# Try to See it My Way:

# The Process and Perils of "Coming Around"

Dan Alexander
University of Rochester
dan.alexander@rochester.edu

Darrian Stacy
Vanderbilt University
darrian.k.stacy@vanderbilt.edu

June 18, 2020

PRELIMINARY

**Abstract**

Can a career civil servant (bureaucrat) achieve long-term policy congruence via strategic reporting to a political appointee who oversees her agency and sets policy? This paper develops a model that explores when bureaucrats will misreport information relevant to an appointee's short-term decision making in order to influence the appointee's belief formation and, thereby, long-term policy setting. We find that when a bureaucrat's short-term and long-term policy goals do not conflict too greatly, she will prefer to provide a truthful report to the appointee – who, in turn, is willing to accept the report as truthful (particularly as it pertains to short-term policy responses). When a bureaucrat's long-term and short-term goals are in conflict, however, she has an incentive to engage in deception in which she reports misinformation to her political appointee in order to influence belief formation. Favorable conditions for this deception to take place include divergent policy preferences between the bureaucrat and the appointee, low likelihood of short-term decisions being necessary, and bureaucratic desire for long-term policy congruence from the appointee.

# 1 Introduction

Within a political system with constant changes in executive oversight, career civil servants (bureaucrats) serve an essential role in the continuity and completion of short-term and long-term goals for the agencies in which they work. Conversely, political appointees supervise the work of bureaucrats, and executives look to their appointees both to manage a given agency or department as well as to promote the administration's goals through the agency's decision-making processes. The inherent differences in the function and position of these two broadly defined roles yields a tension that has been the subject of scholarly study for at least a century.

An abundance of anecdotal evidence depicts political appointees "coming around," which occurs when appointees assimilate to the preferences of and becoming more congruent with the agencies they lead, by the end of their tenure.[1] Yet few studies have examined the direct impact of bureaucrats on the belief formation of political appointees. Our study provides insight into how this process operates, highlighting one possible means of bureaucratic control over political appointees – strategic misreporting of information. Our work demonstrates, though, that if strategic misreporting is employed by bureaucrats, it is constrained by the bureaucrat's desire to maintain some level of trust with the appointee and is not an option to be relied upon in perpetuity.

Specifically, we develop a model to explore when bureaucrats will strategically misrepresent the information they provide to political appointees to affect the belief formation of appointees and, thereby, long-term policy outcomes. The bureaucrat learns the true state at the start of the game and reports on the state, truthfully or not, to the appointee. With this report in-hand, the appointee must make long-term policy and concurrently respond to short-term incidents. The bureaucrat alone knows the true likelihood of each state, and she seeks long-term policy that reflect this likelihood, though the bureaucrat prefers that the appointee act on the true state when setting short-term policy. The appointee has her own beliefs about the likelihood of each state, but she is willing (in a sense to be made precise below) to base her decision-making on the bureaucrat's report to some extent.

When the bureaucrat and the appointee do not disagree too much about the underlying likelihood of each state, and when it is more likely and/or more important that the appointee will need to issue a short-term policy response, the bureaucrat will offer truthful reports which will be accepted by the

[1]Historically, the appointee was said (in rather outdated, colonial terms) to have "gone native." Rather than relying on the appointee to exert influence on the bureaucracy of her behalf, the executive ultimately fears the appointee is exerting the bureaucracy's influence on her. Our first contribution is to suggest alternative terminology for this phenomenon, namely, "coming around." In addition to providing a much-needed update of terminology, it also suggests the process by which the appointee's beliefs change likely shares properties with the processes by which other political actors come around to new perspectives and preferences.

appointee. When there is greater conflict between the views of the bureaucrat and the appointee, however, or when it is less likely and/or less important that a short-term response will be necessary, the bureaucrat may engage in deception by reporting misinformation to her political appointee in order to influence belief formation so that the political appointee "comes around" to the bureaucrat's position more quickly and creates long-term policy accordingly. In cases of extreme disagreement, the bureaucrat issues a single report regardless of the state and the appointee is essentially unresponsive to the bureaucrat's reports. When the bureaucrat's deception occurs only probabilistically, however, the appointee still updates in the direction of the bureaucrat's position. Indeed, over time, the bureaucrat will find herself truthfully reporting to the political appointee because there is less potential gain in misreporting and the appointee's beliefs gradually shift closer to the bureaucrat's beliefs.

In the next section, we place our contributions within a number of related, extant literatures. We then lay out the basics of the model and proceed to analyze the stage game. While we do not yet provide a full analysis of the repeated game, we are nonetheless able to speak to elements of the long-term interaction between the players, including those aspects of stage-game behavior that we expect to persist under dynamic considerations.

## 2 Extant Literature

Motivated by the initial tension between career civil servants and new appointees, much of the prior literature on this relationship has emphasized the principal-agent problem facing executives seeking to use the bureaucracy to further their own ends (Gailmard & Patty 2013). In this paper, we take a different approach to studying the interaction between political appointees and career civil servants. Our work contributes to the literature on the bureaucracy – as well as, more broadly, the literature on how decision makers update beliefs – by illustrating a specific way in which bureaucrats potentially exercise influence over their political appointees.

Gailmard & Patty (2012) lay out a "credibility continuum," in which bureaucrats as agents may provide messages that range from cheap talk to costly signals to fully verifiable (but with the option to issue a report of not). In our setting, the bureaucrat's signals are probabilistically verified, putting this paper in a camp closest to the fully verifiable setting. Unlike much of the existing work, however, the motivation to misreport the state stems not from persistent differences in preferences but rather from the possibility of manipulating appointee beliefs and thus changing appointee preferences.

In light of this mechanism, the work most similar to ours is (Hirsch 2016), which is one of the few papers to the authors' knowledge that also employs uncommon priors among the key actors. Hirsch (2016) models a situation in which the appointee seeks to affect the beliefs of the agent, the opposite of our setting. More importantly, the belief manipulation in his paper arises from experimentation (Bayesian persuasion), while in the model we present the appointee does not require hard evidence to update and is instead willing – conditioned on her own beliefs and equilibrium strategies – to update based on soft information. Certain aspects of the appointee's behavior constitute behavioral assumptions. Reflecting on scholarly accounts of appointee-bureaucrat relations, however, this appears to be the truest way to capture the limited trust offered by a new appointee to civil servants and the advice they proffer.

Our particular approach is informed by two rather distinct literatures. The first is the literature on the public administration of the bureaucracy that examines the interaction between political appointees and bureaucrats. Heclo (1977) provides one of the most extensive treatments of the topic within the American bureaucracy. Drawing on interviews of political appointees and career civil servants, Heclo argues that while appointees often mistrust the bureaucracy when first appointed, they must balance their authority and trust in careerists to be successful in accomplishing their goals. Political appointees are typically at an informational disadvantage when they are appointed to leadership roles within an agency, given their lack of substantive expertise and familiarity with issues that arise within bureaucratic politics. Appointees can help themselves in pursuit of their goals through creating and reinforcing relationships of trust and confidence with their superiors and interest groups outside of an agency. However, Heclo argues that generally this "self-help" is not enough; in order for appointees to accomplish their goals and be effective leaders they must elicit the cooperation of their subordinate civil servants to provide "knowledgeable continuity to warn and propose, and institutionalized responsiveness to help carry out executive decisions" (Heclo 1977, p. 170-171). Career civil servants can potentially offer a number of services–which include, orienting new appointees, providing technical information, forewarning about potential conflicts or opposition, and offering and implementing new policies–to aid political appointees. But, in order for appointees to make effective use of these services, there must be at least a minimum level of appointee-bureaucrat trust.

