Online Appendix for

"At Your Own Risk:
A Model of Delegation with Ambiguous Guidelines"

Table of Contents

A. Proofs to Propositions 1-3

Proof to Proposition 1

Proposition 1 (Ambiguity). *If $q < 0.5$ and review cost is sufficiently low, $\gamma < 4q(1-q)$,*

(1) When the disapproval penalty is mild ($\kappa < \bar{\kappa} \equiv \frac{2(1-q)}{q}$), there is a mixed-strategy PBE where the center reviews $x = R$ with probability $\tilde{\rho}^* = \frac{\kappa}{(1-q)(\kappa+2)} \in (0,1)$ and disapproves it unless verified from review; the competent locality chooses a policy that matches the state, and the incompetent locality chooses policy R with probability $\tilde{\sigma}_i^* = \frac{p(1-q)\{2(1-\alpha)-\gamma\}}{(1-p)\{\gamma-2(1-q)(1-\alpha)\}} \in (0,1)$ and chooses policy Q with probability $1-\tilde{\sigma}_i^*$.

(2) When the disapproval penalty is harsh ($\kappa > \bar{\kappa}$), there is a mixed-strategy PBE where the center reviews $x = R$ with probability $\hat{\rho}^* = \frac{2(1-2q)}{q(\kappa-2)} \in (0,1)$ and approves it unless falsified from review; the competent locality chooses a policy that matches the state, and the incompetent locality chooses policy R with probability $\hat{\sigma}_i^* = \frac{p(1-q)\gamma}{(1-p)\{2q(1+\alpha)-\gamma\}} \in (0,1)$ and chooses policy Q with probability $1-\hat{\sigma}_i^*$.

Proof. Consider an equilibrium where the competent locality chooses a policy that matches the state (i.e., σ_c^*=1), and the incompetent locality chooses policy R with probability $\sigma_i^* \in (0,1)$ and chooses policy Q with probability $1-\sigma_i^*$. The center's belief d_{mix} =Pr$[x = \omega | x = R]$ can be written as:

$$d_{mix} = \frac{p(1-q) + (1-p)\sigma_i^*(1-q)}{p(1-q) + (1-p)\sigma_i^*}.$$

Suppose $d_{mix} < 0.5(1+\alpha)$ such that the center disapproves R unless verified (i.e., $\chi^{-*} = 0$; $\chi^{+*} = 1$ if $\omega = x$, $\chi^{+*} = 0$ if $\omega \neq x$, χ^{+*} is omitted hereafter). For the incompetent locality to mix, it must be indifferent between choosing policy R and policy Q, which is satisfied when $\tilde{\rho}^* = \frac{\kappa}{(\kappa+2)(1-q)}$ from Equation 6. Given $\tilde{\rho}^*$, the competent locality's best response is $\sigma_c^* = 1$ from Equation 4. For the center to mix, it must be indifferent between reviewing and not reviewing, which is satisfied when $\gamma = 2(1-\alpha)d_{mix}$ from Equation 3, or to rearrange, $\sigma_i^* = \tilde{\sigma}_i^* \equiv \frac{p(1-q)\{2(1-\alpha)-\gamma\}}{(1-p)\{\gamma-2(1-q)(1-\alpha)\}}$. Therefore, $(\tilde{\rho}^*, \chi^{+*}, \chi^{-*})$ and $(\sigma_c^*, \tilde{\sigma}_i^*)$ are mutually best responses and sequentially rational. When $2 < \kappa < \bar{\kappa} \equiv \frac{2(1-q)}{q}$, $\tilde{\rho}^* \in (0,1)$ always holds. For any values of γ and q that satisfy $\gamma < 4q(1-q)$, there always exists $\tilde{\sigma}_i^* \in (0,1)$ that satisfies $d_{mix} < 0.5(1+\alpha)$.

The constraint for γ is obtained by the following steps. First, for $\tilde{\sigma}_i^* \in (0,1)$ to hold, it must hold that

$$\tilde{\sigma}_i > 0 \iff 2(1-q)(1-\alpha) < \gamma < 2(1-\alpha),$$
$$\tilde{\sigma}_i < 1 \iff p(1-q)\{2(1-\alpha)-\gamma\} < (1-p)\{\gamma - 2(1-q)(1-\alpha)\}.$$

If $2(1-q)(1-\alpha) < \gamma \iff \alpha > 1 - \frac{\gamma}{2(1-q)}$ holds from the upper equation, holding q and γ constant, the lower equation must hold as p approaches to 0. This implies that there exist $\tilde{\sigma}_i^* \in (0,1)$ for sufficiently low values of p. Second, for $d_{mix} < 0.5(1+\alpha)$ to hold, it must hold that

$$d_{mix} < 0.5(1+\alpha) \iff \frac{\gamma}{2(1-\alpha)} < 0.5(1+\alpha)$$
$$\iff \gamma < (1+\alpha)(1-\alpha)$$
$$\iff \alpha < \sqrt{1-\gamma}.$$

Finally, combining $\alpha > 1 - \frac{\gamma}{2(1-q)}$ and $\alpha < \sqrt{1-\gamma}$ yields $\gamma < 4q(1-q)$.

Now suppose $d_{mix} \geq 0.5(1+\alpha)$ such that the center approves R unless falsified (i.e., $\chi^{-*} = 1$). Given this strategy, the competent locality's best response is always $\sigma_c^* = 1$. For the incompetent locality to mix, it must be indifferent between choosing policy R and policy Q, which is satisfied when $\hat{\rho}^* = \frac{2(1-2q)}{q(\kappa-2)}$ from Equation 5. For the center to mix, it must be indifferent between reviewing and not reviewing, which is satisfied when $\gamma = 2(1+\alpha)(1-d_{mix})$ from Equation 2, or to rearrange, $\sigma_i^* = \hat{\sigma}_i^* \equiv \frac{p(1-q)\gamma}{(1-p)\{2q(1+\alpha)-\gamma\}}$. Therefore, $(\hat{\rho}^*, \chi^{+*}, \chi^{-*})$ and $(\sigma_c^*, \hat{\sigma}_i^*)$ are mutually best responses and sequentially rational. When $\kappa > \bar{\kappa} \equiv \frac{2(1-q)}{q}$, $\hat{\rho}^* \in (0,1)$ always holds. For any values of γ and q that satisfy $\gamma < 4q(1-q)$, there always exists $\hat{\sigma}_i^* \in (0,1)$ that satisfies $d_{mix} \geq 0.5(1+\alpha)$.

