Learning by Lobbying*

Emiel Awad† Gleason Judd‡ Nicolás Riquelme§

Abstract

Effective lobbying requires understanding politicians' preferences, while the act of lobbying itself can reveal those preferences. How does this link between lobbying and learning shape relationships between interest groups and politicians? We develop a game-theoretic model where an interest group can lobby a politician while learning about their ideological alignment. Our analysis highlights strategic tensions where interest groups balance information-gathering against policy influence in their lobbying, while forward-looking politicians manage their reputations to shape future interactions. These forces shape dynamics: policies and transfers shift over time as uncertainty resolves, and early-career politicians can show greater policy variance and extract larger benefits through reputation management than veterans. Politicians with secure positions can receive more favorable treatment due to their stronger incentives to appear less aligned than they truly are. Our results address empirical regularities and provide a theoretical foundation for understanding how lobbying relationships evolve across political careers and institutional contexts.

Forthcoming at *American Journal of Political Science*

*We thank Jordi Blanes i Vidal, Jihwan Do, Daniel Habermacher, Lindsey Gailmard, Heng Liu, Juan S. Pereyra, Patricio Romero, Amna Salam, Keith Schnakenberg, Hye Young You, Emilie Wojcieszynski and audiences at APSA 2022, the 35th Stony Brook International Conference on Game Theory, 2025 Copenhagen Money in Politics Conference, Universidad de Talca, Universidad Católica del Maule, and SECHI 2024 for helpful comments.

†Analyst, Oxera Consulting LLP, (emielawad@gmail.com).

‡Assistant Professor, Department of Politics, Princeton University, (gleason.judd@princeton.edu).

§Assistant Professor, School of Business and Economics, Universidad de los Andes, Chile, (nriquelme@uandes.cl).

Introduction

Information is crucial for effective lobbying. Understanding key policymakers is particularly valuable, even beyond knowledge about policies and procedures.[1] Interest groups gather this information in various ways—researching backgrounds, consulting staff, hiring connected lobbyists—but lobbying itself is particularly useful. Politicians' responses to specific lobbying overtures can reveal preferences that public statements or voting records might not capture.[2] Thus, lobbying and learning are linked: understanding politicians' preferences facilitates effective lobbying, while lobbying can reveal those preferences.

Parsing this link is important for understanding political influence. In lobbying relationships, it creates strategic tensions for interest groups and politicians. Interest groups must decide whether to leverage current information or pursue strategies that yield more information about politicians to lobby them more effectively later on. Meanwhile, politicians may recognize that their policy choices reveal information about their preferences, potentially affecting future lobbying approaches or revolving-door opportunities (McChesney, 1997; Egerod and Tran, 2023).[3] These considerations are especially pronounced with unfamiliar politicians.[4] Additionally, the potential for *learning-by-lobbying* also affects interest groups' broader strategic decisions about which politicians to target, how to learn about them, and how much to invest in relationship-building.

Our central questions confront these fundamental tensions. How does the link between lobbying and learning about politicians shape relationships between interest groups and politicians over time? And how does it interact with other aspects of political influence—such as acquiring information via other means, obtaining access that facilitates lobbying, policymaking constraints, or revolving-door considerations?

We develop a game-theoretic model in which an interest group lobbies and may learn about a politician's preferences over time. In each of two periods, the group lobbies to influence the politician's policy choice. Initially uncertain about the politician's ideology—which might be relatively aligned with the group's goals (*ally*) or more distant (*adversary*)—the group updates its beliefs by observing responses to first-period lobbying before lobbying again in the second period. Meanwhile, the politician strategically manages their reputation,

[1] A seasoned lobbyist emphasizes: "it is not about what you want; it is about what the other person needs" (Levine, 2008, p. 163).

[2] Kerr, Lincoln and Mishra (2014, p. 344) note: "firms may gain from learning about policymakers' private dispositions, which may not be fully reflected in their public positions."

[3] Politicians sometimes attempt to induce more favorable lobbying by posturing via "milker"/"juicer"/"fetcher" bills threatening unfavorable taxes, permitting, price controls, etc. (McChesney, 1997, pgs. 29–32).

[4] Kerr, Lincoln and Mishra (2014, p. 344) note: "the costs of learning and establishing relationships with policymakers are likely to be the highest in a firm's first several years of lobbying."

aware that the group may learn from their choices. Our model highlights the link between learning and lobbying: the first period captures initial interactions under uncertainty, while the second captures later interactions that may evolve with learning.

Our model reflects the frequently granular nature of lobbying, which often addresses specific policy details.[5] The group lobbies by offering a menu of policy-transfer pairs, where transfers represent various forms of political support—campaign contributions, charitable donations, or legislative assistance—that could be relevant in different contexts.[6] The politician sees the offered menu and can either select one pair or set policy independently.

To illustrate a scenario reflected by our model, consider a lobbying interaction involving a trade association and a newly appointed member of a relevant congressional subcommittee. The association must determine how to lobby this politician whose true stance on industry regulation is unclear. It may offer policy packages with varying degrees of regulatory change paired with different levels of industry support. Meanwhile, the politician's response may indicate how closely their stance on this particular issue aligns with the group's. Then, if the regulatory issue were to recur later on with the politician in a relevant position, the trade association would have a clearer understanding of how to lobby the politician.

Our setting highlights two key strategic considerations. First, the interest group can face a tension between influence and learning: foregoing immediate policy gains may help them learn about politicians' true preferences for more effective future lobbying. Second, the politician can have reputational concerns: even when aligned with the interest group, they may strategically act more adversarial to extract better terms later. We analyze how these strategic tensions between influence, learning, and reputation interact over time to shape distinctive patterns in lobbying relationships.

The interest group's equilibrium lobbying strategy varies with the politician's reputation. Under broad conditions, first-period lobbying distinguishes whether the politician is an ally or adversary. The group's particular screening approach can take several distinct forms. If it believes the politician is likely an ally, the first-period menu pairs terms attractive to allies with modest demands for potential adversaries. Conversely, when it suspects an adversary, the first-period menu pairs substantial demands for adversaries with terms that are especially generous for allies.[7] The group foregoes learning, however, if it is not very forward-looking and the politician is likely an adversary—it simply offers terms that would be efficient with

[5] A veteran lobbyist notes, "What matters is getting stuff put in the bill, a line here, a line there... What you're looking to do is put a line in a law, get something tweaked. You're looking to change this line in subsection B. You just need one person to make that change" (Drutman, 2015, p. 31).

[6] For more discussion and examples, see McChesney (1997, pgs. 45–54).

[7] *Generous terms* refer to contracts the politician strictly prefers to accept rather than choose policy independently.

an adversary, which the politician would accept regardless of whether they are an ally or adversary.

The politician's reputation concerns affect the interest group's incentives to distinguish whether they are aligned. Stronger reputational concerns make the politician more inclined to misrepresent their preferences, so screening is costlier. Specifically, an ally would have stronger incentive to appear an adversary, so the group must make the adversary offer less appealing to an ally. Depending on the group's beliefs about the politician, it will either offer increasingly generous terms for allies or make the adversary offer less aggressive. Under some conditions, the group's menu might even include an offer to pay an adversary to shift policy in the opposite direction, in order to deter mimicry by an ally.

Our equilibrium characterization provides insights into lobbying relationships.

First, policies and transfers shift as interest groups learn about politicians or reputational incentives dissipate. If the group initially expects an ally, first-period lobbying includes an overly conservative option intended for adversaries, and if the politician chooses that option, then second-period lobbying is more aggressive (i.e., paying more to shift policy closer to the interest group's preferred outcomes). If the group initially expects an adversary, first-period lobbying includes an option that is overly generous to potential supporters, and if that is chosen then second-period lobbying is less generous. Essentially, initial lobbying that resolves learning and reputational incentives early on will facilitate later lobbying that is more targeted and not distorted by potential posturing. Thus, our learning-by-lobbying mechanism can induce more unpredictable and divergent policy choices by early-career politicians than late-career politicians.

Second, politicians with secure positions can receive more favorable early-career lobbying due to stronger reputational concerns with interest groups. They have stronger incentives to appear less aligned than they truly are to be lobbied more favorably in the future, forcing groups to offer more generous terms to screen them. Thus, different political positions can create distinct patterns in how lobbying relationships evolve. Our results suggest that politicians with more secure offices (committee chairs, favorable constituencies) might receive disproportionately favorable lobbying terms compared to those facing greater electoral uncertainty.

Our extensions study four key aspects of relationships. First, analyzing the value of early-career information reveals that understanding politicians' preferences is more valuable when politicians are patient (as this increases learning costs) and weakly more valuable when interest groups are forward-looking (until they are patient enough to learn the politician's preferences by lobbying). Second, the value of early-career access that provides lobbying opportunities varies with learning and reputational incentives, being lower if politicians are

likely to remain in the position. Third, if the group is not sufficiently confident that the politician is an ally, then more valuable revolving-door opportunities make politicians' reputational incentives less adversarial and learning-by-lobbying is less costly for the group. Fourth, if the politician's scope for setting policy is more constrained, then learning-by-lobbying is costlier and the group is less inclined to pursue it.

Our analysis speaks to empirical patterns in interest group activities. The learning-by-lobbying mechanism helps understand why experience with specific policymakers is valuable, why early interactions are particularly formative, why allies may receive especially favorable terms, and how lobbying relationships evolve. In extensions, we show how this mechanism's role in some of those patterns can vary with other aspects of the broader political landscape through revolving-door considerations or policymaking constraints. Moreover, our extensions also offer implications for lobbying-adjacent activities like background research (for information), campaign contributions (for access), or hiring staffers (for both).

Relationship to Existing Literature

We contribute to understanding lobbying relationships and political influence.[8] Experience with specific policymakers is valuable for lobbying (Kerr, Lincoln and Mishra, 2014; Drutman, 2015), and we illuminate this relationship by modeling how lobbying relationships evolve as interest groups learn about politicians through repeated lobbying. We depart from prior work by combining two key features: incomplete information about politicians and the potential for repeated lobbying. Classic theories with complete information in one-shot exchanges (Hall and Wayman, 1990; Grossman and Helpman, 1994; Besley and Coate, 2001) illuminate static influence but cannot address dynamics or learning. Static models with incomplete information about politicians' preferences (Martimort and Semenov, 2008; Kolotilin et al., 2017; Schnakenberg, 2017; Minaudier, 2022) capture uncertainty but not its evolution. Dynamic models with complete information (Iaryczower and Oliveros, 2017; Chen and Zápal, 2022; Bils, Duggan and Judd, 2021) address temporal aspects but not learning or reputation.

Our *learning-by-lobbying* mechanism provides new perspective on established patterns in lobbying relationships. While policy or procedural expertise and relationship building can each involve information (LaPira and Thomas III, 2017), our analysis emphasizes how interest groups can learn about politicians' preferences through the lobbying process. This sheds new light on why early interactions with politicians are so formative and why lobbying patterns

[8]For empirical work on the dynamics of lobbying, see Kerr, Lincoln and Mishra (2014). For broader overviews of theoretical and empirical work, see Grossman and Helpman (2001), Bombardini and Trebbi (2020), and Schnakenberg and Turner (2024).

evolve systematically as uncertainty resolves (Kerr, Lincoln and Mishra, 2014). Furthermore, we illuminate how learning-by-lobbying can impact other activities by affecting the value of different relationships and how they should be managed.

Our focus on interest groups learning about politicians provides a novel perspective on information flows in lobbying relationships. Existing work has studied how politicians learn about interest groups (Groll and Ellis, 2014, 2017; Ellis and Groll, 2024) or use lobbyists as gatekeepers to screen interest groups (Hirsch et al., 2023), but we address the other direction of learning. Our setup aligns with observations that beliefs about legislators' preferences are updated through lobbying (Austen-Smith and Wright, 1992), but provides a formal account of how this updating affects strategic behavior on both sides. By modeling this two-sided relationship, we capture how interest groups' learning incentives interact with politicians' reputational concerns (McChesney, 1997; Egerod and Tran, 2023).