While Heclo (1977) describes a potential motivation for appointees to maintain a positive working relationship with the bureaucrats that inform them, studies that have followed hence have presented evidence that bureaucrats may seek the same. The primary responsibility of many bureaucrats within the bureaucracy is to explain the merits of policies to their appointee supervisors (Dolan 2000). More-over, bureaucrats desire more influence over policy outcomes than they perceive themselves as having already, but they view appointees as having the most influence (Stehr 1997). Thus, whatever control bureaucrats have over policy, it is often through the channel of educating and informing their inexperienced appointees, who then makes the decisions about policy. We can assume, as some studies have shown, that bureaucrats have an interest in maintaining the trust between themselves and their appointees (Chen & Hsieh 2015, Gailmard, Patty & Gailmard 2007). Yet, other studies acknowledge that bureaucrats also often hold a public service motivation (Berman, Chen, Jan & Huang 2013) that lasts over time and across regimes. Our paper considers the implications of these competing motivations for the quality of information a bureaucrat provides and the extent to which a bureaucrat can influence the belief formation of an appointee.

The question of how much control over policy bureaucrats exercise via their influence in the belief formation of appointees has gone largely unexplored by scholars. The second literature upon which our study draws offers a potential reason for this omission. Specifically, numerous studies demonstrate how difficult it is for one individual to convince another individual to update their beliefs. The "expert problem" is typically framed as some decision maker faced with a problem of how to rationally update their beliefs about the likelihood of some event of interest occurring in light of an expert's opinion (Morris 1974, French 1980, Genest & Schervish 1985). Bayesian decisionmakers can update their prior beliefs after the expert's opinion is revealed by (1) assessing their own prior beliefs and (2) specifying a likelihood function that can be thought of as the decision maker's subjective measure of how credible the expert is (Morris 1974).

In general, a decision maker will update her beliefs in light of an expert's opinion depending on (1) how often the decision maker and expert are in agreement about the likelihood an event is to occur and (2) how far an expert's belief in the likelihood of an event is from the decision maker's prior expectation of the average response from the expert (Morris 1974). If decision makers trusts the expert, they feel confident that the expert's opinion represents a narrow distribution of the likelihood of an event occurring given the extant information, and their prior belief is updated considerably. Conversely, if the decision makers does not trust the expert, they are not confident that the expert's opinion is informative because it is likely invariant regardless of the extant information. Intuitively, in this case of a lack of trust, the updated beliefs of the decision maker will more closely resemble their own prior beliefs.

We build upon these insights to consider not only the belief formation of a particular type of decision maker (a political appointee) but also the strategic behavior of a particular type of expert (a bureaucrat). In what follows, we explore the conditions under which a bureaucrat would engage in deception to influence the belief formation of a political appointee, and we also detail how a bureaucrat would update her beliefs given the bureaucrat's propensity to provide a misleading report. Lastly, we show that the strategic manipulation of an appointee's belief formation by a bureaucrat is possible even with the knowledge of the bureaucrat's incentives to misreport. In the next section, we outline the preliminary features of our model.

# 3 A Model of Appointee-Bureaucrat Interaction

## 3.1 Preliminaries

We consider a political appointee, $A$, and a bureaucrat (career civil servant), $B$. At the start of the game, the bureaucrat learns the true state of the world, $\omega \in \{0, 1\}$ and then issues a report on the state of the world, $r \in \{0, 1\}$, which is not constrained to be $\omega$. Let $q : \{0, 1\} \longrightarrow [0, 1]$ be the probability that $B$ accurately reports the state of the world, i.e., $q(\omega) = \Pr(r = \omega)$.

The appointee has two tasks: setting long-term policy and responding to immediate issues. Immediate issues require a short-term response, $s \in \{0, 1\}$, but only when they arise, which occurs with probability $\psi \in (0, 1)$. Long-term policy, $p \in [0, 1]$, must be set. Long-term policy and short-term responses (if required) are effectively chosen simultaneously, in that the appointee faces the same informational environment when making both choices. We discuss our intended interpretation of long- and short-term policies in the subsection below.

The bureaucrat knows that $\Pr(\omega = 1) = \hat{\pi}$, i.e., an independent Bernoulli trial determines the state of the world. The appointee neither learns the true state of the world at the start of the game nor does she know the distribution of $\omega$. Let $A$'s prior beliefs about the value of $\pi$ be distributed according to the beta distribution with parameters $\alpha, \beta$. Without loss of generality, let $\alpha < \beta$, such that $\frac{\alpha}{\alpha+\beta} < \frac{1}{2}$. Denote the p.d.f. of the appointee's updated beliefs after receiving $r$ by $\nu(\pi)$. Finally, let $A$'s belief that $\omega = 1$ given a report $r$ be given by $\mu(r) \in [0, 1]$. This information structure is common knowledge.

Both players incur quadratic loss in the lack of congruence between $A$'s short-term response and the state of the world and in the extent to which $A$'s long-term policy does not match (what each believes is) the likelihood of the state of the world being 1. The bureaucrat weighs long-term policy congruence by $\lambda \geq 0$.[2] Note that, because the appointee takes both actions simultaneously and independently, we need not consider the relative weight they may place on long-term as opposed to short-term policy responses. The utility functions are given by the rules:

$$u_A(s, p|\omega) = -\psi \cdot (\omega - s)^2 - \int_0^1 (\pi - p)^2 \nu(\pi) d\pi \tag{1}$$

$$u_B(s, p|\omega) = -\psi \cdot (\omega - s)^2 - \lambda(\hat{\pi} - p)^2 \tag{2}$$

The game is one of sequential moves and incomplete information, requiring in equilibrium sequential rationality and consistency of beliefs. While sequential rationality takes its usual form, the conditions that must be met for consistent beliefs require more thought in the presence of uncommon priors. A set of substantive assumption about the information structure dictates the way in which $A$ update their beliefs given their innate disagreement with $B$.

Specifically, suppose the bureaucrat's belief that the true value of the underlying parameter is $\hat{\pi}$ is known, the appointee's prior distribution ($\nu$) over the probability that the state of the world is $\pi$ is also known, and the appointee agrees to update their beliefs about the underlying probability that the state of the world is 1 with three stipulations. First, the appointee does so in as evidence-based a way as possible, i.e., based on hard information about the state of the world or, in the absence of such information, based on the report from the bureaucrat. Second, the appointee takes into account the bureaucrat's belief $\hat{\pi}$ as well as the bureaucrat's equilibrium strategy (i.e., $q(\omega)$) when updating. Third, the appointee uses their own priors over $\pi$ when updating.