Similar to the above, the constraint for γ is obtained by the following steps. First, for $\hat{\sigma}_i^* \in (0,1)$ to hold, it must hold that

$$\hat{\sigma}_i > 0 \iff 2q(1+\alpha) > \gamma,$$
$$\hat{\sigma}_i < 1 \iff p(1-q)\gamma < (1-p)\{2q(1+\alpha)-\gamma\}.$$

If $2q(1+\alpha) > \gamma \iff \alpha > 1 - \frac{\gamma}{2q}$ holds from the upper equation, holding q and γ constant, the lower equation must hold as p approaches 0. This implies that there exist $\hat{\sigma}_i^* \in (0,1)$ for sufficiently low values of p. Second, for $d_{mix} > 0.5(1+\alpha)$ to hold, it must hold that

$$\begin{aligned} d_{mix} > 0.5(1+\alpha) &\iff 1 - \frac{\gamma}{2(1+\alpha)} > 0.5(1+\alpha) \\ &\iff \gamma < (1+\alpha)(1-\alpha) \\ &\iff \alpha < \sqrt{1-\gamma}. \end{aligned}$$

Finally, combining $\alpha > 1 - \frac{\gamma}{2q}$ and $\alpha < \sqrt{1-\gamma}$ yields $\gamma < 4q(1-q)$. □

Proof to Proposition 2

Proposition 2 (Micromanage). *If $q < 0.5$ and review cost is sufficiently low, $\gamma < 4q(1-q)$,*

(1) When the disapproval penalty is mild ($\kappa < \bar{\kappa} \equiv \frac{2(1-q)}{q}$), there are pure-strategy PBE where the center always reviews $x = R$, the competent locality chooses a policy that matches the state, and the incompetent locality chooses policy R. Off the equilibrium path, if the center does not review, it either approves or disapproves policy R.

Proof. First, consider an equilibrium where the center always reviews R and approves it unless falsified (i.e., $\rho^* = 1$, $\chi^{-*} = 1$), the competent locality chooses a policy that matches the state, and the incompetent locality chooses R (i.e., $\sigma_c^* = 1$, $\sigma_i^* = 1$). Given the locality's strategy, the center's belief d_{pure} =Pr$[x = \omega | x = B]$ can be written as:

$$d_{pure} = \frac{1-q}{1-pq}.$$

When $\gamma < 2(1+\alpha)(1-d_{pure})$ and $d_{pure} \geq 0.5(1+\alpha)$ hold, or combined, when $\gamma < (1+\alpha)(1-\alpha)$ holds, the center's best response is $\rho^* = 1$, $\chi^{-*} = 1$. Given the center's strategy, the competent locality's best response is always $\sigma_c^* = 1$, and when $1 \leq \frac{2(1-q)}{p(\kappa-2)}$ the incompetent locality's best response is $\sigma_i^* = 1$ from Equation 5. Therefore, $(\rho^*, \chi^{+*}, \chi^{-*})$ and (σ_c^*, σ_i^*) are mutually best responses and sequentially rational.

Second, consider an equilibrium where the center always reviews R and disapproves it unless verified (i.e., $\rho^* = 1$, $\chi^{-*} = 0$), the competent locality chooses a policy that matches the state, and the incompetent locality chooses R (i.e., $\sigma_c^* = 1$, $\sigma_i^* = 1$). Given the locality's strategy, the center's belief is as before: $d_{pure} = \frac{1-q}{1-pq}$. When $\gamma < 2(1-\alpha)d_{pure}$ and $d_{pure} < 0.5(1+\alpha)$ hold, or combined, when $\gamma < (1+\alpha)(1-\alpha)$ holds, the center's best response is $\rho^* = 1$, $\chi^{-*} = 0$.

Given the center's strategy, when $1 > \frac{\kappa}{(\kappa+2)(1-q)}$ holds, the competent locality's best response is $\sigma_c^* = 1$ from Equation 4, and the incompetent locality's best response is $\sigma_i^* = 1$ from Equation 6. Therefore, $(\rho^*, \chi^{+*}, \chi^{-*})$ and (σ_c^*, σ_i^*) are mutually best responses and sequentially rational.

Note that the above equilibria exist if $q < 0.5$ and $\gamma < (1+\alpha)(1-\alpha)$ hold. From proof to proposition 1, the constraint $\gamma < 4q(1-q)$ is a sufficient condition for the constraint $\gamma < (1+\alpha)(1-\alpha)$. In proposition 2, I used the former constraint for the comparability of propositions 1 and 2. □

Proof to Proposition 3

Proposition 3. *Invoking the ambiguity equilibrium is optimal for the center when α is neither too high nor too low and p is sufficiently low.*

Proof. I prove this proposition in two steps: first, identify the center's equilibrium payoffs, and second, identify when the *ambiguity equilibrium* of each range of disapproval penalty becomes optimal for the center.

Step 1: Identify the center's equilibrium payoffs.

First, suppose that the center designates policy R as safe. The competent locality chooses a policy that matches the state, and the incompetent locality chooses policy R. The center's equilibrium payoff is

$$EU_C(safe) = qp(1+\alpha) + q(1-p)(-1-\alpha) + (1-q)(1-\alpha),$$

which is the probability that the state is Q times the probability that the locality is a competent type times a payoff of correctly implementing policy Q, $1+\alpha$; plus, the probability that the state is Q times the probability that the locality is an incompetent type times a payoff of incorrectly implementing policy R; plus the probability that the state is R times a payoff of correctly implementing policy R by both competent and incompetent types, $1-\alpha$.

Second, suppose that the center designates policy R as prohibited. Both competent and incompetent types of the locality choose policy Q. The center's equilibrium payoff is

$$EU_C(prohibited) = q(1+\alpha) + (1-q)(-1+\alpha),$$

which is the probability that the state is Q times a payoff of correctly implementing policy Q by both competent and incompetent types; plus the probability that the state is R times a payoff of incorrectly implementing policy Q by both incompetent and incompetent types.

Third, suppose that the center designates policy R as gray and always reviews $x = R$. In the *micromanage equilibria*, the competent locality chooses a policy that matches the state, and the incompetent locality chooses policy R when conditions from Proposition 2 are met. The center's equilibrium payoff is

$$EU_C(micromanage) = qp(1+\alpha) + q(1-p)(-\gamma+1+\alpha) + (1-q)(-\gamma+1-\alpha),$$

which is the probability that the state is Q times the probability that the locality is a competent type times a payoff of correctly implementing policy Q; plus the probability that the state is Q times the probability that the locality is an incompetent type times a payoff of reviewing $x = R$ and reverting it to a correct policy Q; plus the probability that the state is R times a payoff of reviewing and approving $x = R$, chosen by both competent and incompetent types.

Fourth, suppose that the center designates policy R as gray and invokes the *ambiguity equilib-*

rium (disapprove-unless-verified, 'duv' for short), which is supported only with mild disapproval penalty, $\kappa < \bar{\kappa}$. The competent locality chooses a policy that matches the state, and the incompetent locality chooses policy R with probability $\tilde{\sigma}_i^*$ and chooses policy Q with probability $1 - \tilde{\sigma}_i^*$. The center's equilibrium payoff is

$$
\begin{aligned}
EU_C(ambiguity, duv) =& qp(1+\alpha)+ \\
& q(1-p)[\tilde{\sigma_i^*}\{\tilde{\rho}^*(-\gamma+1+\alpha) + (1-\tilde{\rho}^*)(1+\alpha)\} + (1-\tilde{\sigma_i^*})(1+\alpha)]+ \\
& (1-q)p\{\tilde{\rho}^*(-\gamma+1-\alpha) + (1-\tilde{\rho}^*)(-1+\alpha)\} \\
& (1-q)(1-p)[\tilde{\sigma_i^*}\{\tilde{\rho}^*(-\gamma+1-\alpha) + (1-\tilde{\rho}^*)(-1+\alpha)\} + (1-\tilde{\sigma_i^*})(-1+\alpha)] \\
=& q(1+\alpha) + (1-q)(-1+\alpha).
\end{aligned}
$$

The first part is the probability that the state is Q times the probability that the locality is a competent type times a payoff of correctly implementing policy Q. The second part is the probability that the state is Q times the probability that the locality is an incompetent type times the sum of expected utilities when the incompetent locality chooses policy R and policy Q, respectively. When the incompetent locality chooses policy R, the center reviews with probability $\tilde{\rho^*}$, and does not review with probability $1 - \tilde{\rho^*}$. If it reviews, it learns that the state is Q and reverts the policy to policy Q. If it does not review, it disapproves $x = R$ and reverts it to policy Q. When the incompetent locality chooses policy Q, it is approved without review. The third and fourth parts follow the similar logic.