We model lobbying as an *exchange* of policy for transfers (Grossman and Helpman, 1994) in the menu-auction tradition (Bernheim and Whinston, 1986; Grossman and Helpman, 2001). Our setup captures the granular nature of real-world influence and encompasses various lobbying tactics. While lobbying can be modeled in various other ways—e.g., legislative subsidies (Hall and Deardorff, 2006) or information transmission (Schnakenberg, 2017; Awad, 2020)—a common theme is that interest groups aim to influence politicians' behavior, with tactics depending on both parties' preferences.[9] Our approach reflects this theme and is particularly suited for analyzing how interest groups adjust their strategies as they learn about politicians' preferences.

Our extensions contribute to broader understanding of interest group politics by parsing how learning-by-lobbying forces interact with other influence-oriented activities. First, our extension on the value of information sheds new light on why revolving-door lobbyists are valuable and how their value can vary depending on which politicians they know about, complementing the role of lobbyists as a certification tool (Hirsch et al., 2023). Second, our extension on value of access highlights how interest groups' desire to cultivate relationships that facilitate lobbying can depend on politicians' job security, career stage, or familiarity, complementing the role of relative ideology (Austen-Smith, 1995; Schnakenberg, 2017; Awad, 2020; Judd, 2023) or proposal power (Judd, 2023). Third, our revolving-door extension clarifies how lobbying and reputation can interact to shape politicians' observed in-office behavior before potentially revolving, complementing politician-centric theories (Che, 1995; Bar-Isaac and Shapiro, 2011; Hübert, Rezaee and Colner, 2023). Finally, our policy-constraints extension addresses how voting rules and collective policymaking can shape how learning and reputation affect lobbying, complementing complete-information theories of lobbying in col-

[9] See Schnakenberg and Turner (2024) for an overview.

lective bodies (Grossman and Helpman, 2001; Judd, 2023).

Finally, our model of learning, reputation, and lobbying influence descends from traditional models of repeated contracting (Hart and Tirole, 1988; Laffont and Tirole, 1990; Salanié, 2005). These models typically feature a *ratchet effect* (Freixas, Guesnerie and Tirole, 1985), where players have incentives to conceal their true characteristics to avoid future exploitation of this information. This incentive also arises in our setting, where appearing adversarial can induce more favorable lobbying. Yet, two key aspects of our lobbying and policymaking setting lead to different equilibrium conditions than existing models of repeated contracting. First, the spatial policy preferences introduce complex interdependencies when offering menus and screening different types that preclude standard single-crossing conditions.[10] Second, the politician can set policy independently if she rejects the group's menu, whereas elsewhere the agent either cannot act after rejection or the game ends. Since these are two prominent political features, our analysis may facilitate future work on political influence.

Model

Players. There are two players: an interest group, G, and a politician, P.

Timing. There are two periods of policymaking. In each period $t \in \{1, 2\}$, P will enact a policy and G can lobby by offering P a menu M_t of policy-transfer pairs, with G choosing both (i) how many pairs to include and (ii) for each pair (x, T), the exact policy $x \in \mathbb{R}$ and transfer $T \geq 0$. Next, P observes M_t and then either selects one pair or rejects all of them. If P selects a pair (x, T) from M_t, then the enacted policy is $x_t = x$ and G transfers T to P. Otherwise, if P rejects the menu, then P chooses x_t freely without receiving a transfer T. Since P's selection determines the realized policy and transfer—that is, we abstract from short-term commitment issues—we follow the literature on menu auctions and refer to each pair in a menu as a *contract*. Accordingly, we denote an arbitrary contract as $c = (x, T)$.

Payoffs. In each period, G's utility function is $\Pi(x, T) := -(x - 1)^2 - T$, where x is the implemented policy, 1 is G's ideal point, and T is the transfer (if accepted). Similarly, P's utility function is $U(x, T) := -(x - \theta)^2 + T$, where θ denotes P's ideal point. Notably, P's ideal point can be either $\underline{\theta}$ or $\overline{\theta}$, so we also refer to θ as P's *type*. We assume $\underline{\theta} < \overline{\theta} < 1$, so that $\underline{\theta}$ is the *adversary* type and $\overline{\theta}$ is the *ally* type, closer to G.

Each player's cumulative payoff sums their utility across both periods. Each player discounts second-period utility with (potentially different) discount factors: $\delta_P, \delta_G \in [0, 1]$.

[10] In seller-buyer or firm-worker environments, this single-crossing condition is typically implicitly assumed, as in, Beccuti and Möller (2018), Gerardi and Maestri (2020) and Breig (2022).

Information. The interest group G does not know P's ideal point, $\theta \in \{\underline{\theta}, \overline{\theta}\}$. At the beginning of the game, G's prior belief puts probability $\mu_0 \in [0,1]$ on $\theta = \overline{\theta}$. After observing P's first-period behavior, G's updated belief about P is denoted μ_1. All other features of the game are common knowledge.

Strategies and Equilibrium Concept. We study Perfect Bayesian Equilibria (PBE) in pure strategies.[11] Thus, strategies are sequentially rational at every information set and Bayes' rule is applied whenever possible. We apply the Never-a-Weak-Best-Response (NWBR) refinement, so off-equilibrium-path beliefs assign higher probability to types more inclined to deviate.[12]

Our results in the main text focus on the players' equilibrium strategies. Formal statements of beliefs and all proofs are provided in the Appendix.

Analysis

Our analysis proceeds in several stages. We first analyze two complete-information benchmark cases: with and without lobbying. We then begin our main analysis, studying lobbying in a static setting with incomplete information, which corresponds to the second period of our model. Working backward, we study the first-period interaction, where lobbying occurs in a dynamic setting with incomplete information. To conclude our main analysis, we analyze how policies and transfers shift over time. Finally, in extensions we study how early-career information, access, policymaking constraints, and revolving-door incentives affect our main insights.

Two benchmarks

To set the stage for our main analysis, we characterize behavior in two benchmark settings: (i) no lobbying and (ii) lobbying with complete information.

Benchmark 1 (No lobbying)**.** If the interest group cannot lobby, then P will set policy at her ideal point, θ, in both periods.

[11] We focus on pure strategies for tractability. Although mixed strategies could yield different equilibrium outcomes (see Bester and Strausz, 2001, for more details), the pure-strategy analysis captures strategic considerations in lobbying relationships. Moreover, our approach follows the precedent set by canonical models using menu auctions (Bernheim and Whinston, 1986).

[12] Either P may reject an on-path offer, leading to an off-path information set, or G may make an off-path offer. In either case, applying NWBR constrains G's off-path beliefs. For an overview of refinements in signaling games, see Fudenberg and Tirole (1991) and Manelli (1997).

Without lobbying, P obtains a payoff of zero in both periods, regardless of type. However, P's type does impact G's payoff, $-(1+\delta_G)(1-\theta)^2$. Specifically, G's payoff is lower if P is an adversary ($\theta = \underline{\theta}$) than if P is an ally ($\theta = \overline{\theta}$).

Benchmark 2 (Lobbying with complete information)**.** If G can lobby and has complete information about P, then in each period P enacts policy $x_\theta := \frac{1}{2}(1+\theta)$ and receives transfer $T_\theta := \frac{1}{4}(1-\theta)^2$.

If G knows P's ideal point, θ, then it can perfectly calibrate its lobbying. Notably, P's ideal point acts as a reservation policy that she will enact if she rejects G's menu. Thus, θ determines both (i) the minimal transfer required for P to deviate from her ideal policy and (ii) G's willingness to provide that minimal transfer. To induce P to enact a policy x, G must offer a transfer of at least $T_\theta(x) := (\theta - x)^2$.

In each period, G optimally balances its marginal benefit of more favorable policy against its marginal cost of increasing the transfer. If P is an ally, then G induces a mild policy shift with a modest transfer. Otherwise, if P is an adversary, G has a larger marginal gain from shifting P's policy and therefore induces a larger shift with a larger transfer.

The politician never receives surplus, since in each period she is indifferent between accepting and rejecting the offered contract. Thus, we refer to $c_{\underline{\theta}} := (x_{\underline{\theta}}, T_{\underline{\theta}})$ as the *adversary-efficient contract*, and analogously define $c_{\overline{\theta}}$ as the *ally-efficient contract*.[13] Accordingly, we define $\pi_{\underline{\theta}} := -(x_{\underline{\theta}}-1)^2 - T_{\underline{\theta}}$ as G's *adversary-efficient payoff* and define $\pi_{\overline{\theta}} := -(x_{\overline{\theta}}-1)^2 - T_{\overline{\theta}}$ as G's *ally-efficient payoff*.

Lobbying with incomplete information

We now begin our main analysis. We study how G's influence is shaped by uncertainty about the legislator's alignment and dynamic strategic considerations. Since we focus on pure strategies, the Revelation Principle implies that it is without loss of generality to focus on menus with at most three options: two 'type-specific' contracts and an 'empty' contract. The type-specific contracts are distinct: an *adversary contract*, $\underline{c} = (\underline{x}, \underline{T})$, intended for an adversary, $\underline{\theta}$, and an *ally contract*, $\overline{c} = (\overline{x}, \overline{T})$, intended for an ally, $\overline{\theta}$. We say a menu is *separating* if $\underline{c} \neq \overline{c}$ and *pooling* if $\underline{c} = \overline{c}$.

Second-period lobbying and policymaking

In the second period, G's sole focus is influence. This parallels Benchmark 2 but now G may not know P's ideal point, θ. Instead, G has updated belief $\mu_1 \in [0,1]$ about θ, where

[13] That is, $c_{\underline{\theta}}$ and $c_{\overline{\theta}}$ are the two possible full-information contracts. Throughout, we refer to equilibrium contracts that differ from these full-information contracts as having *distortions* from the perspective of G.

$\mu_1 := \mu(\theta = \overline{\theta}|h)$ represents the posterior probability that P is an ally given history h.

The interest group wants to induce favorable policy without overpaying or under-lobbying. However, without knowing θ, G's equilibrium menu cannot include both efficient contracts $c_{\underline{\theta}}$ and $c_{\overline{\theta}}$. If such a menu were offered, then an ally would choose the adversary-contract $c_{\underline{\theta}}$, so G would infer they are potentially overpaying an ally to enact *worse* policy: $\underline{x}_2$ instead of $\overline{x}_2 > \underline{x}_2$. More broadly, G's menu must ensure *incentive compatibility* for P to ensure that each type selects its intended contract.

Given P's potential incentive to misrepresent her preference, G will proactively distort the offered menu. Specifically, G's second-period menu will not include both $c_{\underline{\theta}}$ and $c_{\overline{\theta}}$. A distortion—defined as a discrepancy from either the ally-efficient or adversary-efficient contracts—could in principle arise as either a *separating menu* with two type-specific contracts or a *pooling menu* with a single non-empty contract. In principle, a separating menu could become prohibitively expensive for G if excessive distortions are required to induce each type to select the contract intended for them, which would induce G to prefer a pooling menu. However, we show that will not occur in the second period in equilibrium and G will always offer a separating menu.

In G's optimal separating menu, at least one contract will be distorted so that the ally and adversary types choose the contract intended for them. To deter an ally from selecting the adversary-contract, G offers a menu that either (i) overpays an ally to ensure it chooses the ally-efficient policy, $x_{\overline{\theta}}$, or (ii) makes an overly conservative adversary-offer (i.e., $\underline{x}_2 < x_{\underline{\theta}}$ with $\underline{T}_2 = T_{\underline{\theta}}(\underline{x}_2)$), that extracts smaller policy concessions from an adversary but is less appealing to an ally. The particular distortion G chooses depends on μ_1, G's updated belief about P.

If G believes P is probably an ally (i.e., μ_1 is sufficiently high), then G offers a separating menu with only one distortion: an overly conservative adversary-offer. This menu pairs (i) an ally-efficient contract and (ii) an overly conservative adversary-contract. Essentially, G makes the adversary-contract less appealing to an ally to ensure it would choose the ally-efficient contract. Although this sacrifices efficiency in the adversary-contract, that inefficiency is relatively unlikely to materialize. To minimize distortions, G calibrates the menu so that (i) an ally is indifferent between accepting either contract or rejecting them and (ii) an adversary is indifferent between accepting the adversary-contract and rejecting G's menu.

