is the sum of his private rent from office and the citizens' utility U_c^D or U_c^P. Assuming that the policymaker delegates, it is then straightforward to show that this is socially optimal if

$$2\Omega^D < 2\Omega - \frac{\alpha}{2}(e_1^P)^2 - \frac{\alpha}{2}(e_2^P)^2. \tag{20}$$

This condition implies that delegation will always be socially optimal as long as the "rent" the politician obtains by delegating power is smaller than the net rent he extracts when he performs the public task himself. We consider this a sensible condition to impose. The case when delegation occurs because the politician can extract higher private rents from that activity is theoretically uninteresting because, in that case, one could always engineer delegation by making Ω^D arbitrarily high. It is also implausible that a politician would receive higher *private* rents from the public office by giving it to someone else. Indeed, imposing condition (20) makes our theory more interesting—and plausible—because it implies that delegation can be optimal for the politician *even if* it reduces his net private rents.

Proposition 2: *Assume condition (20) holds. When delegation occurs, it is socially optimal.*

Even if delegation is always socially optimal under condition (20), it may well be the case that the politician does not wish to delegate. In other words, we may have $U_c^D > U_c^P$ but $U^D < U^P$. This situation can arise because the politician may receive high rents from the policy task that more than compensate for the (social) gains from delegating. In this case, it would be optimal for society to "force" the politician to delegate by constitutional mandate.

5 Blinder's hypothesis: Insulation

We now turn to another key rationale for delegating policy to a technocrat: the fact that a technocrat is insulated from the whims and pressure of the citizenry while politicians may become "slaves of public opinion." It is indeed well known that politicians have a hard time contradicting or confronting the electorate prior to elections. To be more precise, they have an incentive to ignore the superior private information they have about the policy choice and instead base policy on public opinion. This is always sub-optimal in our model because citizens all share the same preferences and the politician is not maximizing these preferences by following public opinion but instead sacrificing social welfare in order to be reelected.

The setup is the same as in Section 3.1 except for the following changes. We now consider the consequences of the public forecast (ρ_t, $t = 1, 2$) of the shock, which becomes available at the beginning of every period before the decision maker chooses his effort level and observes his

private signal. As discussed in Section 3, the public forecast is formed in accordance with the same forecasting technology as the officeholders forecast, except that, since only the politician is compensated for supplying effort in producing the forecast, citizens' forecasts are based only on average expected ability (i.e., citizens' effort is zero). Note that, in our setup, the realization of the public signal has no effect on the forecasting technology of the officeholder because it already takes all relevant information into account. If the politician follows public opinion, his actions will reveal nothing about his type. We will see that, in equilibrium, he will always be reelected since he has an incumbency advantage. This assumption captures the idea that a politician can always credibly claim that a shock turned out to have been exceptional and that he cannot be blamed for it since the whole population agreed about the forecast.

When the politician can follow public opinion, three equilibria are possible: 1) the politician delegates power to a technocrat (who is, by design, insulated from public opinion); 2) the politician does not delegate but uses his (superior, in expectations) private information to make policy choices; and 3) the politician does not delegate but mimics public opinion (public opinion equilibrium). The first and the second equilibria are exactly the same as we analyzed in the last section, since in these cases the public signal is ignored and the formal analysis is relatively straightforward. We therefore turn directly to the public opinion equilibrium.[13]

In case the politician decides to follow public opinion in period one and to base his decision on the public signal, his maximization problem becomes[14]

$$\max_{e_1} U_1^{PO} = \left\{ 2\Omega - \frac{1}{2}\alpha e_1^2 - \frac{1}{2}\alpha e_2^2 \right\} - \left\{ \frac{1}{2} - \frac{1}{2}\theta_G \right\} - \left\{ \frac{1}{2} - \frac{1}{2}\theta_G - \beta e_1 - e_2 \right\}. \tag{21}$$

Throughout, the "PO" superscript refers to the politician following "public opinion." The first bracket captures, again, the politician's net rents from office. The second bracket is the probability of making the wrong policy decision in period one. When the politician follows public opinion, this probability does not depend on his effort in period one since he will not base his decision on his private signal in that period. The last bracket is the probability of making the wrong policy decision in period two. Note that the probability of making a mistake in period two is also a function of effort in period one because of the learning-by-doing effect. The above problem leads to the following equilibrium effort level:

$$e_1^{PO} = \frac{\beta}{\alpha} > 0, \tag{22}$$

which, depending on whether $\beta > 1$ or not, is higher (or not) than $e^* = \alpha^{-1}$, the period two effort level. Due to the learning-by-doing effect and a positive effort in period one, it is easy to verify that the politician will always be reelected. The reason is that, although his first-period action

[13] It is possible that the politician may choose to go against public opinion even if his own signal is also consistent with public opinion. This type of "anti-herding" or "fake leadership" behavior, as analyzed by Levy (2004) and Canes-Wrone et al. (2001), does not arise in our model, because it yields no benefits for period two.