Finally, suppose that the center designates R as gray and invokes the *ambiguity equilibrium* (approve-unless-falsified, 'auf' for short), which is supported only with harsh disapproval penalty, $\kappa > \bar{\kappa}$. The competent locality chooses a policy that matches the state, and the incompetent locality chooses policy R with probability $\hat{\sigma}_i^*$ and chooses policy Q with probability $1 - \hat{\sigma}_i^*$. The center's equilibrium payoff, following the similar logic from above, is

$$
\begin{aligned}
EU_C(ambiguity, auf) =& qp(1+\alpha)+\\
& q(1-p)[\hat{\sigma_i^*}\{\hat{\rho^*}(-\gamma+1+\alpha)+(1-\hat{\rho^*})(-1-\alpha)\}+(1-\hat{\sigma_i^*})(1+\alpha)]\\
& (1-q)p\{\hat{\rho^*}(-\gamma+1-\alpha)+(1-\hat{\rho^*})(1-\alpha)\}+\\
& (1-q)(1-p)[\hat{\sigma_i^*}\{\hat{\rho^*}(-\gamma+1-\alpha)+(1-\hat{\rho^*})(1-\alpha)\}+(1-\hat{\sigma_i^*})(-1+\alpha)].
\end{aligned}
$$

Step 2: Identify when the *ambiguity equilibrium* is optimal.

First consider the mild disapproval penalty range, $\kappa < \bar{\kappa}$. In this range, the *ambiguity equilibrium* (disapprove unless verified) exists. It yields a higher payoff to the center than designating policy R as safe when

$$
\begin{aligned}
EU_C(ambiguity, duv) \geq EU_C(safe) \iff & \alpha > \underline{\alpha} \equiv \frac{1-2q+pq}{1-pq}\\
\iff & p < \bar{p} \equiv \frac{\alpha+2q-1}{q(1+\alpha)},
\end{aligned}
$$

or when α is sufficiently high and p is sufficiently low. It yields the same payoff as designating policy R as prohibited. It yields a higher payoff to than designating policy R as gray and always reviewing it (*micromanage equilibrium*) when

$$
\begin{aligned}
EU_C(ambiguity, duv) \geq EU_C(micromanage) \iff & \alpha > \underline{\alpha} \equiv 1-\frac{(1-pq)\gamma}{2(1-q)}\\
\iff & p < \bar{p} \equiv \frac{2(1-q)(-1+\alpha)+\gamma}{q\gamma},
\end{aligned}
$$

or when α is sufficiently high and p is sufficiently low. Combining the results yields Proposition 3.

Now consider the harsh disapproval penalty range, $\kappa > \bar{\kappa}$. In this range, the *ambiguity equilibrium* (approve-unless-falsified) exists. It yields a higher payoff than designating policy R

as safe when

$$EU_C(ambiguity, auf) \geq EU_C(safe) \iff \alpha > \bar{\bar{\alpha}} \equiv \frac{(1-pq)\gamma}{2(1-p)q} - 1$$
$$\iff p < \underline{\underline{p}} \equiv \frac{2q(1+\alpha)-\gamma}{2q(1+\alpha)-q\gamma},$$

or when α is sufficiently high and p is sufficiently low. It yields a higher payoff than designating policy R as prohibited when

$$EU_C(ambiguity, auf) \geq EU_C(prohibited) \iff \alpha < \underline{\underline{\alpha}} \equiv \sqrt{1-\gamma},$$

or when α is sufficiently low. Again, combining the results yields Proposition 3. □

B. Relaxing the Assumption that Status Quo Policy is Safe

Since the center designates safe, prohibited, or gray for each policy Q and R, the center has nine possible strategies: {(safe, safe), (safe, prohibited), (safe, gray), (prohibited, safe), $\cdots$, (gray, gray)}, where (x, y) means designating x for policy Q and y for policy R. In the main text, I have analyzed only three of them where policy Q is safe. In this section, I further analyze the remaining six strategies. Again, I restrict attention to $q < 0.5$.

i. When policy Q is prohibited.

Suppose that the center designates policy Q as prohibited: {(prohibited, safe), (prohibited, prohibited), (prohibited, gray)}.

First, consider (prohibited, safe). Then both competent and incompetent types of locality choose policy R, and the locality's policy choice stands. This strategy is strictly dominated by (safe, safe) where the competent locality matches the state while the incompetent locality chooses policy R and the locality's policy choice stands. This is because with (prohibited, safe) the center forgoes the utility from exploiting the competent locality's expertise.

Second, consider (prohibited, prohibited). Since the locality's choice is always disapproved without review and reverted to the other one, the competent locality chooses a policy that does not match the state, and the incompetent locality chooses policy Q, which is less likely to match the state. The center's payoff is the same as under (safe,safe). While this guideline seems odd, it becomes optimal under the same conditions as the (safe,safe) guideline.

Finally, consider (prohibited, gray). Then both competent and incompetent types of locality chooses policy R. Since neither type of the locality uses a mixed-strategy, the center does not either. As not reviewing $x = R$ after designating it as gray is functionally equivalent to designating it safe or prohibited, I only consider a situation where the center always reviews. In this situation, the center makes all the locality choose policy R and then selectively revert it to policy Q through review. The center's payoff is the same as when it investigates the state upfront and dictate the correct policy to the locality. This guideline becomes optimal under the same conditions as the option of investigating upfront as analyzed in Appendix B.

ii. When policy Q is gray.

Suppose that the center designates policy Q as gray; $\{$(gray, safe), (gray, prohibited), (gray, gray)$\}$.

First, consider (gray, safe). Then the incompetent locality will choose policy R. It follows that the center has no incentive to review $x = Q$ because it can infer upon observing $x = Q$ that the policy is chosen by the competent locality. This, in turn, makes the competent locality choose a policy that matches the state. Therefore, under (gray, safe), there is no equilibrium where the center reviews $x = Q$ with positive probability. The center's payoff is the same as under (safe, safe). This guideline becomes optimal under the same conditions as the (safe,safe) guideline.

Second, consider (gray, prohibited). By the same logic as under (prohibited, gray), the center's payoff is the same as when it investigates the state upfront. This guideline becomes optimal under the same conditions as the option of investigating upfront as analyzed in Appendix B.