Otherwise, if μ_1 is lower so an adversary is more likely, G prioritizes favorable terms in the adversary-contract and offers a separating menu with two distortions: an excessive ally-transfer *and* an overly conservative adversary-offer. This menu pairs (i) an ally-contract that overpays for the ally-efficient policy $x_{\overline{\theta}}$ with (ii) an adversary-contract that is again overly

conservative but to a lesser degree. The size of these distortions is inversely related to μ_1. As μ_1 decreases, G increasingly prioritizes efficiency in the adversary-contract while offering an increasingly excessive (but less likely to materialize) ally-payment to ensure incentive compatibility.

Lemma 1 characterizes second-period policy and transfers. A key factor is whether G's belief μ_1 is relatively high or low, distinguished by the cutpoint $\tilde{\mu}_1 := \frac{1-\overline{\theta}}{1-\underline{\theta}} \in (0,1)$.

Lemma 1. *The interest group's second-period menu in equilibrium is separating for every second-period history and every belief μ_1.*

1. *If $\mu_1 \leq \tilde{\mu}_1$, then (i) an ally enacts $\overline{x}_2 = x_{\overline{\theta}}$ and receives $\overline{T}_2 = T_{\overline{\theta}}(\overline{x}_2) + \frac{(\overline{\theta}-\underline{\theta})(1-\overline{\theta}-\mu_1(1-\underline{\theta}))}{1-\mu_1}$, and (ii) an adversary enacts $\underline{x}_2 = x_{\underline{\theta}} - \frac{\mu_1}{1-\mu_1}\frac{\overline{\theta}-\underline{\theta}}{2}$ and receives $\underline{T}_2 = T_{\underline{\theta}}(\underline{x}_2)$.*
2. *If $\mu_1 > \tilde{\mu}_1$, then (i) an ally enacts policy $\overline{x}_2 = x_{\overline{\theta}}$ and receives transfer $\overline{T}_2 = T_{\overline{\theta}}$, and (ii) an adversary enacts $\underline{x}_2 = x_{\underline{\theta}} - \frac{1-\overline{\theta}}{2}$ and receives $\underline{T}_2 = T_{\underline{\theta}}(\underline{x}_2)$.*

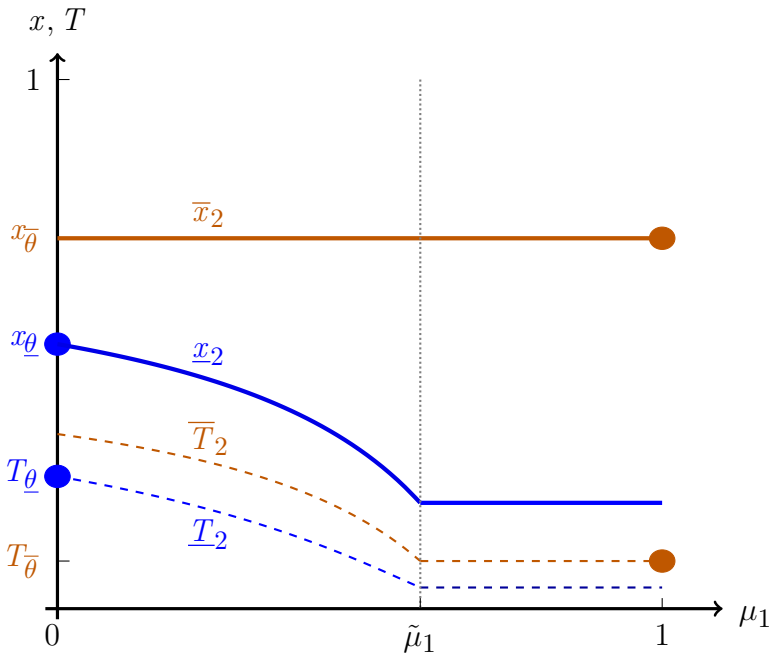
Figure 1 illustrates the findings of Lemma 1 by displaying equilibrium policies and transfers as functions of G's updated belief μ_1. The blue and orange dots illustrate the efficient contracts $c_{\underline{\theta}}$ and $c_{\overline{\theta}}$ that G would offer under complete information ($\mu_1 = 0$ and $\mu_1 = 1$). The ally-contract always contains the same policy offer, but its transfer either decreases with μ_1 or remains constant. In contrast, the adversary's policy- and transfer-offer are always sensitive to G's belief. A key point is that G can only ensure that it offers efficient contracts if it perfectly learns P's alignment. If not, G distorts lobbying, and the degree of distortion generally increases with G's uncertainty about P's alignment. Finally, when comparing the transfers offered to the adversary and ally-type, it is not always the case that the ally (or the adversary) receives a larger transfer. For example, when P is likely an ally, then which type receives a larger transfer depends on the configuration of $\underline{\theta}$ and $\overline{\theta}$, and their proximity to G's ideal point, 1.

First-period lobbying and policymaking

We now analyze how first-period lobbying and policymaking are shaped by the prospect of future lobbying opportunities.

Both the politician and interest group balance first-period incentives against forward-looking considerations about second-period consequences. G wants to influence policy favorably today while also learning about P's preferences to facilitate future lobbying. Meanwhile, P wants to receive favorable terms today while also managing her reputation to induce favorable lobbying later. Each player's static incentive to obtain favorable terms is analogous

Figure 1. Second-Period Policies and Transfers



Note: The figure illustrates equilibrium policies and transfers listed in Lemma 1, while fixing $\underline{\theta} = 0$ *and* $\overline{\theta} = \frac{2}{5}$*. Blue lines represent the equilibrium adversary-contract as a function of* μ_1*, where the dashed line is the adversary-transfer* ($\underline{T}_2$) *and the solid line the adversary-policy* ($\underline{x}_2$)*. Orange lines represent the equilibrium ally-contract, where the dashed line is the ally-transfer* ($\overline{T}_2$) *and the solid line the ally-policy* ($\overline{x}_2$)*. The dots indicate the efficient contracts:* $c_{\underline{\theta}} = (x_{\underline{\theta}}, T_{\underline{\theta}})$ *for* $\mu_1 = 0$ *and* $c_{\overline{\theta}} = (x_{\overline{\theta}}, T_{\overline{\theta}})$ *for* $\mu_1 = 1$*.*

to the second period. That is, all the incentives that are present in the second period (given the same beliefs about P's alignment) are also present in the first period. In the first period, however, their forward-looking incentives—learning and reputation management—introduce new forces above and beyond static incentives. Moreover, these forces are interdependent. On the one hand, G's learning is affected by P's reputational incentives. On the other, P's reputational incentives are affected by P's anticipation of G's learning.

A key factor is how first-period behavior impacts G's updated belief about P, which shapes second-period lobbying and policymaking. How G and P interact in the second period then affects each player's expected future payoffs (i.e., their continuation value following the first period).[14] By learning P's preferences in the first period, G can lobby more effectively in the second compared to lobbying under uncertainty. Remark 1 characterizes how G's continuation value varies with its updated belief, μ_1.

Remark 1. The interest group's continuation value is:

$$V_G(\mu_1) = \begin{cases} (1-\mu_1)(\pi_{\underline{\theta}}) + \mu_1 \left(\pi_{\underline{\theta}} + \frac{(\bar{\theta}-\underline{\theta})^2}{2(1-\mu_1)} \right) & \text{if } \mu_1 \leq \tilde{\mu}_1, \\ (1-\mu_1)\left(\pi_{\underline{\theta}} + \pi_{\bar{\theta}}\right) + \mu_1 \left(\pi_{\bar{\theta}}\right) & \text{if } \mu_1 > \tilde{\mu}_1. \end{cases}$$

Remark 1 implies that G benefits from more precise information. Essentially, G can more confidently tailor its menu of offers towards the politician's likely type (adversary or ally), anticipating that this efficient offer will be chosen. Although the contract for the less likely type is more distorted, it is increasingly unlikely to be chosen. Overall, the net effect makes G better off as its beliefs get more precise. In the limiting case where G has degenerate beliefs, $\mu_1 = 0$ or $\mu_1 = 1$, G's ex ante equilibrium payoff converges to $(1-\mu_0)\,V_G(0) + \mu_0\,V_G(1) = (1-\mu_0)(\pi_{\underline{\theta}}) + \mu_0(\pi_{\bar{\theta}})$ in the second period.

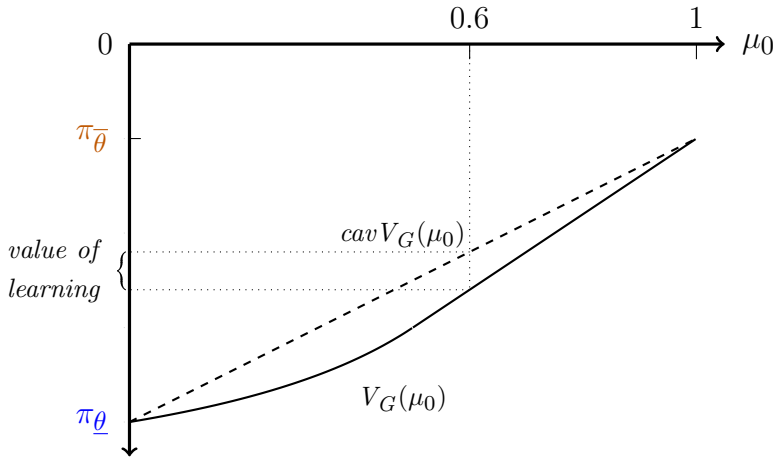
In equilibrium, G's learning is binary: it either learns everything or nothing. It will either offer a separating menu that distinguishes P's preferences ($\mu_1 \in \{0,1\}$); or a pooling menu that both types accept ($\mu_1 = \mu_0$). Thus, G's *value of learning* P's alignment is based on the comparison of (i) G learning P's alignment, and (ii) G lobbying P at the prior belief. Definition 1 formalizes this comparison and the value of learning.

Definition 1. In equilibrium, G's **value of learning** is $W(\mu_0) := (1-\mu_0)\,V_G(0) + \mu_0\,V_G(1) - V_G(\mu_0)$.

G is more inclined to learn P's preferences when G is less certain about P. The value of learning disappears when G is almost certain whether P is an ally or adversary. Figure 2

[14] Formally, these continuation values are the channel for feedback effects through which second-period equilibrium behavior affects first-period incentives.

Figure 2. The Interest Group's Value of Learning



Note: The solid line represents G's continuation value as a function of its prior belief $\mu_0 \in [0,1]$*, while fixing* $\overline{\theta} = \frac{1}{2}$ *and* $\underline{\theta} = 0$*. The dashed line depicts the concavification of* $V_G(\mu_0)$*, which is the upper bound of G's expected continuation value, which it could only obtain if G had complete information about* θ*. For belief* μ_0*, G's value of information is* $cavV_G(\mu_0) - V_G(\mu_0)$*. To illustrate, the value of information at* $\mu_0 = 0.6$ *is shown by the bracketed interval on the y-axis.*

depicts this relationship.

Since G's updated belief affects second-period lobbying, P has incentives to strategically manage her reputation. To clarify key properties, Remark 2 characterizes how P's continuation value depends on μ_1.

Remark 2. In equilibrium, (i) an adversary politician's continuation value is always zero, whereas (ii) an ally politician's continuation value is strictly decreasing over $\mu_1 \leq \tilde{\mu}_1$ and constant over $\mu_1 > \tilde{\mu}_1$. Specifically, $V_P(\underline{\theta}, \mu_1) = 0$ and

$$V_P(\overline{\theta}, \mu_1) = \begin{cases} \frac{(\overline{\theta}-\underline{\theta})(1-\overline{\theta}-\mu_1(1-\underline{\theta}))}{1-\mu_1} & \text{if } \mu_1 \leq \tilde{\mu}_1, \\ 0 & \text{if } \mu_1 > \tilde{\mu}_1. \end{cases}$$

Notably, since an ally's continuation value decreases in μ_1, an ally has incentives to appear adversarial and induce G's belief to be $\mu_1 = 0$.[15] By doing so, they would receive a better lobbying deal in the second period, receiving more than needed to accept.

[15]In contrast, an adversary's continuation value is constant in μ_1, removing any incentive to misrepresent their preferences.