## 3.2 Discussion of the Model's Features

Several of the model's assumptions deserve some additional grounding before proceeding to the analysis, chief among them the nature of and distinctions between the two actions that appointees take. Long-term policy is meant to connote structural policy. In addition to responding to the crisis of the day, a political appointee is also responsible for setting the direction of the agency. In this endeavor, they seek to equip the agency to adeptly respond to the full mix of challenges it is likely to face (in the model's terms, $\omega = 1$ and $\omega = 0$) and in the correct proportion (their perception of $\pi$).

With short-term responses, we intend to evoke the need for agencies to shift resources around quickly to best deal with evolving situations. In the event that such immediate issues arise, the

[2] As it becomes more catastrophic for $B$ to have provided faulty information when a short-term response was ultimately required, $\lambda \to 0$.

appropriateness of the response becomes common knowledge in a way that the appropriateness of long-term policy does not. Thus, the appointee's preferences around short-term responses depend on the actual state of the world, while their preferences around long-term policy depend on their (updated) perception of the distribution of $\pi$. The parameter capturing the probability that short-term responses become necessary, $\psi$, may just as well be taken to be a weighting parameter. While ultimately only the ratio of $\psi$ to the long-term weighting parameter is all that matters, it is instructive to maintain a reminder that $\psi$ may represent more than a ratio of weights, but also the probability that an agency faces a situation of some urgency and must make public and material use of the bureaucrat's response. For instance, certain agencies may have responsibilities that entail more frequent emergencies (e.g., the Department of Defense) while others may rarely face the need to appear at a press conference to address an urgent situation (e.g., the Department of the Interior).

The decision to employ uncommon priors was not undertaken lightly, yet we believe it best captures the different bodies of evidence to which the appointee and the bureaucrat have been privy (Aumann 1976). One might ask why the appointee should listen to the bureaucrat at all, and then only with the report? The appointee has no reason to update their genuinely held beliefs about the distribution of $\pi$ in response to the bureaucrat's assertion that the true value is $\hat{\pi}$ because the bureaucrat has risked nothing – and thus presented no new evidence – to offer an opinion the appointee. Because the appointee and bureaucrat find their payoffs linked once engaged in the sort of interaction modeled herein, the appointee may no longer dismiss the bureaucrat's statements as cheap talk. It is in this sense that uncommon priors may gradually be reconciled. This logic also explains Why simply allowing the bureaucrat to issue a short-term report along with a long-term report would be inconsequential; only the short-term report would carry any information for the appointee, as it is the only one in which the bureaucrat may have risked their payoffs to issue.

The partial reconciliation of priors takes the form of convergence of beliefs about $\pi$. Because payoffs from long-term policy are based on each player's beliefs about $\pi$, there is a sense in which the appointee's preferences are converging to those of the bureaucrat. The bureaucrat's and appointee's disagreement stems not from different preferences, *per se*. In this model, both players want short-term responses to match the true state of the world and long-term policy to reflect the underlying data-generating process for the state of the world. Their disagreement is a product of the different information and thus different beliefs they bring to bear in determining what the long-term policy ought to be.

The interaction in the model captures the idea that the appointee learns about the best long-term policy through reports that also serve to best prepare for the possibility of an urgent, short-term response. Yet there is only a single round of such learning. One potential interpretation might be that the bureaucrat has an initial opportunity to shift appointee beliefs upon their arrival at the agency – a time at which appointees take a number of actions to set long-term agency direction. Of course, if the bureaucrat provides faulty information right off the bat, and it becomes known through the need for a short-term response, the bureaucrat has jeopardized their relationship with the appointee.

A broader, less literal interpretation, however, might be that the interaction between the appointee and the bureaucrat is periodic, recurring at regular intervals. Perhaps the bureaucrat gives the appointee a weekly report of the state of the world – reflecting important information should any of a number of possible situations escalate. While this suggests the game outlined above be repeated, with the beliefs produced in one round carrying over as the priors for the next round, we have no reason to suspect the cost in terms of parsimony would produce benefit in terms of insight that the one-shot game does not. Further, this allows us to avoid complications, such as the question of whether the bureaucrat would update their beliefs in response to a history that runs counter to expectations given by $\hat{\pi}$.

Finally, the process and perils of "coming around" are certainly not limited to political appointees and career civil servants. The model speaks to nearly any environment in which: i) group A has mistaken beliefs about the reality facing a group B, ii) group B self-reports messages about specific incidents – which are occasionally receive external verification – with the knowledge that group A's beliefs will update in response, and iii) group A has explicitly or implicitly institutionalized power over group B. The United States in particular has in the past few years undergone at least two major reckonings in which an almost entirely white/male/straight/cis establishment has encountered impossible-to-deny evidence, preceded and accompanied by an avalanche of self-reports, leading it to update its beliefs about the reality facing non-white/female/LGBTQ+ groups. What accounts for the persistence of different beliefs? How did the beliefs begin to converge at long last? What characterized trust of self-reports historically and once convergence began to occur? Without belaboring the point, we turn to analyze the model, which we believe speaks to exactly these questions.

## 4 Analysis of the Model's Equilibria

We employ the concept of weak perfect Bayesian equilibrium and begin the analysis with the appointee's decisions, given a report from the bureaucrat.

### 4.1 The Appointee's Problem

The analysis begins with the appointee taking two actions in response to a report from the bureaucrat, $r$, and specifically her updated beliefs about the state of the world and the underlying data-generating process. Recall the appointee's updated beliefs about the state of the world are given by $\mu(r) = \Pr(\omega = 1|r)$. The appointee's updated beliefs about the parameter $\pi$ (the unconditional probability that $\omega = 1$) remain within the Beta-distribution family.

The two results immediately below address the two elements of $A$'s equilibrium behavior following from these updated beliefs. Both results take as given the quantity $\mu(r)$. The characterization of $\mu(r)$ is not trivial, and it is precisely where the assumptions above about the way in which $A$ processes reports from $B$ and updates her beliefs become most relevant. This aspect of equilibrium receives greater attention below, though in the meantime it is reasonable to assume $\Pr(\omega = 1|r = 1) := \mu(1) \geq \mu(0) =: \Pr(\omega = 1|r = 0)$.

**Lemma 1.** *If $A$'s belief is that $\pi \sim Beta(\alpha, \beta)$, and $A$ receives information that leads her to believe that the state of the world is 1 with probability $\mu$, her updated belief is that $\pi \sim Beta(\alpha+\mu, \beta+(1-\mu))$.*

**Lemma 2.** *$A$'s optimal short-term response, given a report from $B$ of $r$, is $s^* = 1$ if $\mu(r) \geq 1/2$ and $s^* = 0$ if $\mu(r) \leq 1/2$. $A$'s optimal long-term policy is $p^* = \frac{\alpha+\mu(r)}{\alpha+\beta+1}$.*

The first result merely recapitulates the conjugacy of a beta-distributed prior over the probability parameter in a Bernoulli trial. The second result specifies that for short-term responses to immediate issues, $A$ will choose $s \in \{0, 1\}$ in accordance with the state she thinks most likely given the report she received. For long-term policy, however, she will minimize squared distance from her policy choice to the actual value of $\pi$ by choosing the mean of the distribution describing her updated beliefs about $\pi$.