Finally, consider (gray, gray). Let us focus on situations where the center reviews $x = Q$ and $x = R$ with positive probability, because otherwise the strategy becomes functionally equivalent to one of the other strategies. For the center to review both $x = Q$ and $x = R$ with positive probability, it must be that the incompetent mixes between the two; otherwise, the center has no incentive to review the policy not chosen by the incompetent locality. For the incompetent locality to mix, the center must also mix between review and not review for at least one policy. While the center can mix between review and not review for both policies, I do not consider these equilibria, for they exist only for a knife-edge parameter value. Thus, there are four review and approval strategies for the center under (gray, gray):

(1) always review $x = R$, and randomly review $x = Q$ (approve unless verified),

(2) always review $x = R$, and randomly review $x = Q$ (disapprove unless falsified),

(3) always review $x = Q$, and randomly review $x = R$ (approve unless verified), and

(4) always review $x = Q$, and randomly review $x = R$ (disapprove unless falsified).

Only the first strategy may be supported as an equilibrium under the (gray, gray) guideline, though we need to check when it becomes optimal for the center. The second strategy cannot be supported because the center has incentives to mix between review not review if and only if the incompetent locality chooses policy R with probability $\sigma_{i(2)} = \frac{(1-p+pq)\gamma-2(1+\alpha)q}{(1-p)\gamma-2(1-p)(1+\alpha)q}$ and chooses policy Q with probability $1-\sigma_{i(2)}$, but $\sigma_{i(2)} > 1$ for any α and p. The third strategy cannot be supported because the incompetent locality mixes between policy Q and policy R if and only if the center reviews $x = R$ with probability $\rho_{(3)} = \frac{\kappa-(\kappa+2)q}{(\kappa-2)q}$ and does not review with probability $1-\rho_{(3)}$, but $\rho_{(3)} > 1$ for any $\kappa > 2$ and $q < 0.5$. The fourth strategy cannot be supported because the incompetent locality mixes between policy Q and policy R if and only if the center reviews $x = R$ with probability $\rho_{(4)} = \frac{2-q-\kappa}{2-3q}$ and does not review with probability $1-\rho_{(4)}$, but $\rho_{(4)} < 0$ for any $\kappa > 2$ and $q < 0.5$.

For the first strategy, there exist mixed-strategy equilibrium as follows:

Proposition 4. *If $q < 0.5$ and review cost is sufficiently low, $\gamma < 4q(1-q)$,*

(1) When the disapproval penalty is mild ($\kappa < \bar{\kappa} \equiv \frac{2(1-q)}{q}$), there is a mixed-strategy PBE where the center reviews $x = Q$ with probability $\rho^*_{(1)} = \frac{2-2q-q\kappa}{2-2q+q\kappa} \in (0,1)$ and approves it unless falsified from review, and reviews $x = R$ with certainty; the competent locality chooses a policy that matches the state, and the incompetent locality chooses policy R with probability $\sigma^*_{i(1)} = \frac{(1-p+pq)\gamma-2(1-p)(1-\alpha)(1-q)}{(1-p)\gamma-2(1-p)(1-\alpha)(1-q)} \in (0,1)$ and chooses policy Q with probability $1-\sigma^*_{i(1)}$.

Proof. Consider an equilibrium where the competent locality chooses a policy that matches the state (i.e., σ^*_c=1), and the incompetent locality chooses policy R with probability $\sigma^*_{i(1)} \in (0,1)$ and chooses policy Q with probability $1-\sigma^*_{i(1)}$. The center's belief d_{mixQ} =Pr$[x = \omega | x = Q]$ can be written as:

$$d_{mixQ} = \frac{pq + (1-p)(1-\sigma^*_{i(1)})q}{pq + (1-p)(1-\sigma^*_{i(1)})}.$$

Upon observing $x = Q$ but not knowing the state, the center approves $x = Q$ if and only if $d_{mixQ} \geq 0.5(1 - \alpha)$. Suppose it holds, such that the center approves $x = Q$ unless falsified (i.e., $\chi^{-*} = 1$; $\chi^{+*} = 1$ if $\omega = x$, $\chi^{+*} = 0$ if $\omega \neq x$). For the incompetent locality to mix, it must be indifferent between choosing policy R and choosing policy Q, which is satisfied when $\rho^*_{(1)} = \frac{2-2q-q\kappa}{2-2q+q\kappa}$. Given this strategy, the competent locality's best response is always $\sigma^*_c = 1$. For the incompetent locality to mix, it must be indifferent between choosing policy R and choosing policy Q, which is satisfied when $\rho^*_{(1)} = \frac{2-2q-q\kappa}{2-2q+q\kappa}$. For the center to mix, it must be indifferent between reviewing and not reviewing, which is satisfied when $\gamma = 2(1-\alpha)(1-d_{mixQ})$, or to rearrange, $\sigma^*_{i(1)} = \frac{(1-p+pq)\gamma-2(1-p)(1-\alpha)(1-q)}{(1-p)\gamma-2(1-p)(1-\alpha)(1-q)}$. Therefore, $(\rho^*_{(1)}, \chi^{+*}, \chi^{-*})$ and $(\sigma^*_c, \sigma_{i(1)})$ are mutually best responses and sequentially rational. When $\kappa > \bar{\kappa} \equiv \frac{2(1-q)}{q}$, $\rho^*_{(1)} \in (0, 1)$ always holds. For any values of γ and q that satisfy $\gamma < 4q(1-q)$, there always exists $\sigma^*_{i(1)} \in (0, 1)$ that satisfies $d_{mixQ} \geq 0.5(1 - \alpha)$. □

This guideline, however, is never optimal for the center because it is strictly dominated by the (safe, gray) guideline in which the center always reviews $x = R$ (the *micromanage equilibrium* in the main text). Therefore, the center has no incentive to provide the (gray, gray) guideline for any possible review and approval strategies.

From above, relaxing the assumption that Q is always safe adds two observations. First, the center receives the same payoff as the (safe, safe) guideline when it provides the (prohibited, prohibited) or (gray, safe) guidelines. These guidelines become optimal under the same conditions as the (safe, safe) guidelines, which does not alter the parameter space for the *ambiguity equilibrium* under the (safe, gray) guideline in the main text. Second, the center receives the same payoff as the option of investigating upfront as analyzed in Appendix B when it provides the (prohibited, gray) or (gray, prohibited) guideline. These guidelines become optimal under the same conditions as the option of investigating upfront, which shrinks the parameter space for the *ambiguity equilibrium* but does not alter the main result stated in Proposition 3.

C. The Option to Investigate Upfront

Suppose that the center has the option to investigate the state of the world upfront and dictate the correct policy to the locality. The center's expected utility for this option is

$$EU_C(\textit{investigation upfront}) = -\gamma + \{q(1+\alpha) + (1-q)(1-\alpha)\}.$$

First consider the mild disapproval penalty range, $\kappa < \bar{\kappa}$. In this range, the option of investigating upfront is strictly dominated by not investigating, designating policy R as gray, and invoking the *micromanage equilibrium*, for any values of p and α. Intuitively, both options guarantee that the center adopts the correct policy. If the center investigates upfront, it pays an extra cost of investigation for the occasion where the competent locality chooses policy Q, which could be saved under the *micromanage equilibrium*. Therefore, adding the option of investigating upfront does not alter Proposition 3.