P's potential reputational incentives make learning more costly for G. The prospect of future lobbying complicates G's efforts to induce P to reveal her preferences in the first period.[16] The interplay between learning and reputational incentives that shape first-period lobbying is determined by how much each player values the future. A more patient *interest group* is more willing to concede favorable first-period terms to learn about P and obtain more favorable second-period terms. Conversely, a more patient *politician* is more inclined to forego favorable first-period terms to misrepresent its preferences and receive more favorable second-period terms. Thus, a higher δ_G results in G having a higher willingness to learn P's preferences, while a higher δ_P results in G facing greater costs of learning about P's alignment.[17]

If P is likely to be an ally, then G's first-period menu is separating. Moreover, G prioritizes ally-efficiency and supports it with an overly conservative adversary-contract. Despite the low value of learning, G chooses to learn through offering separating contracts as it is cheap in expectation.[18] Indeed, learning is cheap enough that it is always worthwhile and G offers a separating menu for all $\delta_G \in [0, 1]$.

Conversely, if P is unlikely to be an ally, then G's first-period menu is not separating if G is sufficiently impatient. Then, G offers a pooling menu containing only the adversary-efficient contract, preventing any learning.[19] If G is more patient, however, then G is more inclined to learn and offers a separating menu that overpays an ally politician while being overly conservative towards an adversary politician. Although these distortions are qualitatively similar to second-period lobbying (at equivalent beliefs), their magnitude is amplified due to P's reputation incentive. In turn, the magnitude of distortions increases with P's patience.

First-period distortions can take forms that would not occur in the second period. If G is sufficiently patient and believes P is relatively unlikely to be an ally, then G will distort the ally-policy upward so that it is overly aggressive in both transfer *and* policy. This additional distortion occurs because the substantial ally-transfer makes the ally-contract appealing to an adversary. The incentive compatibility constraints bind for both types, so G adjusts the ally-policy to deter an adversary from accepting the ally-contract.

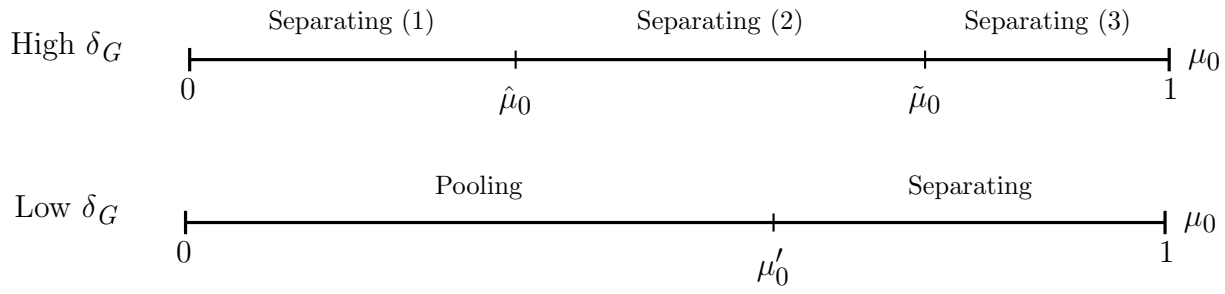
[16] That is, inducing P to reveal her preferences in the second period is already costly (even without dynamic considerations), but this is even more the case in the first. Fixing a particular belief—i.e., $\mu_0 = \mu_1$—a separating menu in the first period must be more distorted than it would be in the second period.

[17] An alternative interpretation is to compare a politician P with a low δ_P and a high δ_P. The same applies to other parameters and players.

[18] This is because distortions are unlikely to materialize in this case, namely only in the unlikely event that P is an adversary.

[19] In the appendix, we formally define the threshold for δ_G (namely $\delta_G^\dagger$) that determines whether pooling can occur.

Figure 3. Informativeness of First-period Lobbying



Note: The figure displays how G's equilibrium menu varies with its patience δ_G and prior belief μ_0. As characterized in Proposition 1, there are three qualitatively distinct separating cases and a pooling case.

Proposition 1 characterizes first-period equilibrium behavior. To state the result, we define several belief thresholds. Let $\hat{\mu}_0 := \max\left\{0, \frac{\delta_P(1-\overline{\theta})-(\overline{\theta}-\underline{\theta})}{\delta_P(1-\overline{\theta})}\right\}$ and $\tilde{\mu}_0 := \frac{(1+\delta_P)(1-\overline{\theta})}{1-\underline{\theta}+\delta_P(1-\overline{\theta})}$.[20]

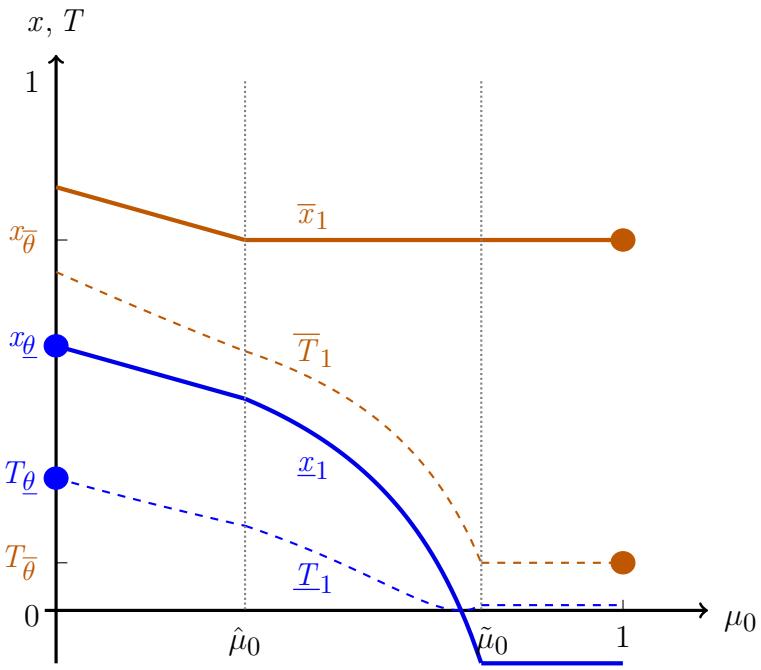
Proposition 1. *In any equilibrium, first-period behavior is as follows. If $\mu_0 < \tilde{\mu}_0$ and $\delta_G < \delta_G^\dagger$, then G offers a pooling menu and P enacts policy $x_1 = x_{\underline{\theta}}$ to receive transfer $T_1 = T_{\underline{\theta}}(x_1)$. Otherwise, G offers a separating menu in which:*

1. *if $\mu_0 \leq \hat{\mu}_0$, then (i) an ally will enact $\overline{x}_1 = x_{\overline{\theta}} + \delta_P(1-\mu_0)\frac{1-\overline{\theta}}{2} - \frac{\overline{\theta}-\underline{\theta}}{2}$ to receive $\overline{T}_1 = T_{\overline{\theta}}(\overline{x}_1) + (\overline{\theta}-\underline{\theta})\Big((1-\overline{\theta})(1+\delta_P(1-\mu_0))\Big)$, and (ii) an adversary will enact $\underline{x}_1 = x_{\underline{\theta}} - \delta_P\mu_0\frac{1-\overline{\theta}}{2}$ to receive $\underline{T}_1 = T_{\underline{\theta}}(\underline{x}_1)$;*
2. *if $\mu_0 \in (\hat{\mu}_0, \tilde{\mu}_0]$, then (i) an ally will enact $\overline{x}_1 = x_{\overline{\theta}}$ to receive $\overline{T}_1 = T_{\overline{\theta}}(\overline{x}_1) + \frac{\overline{\theta}-\underline{\theta}}{1-\mu_0}\Big((1-\overline{\theta})(1+\delta_P(1-\mu_0)) - (1-\underline{\theta})\mu_0\Big)$, and (ii) an adversary will enact $\underline{x}_1 = x_{\underline{\theta}} - \frac{\mu_0}{1-\mu_0}\frac{\overline{\theta}-\underline{\theta}}{2}$ to receive $\underline{T}_1 = T_{\underline{\theta}}(\underline{x}_1)$; and*
3. *if $\mu_0 > \tilde{\mu}_0$, then (i) an ally will enact $\overline{x}_1 = x_{\overline{\theta}}$ to receive $\overline{T}_1 = T_{\overline{\theta}}(x_{\overline{\theta}})$, and (ii) an adversary will enact $\underline{x}_1 = x_{\underline{\theta}} - (1+\delta_P)\frac{1-\overline{\theta}}{2}$ to receive $\underline{T}_1 = T_{\underline{\theta}}(\underline{x}_1)$.*

Proposition 1 reveals the impact of learning and reputation considerations. Figure 3 displays the qualitatively distinct equilibrium menus depending on G's belief μ_0 and patience

[20]Note that $\hat{\mu}_0 < \tilde{\mu}_0$ and $\tilde{\mu}_1 \in (0, \tilde{\mu}_0)$ always hold. Moreover, $\hat{\mu}_0 > 0$ if and only if $\delta_P > \frac{\overline{\theta}-\underline{\theta}}{1-\overline{\theta}}$. Therefore, if $\delta_P \leq \frac{\overline{\theta}-\underline{\theta}}{1-\overline{\theta}}$, then the first case in Proposition 1 does not materialize.

Figure 4. First-period Policies and Transfers



Note: The figure depicts equilibrium policies and transfers from Proposition 1, while fixing $\delta_P = 1$, $\underline{\theta} = 0$, *and* $\overline{\theta} = \frac{2}{5}$. *Blue lines represent the equilibrium adversary-transfer (dashed line) and adversary-policy (solid line). Orange lines represent the equilibrium ally-transfer (dashed line) and ally-policy (solid line). The dots indicate the efficient contracts:* $c_{\underline{\theta}} = (x_{\underline{\theta}}, T_{\underline{\theta}})$ *for* $\mu_0 = 0$ *and* $c_{\overline{\theta}} = (x_{\overline{\theta}}, T_{\overline{\theta}})$ *for* $\mu_0 = 1$.

δ_G. Figure 4 illustrates the separating contracts as a function of G's belief μ_0 when G is sufficiently patient. Without reputation concerns ($\delta_P = 0$), the characterization mirrors Lemma 1. Otherwise, G alters its menu by either adjusting the terms it offers to learn P's preferences or by forgoing learning in favor of a pooling menu.

The politician's reputation incentives have several consequences. First, δ_P affects the conditions under which G offers a separating menu. As δ_P increases, P (if she is an ally) increasingly values the second-period gains from appearing adversarial, thereby forcing G to make an increasingly generous ally-contract to screen P. Therefore learning becomes more costly. Thus, G is less inclined to learn about P's preferences and will do so under fewer conditions.

Second, δ_P also affects the equilibrium contracts that G offers in its separating menu. Broadly, as δ_P increases, G's offers become more distorted. Moreover, depending on the belief μ_0, G will adjust different aspects of the menu in ways that minimize the associated policy or monetary costs. Roughly, (i) if the ally is unlikely, G increases the ally-contract while decreasing the adversary-contract; (ii) for intermediate beliefs, G only increases the ally-transfer; and (iii) if the ally is likely, G decreases the adversary-policy but may increase or decrease its transfer. Proposition 2 makes these observations precise by characterizing how the first-period menu changes as δ_P increases.

Proposition 2. *Suppose P's patience increases from δ_P to δ_P' and fix $\delta_G > \delta_G^{\dagger\prime}$.*

1. *If $\mu_0 < \hat{\mu}_0'$, then the adversary-contract $(\underline{x}_1, \underline{T}_1)$ decreases and the ally-contract $(\overline{x}_1, \overline{T}_1)$ increases.*

2. *If $\mu_0 \in (\hat{\mu}_0', \tilde{\mu}_0)$, then $\overline{T}_1$ increases while $\overline{x}_1$ and $(\underline{x}_1, \underline{T}_1)$ are both constant.*

3. *If $\mu_0 > \tilde{\mu}_0$, then $\overline{x}_1$ is constant and $\overline{T}_1$ increases, while $\underline{x}_1$ decreases and $\underline{T}_1$ decreases for $\underline{x}_1 \geq \underline{\theta}$ and increases for $\underline{x}_1 < \underline{\theta}$.*[21]

How are equilibrium behaviors affected by G's beliefs about P's preferences? First-period lobbying varies with G's belief about P's preference, in several ways. Broadly, as P is more likely to be an ally, G requests smaller policy concessions. Yet, G adjusts fewer aspects of its offers as μ_0 increases: over low μ_0, G decreases policies and transfers for both types; over intermediate μ_0, G does not adjust the ally-policy; and over high μ_0, G's offer is constant. Proposition 3 formally states these observations, comparing lobbying outcomes with a politician who is likely more aligned to outcomes when the politician is relatively unlikely to be an ally.