[14] We assume parameters such that the politician would never want to resign voluntarily and would never want to follow public opinion in period two. It can be verified that we can always assume parameters for the rent and β so that this is the case.

does not reveal any information to the public (because it is conditioned on the public signal), since the politician puts in (strictly) positive effort in period one, he is expected to do better than anyone who runs for office against him (an incumbency advantage).

Proposition 3: *If the politician follows public opinion, he is always reelected.*

Using this proposition, we can then express the utility of the politician if he follows public opinion. It is obtained after substituting the equilibrium effort values (6) and (22) into (21), i.e.,

$$U^{PO} = 2\Omega - 1 + \theta_G + \frac{1+\beta^2}{2\alpha}. \tag{23}$$

When the politician follows the whims of public opinion, utility takes exactly the same form as welfare under delegation, i.e., the line has a slope 1 in (U, θ_G) space. To investigate whether a politician will find it individually rational to follow public opinion and ignore his private information, we can perform the same analysis as when analyzing whether or not the politician delegates (Section 3.2). Drawing up a diagram in (U, θ_G) space, we get the same picture as shown in Figure 2, with the line for the technocrat replaced by the politician that is not insulated from public opinion. Thus the politician may or may not choose to follow public opinion, depending on parameter values.

It is easy to verify the proposition below following the steps shown in the previous section. In Proposition 4, we denote the equilibrium when the politician bases his decisions on his private information rather than public information by "P," and it is the same equilibrium as analyzed in last section where "P" stood for "politician."

Proposition 4: *The politician always follows public opinion if $U^{PO} > U^{P}$. Both equilibria exist, and which arises depends on the value of the parameters $(\theta_G, \Omega, \alpha, \beta)$. For the policy task, the following holds at the margin:*

i) An increase in the complexity of the task (β) reduces the incentive to follow public opinion.

ii) An increase in the rent from office (Ω) increases the incentive to follow public opinion.

iii) An increase in the uncertainty of the officeholder's ability ($\theta_G - \theta_B$) reduces the incentive to follow public opinion.

We now investigate whether a politician who always follows public opinion would ever choose to delegate policy to a technocrat so as to insulate policies from public perceptions.

Consider a range of parameters $(\theta_G, \Omega, \alpha, \beta)$ such that the politician prefers to follow public opinion (i.e., $U^D < U^{PO}$). Note that this inequality does not depend on the value of θ_G. This is because the politician will never be fired. The key trade-off in this case comes through variations in the parameter α. The reason is that, for low values of α, the technocrat invests more effort into policymaking and thus has superior information about the policy task compared to the public. The politician who follows public opinion, however, invests little effort since he does not take advantage of his private information. This means that the cost of following public opinion is higher the lower is α. Figure 4 highlights the basic trade-offs facing a politician who follows public opinion. His welfare increases more slowly in α^{-1}—which is a measure of his private information—if he retains the policy task than if he delegates it to the technocrat. This can be seen by taking a partial derivative of U^D with respect to α^{-1}, i.e.,

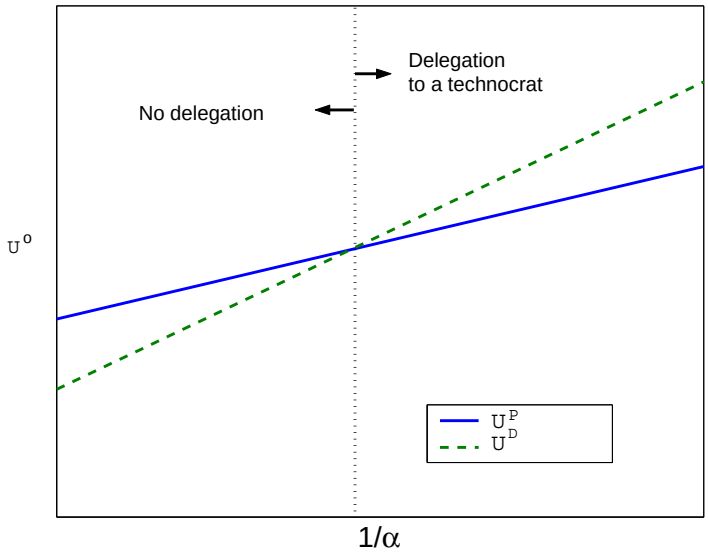

Figure 4: Basic insulation/mimicking trade-off

$$\frac{\partial U^D}{\partial \alpha^{-1}} = 1 + \beta^2 + 1 + 2\beta > \frac{\partial U^{PO}}{\partial \alpha^{-1}} = \frac{1+\beta^2}{2}.$$