Now consider the harsh disapproval penalty range, $\kappa > \bar{\kappa}$. In this range, the *micromanage equilibrium* does not exist. Comparing the center's possible equilibrium payoffs yield the following figure, a modification of Figure 3(b) in the main text.

Figure A1 suggests that the option of investigating upfront shrinks the parameter space for the *ambiguity equilibrium*. However, this does not alter Proposition 3. Notice that this new option is optimal for the center when α and p are sufficiently low. Intuitively, the center has more incentives to investigate upfront when its stakes for correctly implementing a policy (normalized to 1 in the model) relative to α is sufficiently high. Also, the center has more incentives to investigate upfront and dictate the correct policy—that is, not delegating the authority to the locality—when the proportion of competent localities is sufficiently low.

Figure A1: The optimal guidelines for the center with the option to investigate upfront

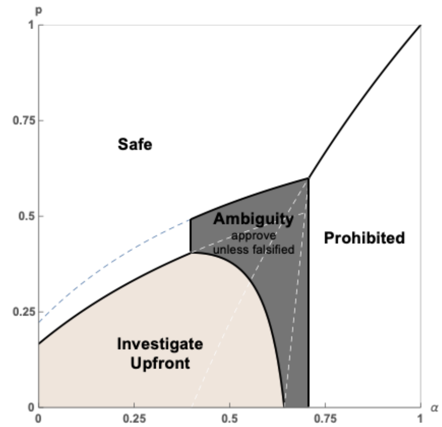

Notes: The x-axis, α, is the center's ideological bias against privatization. The y-axis, p, is the proportion of competent localities who are capable of adopting an appropriate policy.

D. Imperfect Review

In the main text, I assume that the center's review is costly but perfect. That is, whenever the center pays the cost to review, it learns the state of the world ω with certainty. Here I relax the assumption and analyze a situation where the center's review is not perfect. Suppose that when the center reviews, it learns the state of the world ω with probability $\lambda \in (0, 1)$ and fails to learn the state with probability $1 - \lambda$.

Analysis

The center's intervention

First, approval decision. When the center reviews and learns the state, it approves if and only if $\omega = R$. When the center does not review *or reviews but fails to learn*, it approves $x = R$ if and only if $d \geq 0.5(1+\alpha)$, where d is defined the same as the Equation 1 in the main text.

Second, review decision. If the center will approve $x = R$ unless it is falsified from review (i.e., approve if and only if $\omega = R$ when it learns the state; approve $x = R$ when it does not review or reviews but fails to learn), the center reviews if and only if

$$\gamma \leq 2\lambda(1+\alpha)(1-d).$$

If the center will disapprove $x = R$ unless verified from review (i.e., approve if and only if $\omega = R$ when it learns the state; disapprove $x = R$ when it does not review or reviews but fails to learn), the center reviews if and only if

$$\gamma \leq 2\lambda(1-\alpha)d.$$

The locality's policy choice

Let us begin by the competent locality's strategy. We know that it chooses $x = Q$ whenever it observes $\omega = Q$. The only question about its decision is when it observes $\omega = R$. If the center will approve R unless it is falsified from review, the competent locality always chooses policy R. If the center will disapprove $x = R$ unless verified from review, the competent locality chooses policy R if and only if

$$\rho\lambda \geq \frac{\kappa}{\kappa + 2}.$$

The incompetent locality cannot condition its policy choice on the state of the world ω. When the center will approve $x = R$ unless falsified from review, the incompetent locality chooses policy R if and only if

$$\rho\lambda \leq \frac{2(1 - 2q)}{q(\kappa - 2)}.$$

When the center will disapprove $x = R$ unless verified from review, the incompetent locality chooses policy R if and only if

$$\rho\lambda \geq \frac{\kappa}{(\kappa + 2)(1 - q)}.$$

Equilibrium

Consequences of Ambiguous Guidelines

The same as in the main text, I first specify various equilibria under ambiguous guidelines and identify the optimal guidelines for the center. In what follows, I update the results with consideration of λ and discuss changes from the model without this parameter.

Proposition 5 (Ambiguity). *If $q < 0.5$ and review cost is sufficiently low, $\gamma < 4\lambda q(1 - q)$,*

(1) When disapproval penalty is mild ($\kappa < \underline{\kappa} \equiv \frac{2\lambda(1-q)}{1-\lambda(1-q)}$), there is a mixed-strategy PBE where the center reviews $x = R$ with probability $\tilde{\rho}^* = \frac{\kappa}{\lambda(1-q)(\kappa+2)} \in (0, 1)$ and disapproves it unless

verified from review; the competent locality chooses a policy that matches the state, and the incompetent locality chooses policy R with probability $\tilde{\sigma}_i^* = \frac{p(1-q)\{2\lambda(1-\alpha)-\gamma\}}{(1-p)\{\gamma-2\lambda(1-q)(1-\alpha)\}} \in (0,1)$ and chooses policy Q with probability $1-\tilde{\sigma}_i^*$.

(2) *When disapproval penalty is harsh ($\kappa > \bar{\kappa} \equiv 2 + \frac{2(1-2q)}{\lambda q}$), there is a mixed-strategy PBE where the center reviews $x = R$ with probability $\hat{\rho}^* = \frac{2(1-2q)}{\lambda q(\kappa-2)} \in (0,1)$ and approves it unless falsified from review; the competent locality chooses a policy that matches the state, and the incompetent locality chooses policy R with probability $\hat{\sigma}_i^* = \frac{p(1-q)\gamma}{(1-p)\{2q\lambda(1+\alpha)-\gamma\}} \in (0,1)$, and chooses policy Q with probability $1-\hat{\sigma}_i^*$.*

Proof. Proofs to propositions 4-6 are omitted. □

Figure A2: The center's review and approval strategy in the *ambiguity equilibria* with imperfect review

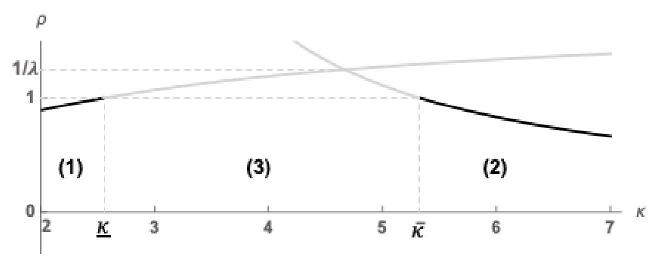

Notes: The x-axis, $\kappa > 2$, is the disapproval penalty. The y-axis, ρ, is the probability of review by the center. The threshold values are $\underline{\kappa} \equiv \frac{2\lambda(1-q)}{1-\lambda(1-q)}$ and $\bar{\kappa} \equiv 2 + \frac{2(1-2q)}{\lambda q}$. For this figure, I set $q = 0.3$ and $\lambda = 0.8$. For any $q < 0.5$ and λ, $\bar{\kappa} > \underline{\kappa}$ holds. Range (1) exists only if $\lambda(1-q) > 0.5$, given $\kappa > 2$.