[21] Because $\underline{T}_1 = (\underline{x}_1 - \underline{\theta})^2$, the adversary-transfer decreases as the adversary-policy approaches $\underline{\theta}$ and increases once the policy passes below it.

Proposition 3. *Fix $\delta_G > \delta_G^\dagger$. As μ_0 increases:*

(i) if $\mu_0 < \hat{\mu}_0$, then the ally-contract $(\overline{x}_1, \overline{T}_1)$ and adversary-contract $(\underline{x}_1, \underline{T}_1)$ will decrease;

(ii) if $\mu_0 \in (\hat{\mu}_0, \tilde{\mu}_0)$, then $\underline{x}_1$ will decrease while $\underline{T}_1$ will decrease or increase, and $\overline{T}_1$ will decrease while $\overline{x}_1$ is constant;

(iii) if $\mu_0 > \tilde{\mu}_0$, then $(\overline{x}_1, \overline{T}_1)$ and $(\underline{x}_1, \underline{T}_1)$ are constant.

As μ_0 increases, G is more likely to face an ally. If an ally is unlikely ($\mu_0 < \hat{\mu}_0$), this change leads G to decrease the ally-contract towards the ally-efficient contract while simultaneously offering a less aggressive adversary offer to distinguish P's alignment. If G's belief is intermediate ($\mu_0 \in (\hat{\mu}_0, \tilde{\mu}_0)$), then G places more emphasis on offering an efficient ally-contract and distorting the adversary-contract is less costly in expectation. Thus, it reduces its excessive transfer for the ally-efficient policy while also making the adversary-contract less aggressive or even paying the adversary to shift policy leftward. Finally, if an ally is likely ($\mu_0 > \tilde{\mu}_0$), the ally-contract is always efficient and no adjustments are necessary to maintain incentive compatibility.

Dynamics of Policymaking and Lobbying

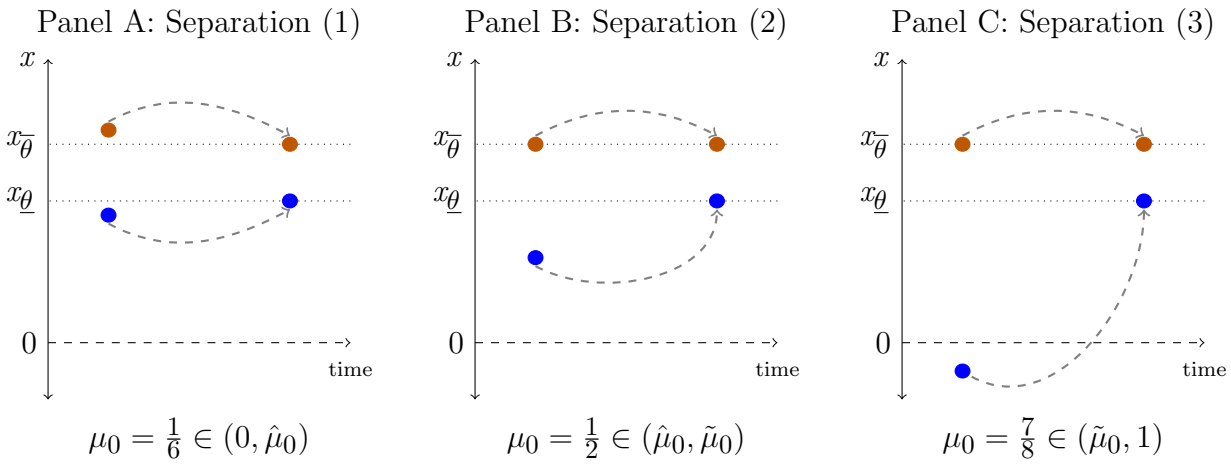
We now trace the dynamics of policies and transfers over time. These dynamics depend on G's prior belief about P's preferences, as well as each player's discount factor.

Proposition 4 characterizes the trajectories of policies and transfers. In equilibrium, an increasing transfer always coincides with increasing policy, and decreasing policy always coincides with a decreasing transfer. Under broad conditions, the trajectories are ambiguous and can go in either direction. They are not ambiguous if P is likely an ally, however, and in that case they will weakly increase over time.

Proposition 4. *In equilibrium, $T_1 < T_2$ implies $x_1 < x_2$, and $x_1 > x_2$ implies $T_1 > T_2$. Furthermore, $\mu_0 > \tilde{\mu}_0$ implies $x_1 \leq x_2$ and $T_1 \leq T_2$, but otherwise $x_1 \geq x_2$ and $T_1 \geq T_2$ is also possible.*

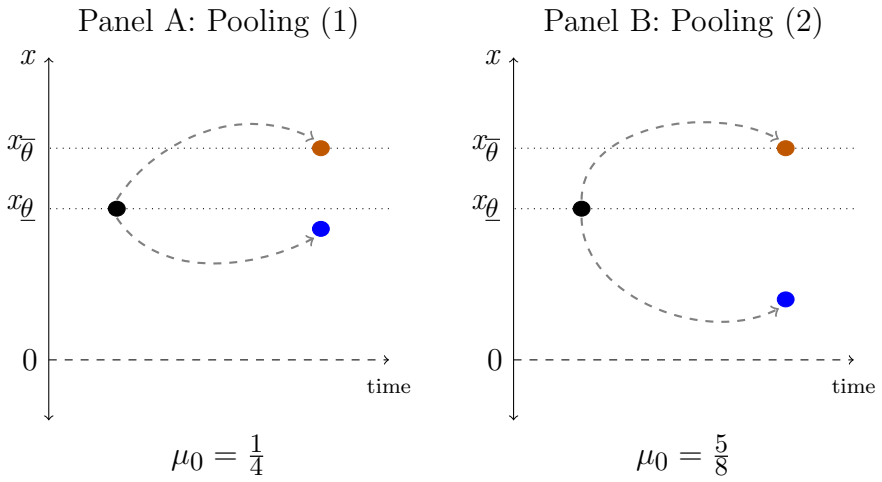
The trajectory of observed policies and transfers is primarily driven by G's learning. If G offers a separating menu to learn P's preference in the first period, that information facilitates an efficient second-period offer. Thus, if G is relatively patient, the observed policy and transfer will either (i) start low and then increase if P is an adversary, or (ii) start high

Figure 5. Dynamics of Equilibrium Policy (First-period Separation)



Note: In each panel, the left dots depict first-period policies while the right dots depict second-period policies. The blue dots indicate adversary-policies and the orange dots indicate ally-policies in equilibrium. In every scenario, equilibrium policies converge to the efficient policies $x_{\underline{\theta}}$ and $x_{\overline{\theta}}$. Each panel depicts the case with $\underline{\theta} = 0$, $\overline{\theta} = \frac{2}{5}$, $\delta_P = 1$, and δ_G sufficiently large.

Figure 6. Dynamics of Equilibrium Policy (First-period Pooling)



Note: In each panel, the left dot depicts first-period policies while the right dots depict second-period policies. The black dot is the pooled policy for both the adversary and ally, the orange dot is the ally-policy and the blue dot is the adversary-policy. Each panel depicts the case with $\underline{\theta} = 0$, $\overline{\theta} = \frac{2}{5}$, $\delta_P = 1$, *and* δ_G *sufficiently small.*

and then decrease or stay constant if P is an ally. Figure 5 displays the three qualitatively different possibilities.

If G's first-period menu is pooling, then learning is delayed and G will instead make a separating offer in the second period. Under these conditions, the observed policies will initially be low before shifting upward or downward depending on P's alignment. Figure 6 displays the two possible scenarios in which G offers a pooling menu in the first period and separating menu in the second period. In this case, second-period policies do not converge to the efficient ones, and the adversary offer is distorted downwards.

Several factors impact the magnitude of changes in observed policies and transfers over time. As δ_P increases, policy shifts become larger for either type. Transfer shifts also become larger while the adversary-policy remains above $\underline{\theta}$. If it is below, however, then the adversary-transfer shift can shrink. This relationship emerges because G must offer more distorted contracts to effectively learn about P's preferences. In contrast, the impact of G's beliefs depends on P's preference: for higher μ_0, the ally-contract shifts less, while the adversary-contract shifts more. In the limiting case where P is entirely short-sighted ($\delta_P = 0$), P's reputational incentive disappears, and therefore its incentives to misrepresent are constant across periods. Thus, G can offer contracts closer to the efficient ally and adversary contracts, while still accounting for the inherent screening obstacles that arise even

without reputational considerations. Proposition 5 formally characterizes these comparative statics.

Proposition 5. *If the interest group is sufficiently patient with $\delta_G > \delta_G^\dagger$, then:*

1. *increasing δ_P will increase $|x_1 - x_2|$ and, provided that the adversary-policy remains above $\underline{\theta}$, also increase $|T_1 - T_2|$; and*

2. *increasing μ_0 will (i) decrease $|\overline{x}_1 - \overline{x}_2|$ and $|\overline{T}_1 - \overline{T}_2|$, but (ii) increase $|\underline{x}_1 - \underline{x}_2|$ and, provided that the adversary-policy remains above $\underline{\theta}$, also increase $|\underline{T}_1 - \underline{T}_2|$.*[22]

Since policies and transfers can shift in either direction depending on P's type, forecasts about the dynamics of policymaking and lobbying will depend on the probability of facing an ally politician. Our focus below is on the expected policy and transfer, analyzing a weighted average of these equilibrium objects for the adversary and ally types.[23]

First, we focus on how transfers evolve over time through lobbying. When the probability of an adversary politician is high, the expected transfer to P can decrease over time. This is because learning is relatively costly since it requires large transfers. Otherwise, if an ally is sufficiently likely, then P receives a higher expected transfer in the second period than the first period. A different effect dominates in this case—as an ally is more likely, G is overly conservative in its adversary-contract. Once G learns that P is an adversary, however, it can make a more aggressive offer without having to worry about an ally's incentive constraint.

Second, we show that if the first-period menu is separating, then the expected policy shift over time is positive. Thus, ex ante, P's expected policy choice shifts towards G's ideal point over time. As time evolves, the expected policy becomes more in line with G's preferences. The primary policy distortion in any such menu is an overly conservative adversary policy to deter an ally from mimicking an adversary. Thus, although policy can shift weakly downward if P is an ally, there is always a larger increase in the adversary policy over time and that trajectory is likely enough that the expected policy shift is positive. This is visible in Figure 5: although the ally-policy decreases or stays constant over time in each case, the expected policy always increases.

Alternatively, if G does not aim to learn the politician's preferences, it simply treats P like an adversary. In this case, policy shifts upward over time if P is an ally but downward if P is an adversary. When assessing the expected policy change, it can either be the case

[22] When reputation concerns are strong enough that the adversary-policy falls below $\underline{\theta}$, the adversary-transfer gap can move in the opposite direction.

[23] There is a direct effect determined by the weights of these objects that vary with μ_0. There is also an indirect effect, however, since policies and transfers can depend on μ_0.

that the policy shift is zero over time ($\mu_0 < \tilde{\mu}_1$) or positive ($\mu_0 > \tilde{\mu}_1$).[24]

Finally, policy variance simply depends on whether G provides a separating or pooling offer to P. As displayed in Figure 5, first-period separating policy offers are more distinct than second-period policy offers. On the other hand, however, given first-period pooling, Figure 6 shows that the lack of variance in the first period is followed by more uncertain policy predictions in the second, where the level of uncertainty depends on G's belief, μ_0.

Extensions

We study four extensions capturing key aspects of other political features such as hiring lobbyists, making campaign contributions, veto players, voting rules, or revolving-door hiring. We extend our main analysis to analyze: (i) early-career information, (ii) early-career access, (iii) policymaking constraints, and (iv) revolving-door incentives. For each, we highlight the impact on policies and transfers, their dynamics, as well as the degree of G's policy influence.