Furthermore, the higher the value of β, the higher the incentive of a politician who follows public opinion to delegate to an insulated technocrat. To see this, note that

$$\frac{\partial U^D}{\partial \beta} = e_1^T = \frac{1+\beta}{\alpha} > \frac{\partial U^{PO}}{\partial \beta} = e_1^{PO} = \frac{\beta}{\alpha}.$$

The intuition for why a higher β increases the incentive for delegation is straightforward. Since the insulated technocrat supplies a higher effort in period one than the politician (because he uses this effort to retrieve a private signal for policymaking), β, the learning-by-doing parameter, has a stronger impact on utility under delegation.

Interestingly, we can interpret both these parameters (α and β) as the degree of complexity of a task. Presumably, the more complex the task, the more important the learning effect. One may also argue that the more important private information is for a given task—namely, information that is available only to the technocrat—the more complex and "specialized" that task is. Thus these two effects reinforce each other: The more complex the task, the more important it is to delegate it.

Proposition 5: *The more complex the policy task (as measured by an increase in β and α^{-1}), the more likely a politician enslaved to public opinion is to delegate it, insulating the policymaker from the whims of public opinion.*

One implication of our analysis is that if an equilibrium takes place in which the politician ignores his own information in favor of public opinion, it is always socially optimal to delegate to an insulated technocrat. To see this, note that when the politician follows public opinion he does not reveal anything about his type, which indicates that there is no advantage of elections.

In contrast, there is always a benefit to delegation since the technocrat is insulated from public opinion and therefore has an incentive to use his (superior) private information rather than the public signal. It is not guaranteed, however, that a politician who would follow public opinion would want to delegate even if it were socially optimal to do so. This depends on the private rents he extracts from office and the other parameters of the model.

Proposition 6: *If the politician ignores his private information and follows public opinion in equilibrium, it is always socially optimal to delegate power to an insulated technocrat. However, it may not be in the interest of a politician to delegate.*

Another implication of the analysis is that the public signal makes it less likely that a politician will choose to delegate. This should be obvious because the public signal can only make a politician who chooses to retain power better off (since he is always free to ignore the signal). For a public signal to increase the incentive to delegate, we would need to introduce some altruistic motives for the politician (e.g., a legacy motive) or a belief that delegation today can commit future politicians to leave the policy task in the hands of technocrats. These extensions seem relatively straightforward, and we leave them to future research.

6 Example: Central Bank Independence

In this section, we embed our simple political economy model into a small but standard macro model. In the last sections we have shown two key motivations for delegation to a bureaucrat with a long employment contract that is independent of the elected politician. This example illustrates that the theory provides natural new foundation for central bank independence.

The standard theory explains central bank independence as a mechanism to eliminate an "inflation bias". A discretionary policymaker that maximizes social welfare has an incentive to exploit the short-run trade-off between inflation and output to increase output through unexpected inflation. In equilibrium, the public anticipates this and there is no gain in output and only excess inflation – an "inflation bias". The solution to this, according to the literature (see, e.g., Rogoff (1985)), is to appoint a "conservative" and "independent" central banker that does not have an objective function that gives rise to an inflation bias, or endow an independent central bank with a "contract" that eliminates this bias (Walsh (1995)). Here we abstract from the inflation bias to emphasize a new motivation for central bank independence. Our theory therefore complements the standard theory, rather than contradicting it. In our example, the shocks and the policy objective we consider create no inherent trade-off between inflation and output, and hence no inflation bias. The superiority of the independent central banker only comes from his ability to be in a better position to judge the economy, due to a longer-term perspective (as suggested by Hamilton), and through his stronger incentive to withstand short-term electoral pressures, i.e. refrain from "pandering to public opinion" (as suggested by Blinder).