Compared to the model in the main text, there exist two thresholds $\underline{\kappa}$ and $\bar{\kappa}$ that create three ranges in the disapproval penalty: mild, moderate, and harsh. The *ambiguity equilibria* exist only when the disapproval penalty is mild or harsh.

Proposition 6 (Micromanage). *If $q < 0.5$ and review cost is sufficiently low, $\gamma < 4\lambda q(1-q)$,*

(1) When disapproval penalty is mild ($\kappa < \underline{\kappa}$), there is a pure-strategy PBE where the center always reviews $x = R$ and disapproves it unless verified, the competent locality chooses a policy that matches the state, and the incompetent locality chooses policy R.

(2) When disapproval penalty is moderate ($\kappa < \underline{\kappa}$) or moderate ($\underline{\kappa} < \kappa < \bar{\kappa}$), there is a pure-strategy PBE where the center always reviews $x = R$ and approves it unless falsified, the competent locality chooses a policy that matches the state, and the incompetent locality chooses policy R.

(3) When disapproval penalty is harsh ($\kappa > \bar{\kappa}$), there is a pure-strategy PBE where the center always reviews $x = R$ and disapproves it unless verified, and both competent and incompetent types of locality choose policy Q.

Unlike the model in the main text, two kinds of micromanage equilibrium, approve-unless-falsified and disapprove-unless-verified, exist in the different ranges of the disapproval penalty. Also, there exists a new equilibrium where the center always reviews and the locality never chooses R. Combining Propositions 1 and 2 yields the following result:

Corollary 2. *If $q < 0.5$ and review cost is sufficiently low, $\gamma < 4q(1-q)$,*

(1) When the disapproval penalty is mild ($\kappa < \underline{\kappa}$), the ambiguity equilibrium (disapprove unless verified) and the micromanage equilibria (both approve unless falsified and disapprove unless verified) exist.

(2) When the disapproval penalty is moderate ($\underline{\kappa} < \kappa < \bar{\kappa}$), the micromanage equilibrium (approve unless falsified) exists.

(3) When the disapproval penalty is harsh ($\kappa > \bar{\kappa}$), the ambiguity equilibrium (approve unless falsified) and the micromanage equilibrium (disapprove unless verified; the locality never chooses R) exist.

Desirability of Ambiguous Guidelines

The center's equilibrium payoffs under clear guidelines (i.e., designating R as safe or prohibited) remain the same, but those under gray guidelines are updated as follows. In this section, I omit the micromanage equilibrium under harsh penalty where the locality never chooses R, which yields the same payoff as when R is prohibited.

$$EU'_C(micromanage, auf) = qp(1+\alpha) + q(1-p)(-\gamma + (1+\alpha)(2\lambda - 1)) + (1-q)(-\gamma + 1 - \alpha),$$

$$EU'_C(micromanage, duv) = qp(1+\alpha) + q(1-p)(-\gamma + 1 + \alpha) + (1-q)(-\gamma + (1-\alpha)(2\lambda - 1)),$$

$$\begin{aligned} EU'_C(ambiguity, duv) =& q(1+\alpha) + (1-q)(-1+\alpha), \\ EU'_C(ambiguity, auf) =& qp(1+\alpha) + \\ & q(1-p)[\hat{\sigma_i^*}\{\hat{\rho^*}\lambda(1+\alpha) + (1-\hat{\rho^*}\lambda)(-1-\alpha) - \hat{\rho^*}\gamma\} + (1-\hat{\sigma_i^*})(1+\alpha)] \\ & (1-q)p(1-\alpha-\hat{\rho^*}\gamma) + \\ & (1-q)(1-p)[\hat{\sigma_i^*}(1-\alpha-\hat{\rho^*}\gamma) + (1-\hat{\sigma_i^*})(-1+\alpha)]. \end{aligned}$$

In Figure A3, shaded areas shrink compared to Figure 3 in the main text. This is because the value of review decreases, as the center's capacity to learn from review is not perfect (i.e., $\lambda < 1$). Nevertheless, the main result on the desirability of ambiguous guidelines (Proposition 3 in the main text) does not change.

Figure A3: The optimal guidelines for the center with imperfect review

(a) Mild penalty: $1 < \kappa < \underline{\kappa}$

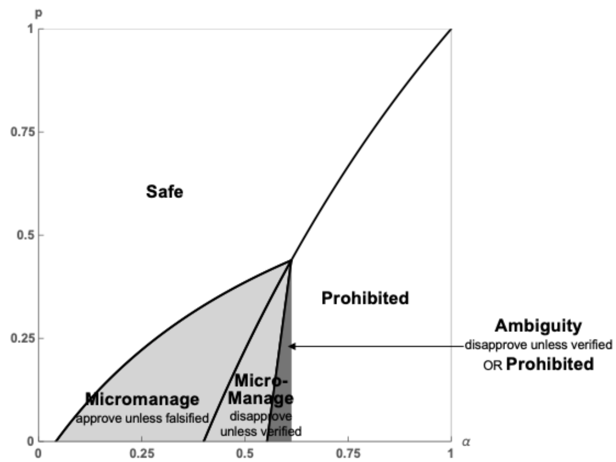

(b) Harsh penalty: $\kappa > \bar{\kappa}$

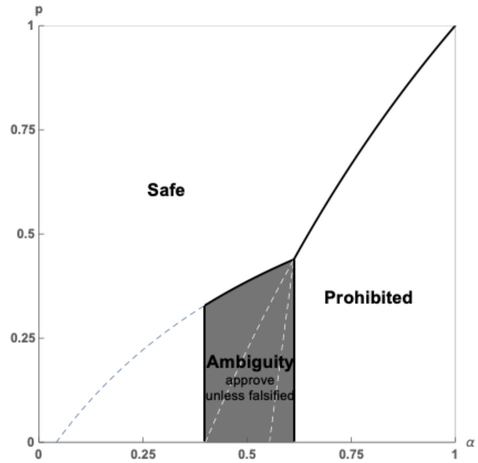

(c) Moderate penalty: $\underline{\kappa} < \kappa < \bar{\kappa}$

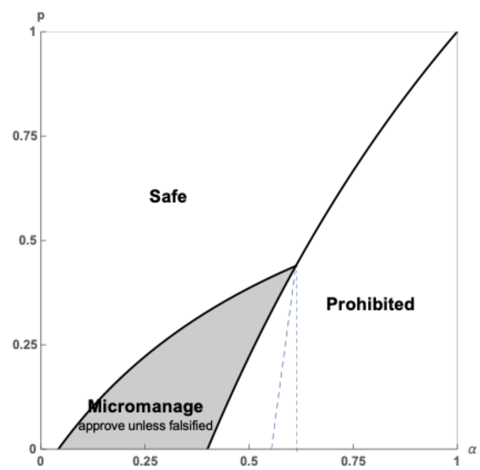

Notes: The x-axis, α, is the center's ideological bias against policy R. The y-axis, p, is the center's prior probability that the locality is competent. Shaded areas represent designating policy R as gray. For all figures, I set $q = 0.3$, $\lambda = 0.8$, and $\gamma = 0.5$.