### 4.2 The Bureaucrat's Decision to (Mis-)Report the State

Analyzing the bureaucrat's state-contingent decision to tell the truth or lie in her report, $r$, is informative of the set of potential equilibrium cases. In describing $B$'s decision problem, it is useful to note that, generically, $A$ will not mix over short-term responses in equilibrium, which would require $\mu = 1/2$. As such, $A$ will either choose $s = r$, or instead always choose $s = 1$ or $s = 0$ regardless of whether $r = 1$ or $r = 0$. In the first case, $A$ may be said to be responsive to $B$'s report. The appointee is unresponsive in the latter case.

Supposing momentarily that the appointee is responsive to the bureaucrat's reports, i.e., $\mu(1) > 1/2 > \mu(0)$, $B$ reports truthfully when the $\omega = 1$ if:

$$u_B(r=1|\omega=1) \geq u_B(r=0|\omega=1)$$
$$-\psi \cdot (1-\lambda)(1-1)^2 - \lambda \left(\hat{\pi} - \frac{\alpha+\mu(1)}{\alpha+\beta+1}\right)^2 \geq -\psi \cdot (0-1)^2 - \lambda \left(\hat{\pi} - \frac{\alpha+\mu(0)}{\alpha+\beta+1}\right)^2.$$

Combining terms, $B$ reports $r = 1$ when $\omega = 1$ when:

$$\frac{\psi}{\lambda} \geq \left(\hat{\pi} - \frac{\alpha+\mu(1)}{\alpha+\beta+1}\right)^2 - \left(\hat{\pi} - \frac{\alpha+\mu(0)}{\alpha+\beta+1}\right)^2. \tag{3}$$

The left-hand side of (3) increases as the probability of a short-term response increases or the weight $B$ places on the appropriateness of long-term policy decreases. The terms on the right-hand side represent the squared distance from the true value of $\pi$ to the posterior mean after a report of $r = 1$ and the squared distance from the true value of $\pi$ to the posterior mean after a report of $r = 0$. The whole right-hand side is positive if reporting 0 would bring the posterior mean (the appointee's optimal long-term policy, $p^*$) closer to $\hat{\pi}$ than reporting 1 would. Truth-telling would then be optimal only if the benefit of misreporting in terms of belief manipulation did not overwhelm short-term benefit of truth-telling captured by the left-hand side. Conversely, if reporting 1 would have a more desirable effect on the posterior mean, the right-hand side will be negative, (3) will hold regardless of the value of the left-hand side, and truth-telling will be optimal for $B$ when the state is $\omega = 1$.

The analogous expressions for $B$'s payoffs when $\omega = 0$ are as follows:

$$u_B(r=0|\omega=0) \geq u_B(r=1|\omega=0)$$
$$-\psi \cdot (1-\lambda)(1-1)^2 - \lambda \left(\hat{\pi} - \frac{\alpha+\mu(0)}{\alpha+\beta+1}\right)^2 \geq -\psi \cdot (0-1)^2 - \lambda \left(\hat{\pi} - \frac{\alpha+\mu(1)}{\alpha+\beta+1}\right)^2.$$

Rearranging, we see that the greater the disagreement between $B$ and $A$, the greater $B$'s motivation to misrepresent a state of 0 with a report of 1:

$$\frac{\psi}{\lambda} \geq \left(\hat{\pi} - \frac{\alpha + \mu(0)}{\alpha + \beta + 1}\right)^2 - \left(\hat{\pi} - \frac{\alpha + \mu(1)}{\alpha + \beta + 1}\right)^2. \tag{4}$$

Specifically the higher $\hat{\pi}$ is than the prior mean $\frac{\alpha}{\alpha+\beta}$, the better the conditions for $B$ to optimally report $r = 1$ when $\omega = 1$.

The condition in (4) is quite similar to that in (3), except for the sign of the expression on the right-hand side. Indeed, if $B$ misreports when the state is 0, she will certainly tell the truth when the state is 1. Conversely, if she finds it worthwhile to misreport when the state is 1, then she will truthfully report when the state is 0. In equilibrium, then, $B$ will lie in at most one state of the world.

If $B$ always reports truthfully in both states, then $\Pr(\omega = 1|r = 0) := \mu(0) = 0$ and $\Pr(\omega = 1|r = 1) := \mu(1) = 1$. If $\mu(1) < 1$ (resp., $\mu(0) > 0$), $B$ must in equilibrium be mixing over reports when $\omega = 0$ (resp., $\omega = 1$). Mixing requires indifference, and an examination of the conditions for truthful reporting above indicated that $B$ can be indifferent between reporting 1 and 0 in at most one state. The next result summarizes the insights thus far about the possible equilibrium cases.

**Lemma 3.** *If A's short-term policies are responsive to B's reports, then one of the following cases will obtain:*

*(i)* $B$ *always reports* $r = 0$

*(ii)* $B$ *mixes between* $r = 1$ *and* $r = 0$ *if* $\omega = 1$ *and reports* $r = 0$ *if* $\omega = 0$

*(iii)* $B$ *truthfully reports the state, such that* $r = \omega$

*(iv)* $B$ *reports* $r = 1$ *if* $\omega = 1$ *and mixes between* $r = 1$ *and* $r = 0$ *if* $\omega = 0$

*(v)* $B$ *always reports* $r = 1$

Note that in cases (i), the inequality in (3) is reversed, strictly; in (ii), (3) holds with equality; and in (iii-v), the inequality in (3) holds strictly. In cases (i-iii), the inequality in (4) holds strictly; in (iv), (4) holds with equality; and in (v), the inequality in (4) is reversed, strictly. These are the only possible cases if $A$ chooses short-term policies responsively, for reasons noted previously about the relationship between the two inequalities.

The result above presumes that $A$ is responsive to $B$. If $A$ instead selects a given short-term response regardless of $B$'s report, then $B$'s only influence over $A$ is on long-term policy. The following lemma states, as expected, that $B$'s best response given unresponsive short-term choices by $A$ is to report state 1 if she wishes to shift $A$'s belief about the distribution of $\pi$ upwards and 0 if she prefers a downward shift.

**Lemma 4.** *Suppose $\mu(1) \geq \mu(0) > 1/2$ or $1/2 > \mu(1) \geq \mu(0)$ such that the appointee best short-term responses are such that $s^*(1) = s^*(0)$. Then $B$ issues $r = 1$ in all states $\omega$ if*

$$\hat{\pi} > \frac{1}{2}\left(\frac{\alpha + \mu(1)}{\alpha + \beta + 1} + \frac{\alpha + \mu(0)}{\alpha + \beta + 1}\right), \tag{5}$$

*and $r = 0$ in all states $\omega$ if the inequality is reversed.*

The condition in (5) determines the sign of the bracketed expressions in (3) and (4). Before proceeding further in characterizing the equilibrium cases, the formation of the appointee's equilibrium beliefs about the probability that the state is 1 given a report $r$ requires careful attention. It is the formation of this belief in which the behavioral assumptions underlying this model become most consequential.