The value of early-career information

Interest groups often have access to various tools for learning about politicians' motivations and interests before lobbying: conducting interviews with staff members, researching politicians' lawmaking and voting records, or hiring lobbyists with established connections.[25] Given these avenues for acquiring information, we address two key questions. First, what is the value of obtaining such information for interest groups compared to *learning-by-lobbying*? And, second, which factors influence the value of this information?

We analyze the value of early-career information by comparing G's payoff under full information to G's payoff in the main model.[26]

Under full information, G receives either $\pi_{\underline{\theta}}$ if P is an adversary or $\pi_{\overline{\theta}}$ if P is an ally. From G's first-period perspective, its expected payoff is:

$$V_G^{informed} = (1 + \delta_G)\left((1 - \mu_0)\pi_{\underline{\theta}} + \mu_0 \pi_{\overline{\theta}}\right).$$

There are two main cases. In the first, G offers a separating menu and learns P's preferences in the first period. Here, information is only valuable in the first period, since

[24] The stated condition here focuses on second-period equilibrium policies, while implicitly assuming that G still finds it optimal to give a pooling offer to P.

[25] For instance, a primary motive for revolving-door hiring is "buying advice on who is likely to be sympathetic to them on a particular issue, how best to win the support of particular members" (Drutman, 2015, p. 163).

[26] See Lemma 1 and Proposition 1.

G will be fully informed in the second period regardless of having access to early-career information. The value of information is proportionate to the distortions G induces when lobbying.

In the second case, G makes a pooling offer in the first period, and then a separating offer in the second. Here, information is valuable in both periods, since it allows for efficient lobbying compared to the lack of efficient contracts absent information. In the first period, the adversary-contract is efficient, while the ally-contract is distorted. Thus, the first-period value of information equals $\mu_0\left(\pi_{\overline{\theta}} - \pi_{\underline{\theta}}\right)$. In the second period, the value of information also depends on the prior, μ_0. If $\mu_0 \leq \tilde{\mu}_1$, then G's second-period value of information is $\mu_0\left(\pi_{\overline{\theta}} - \frac{(\overline{\theta}-\underline{\theta})^2}{2(1-\mu_0)} - \pi_{\underline{\theta}}\right)$. Otherwise, it is $(1-\mu_0)(-\pi_{\overline{\theta}})$.[27]

Proposition 6 summarizes key properties about G's value of early-career information.

Proposition 6. *The interest group's value of early-career information is positive and weakly increasing in δ_P and δ_G.*

Although G always benefits from more information about P's preferences, its willingness to pay for that information in the first period varies with the political context. When P is more forward-looking, screening becomes more expensive, so early-career information about P is more valuable to G. When G is more patient but would otherwise not learn by lobbying, early-career information is weakly more valuable. Once G is patient enough to learn by lobbying, however, its value is unchanged.

The value of early-career access

To lobby particular politicians, interest groups typically need access (Hansen, 1991). This access often requires strategic investments to establish relationships, such as hiring lobbyists (Blanes i Vidal, Draca and Fons-Rosen, 2012; Bertrand, Bombardini and Trebbi, 2014) or providing campaign contributions (Fouirnaies, 2018; **?**). To quantify the importance of early-career access, we compare two settings: (i) *full access*, where G can lobby in both periods (as in our main model), and (ii) *late-career access*, where G can only lobby in the second period.

We define the *value of early-career access* as the difference between G's equilibrium payoff with full access versus its payoff with only late-career access. This quantity measures the strategic importance of early engagement with politicians.

We show that politicians may have incentives to misrepresent their alignment even when G cannot lobby in the first period. These incentives stem from the anticipation of lobbying

[27] Recall that $\pi_{\overline{\theta}}$ is negative, implying a positive value of information.

efforts in the second period. Similar to our main model, ally politicians may be motivated to feign disagreement to secure more favorable terms from interest groups in subsequent interactions. The strength of these incentives can be substantial. In some cases, ally politicians might moderate their chosen policies without receiving any immediate transfer, solely based on the prospect of future lobbying.

In equilibrium, P's first-period policy choice may or may not reveal its alignment. First, each politician type may choose its ideal policy, θ, and G perfectly learns P's alignment. Second, an ally may mimic an adversary and choose policy $x_1 = \underline{\theta}$, leaving G without any new information about P.[28]

Proposition 7 characterizes key properties of G's value of early-career access.

Proposition 7. *Suppose $\delta_G > \delta_G^\dagger$. The interest group's value of early-career access is positive, decreasing in δ_P, and weakly increasing in δ_G.*

Early access is less valuable if P is more patient. Then, P cares more about her future reputation, so G must incur higher costs to learn, which reduces the appeal of learning-by-lobbying in the first period. The more patient G is, however, the more valuable access becomes when P's policy without influence would not reveal P's alignment. A far-sighted G is more willing to learn-by-lobbying because it gets more out of eventually using the information it learns.

This extension highlights how politicians' job security, career stage, or familiarity can affect interest groups' desire to cultivate relationships with them. Our learning-by-lobbying mechanism sheds new light on why the value of access varies across politicians based on factors like their procedural rights (Fouirnaies and Hall, 2018; Berry and Fowler, 2018).

Lobbying with policymaking constraints

Politicians have different degrees of policy influence, with some playing more active roles in developing and drafting proposals on particular issues than others. Yet, even the most powerful politicians often face constraints when crafting proposals due to, e.g., voting rules or veto players. How do such policymaking constraints affect lobbying, influence, and the value of access?

To study this question, we extend our model by restricting the set of policies that P can enact.[29] Formally, we introduce a maximum policy, denoted $\overline{y} \in [\max\{x_{\underline{\theta}}, \overline{\theta}\}, x_{\overline{\theta}})$,

[28] Mixed strategy equilibria also exist, causing G to only partially learn what P prefers. For the sake of presentation, we omit such equilibria given our focus on pure strategy equilibria.

[29] Alternatively, suppose some politicians face higher costs than others when changing policy. Then, lobbying lower-cost politicians is more fruitful, so they are more valuable to learn about. This should produce qualitatively similar results.

which P's policy cannot exceed. This extension yields two key strategic implications. First, the constraint limits G's scope for shifting the ally's policy, which lowers G's payoff from influencing an ally—and the associated transfer—even under complete information. Second, the constraint raises the first-period cost of screening P, while leaving the second-period value of learning unchanged.

Qualitatively, equilibrium strategies remain similar to those in the main model. Whenever an ally-contract is offered, the cap binds its policy at $\overline{y}$, and its transfer adjusts to satisfy the relevant incentive and participation constraints. The first-period pooling contract remains feasible, whereas the unconstrained separating menu does not. A tighter cap therefore lowers the group's first-period payoff from screening relative to pooling, while leaving the second-period value of learning unchanged. Proposition 8 formalizes these observations.

Proposition 8. *A tighter policy constraint (lower $\overline{y}$) makes G less inclined to learn P's preferences by lobbying, and G's equilibrium payoff is increasing in $\overline{y}$.*

Interest groups have more to gain from influencing politicians with greater policy discretion, and screening those politicians by lobbying is also less costly. Both forces can strengthen groups' incentives to target powerful committee leaders.

Revolving-door incentives

In our final extension, we study how politicians' revolving-door incentives can affect lobbying and policymaking dynamics. The revolving door is a lucrative post-government option (Blanes i Vidal, Draca and Fons-Rosen, 2012; Bertrand, Bombardini and Trebbi, 2014; McCrain, 2018) that can encourage politicians to signal their alignment with potential future employers while still in office (Cornaggia, Cornaggia and Xia, 2016; Tabakovic and Wollmann, 2018). We analyze how this incentive interacts with our core learning-by-lobbying forces to shape lobbying and politicians' observed in-office behavior.

We introduce a new parameter, $R > 0$, representing the additional payoff P receives if hired by G after the second period. Specifically, G commits to hire P if and only if P accepts a first-period contract whose acceptance makes G sufficiently confident that P is an ally (a posterior above $\tilde{\mu}_1$), whereas a politician who rejects the menu is not hired.[30]

The benefit of appearing aligned has two key strategic implications. First, P's signaling incentives become more complex. Previously, an ally was motivated to signal misalignment solely to secure better second-period lobbying. Now, this incentive is counterbalanced by the desire to appear aligned to get hired as a revolving-door lobbyist. Second, in turn, provided

[30] Conditioning hiring on acceptance ensures that rejecting is not a route to the hiring rent, so the baseline analysis of rejection carries over. To ensure non-negative equilibrium transfers, we focus on $\delta_P R \leq T_{\overline{\theta}}$.

P is not too likely an ally ($\mu_0 \leq \tilde{\mu}_1$), G can learn P's preferences at lower cost, since an ally who postures as an adversary forfeits the hiring rent.[31] These strategic implications lead to the main empirical implication of this extension.

Proposition 9. *Suppose $\mu_0 \leq \tilde{\mu}_1$. Increasing the value of the revolving door expands the set of parameters under which G successfully learns P's preferences.*

Substantively, this suggests a novel way that interest groups can benefit from situations in which politicians strongly value becoming a revolving-door lobbyist. For instance, larger firms that can promise and offer higher future wages and better careers to incumbent politicians may also learn more easily whether they are dealing with an ally or adversary via learning-by-lobbying. Proposition 10 establishes this effect on G's expected payoff, as well as how it is amplified by P's patience, δ_P, and G's belief that it faces an ally politician, μ_0.

Proposition 10. *Suppose $\mu_0 \leq \tilde{\mu}_1$ and $\delta_G > \delta_G^{\dagger}$. The interest group's equilibrium payoff increases in the politician's value of the revolving door, R, and this effect is amplified by δ_P and μ_0.*

This result highlights how strategic links between lobbying and revolving door opportunities can benefit powerful and wealthy interest groups. By providing the prospect of appealing revolving door opportunities, these groups can get more bang for their buck from lobbying. Not only do they have more resources for lobbying politicians, but their capacity to provide lucrative post-government compensation to ally politicians can also help them lobby more efficiently by weakening politicians' reputational incentives to appear adversarial. This facilitates *learning-by-lobbying* since these groups do not have to distort their initial lobbying overtures as much to learn about politicians.

Discussion and Conclusion

We study how lobbying relationships evolve when interest groups may learn about politicians by lobbying them. Our game-theoretic analysis clarifies how *learning-by-lobbying* can impact patterns of lobbying behavior throughout political careers and across institutional contexts. This fundamental mechanism complements previous insights about information transmission, legislative subsidies, or financial incentives. In doing so, we sharpen theoretical understanding of information, influence, and policy outcomes.

Our main analysis has implications for empirical patterns in lobbying. Early-career politicians with secure positions can attract intense lobbying partly due to their stronger

[31] If instead $\mu_0 > \tilde{\mu}_1$, pooling keeps both types hireable and lets G extract the rent without screening, so the revolving door can instead discourage learning.

reputational considerations, which force groups to offer more generous terms for effective learning. Additionally, our learning-by-lobbying mechanism may contribute to observed patterns of generous treatment toward allies, especially early in politicians' careers—a pattern that standard exchange models struggle to explain. When the group screens early on, our model also predicts systematic variation in policy uncertainty across career stages, with more dispersed policy choices by newcomers and a convergence in late-career lobbying as learning and reputation forces diminish. Politicians receiving aggressive lobbying early might face milder approaches later, and vice versa—not because their preferences changed, but because interest groups learned about them.

Our extensions offer additional empirical insights into broader interest group activities. Our early-career information extension suggests interest groups are especially inclined to solicit information about patient politicians with secure positions. In contrast, our early-career access extension highlights potential downsides to acquiring access to those same politicians. Our policymaking constraint extension suggests they are less inclined to screen politicians who have limited discretion by lobbying them. Notably, some interest group activities—like hiring former staffers—serve both information and access functions, complicating empirical expectations. Finally, when groups initially doubt that a politician is an ally, our revolving-door extension suggests that learning-by-lobbying forces can encourage wealthy industries with lucrative post-career opportunities to prioritize access over background research and enjoy lower overall lobbying costs.