Remark. The same as the model in the main text, between the two strategies for the *ambiguity equilibria*, the approve-unless-falsified strategy combined with the harsh disapproval penalty is a more effective tool for the center than the disapprove-unless-verified strategy combined with mild disapproval penalty. Furthermore, notice that latter option is not always available, given that the mild disapproval penalty range exists only if $\lambda(1-q) > 0.5$, or when λ is sufficiently high. In contrast, the former option is always available as long as κ is sufficiently high. Although κ is an exogenous parameter in the model, *if the center has leverage in choosing the size of the disapproval penalty* κ, it can always invoke the this equilibrium even if its learning capacity λ is limited. This observation lends additional support to the idea that ambiguous guidelines are most desirable when combined with high disapproval penalty, making the locality choose the reform policy at their own risk.

E. Political Selection and Career Incentives

This section extends the model to incorporate the center's incentives to recruit a competent locality and the locality's incentives to stay in office, and check whether ambiguity is still a useful screening tool for the center. Suppose that the center has the option to retain or replace the locality after policy outcomes are realized—that is, after the center learns about whether the final policy matches the state ω or not. If the center retains the locality, the same locality will stay in office for future periods (which I do not model explicitly) with its type unchanged. If replaces, a new locality, whose type is competent with probability $p \in (0,1)$ and incompetent with probability $1-p$, will come into office for future periods. The center receives an additional payoff of $\pi(\cdot) \geq 0$ that increases as a function of the expected competence of the locality in future periods, $\dot{\theta}$; for simplicity, let $\pi(\cdot)$ be $\pi\dot{\theta}$. The locality receives an additional payoff of $\beta \geq 0$ if it stays in office, and 0 otherwise.

Analysis

The Center's Recruitment

By backward induction, I start from the center's decision to retain or replace the locality. If retains, the expected competence of the locality in future periods $\dot{\theta}$ equals to the center's updated belief about the locality's type in the current period, which are formed as below. If the center replaces the locality, the expected competence of the locality in future periods is simply p. Then,

Lemma 1. *The center retains the current locality if and only if the center's updated belief that the current locality is competent is greater than or equal to* p*.*

Proof. It directly follows that $\pi(\cdot)$ strictly increases as a function of the expected competence of the locality in future periods. □

First, suppose that the center designates policy R as safe. The competent locality chooses a policy that matches the state, and the incompetent locality chooses policy R. The center observes the locality's policy choice x, and once the outcomes (i.e., whether the final policy matches the state) are realized, the center can directly infer the state ω. By the Bayes's rule, the center's beliefs about the locality's type are formed as follows:

$$Pr(\theta = competent|x = R, \omega = Q) = 0,$$

$$Pr(\theta = competent|x = R, \omega = R) = p,$$

$$Pr(\theta = competent|x = Q, \omega = Q) = 1,$$

$$Pr(\theta = competent|x = Q, \omega = R) = \ddot{\theta} \text{ (unconstrained)}.$$

Second, suppose that the center designates policy R as prohibited. Both competent and incompetent types of the locality choose policy Q. By the Bayes's rule, the center's beliefs are (excluding the beliefs that are unconstrained)

$$Pr(\theta = competent|x = Q, \omega = Q) = p,$$
$$Pr(\theta = competent|x = Q, \omega = R) = p.$$

Third, suppose that the center designates policy R as gray invokes the *micromanage equilibria*. The competent locality chooses a policy that matches the state, and the incompetent locality chooses policy R. The center's beliefs are the same as when the center designates policy R as safe.

Finally, suppose that the center designates policy R as gray and invokes the *ambiguity equilibria*. For the "disapprove-unless-verified" kind, the competent locality chooses a policy that matches the state, and the incompetent locality chooses policy R with probability $\tilde{\sigma}_i^*$ and chooses policy Q with probability $\tilde{\sigma}_i^*$. For the "approve-unless-falsified" kind, everything is the same except that $\hat{\sigma}_i^*$ replaces $\tilde{\sigma}_i^*$. By the Bayes's rule, the center's beliefs are

$$Pr(\theta = competent|x = R, \omega = Q) = 0,$$
$$Pr(\theta = competent|x = R, \omega = R) = \frac{p}{p + (1-p)\sigma^*} > p,$$
$$Pr(\theta = competent|x = Q, \omega = Q) = \frac{p}{p + (1-p)(1-\sigma^*)} > p,$$
$$Pr(\theta = competent|x = Q, \omega = R) = 0,$$

where σ^* denotes $\tilde{\sigma}_i^*$ and $\hat{\sigma}_i^*$, respectively, for each kind.

Lemma 2. *It is sequentially rational for the center to retain the locality if and only if the policy choice x matches the state ω, and replaces it otherwise.*

Proof. It directly follows Lemma 1 and the center's beliefs under each scenario. □

Given the center's retention and replacement strategy, the locality may expect to receive the payoff of β whenever it chooses a policy x that matches the state ω. The locality's strategies under R being safe or prohibited are the same as the main model. Here I focus on analyzing the case where the center designates R as gray.

The Center's Intervention

The center's review and approval decisions are the same as the main model. Although in the new setup the center has an additional payoff from recruiting a competent locality in future periods, the center's expected additional payoff is the same regardless of its review and approval decisions.

The Locality's Policy Choice.

The locality's utility can be written as:

$$U_L = \Phi + (1 - a)(-\kappa) + \beta \times \mathbf{1}_{x=\omega},$$

where $\Phi \in \{-1, 1\}$ denotes whether the final policy matches the state ω: 1 if matches, and -1 otherwise. It receives β if and only if $x = \omega$ by Lemma 2.

First, consider the competent locality. We know that it always chooses policy Q when it observes $\omega = Q$. The only question about its decision is when it observes $\omega = R$. If $d \geq 0.5(1+\alpha)$, or the center approves unless falsified, the competent locality always chooses policy R because regardless of whether the center reviews or not its policy choice will never be disapproved. If $d < 0.5(1 + \alpha)$, or the center disapproves unless verified, the competent locality now needs to consider the possibility of disapproval in choosing policy R. It chooses policy R if and only if its expected utility for choosing policy R, which is $\rho(1+\beta)+(1-\rho)(-1-\kappa+\beta)$, is greater than its utility for choosing policy Q, which is -1. Therefore, when the center disapproves unless verified and $\omega = R$, the competent locality chooses policy R if and only if

$$\rho \geq \frac{\kappa-\beta}{\kappa+2}.$$

As β increases, the competent locality is more inclined to choose R, the correct policy.