Consider the standard "Phelps problem," whereby the government minimizes deviation of inflation and the output gap from zero subject to an expectation-augmented Phillips curve, as in Kydland and Prescott (1977) and Barro and Gordon (1983). Unlike these authors, we do

not assume any inflation bias. The period utility function of a (representative) citizen is $v_t^c = -\frac{1}{2}\pi_t^2 - \frac{1}{2}\lambda x_t^2$, where π_t is inflation and x_t is the output gap. The supply side is $\pi_t = \kappa x_t + \pi_t^e$, where π_t^e denotes expectations of inflation that are formed before monetary policy is set in each period. The output gap is related to monetary policy by $x_t = -\varepsilon(i_t - \bar{r}_t)$, where i_t is the nominal interest rate and $\bar{r}_t$ is the natural rate of interest, which is exogenous. We assume that the central bank chooses i_t in each period before the shocks are realized so that the policy problem is exactly the same as discussed in Section 3 (but here the loss function is *endogenous* since it is a function of π_t and x_t).

Consider first the exogenous shock $\bar{r}_t$. As before, we assume that $\bar{r}_t$ is equal to r^H and r^L with equal probability. If the central bank could perfectly forecast the natural rate of interest, it would set $i_t = \bar{r}_t$, resulting in zero inflation and a zero output gap. This minimizes the bank's loss function, and hence there is no trade-off between inflation and output. If the central bank cannot perfectly forecast the natural rate of output, this equilibrium may not be feasible. Since the central bank sets the nominal interest rate before observing $\bar{r}_t$, its problem is to predict the future value of $\bar{r}_t$ in order to minimize its losses. For now we only consider the case in which there is no public information about the shock, and hence abstract from the insulation effect.

In the last (second) period, the maximization problem of the officeholder is to $\max_{\{e,i,\pi\}} E\left[v^o\right]$. We analyze optimal policy under discretion. The government treats expectations as constants and this reduces the problem to a one-period model. Substituting the constraints into the objective of the officeholder, we can write the maximization problem as

$$\max_{i,e,\pi} \sum_{s\in\{HR,LR,HW,LW\}} p(s)\{-(\frac{1}{2}(\pi^s)^2 + \frac{1}{2}\lambda(x^s)^2) + \Omega - \frac{\alpha}{2}(e)^2\} \tag{24}$$

$$\sum_{\sigma=H,L}\sum_{r=H,L} p(s(r,\sigma))\{\frac{1}{2}(-\varepsilon\kappa(i^\sigma - \bar{r}^s) + \pi^e)^2 + \frac{1}{2}\lambda\left(-\varepsilon(i^\sigma - \bar{r}^s)\right)^2 + \Omega - \frac{\alpha}{2}(e)^2 \tag{25}$$

and as in Section 3 the superscript $s \in \{HR, LR, HW, LW\}$ refers to the combination of the shock $r \in \{H, L\}$ and the forecast $\sigma \in \{H, L\}$, i.e., whether σ turns out to be a correct (R) forecast about r or not (W). Hence we can write s as a function $s(r, \sigma)$. Note that the quality of the signal σ is again determined by the ability (θ_G) and effort (e) of the policymaker and we assume the same forecasting technology as before. In the second line we have used the structural equations to substitute out for π^s and x^s. Taking first-order conditions with respect to e, i^H, i^L, π_{HR}, π_{HW}, π_{LR}, and π_{LW}, we can derive equilibrium values of i^H, i^L, π_{HR}, π_{HW}, π_{LR}, and π_{LW}, shown in the appendix (note that to derive these we will need to substitute for the conditional probabilities of being right or wrong given by $p(s(r, \sigma))$ as function of the effort and ability as in sections 3). To find the value for e, we again obtain condition (6); we can then substitute the endogenous values of π_s and x_s into the loss functions and obtain the solution $e_2^s = \frac{\phi}{\alpha-\phi}E_2^s\theta$ where we assume that the cost function is $c(e) = \frac{1}{2}\alpha e^2$ and define the following coefficient $\phi = 4(1 + \frac{\lambda}{\kappa^2})\varepsilon^2 k^2 \left(r^H - r^L\right)^2$.[15]

[15] Note that we need $\alpha > \phi$ for an equilibrium with positive effort to exist. The effort choice depends on the realization of s_1 in period one since this has an effect on the expected ability of the officeholder. This implies that,