### 4.3 Appointee Updating from Bureaucrat Reports

When choosing a short-term policy response or a long-term policy response when no immediate response arose, the appointee only has the bureaucrat's unverified report from which to update her beliefs. Yet by the assumptions laid out above, the appointee will still learn what she can from this report. She will do so, however, taking into account the strategy of the bureaucrat (viz., the truthfulness with which $B$ reports each state) and the bureaucrat's known belief ($\hat{pi}$), as well as according to her own prior beliefs ($\nu$).

To reflect these assumptions, and recalling that $B$'s strategy is given by $q(\omega) = \Pr(r = \omega)$, set $\mu(1) := \mathbb{E}(Pr(\omega = 1|r = 1)|q(1), q(0), \nu(\pi))$, and analogously for $\mu(0) := \mathbb{E}(Pr(\omega = 1|r = 0)|q(1), q(0), \nu(\pi))$.

$$\begin{aligned}
\mu(1) =& \mathbb{E}\left(\frac{\Pr(r = 1|\omega = 1)\Pr(\omega = 1)}{\Pr(r = 1|\omega = 1)\Pr(\omega = 1) + \Pr(r = 1|\omega = 0)\Pr(\omega = 0)}\right) \\
=& \mathbb{E}\left(\frac{q(1)\pi}{q(1)\pi + (1 - q(0))(1 - \pi)}\right) \\
=& \int_0^1 \frac{q(1)\pi}{q(1)\pi + (1 - q(0))(1 - \pi)}\nu(\pi)d\pi
\end{aligned}$$

$$\begin{aligned}\mu(0) &= \mathbb{E}\left(\frac{\Pr(r=0|\omega=1)\Pr(\omega=1)}{\Pr(r=0|\omega=1)\Pr(\omega=1)+\Pr(r=0|\omega=0)\Pr(\omega=0)}\right)\\ &= \mathbb{E}\left(\frac{(1-q(1))\pi}{(1-q(1))\pi+q(0)(1-\pi)}\right)\\ &= \int_0^1 \frac{(1-q(1))\pi}{(1-q(1))\pi+q(0)(1-\pi)}\nu(\pi)d\pi\end{aligned}$$

It will not in general be possible to obtain a closed form solution for interior values of $q(1), q(0)$.[3] By Lemma 3, however, at most one of $q(1), q(0)$ will be interior in any equilibrium, and in three cases neither will be interior. As such, it is possible to pin down exact values for $q(1), q(0)$ in three of the cases and to develop bounds for the remaining two cases in Lemma 3, which the next result does. Note that in cases $(i)$ and $(v)$, because $B$ is pooling on a single report, $A$'s posterior beliefs on the equilibrium path are simply her priors.

**Lemma 5.** *$B$'s probability of reporting the state truthfully $q(\omega)$ and $A$'s corresponding beliefs that $\omega = 1$ given a report of $r = 1$ ($\mu(1)$) or given a report of $r = 0$ ($\mu(0)$) in each of the cases from Lemma 3 are as follows:*

*(i)* $q(1) = 0, q(0) = 1 \Longrightarrow \mu(1) = 1,$[4] $\mu(0) = \frac{\alpha}{\alpha+\beta}$

*(ii)* $q(1) \in (0,1), q(0) = 1 \Longrightarrow \mu(1) = 1, \mu(0) \in \left(0, \frac{\alpha}{\alpha+\beta}\right)$

*(iii)* $q(1) = 1, q(0) = 1 \Longrightarrow \mu(1) = 1, \mu(0) = 0$

*(iv)* $q(1) = 1, q(0) \in (0,1) \Longrightarrow \mu(1) \in \left(\frac{\alpha}{\alpha+\beta}, 1\right), \mu(0) = 0$

*(v)* $q(1) = 1, q(0) = 0 \Longrightarrow \mu(1) = \frac{\alpha}{\alpha+\beta}, \mu(0) = 0$[5]

*Further, $\mu(1)$ is monotonically increasing in $q(0)$, and $\mu(0)$ is monotonically decreasing in $q(1)$.*

The final statement of this result says that the more likely $B$ is to truthfully report when $\omega = 0$, the more likely $A$ is to believe the state is $\omega = 1$ based on a report $r = 0$. Analogously, the more likely

[3] It is possible to obtain a closed-form solution (and somewhat instructive) if $\nu(\pi) = 1$, such that $A$ believes $\pi \sim U(0,1)$, and if $\hat{\pi}$ is sufficiently large that $q(1) = \Pr(r = 1|\omega = 1) = 1$, then

$$\mu(1; q(1) = 1, q(0)) = \int_0^1 \frac{\pi}{\pi + (1-q(0))(1-\pi)}d\pi = \frac{q + (1-q(0))\ln(1-q(0))}{q(0)^2}, q(0) \in (0,1).$$

When $q(0) = 0$, such that $B$ always reports $\omega$ truthfully, then $\mu(1) = 1$, as expected, and this is the limit of the function above as $q(0) \to 0$. When $q(0) = 1$, such that $B$ always reports 1 regardless of the state, then $\mu(1) = \int_0^1 \pi d\pi = 1/2$, the limit of the above function as $q(0) \to 1$; in this case, where $B$'s signal is entirely uninformative, $A$'s posteriors are equal to her priors. For interior values of $q(0)$, $\mu(1)$ interpolates monotonically between 1 and 1/2.

[4] Here, $r = 1$ is off the equilibrium path, but we impose a belief that $\omega = 1$ with certainty in response to $r = 1$.

[5] While $r = 0$ is off the equilibrium path, we impose a belief that $\omega = 0$ with certainty in response to $r = 0$.

$B$ is to truthfully report when $\omega = 1$, the less likely $A$ is to believe the state is $\omega = 0$ based on a report $r = 1$. Simply, the more truthful $B$'s reports are, the more informative $A$ finds them about the state.

### 4.4 Characterizing Equilibrium Cases, with Refinements

Putting the results of the preceding subsections together, the following proposition describes the equilibrium cases. The distance from the mean of $A$'s distribution of beliefs about $\pi$ to $\hat{\pi}$ (the true underlying probability according to $B$) is one determinant of the number and nature of the equilibria. The quantity $\frac{\psi}{\lambda}$, which captures the weight that $B$ places on short-term relative to long-term considerations, also plays a key role in delimiting the various equilibrium cases. Indeed, the proposition refers to specific threshold values of $\frac{\psi}{\lambda}$.