Now consider the incompetent locality. When $d \geq 0.5(1+\alpha)$, or the center approves unless falsified, the incompetent locality chooses policy R if and only if its expected utility for choosing policy R, which is $\rho\{q(1-\kappa)+(1-q)(1+\beta)\}+(1-\rho)\{q(-1)+(1-q)(1+\beta)\}$, is greater than its expected utility for choosing policy Q, which is $q(1+\beta)+(1-q)(-1)$. Therefore, the incompetent locality chooses policy R if and only if

$$\rho \leq \frac{(2+\beta)(1-2q)}{q(\kappa-2)}.$$

When $d < 0.5(1+\alpha)$, or the center disapproves unless verified, the incompetent locality chooses policy R if and only if its expected utility for choosing policy R, which is $\rho\{q(1-\kappa)+(1-q)(1+\beta)\}+(1-\rho)(q(1-\kappa)+(1-q)(-1-\kappa+\beta))$, is greater than its expected utility for choosing policy Q, which is $q(1+\beta)+(1-q)(-1)$. Therefore, the incompetent locality chooses policy R if and only if

$$\rho \geq \frac{\kappa-\beta(1-2q)}{(\kappa+2)(1-q)}.$$

Above constraints suggest that when $q > 0.5$ the incompetent locality never chooses R, and when $q < 0.5$ as β increases it is more inclined to choose R, the more likely correct policy.

Equilibrium

Consequences of Ambiguous Guidelines

The same as in the main text, I first specify various equilibria under ambiguous guidelines and identify the optimal guidelines for the center. In what follows, I update the results with consideration of π and β and discuss changes from the model without these parameters.

Proposition 7 (Ambiguity). *If $q < 0.5$ and review cost is sufficiently low, $\gamma < 4q(1-q)$,*

(1) When the disapproval penalty is mild ($\kappa < \bar{\kappa} \equiv \frac{2(1-q)+\beta(1-2q)}{q}$), there is a mixed-strategy PBE where the center reviews $x = R$ with probability $\tilde{\rho}^* = \frac{\kappa-\beta(1-2q)}{(1-q)(\kappa+2)} \in (0,1)$ and disapproves it unless verified from review; the competent locality chooses a policy that matches the state, and the incompetent locality chooses policy R with probability $\tilde{\sigma}_i^* = \frac{p(1-q)\{2(1-\alpha)-\gamma\}}{(1-p)\{\gamma-2(1-q)(1-\alpha)\}} \in (0,1)$ and chooses policy Q with probability $1-\tilde{\sigma}_i^*$.

(2) When the disapproval penalty is harsh ($\kappa > \bar{\kappa}$), there is a mixed-strategy PBE where the center reviews $x = R$ with probability $\hat{\rho}^* = \frac{(2+\beta)(1-2q)}{q(\kappa-2)} \in (0,1)$ and approves it unless falsified from review; the competent locality chooses a policy that matches the state, and the incompetent locality chooses policy R with probability $\hat{\sigma}_i^* = \frac{p(1-q)\gamma}{(1-p)\{2q(1+\alpha)-\gamma\}} \in (0,1)$ and chooses policy Q with probability $1-\hat{\sigma}_i^*$.

Compared to Proposition 1 in the main text, the incompetent locality's strategies $\tilde{\sigma}_i^*$ and $\hat{\sigma}_i^*$ remain the same, while the center's strategies $\tilde{\rho}^*$ and $\hat{\rho}^*$ change. This is because only the locality's payoffs and decision-making have been affected by the new parameters. Proposition 2 (Micromanage) and Corollary 1 in the main text apply the same, except for the γ constraint and the value of $\bar{\kappa}$, which should be updated to as presented in Proposition 7.

Desirability of Ambiguous Guidelines

The center's new equilibrium payoffs (marked with apostrophes) are updated as follows:

$$EU'_C(safe) = EU_C(safe) + p\pi\{1 + (1-p)q\},$$

$$EU'_C(prohibited) = EU_C(prohibited) + p\pi,$$

$$EU'_C(micromanage) = EU_C(micromanage) + p\pi\{1 + (1-p)q\},$$

$$EU'_C(ambiguity, duv) = EU_C(ambiguity, duv) + p\pi\{1 + (1-p)q + (1-p)(1-\tilde{\sigma_i^*})(1-2q)\},$$

$$EU'_C(ambiguity, auf) = EU_C(ambiguity, auf) + p\pi\{1 + (1-p)q + (1-p)(1-\hat{\sigma_i^*})(1-2q)\}.$$

In Figure A4, shaded areas expand compared to Figure 3 in the main text (the higher the value of π, the greater the difference). This is because the center's additional payoffs from retaining or replacing the locality, as indicated above, are the highest when it designates policy R as gray and invokes the *ambiguity* equilibrium (given $q < 0.5$), the next highest when it designates it as gray and invokes the *micromanage* equilibrium or designates it as safe, and the lowest when it designates it as prohibited. With the mild penalty (Figure A4(a)), the *ambiguity* equilibrium now yields the uniquely highest payoff for the center unlike in the main text, while the areas (i.e., the dark-gray triangle) remain exactly the same. With the harsh penalty (Figure A4(b)), the dark-gray areas expand toward lower α values. The main result on the desirability of ambiguous guidelines (Proposition 3 in the main text) does not change. Overall, this extension suggests that ambiguity becomes an even more useful screening tool for the center, once incorporating the issue of political selection and career incentives.

Figure A4: The optimal guidelines for the center with the option to retain or replace the locality

(a) Mild penalty: $\kappa < \bar{\kappa}$

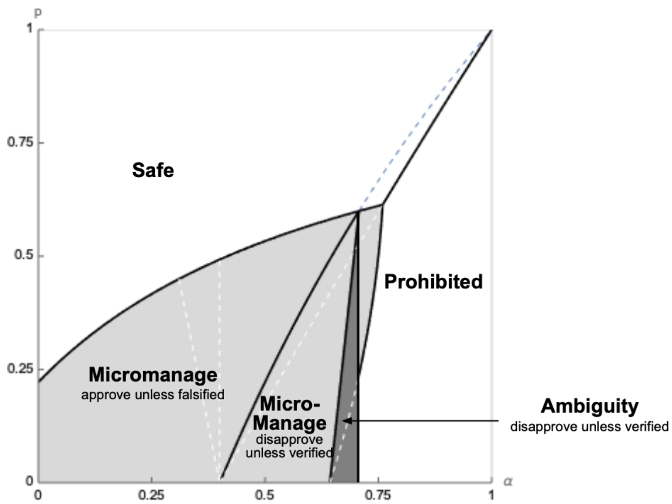

(b) Harsh penalty: $\kappa > \bar{\kappa}$

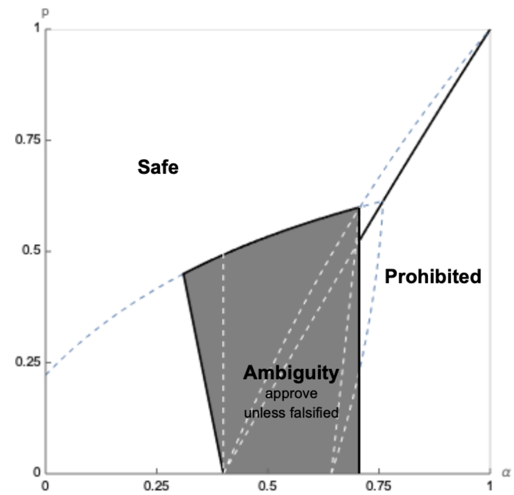

Notes: The x-axis, α, is the center's ideological bias against policy R. The y-axis, p, is the center's prior probability that the locality is competent. Shaded areas represent designating policy R gray. For both figures, I set $q = 0.3$, $\gamma = 0.5$, $\pi = 1$, and $\beta = 1$.