The first-period allocation, as far as the macroeconomic variables are concerned (i.e., i^H, i^L, π_s, x_s), will be the same as those derived in the appendix for period two except that we replace $E_2[\theta]$ with $E_1[\theta]$ and, in this case, e_1 does not appear in the expression for $E_1[\theta]$. The derivation of the first-period equilibrium effort level follows directly from the analysis in the previous sections with the added analytical complication that the loss function is endogenous. Note that the derivations (26)-(32) imply that, for given ability, a higher level of effort reduces output and inflation variability. Since delegation implies higher effort, this implies that if a politician delegates power, then central bank independence leads to a reduction in both output and inflation variability, overcoming what has sometimes been considered a weakness of the earlier literature. Hence this simple example establishes that our theory can account for central bank independence. Furthermore, the theory predicts that an independent central bank is associated with lower inflation *and* output variability. In the discussion above, we have not included the public signal which gives rise to the insulation effect. That analysis follows in exactly the same fashion as before. Another extension is to add cost push shock u_t to the Phillips curve. In this case, there is a trade-off between inflation and output. Exactly the same analysis applies, however.

7 Conclusion

In this paper, we have studied the costs and benefits of delegating public policy to an independent agency free from election pressures. To get at the heart of the issue—the role of election pressures—we considered the most simple framework possible: a two-period model with the possibility of elections at the end of period one. While this framework is extremely simple, we hope that it gets at some basic issues likely to remain important in a more general setting. In particular, the dynamic incentive structure that arises because of elections enables us to formalize both Hamilton's (1788) and Blinder's (1998) rationales for delegation. One limitation of our model is that it has only two "periods"; however, we conjecture that our results would extend to a more general setting. An extension of this kind could answer several questions quite unrelated to the issue of optimal delegation of political power, because it would allow us to analyze the *optimal degree of election pressures* over time for a given policymaker. Our basic framework suggests that the role of learning on the job, resistance to short-term political pressure, and the danger of ending up with incompetent policymakers would all be important to understanding the optimal task allocation. Our hope is that the present framework will be helpful in addressing these important questions in future work.

in the macroeconomic setup, the effort choice can be higher under delegation in both period one and two, whereas in our previous section the effort choice was the same in period two across the two regimes.

A Appendix

A.1 Proof that $\alpha - 2\beta > 0$

Since the probability has to be between zero and one, we have $\frac{1}{2} + \frac{1}{2}\theta_G + e_1 + \beta e_2 \leq 1$ and, using the equilibrium values for e_1 and e_2 under delegation, we have $\frac{1}{2} + \frac{1+\beta}{\alpha} + \frac{\beta}{\alpha} \leq 1$; so that $\frac{1+2\beta}{\alpha} \leq \frac{1}{2}$ and $2(1+2\beta) \leq \alpha$, which gives $\alpha - 2\beta \geq 2(1+\beta) > 0$.

A.2 Proof that the coefficient γ_3 of $U^D = \gamma_1 + \gamma_2\theta_G + \gamma_3\theta_G^2$ is strictly positive

The coefficient γ_3 is strictly positive if

$$\frac{\beta^2}{\alpha - 2\beta} + \frac{\beta^3}{(\alpha - 2\beta)^2} > \frac{1}{2}\alpha\frac{\beta^2}{(\alpha - 2\beta)^2},$$

i.e., if $(\alpha - 2\beta)\beta^2 + \beta^3 > \frac{1}{2}\alpha\beta^2 \iff \alpha\beta^2 - \beta^3 > \frac{1}{2}\alpha\beta^2 \iff \alpha - \beta > \frac{1}{2}\alpha \iff \frac{1}{2}\alpha - \beta > 0 \iff \alpha - 2\beta > 0$, which we know is strictly positive from (**??**).□

A.3 Equilibrium values in example

After rearranging the first-order conditions (w.r.t. e, i^H, i^L, π_{HR}, π_{HW}, π_{LR}, π_{LW}), the equilibrium values are:

$$i^H = \left(\frac{1}{2} + E_2\left[\theta\right]e\right)\bar{r}^H + \left(\frac{1}{2} - E_2\left[\theta\right] - e\right)\bar{r}^L \tag{26}$$

$$i^L = \left(\frac{1}{2} - E_2\left[\theta\right]e\right)\bar{r}^H + \left(\frac{1}{2} + E_2\left[\theta\right] + e\right)\bar{r}^L \tag{27}$$

$$\pi_{HR} = \varepsilon k\left(\frac{1}{2} - E_2\left[\theta\right] - e\right)\left(\bar{r}^H - \bar{r}^L\right) \tag{28}$$