Our work has broader implications for democratic representation and special interest influence. Since new politicians increasingly lack extensive public-service track records (Porter and Treul, 2025), the ability to understand and anticipate policymakers' preferences is increasingly valuable for shaping legislation. We set aside factors such as competing interest groups, multiple politicians, and electoral considerations that have been emphasized elsewhere (Austen-Smith and Wright, 1994; Bils, Duggan and Judd, 2021; Groseclose and Snyder, 1996). Although relevant in certain contexts, such factors are not central to our core insights about learning and influence. Analyzing how they interact with learning-by-lobbying is a promising avenue for future work.

References

Austen-Smith, David. 1995. "Campaign Contributions and Access." *American Political Science Review* 89(3):566–581.

Austen-Smith, David and John R. Wright. 1992. "Competitive Lobbying for a Legislator's Vote." *Social Choice and Welfare* 9(3):229–257.

Austen-Smith, David and John R. Wright. 1994. "Counteractive Lobbying." *American Journal of Political Science* 38(1):25–44.

Awad, Emiel. 2020. "Persuasive Lobbying with Allied Legislators." *American Journal of Political Science* 64(4):938–951.

Bar-Isaac, Heski and Joel Shapiro. 2011. "Credit Ratings Accuracy and Analyst Incentives." *American Economic Review: Papers and Proceedings* 101(3):120–124.

Beccuti, Juan and Marc Möller. 2018. "Dynamic Adverse Selection with a Patient Seller." *Journal of Economic Theory* 173:95–117.

Bernheim, B. Douglas and Michael D. Whinston. 1986. "Menu Auctions, Resource Allocation, and Economic Influence." *Quarterly Journal of Economics* 101(1):1–31.

Berry, Christopher R. and Anthony Fowler. 2018. "Congressional Committees, Legislative Influence, and the Hegemony of Chairs." *Journal of Public Economics* 158:1–11.

Bertrand, Marianne, Matilde Bombardini and Francesco Trebbi. 2014. "Is It Whom You Know or What You Know? An Empirical Assessment of the Lobbying Process." *American Economic Review* 104(12):3885–3920.

Besley, Timothy and Stephen Coate. 2001. "Lobbying and Welfare in a Representative Democracy." *The Review of Economic Studies* 68(1):67–82.

Bester, Helmut and Roland Strausz. 2001. "Contracting with Imperfect Commitment and the Revelation Principle: The Single Agent Case." *Econometrica* 69(4):1077–1098.

Bils, Peter, John Duggan and Gleason Judd. 2021. "Lobbying and Policy Extremism in Repeated Elections." *Journal of Economic Theory* 193:105223.

Blanes i Vidal, Jordi, Mirko Draca and Christian Fons-Rosen. 2012. "Revolving Door Lobbyists." *American Economic Review* 102(7):3731–3748.

Bombardini, Matilde and Francesco Trebbi. 2020. "Empirical Models of Lobbying." *Annual Review of Economics* 12:391–413.

Breig, Zachary. 2022. "Repeated Contracting without Commitment." *Journal of Economic Theory* 204:105514.

Che, Yeon-Koo. 1995. "Revolving Doors and the Optimal Tolerance for Agency Collusion." *RAND Journal of Economics* 26(3):378–397.

Chen, Ying and Jan Zápal. 2022. "Sequential Vote Buying." *Journal of Economic Theory* 205:105529.

Cornaggia, Jess, Kimberly J. Cornaggia and Han Xia. 2016. "Revolving Doors on Wall Street." *Journal of Financial Economics* 120(2):400–419.

Drutman, Lee. 2015. *The Business of America is Lobbying*. Oxford University Press New York.

Egerod, Benjamin C.K. and Hai Tran. 2023. "Extremists Not on Board: Labor Market Costs to Radical Behavior in Elected Office." *Journal of Politics* 85(3):1161–1165.

Ellis, Christopher J. and Thomas Groll. 2024. "Who Lobbies Whom? Special Interests and Hired Guns." Unpublished manuscript.

Fouirnaies, Alexander. 2018. "When Are Agenda Setters Valuable?" *American Journal of Political Science* 62(1):176–191.

Fouirnaies, Alexander and Andrew B. Hall. 2018. "How Do Interest Groups Seek Access to Committees?" *American Journal of Political Science* 62(1):132–147.

Freixas, Xavier, Roger Guesnerie and Jean Tirole. 1985. "Planning under Incomplete Information and the Ratchet Effect." *The Review of Economic Studies* 52(2):173–191.

Fudenberg, Drew and Jean Tirole. 1991. *Game Theory*. MIT Press.

Gerardi, Dino and Lucas Maestri. 2020. "Dynamic Contracting with Limited Commitment and the Ratchet Effect." *Theoretical Economics* 15(2):583–623.

Groll, Thomas and Christopher J. Ellis. 2014. "A Simple Model of the Commercial Lobbying Industry." *European Economic Review* 70:299–316.

Groll, Thomas and Christopher J. Ellis. 2017. "Repeated Lobbying by Commercial Lobbyists and Special Interests." *Economic Inquiry* 55(4):1868–1897.

Groseclose, Tim and James M. Snyder. 1996. "Buying Supermajorities." *American Political Science Review* 90(2):303–315.

Grossman, Gene M. and Elhanan Helpman. 1994. "Protection for Sale." *American Economic Review* 84(4):833–850.

Grossman, Gene M. and Elhanan Helpman. 2001. *Special Interest Politics*. MIT Press.

Hall, Richard L. and Alan V. Deardorff. 2006. "Lobbying as Legislative Subsidy." *American Political Science Review* 100(1):69–84.

Hall, Richard L and Frank W Wayman. 1990. "Buying Time: Moneyed Interests and the Mobilization of Bias in Congressional Committees." *American Political Science Review* 84(3):797–820.

Hansen, John Mark. 1991. *Gaining Access: Congress and the Farm Lobby, 1919-1981*. University of Chicago Press.

Hart, Oliver D and Jean Tirole. 1988. "Contract Renegotiation and Coasian Dynamics." *The Review of Economic Studies* 55(4):509–540.

Hirsch, Alexander V., Karam Kang, B. Pablo Montagnes and Hye Young You. 2023. "Lobbyists as Gatekeepers: Theory and Evidence." *The Journal of Politics* 85(2):731–748.

Hübert, Ryan, Janna King Rezaee and Jonathan Colner. 2023. "Going into Government: How Hiring from Special Interests Reduces Their Influence." *American Journal of Political Science* 67(2):485–498.

Iaryczower, Matías and Santiago Oliveros. 2017. "Competing for Loyalty: The Dynamics of Rallying Support." *American Economic Review* 107(10):2990–3005.

Judd, Gleason. 2023. "Access to Proposers and Influence in Collective Policy Making." *The Journal of Politics* 85(4):1430–1443.

Kerr, William R., William F. Lincoln and Prachi Mishra. 2014. "The Dynamics of Firm Lobbying." *American Economic Journal: Economic Policy* 6(4):343–379.

Kolotilin, Anton, Tymofiy Mylovanov, Andriy Zapechelnyuk and Ming Li. 2017. "Persuasion of a Privately Informed Receiver." *Econometrica* 85(6):1949–1964.

Laffont, Jean-Jacques and Jean Tirole. 1990. "Adverse Selection and Renegotiation in Procurement." *The Review of Economic Studies* 57(4):597–625.

LaPira, Timothy and Herschel F. Thomas III. 2017. *Revolving Door Lobbying: Public Service, Private Influence, and the Unequal Representation of Interests*. University Press of Kansas.

Levine, Bertram J. 2008. *The Art of Lobbying: Building Trust and Selling Policy*. CQ Press.

Manelli, Alejandro M. 1997. "The Never-a-weak-best-response Test in Infinite Signaling Games." *Journal of Economic Theory* 74(1):152–173.

Martimort, David and Aggey Semenov. 2008. "Ideological Uncertainty and Lobbying Competition." *Journal of Public Economics* 92(3-4):456–481.

McChesney, Fred S. 1997. *Money for Nothing: Politicians, Rent Extraction, and Political Extortion*. Harvard University Press.

McCrain, Joshua. 2018. "Revolving Door Lobbyists and the Value of Congressional Staff Connections." *The Journal of Politics* 80(4):1369–1383.

Minaudier, Clement. 2022. "The Value of Confidential Policy Information: Persuasion, Transparency, and Influence." *The Journal of Law, Economics, and Organization* 38(2):570–612.

Porter, Rachel and Sarah A. Treul. 2025. "Evaluating (in)experience in Congressional Elections." *American Journal of Political Science* 69(1):284–298.

Salanié, Bernard. 2005. *The Economics of Contracts: A Primer*. MIT Press.

Schnakenberg, Keith E. 2017. "Informational Lobbying and Legislative Voting." *American Journal of Political Science* 61(1):129–145.

Schnakenberg, Keith E. and Ian R. Turner. 2024. "Formal Theories of Special Interest Influence." *Annual Review of Political Science* 27(1):401–421.

Tabakovic, Haris and Thomas G. Wollmann. 2018. From Revolving Doors to Regulatory Capture? Evidence from Patent Examiners. Technical report National Bureau of Economic Research.

Online Appendix

Contents

A Proofs of Main Results **A-1**

A.1 Proof of Lemma 1 . . . . . A-1

A.2 Proofs of Remarks 1 and 2 . . . . . A-5

A.3 Proof of Proposition 1 . . . . . A-5

A.4 Proofs of Propositions 2–5 . . . . . A-11

B Proofs of Extensions **A-15**

B.1 Proof of Proposition 6 . . . . . A-16

B.2 Proof of Proposition 7 . . . . . A-17

B.3 Proof of Proposition 8 . . . . . A-18

B.4 Proofs of Propositions 9 and 10 . . . . . A-20

A Proofs of Main Results

Throughout the appendix, $\Delta := \overline{\theta}-\underline{\theta}$, $\delta_P^* := \frac{\Delta}{1-\overline{\theta}}$, $\mu_0^* := \frac{1-\underline{\theta}}{1+\overline{\theta}-2\underline{\theta}}$, and $S(\mu_1) := \frac{(\overline{\theta}-\underline{\theta})(1-\overline{\theta}-\mu_1(1-\underline{\theta}))}{1-\mu_1}$, so that an ally's continuation value is $V_P(\overline{\theta},\mu_1) = \max\{S(\mu_1),0\}$ for $\mu_1 < 1$ (Remark 2).

A.1 Proof of Lemma 1

Fix $\mu := \mu_1 \in (0,1)$.In the second period, the interest group has the following constrained maximization problem,

$$\begin{aligned}\max_{\underline{x},\overline{x},\underline{T},\overline{T}\in\mathbb{R}} \; & \mu\left(-\left(\overline{x}-1\right)^2-\overline{T}\right)+(1-\mu)\left(-\left(\underline{x}-1\right)^2-\underline{T}\right)\\ s.t. \quad & (IC_{\overline{\theta}}),(IC_{\underline{\theta}}),(P_{\overline{\theta}}) \text{ and } (P_{\underline{\theta}}).\end{aligned}$$

The four restrictions include two incentive compatibility constraints and two participation constraints:

$$-\left(\overline{x}-\overline{\theta}\right)^2+\overline{T} \geq -\left(\underline{x}-\overline{\theta}\right)^2+\underline{T}, \tag{$IC_{\overline{\theta}}$}$$
$$-\left(\underline{x}-\underline{\theta}\right)^2+\underline{T} \geq -\left(\overline{x}-\underline{\theta}\right)^2+\overline{T}, \tag{$IC_{\underline{\theta}}$}$$
$$-\left(\overline{x}-\overline{\theta}\right)^2+\overline{T} \geq 0, \tag{$P_{\overline{\theta}}$}$$
$$-\left(\underline{x}-\underline{\theta}\right)^2+\underline{T} \geq 0. \tag{$P_{\underline{\theta}}$}$$

We study the following relaxed problem and then verify that the solution satisfies the constraint $(IC_{\underline{\theta}})$:

$$\begin{aligned}\max_{\underline{x},\overline{x},\underline{T},\overline{T}\in\mathbb{R}} \; & \mu\left(-\left(\overline{x}-1\right)^2-\overline{T}\right)+(1-\mu)\left(-\left(\underline{x}-1\right)^2-\underline{T}\right)\\ s.t. \quad & (IC_{\overline{\theta}}),(P_{\overline{\theta}}) \text{ and } (P_{\underline{\theta}}).\end{aligned}$$

We begin our analysis by setting up the Lagrangian for the relaxed problem:

$$\begin{aligned}\mathcal{L} = & \mu\left(-\left(\overline{x}-1\right)^2-\overline{T}\right)+(1-\mu)\left(-\left(\underline{x}-1\right)^2-\underline{T}\right)\\ & +\lambda_1\left(-\left(\overline{x}-\overline{\theta}\right)^2+\overline{T}+\left(\underline{x}-\overline{\theta}\right)^2-\underline{T}\right)\\ & +\lambda_2\left(-\left(\overline{x}-\overline{\theta}\right)^2+\overline{T}\right)\\ & +\lambda_3\left(-\left(\underline{x}-\underline{\theta}\right)^2+\underline{T}\right),\end{aligned}$$

where λ_1, λ_2 and λ_3 are the multipliers for $(IC_{\overline{\theta}})$, $(P_{\overline{\theta}})$ and $(P_{\underline{\theta}})$ respectively. Using this notation, the Kuhn-Tucker first-order necessary conditions are given as follows.