Referring to the cases in Lemmas 3 and 5, the proposition refers to "truth-telling" equilibria (case (iii), in which $B$ accurately reports the state, and $A$ chooses $s, p = r$), "mixing-if-$\omega$" equilibria (cases (ii) and (iv), in which $B$ accurately reports the state in $1-\omega$ but randomizes between reporting truthfully and not when the state is $\omega$, and $A$ chooses $s = r$), and "pooling-on-$r$" equilibria (cases (i) and (v), in which $B$ reports $r$ regardless of the state). Within this last type of equilibria, $A$'s posterior distribution of beliefs about $\pi$ are the same as her priors, so her action is determined by the mean of that prior distribution, i.e., $\frac{\alpha}{\alpha+\beta}$; as per Lemmas 2 and 4, the appointee's response is governed by whether $\frac{\alpha}{\alpha+\beta}$ is greater or less than 1/2.

One final note is in order before stating the equilibrium cases. If $\hat{\pi} > \frac{\alpha+1}{\alpha+\beta+1}$, the equilibria are unique within each case. If not, however, there is a potential multiplicity of equilibria. The equilibria we list below are Pareto optimal, display continuity with the cases when equilibria are unique, and possess other intuitive properties discussed further in the appendix. Specifically, we assume that when truth-telling is a feasible equilibrium, it is the equilibrium players will coordinate on. If not, players will coordinate on a mixed strategy for $B$, if feasible. Otherwise, $B$ will pool on a single report.

**Proposition 1.** *Let $\frac{\alpha}{\alpha+\beta} < \frac{1}{2}$. In all cases, $A$ optimally sets $p^*(r) = \frac{\alpha+\mu(r)}{\alpha+\beta+1}$; $A$ optimally chooses $s^*(r) = r$ unless otherwise noted.*

- *If $\frac{\psi}{\lambda} \geq \left| \left(\hat{\pi} - \frac{\alpha+1}{\alpha+\beta+1}\right)^2 - \left(\hat{\pi} - \frac{\alpha}{\alpha+\beta+1}\right)^2 \right|$, then a separating (truth-telling) equilibrium obtains, with $r^*(\omega) = \omega$, and $\mu(1) = 1$, $\mu(0) = 0$.*[6]

- *If $\frac{\psi}{\lambda} \in \left[ \left(\hat{\pi} - \frac{\alpha+1}{\alpha+\beta+1}\right)^2 - \left(\hat{\pi} - \frac{\alpha}{\alpha+\beta}\right)^2, \left(\hat{\pi} - \frac{\alpha+1}{\alpha+\beta+1}\right)^2 - \left(\hat{\pi} - \frac{\alpha}{\alpha+\beta+1}\right)^2 \right)$, then $r^*(0) = 0$, and*

[6]Case (*iii*) in Figure 1.

*B randomizes when* $\omega = 1$*, such that* $\mu(1) = 1$*, and* $\mu(0) = \frac{\left(\hat{\pi} - \frac{\alpha}{\alpha+\beta+1}\right) + \sqrt{\left(\hat{\pi} - \frac{\alpha+1}{\alpha+\beta+1}\right)^2 - \psi/\lambda}}{1/(\alpha+\beta+1)}$.[7]

- *If* $\frac{\psi}{\lambda} \in \left[\left(\hat{\pi} - \frac{\alpha}{\alpha+\beta}\right)^2 - \left(\hat{\pi} - \frac{\alpha+1}{\alpha+\beta+1}\right)^2, \left(\hat{\pi} - \frac{\alpha}{\alpha+\beta}\right)^2 - \left(\hat{\pi} - \frac{\alpha+1/2}{\alpha+\beta+1}\right)^2\right)$*, then* $r^*(1) = 1$*, and B randomizes when* $\omega = 0$*, such that* $\mu(1) = \frac{\left(\hat{\pi} - \frac{\alpha}{\alpha+\beta+1}\right) - \sqrt{\left(\hat{\pi} - \frac{\alpha}{\alpha+\beta+1}\right)^2 - \psi/\lambda}}{1/(\alpha+\beta+1)}$*, and* $\mu(0) = 0$.[8]

- *Otherwise, one of two cases of a pooling equilibrium obtains:*

  - *If* $\hat{\pi} < \frac{\alpha+1/2}{\alpha+\beta+1}$*, then* $r^*(\omega) = 0$*,* $\mu(1) = 1$*,* $\mu(0) = \frac{\alpha}{\alpha+\beta}$.[9]

  - *If* $\hat{\pi} > \frac{\alpha+1/2}{\alpha+\beta+1}$*, then* $r^*(\omega) = 1$*, and* $\mu(1) = \frac{\alpha}{\alpha+\beta}$*,* $\mu(0) = 0$*. Because* $\frac{\alpha}{\alpha+\beta} < \frac{1}{2}$*,* $s^*(0) = 0$ *and* $s^*(1) = 0$.[10]

Figure 1 depicts these cases. We may observe that, given $A$'s priors, the range of $\hat{\pi}$'s for which there exists a truth-telling equilibrium increases as the bureaucrat's short-to-long-term weighting increases. The extent of disagreement between $A$ and $B$ may be quite large and a truth-telling equilibrium still exist, as long as $B$ weights short-term considerations enough relative to long-term considerations. For intermediate values of this weighting, and $\hat{\pi}$ sufficiently larger or smaller than $\alpha/(\alpha+\beta)$, an equilibrium in which $B$ mixes in one state of the world and truthfully reports in the other is possible. For smaller short-to-long-term weightings, $B$ will pool on the report that would bring $A$'s beliefs towards $\hat{\pi}$, though no such convergence occurs in such cases. The range of weightings that support pooling equilibria increases as the disagreement between $B$ and $A$ increases. The next section discusses the mechanics and significance of each type of equilibrium case in greater detail.

# 5 Discussion of the Different Cases

The three types of equilibrium cases represent different stages of the process of coming around. Even though the game is not modeled as a repeated interaction, over time $\alpha/(\alpha+\beta)$ will approach $\hat{\pi}$, either through updating of beliefs from reports from the bureaucrat or realizing the congruence, or lack thereof, between short-term responses and the state of the world. As the disagreement in beliefs (and thus policy preferences) between bureaucrat and appointee decreases, we move from pooling to mixing to separating equilibria.

[7] Case (*ii*) in Figure 1.
[8] Case (*iv*) in Figure 1.
[9] Case (*i*) in Figure 1.
[10] Case (*v*) in Figure 1.