$$\pi_{HW} = -\varepsilon k\left(\frac{1}{2} + E_2\left[\theta\right] + e\right)\left(\bar{r}^H - \bar{r}^L\right) \tag{29}$$

$$\pi_{LR} = -\varepsilon k\left(\frac{1}{2} - E_2\left[\theta\right] - e\right)\left(\bar{r}^H - \bar{r}^L\right) \tag{30}$$

$$\pi_{LW} = \varepsilon k\left(\frac{1}{2} + E_2\left[\theta\right] + e\right)\left(\bar{r}^H - \bar{r}^L\right) \tag{31}$$

$$x_s = k^{-1}\pi_s; \quad s \in \{HR, LR, HW, LW\} \tag{32}$$

where $E_2[\theta] = q\theta_G + (1-q)\,\theta_B + e_1$. Note that we define the notation $E_2\theta$ as expected ability, which includes the effort in the previous period. However, as in Section 3, under democracy, if the incumbent politician is not reelected, the e_1 term drops out. With an endogenous loss function, all the equilibrium values are now functions of the expected ability of the officeholder, i.e., $e = e^*\left(q\right) > 0$, $i^H = i^{H*}\left(q\right)$, $i^L = i^{L*}\left(q\right)$, $\pi_s = \pi_s^*\left(q\right)$, and $x_s = x_s^*\left(q\right)$.

B References

Alesina, A., and G. Tabellini, 2007, "Bureaucrats or Politicians? Part I: A Single Policy Task," *American Economic Review*, Vol. 97, pp. 169-79.

Arrow, K., 1962, "The Economic Implications of Learning by Doing," *Review of Economic Studies*, Vol. 29, pp. 155-73.

Barro, R.J., and D.B. Gordon, 1983, "A Positive Theory of Monetary Policy in a Natural Rate Model," *Journal of Political Economy*, Vol. 91, pp. 589-610.

Besley, T., and S. Coate, 2003, "Elected Versus Appointed Regulators: Theory and Evidence," *Journal of the European Economics Association*, Vol. 1, pp. 1176-1206.

Blinder, A. 1998, "Central Banking in Theory and Practice," MIT Press, Cambridge.

Canes-Wrone, B., Herron, M., and K.W. Shotts, 2001, "Leadership and Pandering: A Theory of Executive Policymaking," *American Journal of Political Science*, Vol. 45, pp. 532-50.

Drazen, Allan, 2002, "'Laying Low' During Elections: Political Pressure and Monetary Accommodation," mimeo, University of Maryland.

———, 2000, *Political Economy in Macroeconomics*, Princeton University Press, Princeton, NJ.

Elster, J., 1995, "The Impact of Constitutions on Economic Performance," *Proceedings of the Annual World Bank Conference on Economic Development*, Washington: The World Bank, pp. 209-226.

Hamilton, A. 1788, "The Duration of Office in the Executive", Federalist Papers No. 71.

Holmström, B., 1999, "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies*, Vol. 66, pp. 169-82. Originally, 1982, in *Essays in Economics and Management in Honor of Lars Wahlbeck* (Stockholm: Swedish School of Economics).

Kydland, F.E., and E.C. Prescott, 1977, "Rules Rather than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy,*.Vol. 85, pp. 473-91.

Le Borgne, E., and B. Lockwood, 2006, "Do Elections Always Motivate Incumbents? Experimentation vs. Career Concerns," *Public Choice*, Vol. 129, pp. 41-60.

Levy, G., 2004, "Anti-herding and Strategic Consultation," *European Economic Review*, Vol. 48, pp. 503-25.

Maskin, E., and J. Tirole, 2004, "The Politician and the Judge: Accountability in Government," *American Economic Review*, Vol. 94, pp. 1034-54.

Persson, T., and G. Tabellini, 2000, *Political Economics: Explaining Economic Policy*, MIT Press, Cambridge, MA.

Rogoff, K., 1985, "The Optimal Degree of Commitment to an Intermediate Monetary Target," *Quarterly Journal of Economics*, Vol. 100, pp. 1169-89.

———, and A. Sibert, 1988, "Elections and Macroeconomic Policy Cycles," *Review of Economic Studies*, Vol. 55, pp. 1-16.

Sharfstein, D. and J.C. Stein„ 1990, "Herd Behavior and Investment," *American Economic Review*, Vol. 80, pp. 465-79

Walsh, C.E., 1995, "Optimal Contracts for Central Bankers," *American Economic Review*, Vol. 89, pp. 150-67