Figure 1: Equilibrium Cases

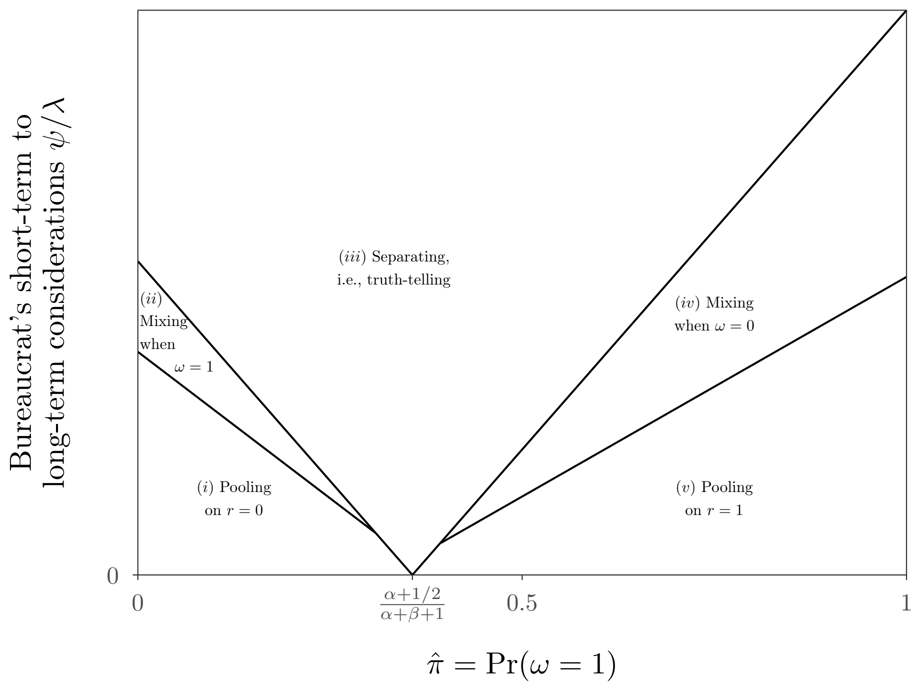


When disagreement between the appointee and the bureaucrat is quite large and the bureaucrat's short-to-long-term weights are sufficiently small, both sides will seem to entrench. The bureaucrat will issue the report closest to $\hat{\pi}$ regardless of the state. The appointee does not update in the direction of $\hat{\pi}$ from the reports, however, which are uninformative. The appointee's belief (i.e., the mean of the distribution $\nu(\pi)$) gradually shifts towards $\hat{\pi}$ only when verified by the congruence (or lack thereof) of a short-term response. This type of equilibrium is possible only for low-enough values of $\frac{\psi}{\lambda}$, and if this quantity is small on account of a low probability of a short-term response being necessary, the updating of beliefs towards $\hat{\pi}$ could occur quite slowly.

If the appointee is so predisposed against the bureaucrat's view that they each deem a different state to be more likely than the other (i.e., $\alpha/(\alpha+\beta) < 1/2$ and $\hat{\pi} > (\alpha+1/2)/(\alpha+\beta+1)$), then appointee will choose policies that directly contradict the bureaucrat's report. This constitutes one explanation for the observation of parties in a disagreement seeming to cede less ground the more polarized their positions relative to one another. If the appointee and the bureaucrat agree on which state is more likely, however, the appointee partakes in a different sort of entrenchment. In this case, the appointee will follow the bureaucrat's advice to an extent that overstates her confidence that the state is as the bureaucrat says. This equilibrium behavior may provide insight into notions of a vanishing middle, in which even moderates tend to behave similarly to ideological purists.

As disagreement abates and/or short-term considerations carry increasing weight relative to longer-term alignment, mixing equilibrium cases become possible, in which the bureaucrat honestly reports one state while occasionally lying in the other state. While the appointee knows the bureaucrat's strategy and belief that $\hat{\pi}$ is the true probability the state is 1, and though she uses her prior distribution to process information, somewhat surprisingly the appointee's belief still updates in the direction of $\hat{\pi}$ in a mixing equilibrium. The mixing equilibria represent the most subtle of the perils of bringing someone around. The bureaucrat manipulates the appointee to expedite the process of her coming around, but not so much that the appointee loses trust in the bureaucrat. The motivation to do so persists despite the expectation that the appointee will occasionally learn the true state because the bureaucrat would, all things equal, prefer the appointee update her beliefs in the direction of $\hat{\pi}$ as quickly as possible. It is worth noting that this is only possible when the appointee and bureaucrat both believe the same state is the more likely of the two.

In the third type of equilibrium case, the bureaucrat always accurately reports the state, and the appointee follows this report when choosing policies. This case is possible when disagreement between the appointee's priors hew towards $\hat{\pi}$ and the bureaucrat's short-term considerations are strong enough relative to her long-term considerations. If the bureaucrat's incentives to misrepresent the state are too high, a strategy of truth-telling would not be credible. Such a strategy does highlight the counterfactual to the less truthful strategies, namely, that if the bureaucrat reported truthfully, the appointee would gradually come around. It is the desire to bring the appointee around more quickly than she might otherwise come around that prevents truth-telling equilibria from always being possible.

# 6 Conclusion

The model provides a basic intuition into the process of coming around. A bureaucrat may strategically misrepresent the information she provides to her political appointee when there is significant disagreement between the two or a low likelihood of a need to respond to an immediate issue. Unless there is extreme disagreement between the appointee and the bureaucrat, the appointee updates closer to the bureaucrat's position. This outcome is consistent with the bureaucrat's goal of having the appointee come around more quickly to her position than the appointee might have otherwise done if left alone to update according to confirmations (or refutations) of prior beliefs over time.

Perhaps more fundamentally, the model also highlights the role of trust within the professional relationship between appointees and bureaucrats. Although bureaucrats may engage in deception, there is little reason to expect them to be completely uninformative at all times. Consistent with the literature, in the cases we examined, a bureaucrat will tilt towards telling the truth more often than not.

For an appointee, maintaining trust with a bureaucrat who provides information is relevant in order for the appointee to be effective. As Heclo (1977) notes, political appointees who arrive at an agency skeptical of the motivation and expertise of bureaucrats are at a disadvantage to their more open-minded peers; to be effective, a political appointee "does not have to trust everyone in the bureaucracy, but they need to trust some" (189). Our model shows why this might be the case. Specifically, an appointee who remains skeptical and treats the reports they receive from a bureaucrat as uninformative during their tenure will update their beliefs more gradually, according to trial-and-error, than their less extreme peers. In practice, this strategy requires appointees to invest a lot of time to arrive at accurate beliefs, despite it mitigating the risk of being deceived by their bureaucrat. In the appointee's position, however, time is a resource that is limited and uncertain.

Lastly, although more work is needed to extend our analysis to infinitely-repeated contexts, the model suggests that truth-telling by a bureaucrat is increasingly optimal over time as potential gains from misreporting the true state becomes smaller. The appointee's beliefs shift closer to the bureaucrat's position over time as she learns from prior observations (and reports, potentially). Taken collectively, our work shows that coming around is not exactly a process by which a bureaucrat manipulates a political appointee to become a pawn of their agency. Instead, coming around is about how a political appointee accepts and updates what they believe is true, which is already known by bureaucrats and occurs (albeit gradually) regardless of strategic manipulation. From this perspective, there is little wonder why most political appointees "learn to qualify their preliminary distrust of the bureaucracy" and leave their positions (often publicly) "regretting the bulk of their early suspicions" (Heclo 1977, 187).

# References


Aumann, Robert J. 1976. "Agreeing to Disagree." *The Annals of Statistics* 4(6):1236–1239.

Berman, Evan, Don-Yun Chen, Chung-Yuang Jan & Tong-Yi Huang. 2013. "Public Agency Leadership: The Impact of Informal Understandings With Political Appointees on Perceived Agency Innovation in Taiwan." *Public Administration* 91(2):303–324.

Chen, Chung-An & Chih-Wei Hsieh. 2015. "Knowledge sharing motivation in the public sector: the role of public service motivation." *International Review of Administrative Sciences* 81(4):812–832.

Dolan, Julie. 2000. "Influencing Policy at the Top of the Federal Bureaucracy: A Comparison of Career and Political Senior Executives." *Public Administration Review* 60(6):573–581.

French, Simon. 1980. "Updating of Belief in the Light of Someone Else's Opinion." *Journal of the Royal Statistical Society. Series A (General)* 143(1):43.

Gailmard, Sean & John W. Patty. 2012. "Formal Models of Bureaucracy." *Annual Review of Political Science* 15(1):353–377.

Gailmard, Sean & John W. Patty. 2013. *Learning While Governing.* Chicago: The University of Chicago Press.

Gailmard, Sean, John W Patty & Sean Gailmard. 2007. "Slackers and Zealots: Civil Policy and Bureaucratic Expertise." *American Journal of Political Science* 51(4):873–889.

Genest, Christian & Mark J. Schervish. 1985. "Modeling Expert Judgments for Bayesian Updating." *The Annals of Statistics* 13(3):1198–1212.

Heclo, Hugh. 1977. *A government of strangers : executive politics in Washington.* Brookings Institution.

Hirsch, Alexander V. 2016. "Experimentation and Persuasion in Political Organizations." *American Political Science Review* 110(1):68–84.

Morris, Peter A. 1974. "Decision Analysis Expert Use." *Management Science* 20(9):1233–1241.

Stehr, Steven D. 1997. "Top Bureaucrats and the Distribution of Influence in Reagan's Executive Branch." *Public Administration Review* 57(1):75.

## Proofs

*Proof of Lemma 1.* The beta distribution is a conjugate prior for an unknown parameter $\pi$ in a data-generating process $Bern(\pi)$, a fact that holds even for fractional "observations" such as $\mu(r)$. ■

*Proof of Lemma 2.* For short-term responses, $A$ solves for $s^* = \arg\max_{s \in \{0,1\}} -\mu(r)(s-1)^2 - (1-\mu(r))(s-0)^2$. Setting $s = 1$ achieves a strictly higher (resp., lower) value than $s = 0$ if $\mu(r) > 1/2$ (resp., $<$).

For long-term policy, $A$ solves for $p^* = \arg\max_{p \in [0,1]} \int_0^1 -(\pi - p)^2 \nu(\pi) d\pi$. Setting $p$ equal to the mean of the distribution over $\pi$ minimizes the squared difference. ■

Explicit proofs of 3-5 forthcoming, but they entail some trivial algebra following from the inequalities (3) and (4).

We employ the following refinements:

- If $B$ misreports the state of the world, it is to bring $A$'s posterior belief closer to $\hat{\pi}$, implying that if $\hat{\pi} > \frac{1}{2}\left(\frac{\alpha+1}{\alpha+\beta+1} + \frac{\alpha}{\alpha+\beta+1}\right)$, $B$ will misreport only if (though not necessarily) $\omega = 0$.
- Given the above restriction, if multiple equilibria exist, separating is chosen if possible, then mixed-strategy if possible, and only otherwise pooling.

***Please excuse the notes to ourselves in the proofs below.

*Proof of Proposition 1.* Lemma 2 provides $p^*$ and $s^*$ for all cases.

For pooling equilibria, posteriors equal priors on the equilibrium path. If $B$ pools on $r = 1$, $\mu(1) = \frac{\alpha}{\alpha+\beta}$ by Lemma 1, and $\mu(0) = 0$ by Lemma 3; when $B$ pools on $r = 0$, $\mu(0) = \frac{\alpha}{\alpha+\beta}$ by Lemma 1, and $\mu(1) = 1$ by Lemma 3. In the latter case, since $\mu(0) < 1/2 < \mu(1)$, $A$'s short-term responses are responsive to $B$'s reports, even off the equilibrium path. In the former case, since $\mu(0) < \mu(1) < 1/2$, $A$ selects $s^* = 0$ even off the equilibrium path. More importantly, $1/2$ is the lower bound for $\mu(1)$ in a mixing-when-$\omega = 0$ equilibrium, while $\alpha/(\alpha + \beta)$ is the upper bound for $\mu(0)$ in a mixing-when-$\omega = 1$ equilibrium. This would be reversed if $\alpha > \beta$ such that $\alpha/(\alpha + \beta) > 1/2$.

For truth-telling, i.e., separating equilibria, we need the inequalities in (3) and (4) to be satisfied at $\mu(0) = 0$ and $\mu(1) = 1$, which leads to

$$\frac{\psi}{\lambda} \geq \left| \left(\hat{\pi} - \frac{\alpha+1}{\alpha+\beta+1}\right)^2 - \left(\hat{\pi} - \frac{\alpha}{\alpha+\beta+1}\right)^2 \right|.$$

The beliefs and actions follow from Lemmas 3 and 5.

Equality of the (3) or (4) define quadratic equations; with respect to $\mu(0)$ for the former, holding fixed $\mu(1) = 1$, and with respect to $\mu(1)$ for the latter, holding fixed $\mu(0) = 0$, both strictly concave. Denote the solutions by $\mu(0)_* < \mu(0)^*$ and $\mu(1)_* < \mu(1)^*$, where each of these solutions could represent a mixed-strategy equilibrium. In the case of an equilibrium based on $\mu(0)_*$ or $\mu(1)^*$, there also exists a separating equilibrium,[11] so they are disqualified given our refinements. For the other solutions to be valid, by Lemmas 4 and 3, $\mu(0)$ must be less than $\min\{1/2, \alpha/(\alpha+\beta)\}$ if mixing is occurring when $\omega = 1$ and $\mu(1)$ greater than $\max\{1/2, \alpha/(\alpha+\beta)\}$ when mixing is occurring on $\omega = 0$. Given our assumption that $\alpha < \beta$, we know the first bound is $\alpha/(\alpha+\beta)$ and the latter is $1/2$.[12] These produce the following inequalities that must be satisfied:

$$\frac{\alpha}{\alpha+\beta} > \mu(0)^* = \frac{\left(\hat{\pi} - \frac{\alpha}{\alpha+\beta+1}\right) + \sqrt{\left(\hat{\pi} - \frac{\alpha+1}{\alpha+\beta+1}\right)^2 - \psi/\lambda}}{1/(\alpha+\beta+1)} \tag{6}$$

$$\frac{1}{2} < \mu(1)_* = \frac{\left(\hat{\pi} - \frac{\alpha}{\alpha+\beta+1}\right) - \sqrt{\left(\hat{\pi} - \frac{\alpha}{\alpha+\beta+1}\right)^2 - \psi/\lambda}}{1/(\alpha+\beta+1)} \tag{7}$$

Finally, the first of the refinements dictates which mixing or pooling equilibrium may be played, depending on the location of $\hat{\pi}$ relative to $\frac{\alpha+1/2}{\alpha+\beta+1}$. ■

[11] Need to show in separate lemma.
[12] Explain why in greater detail.