The first-order conditions with respect to the tuple $(\underline{x}, \overline{x}, \underline{T}, \overline{T})$ are:

$$\frac{\partial \mathcal{L}}{\partial \underline{x}} = -2(1-\mu)(\underline{x}-1) + 2\lambda_1(\underline{x}-\overline{\theta}) - 2\lambda_3(\underline{x}-\underline{\theta}) = 0, \tag{A.1}$$

$$\frac{\partial \mathcal{L}}{\partial \overline{x}} = -2\mu(\overline{x}-1) - 2\lambda_1(\overline{x}-\overline{\theta}) - 2\lambda_2(\overline{x}-\overline{\theta}) = 0, \tag{A.2}$$

$$\frac{\partial \mathcal{L}}{\partial \underline{T}} = -(1-\mu) - \lambda_1 + \lambda_3 = 0, \tag{A.3}$$

$$\frac{\partial \mathcal{L}}{\partial \overline{T}} = -\mu + \lambda_1 + \lambda_2 = 0. \tag{A.4}$$

The complementary slackness conditions for $(IC_{\overline{\theta}}), (P_{\overline{\theta}})$, and $(P_{\underline{\theta}})$ respectively are:

$$\lambda_1 \left(-\left(\overline{x}-\overline{\theta}\right)^2 + \overline{T} + \left(\underline{x}-\overline{\theta}\right)^2 - \underline{T} \right) = 0, \tag{A.5}$$

$$\lambda_2 \left(-\left(\overline{x}-\overline{\theta}\right)^2 + \overline{T} \right) = 0, \tag{A.6}$$

$$\lambda_3 \left(-\left(\underline{x}-\underline{\theta}\right)^2 + \underline{T} \right) = 0. \tag{A.7}$$

The non-negative Lagrangian multipliers and the constraints are:

$$\lambda_1, \lambda_2, \lambda_3 \geq 0, \tag{A.8}$$

$$(IC_{\overline{\theta}}), (P_{\overline{\theta}}) \text{ and } (P_{\underline{\theta}}). \tag{A.9}$$

From (A.3), we deduce that $\lambda_3 = \lambda_1 + (1-\mu) > 0$. Condition (A.7) then implies $\left(-\left(\underline{x}-\underline{\theta}\right)^2 + \underline{T} \right) = 0$. From (A.4), we deduce that $\lambda_1 + \lambda_2 = \mu$. We have three possible cases: (i) $\lambda_1 > 0$ and $\lambda_2 > 0$, (ii) $\lambda_1 = \mu$, $\lambda_2 = 0$, and (iii) $\lambda_1 = 0$, $\lambda_2 = \mu$.

Case (i). The fact that $\lambda_1 > 0$ and $\lambda_2 > 0$ implies that constraints $(IC_{\overline{\theta}}), (P_{\overline{\theta}})$ and $(P_{\underline{\theta}})$ are all binding. These binding constraints plus the first-order conditions create a system of 7 equations and 7 unknowns with the following solution:

$$\underline{x}^* = \frac{\underline{\theta}+\overline{\theta}}{2},\ \overline{x}^* = \frac{1+\overline{\theta}}{2},\ \underline{T}^* = \frac{1}{4}\Delta^2, \overline{T}^* = \frac{1}{4}\left(1-\overline{\theta}\right)^2,$$
$$\lambda_1^* = \frac{(1-\mu)(1-\overline{\theta})}{\Delta},\ \lambda_2^* = \frac{\mu(1-\underline{\theta}) - (1-\overline{\theta})}{\Delta},\ \lambda_3^* = \frac{(1-\mu)(1-\underline{\theta})}{\Delta}.$$

Note that $\lambda_1^*, \lambda_3^* > 0$. Also, $\lambda_2^* > 0$ if $\mu > \tilde{\mu}_1 := \frac{1-\overline{\theta}}{1-\underline{\theta}}$. Thus, if $\mu > \tilde{\mu}_1$, the vector $\left(\underline{x}^*, \overline{x}^*, \underline{T}^*, \overline{T}^*, \lambda_1^*, \lambda_2^*, \lambda_3^*\right)$ satisfies the Kuhn-Tucker first-order necessary conditions of the

relaxed problem.

Case (ii). Now $(IC_{\overline{\theta}})$ and $(P_{\underline{\theta}})$ are binding and, together with the first-order conditions, create a system of 6 equations and 6 unknowns with the following solution:

$$\underline{x}' = \frac{1+\underline{\theta}-\mu(1+\overline{\theta})}{2(1-\mu)},\ \overline{x}' = \frac{1+\overline{\theta}}{2},$$
$$\underline{T}' = \frac{1}{4}\frac{(1-\underline{\theta}-\mu(1+\overline{\theta}-2\underline{\theta}))^2}{(1-\mu)^2},\ \overline{T}' = \frac{1}{4}\frac{(1-\overline{\theta})(1+3\overline{\theta}-4\underline{\theta})-\mu(1+\overline{\theta}-2\underline{\theta})^2}{(1-\mu)},$$
$$\lambda_1' = \mu,\ \lambda_2' = 0,\ \lambda_3' = 1.$$

Replacing these values on the constraint $(P_{\overline{\theta}})$, we obtain that it must be that $\frac{\Delta(1-\overline{\theta}-\mu(1-\underline{\theta}))}{(1-\mu)} \geq 0$, which is satisfied if $\mu \leq \tilde{\mu}_1$. Thus, if $\mu \leq \tilde{\mu}_1$, the vector $\left(\underline{x}', \overline{x}', \underline{T}', \overline{T}', \lambda_1', \lambda_2', \lambda_3'\right)$ satisfies the Kuhn-Tucker first-order necessary conditions of the relaxed problem.

Case (iii). Now $(P_{\overline{\theta}})$ and $(P_{\underline{\theta}})$ are binding and, together with the first-order conditions, create a system of 6 equations and 6 unknowns with the following solution:

$$\underline{x}'' = \frac{\underline{\theta}+1}{2},\ \overline{x}'' = \frac{1+\overline{\theta}}{2},\ \underline{T}'' = \frac{1}{4}(1-\underline{\theta})^2,\ \overline{T}'' = \frac{1}{4}\left(1-\overline{\theta}\right)^2,$$
$$\lambda_1'' = 0,\ \lambda_2'' = \mu,\ \lambda_3'' = (1-\mu).$$

Replacing these values on the constraint $(IC_{\overline{\theta}})$ we obtain that it must be that $-(1-\overline{\theta})(\overline{\theta}-\underline{\theta}) \geq 0$, which is never satisfied. Thus, there are no values that satisfy the Kuhn-Tucker first-order necessary conditions in case (iii).

Note that $\frac{\partial^2 \mathcal{L}}{\partial \underline{x}^2} = -2(1-\mu) + 2(\lambda_1 - \lambda_3) = -4(1-\mu) < 0$, using $(\lambda_1 - \lambda_3) = -(1-\mu)$ from (A.3). Also, $\frac{\partial^2 \mathcal{L}}{\partial \overline{x}^2} = -2\mu - 2\lambda_1 - 2\lambda_2 < 0$. Since $\mathcal{L}$ is linear in $(\underline{T}, \overline{T})$, as well as additively separable and strictly concave in $\underline{x}$ and $\overline{x}$, the function $\mathcal{L}(\cdot;\lambda)$ is concave. Thus, the Kuhn-Tucker first-order conditions are sufficient for a global maximum of the relaxed problem.

In sum, the solution of the relaxed problem is the following:

$$(\underline{x}, \overline{x}) = \begin{cases} \left(\dfrac{\underline{\theta}+\overline{\theta}}{2}, \dfrac{1+\overline{\theta}}{2} \right), & \text{if } \mu > \tilde{\mu}_1, \\ \left(\dfrac{1+\underline{\theta}-\mu(1+\overline{\theta})}{2(1-\mu)}, \dfrac{1+\overline{\theta}}{2} \right), & \text{if } \mu \leq \tilde{\mu}_1, \end{cases}$$

$$\underline{T} = \begin{cases} \dfrac{\Delta^2}{4}, & \text{if } \mu > \tilde{\mu}_1, \\ \dfrac{\left(1-\underline{\theta}-\mu(1+\overline{\theta}-2\underline{\theta})\right)^2}{4(1-\mu)^2}, & \text{if } \mu \leq \tilde{\mu}_1, \end{cases}$$

$$\overline{T} = \begin{cases} \dfrac{\left(1-\overline{\theta}\right)^2}{4}, & \text{if } \mu > \tilde{\mu}_1, \\ \dfrac{(1-\overline{\theta})(1+3\overline{\theta}-4\underline{\theta})-\mu(1+\overline{\theta}-2\underline{\theta})^2}{4(1-\mu)}, & \text{if } \mu \leq \tilde{\mu}_1. \end{cases}$$

After some calculations, it is direct to see that the solution of the relaxed problem strictly satisfies ($IC_{\underline{\theta}}$). Thus, it is a solution of the original problem.

The interest group's expected payoff in the second period if it offers two contracts can be simplified to

$$V_2 = \begin{cases} (1-\mu)\left(-\frac{(1-\overline{\theta})^2}{2} - \frac{(1-\underline{\theta})^2}{2}\right) + \mu\left(-\frac{(1-\overline{\theta})^2}{2}\right) & \text{if } \mu > \tilde{\mu}_1, \\ \frac{1}{2(1-\mu)}\left((1-\mu)(-(1-\underline{\theta})^2) + \mu(\overline{\theta}-\underline{\theta})^2\right) & \text{if } \mu \leq \tilde{\mu}_1. \end{cases}$$

It remains to rule out menus in which only one type accepts, no type accepts, or a single nonempty contract is accepted by both types. By the Revelation Principle each such menu induces an allocation feasible in the program above: a rejection by type θ is reproduced by the outside-option contract $(\theta, 0)$, and acceptance of one common contract by both types is reproduced by setting the two type-specific contracts equal. We show each is strictly suboptimal. Let $z^* = (\underline{x}^*, \overline{x}^*, \underline{T}^*, \overline{T}^*)$ be the solution above and $\lambda^* \geq 0$ its multipliers, and let z be any feasible allocation. Because z satisfies the constraints and $\lambda^* \geq 0$, G's objective at z is at most $\mathcal{L}(z; \lambda^*)$, by concavity of $\mathcal{L}(\cdot; \lambda^*)$ and optimality of z^*, $\mathcal{L}(z; \lambda^*) \leq \mathcal{L}(z^*; \lambda^*)$, which equals G's objective at z^* by complementary slackness.

Since $\mathcal{L}(\cdot; \lambda^*)$ is strictly concave in $(\underline{x}, \overline{x})$, the inequality is strict if z and z^* differ in the policy coordinates. Every pooling allocation has $\underline{x} = \overline{x}$, whereas $\underline{x}^* \neq \overline{x}^*$ at every interior belief. Specifically, $\overline{x}^* = \frac{1+\overline{\theta}}{2}$ for all interior μ, while $\underline{x}^* = \frac{\theta+\overline{\theta}}{2}$ for $\mu > \tilde{\mu}_1$ and $\underline{x}^* \leq x_{\underline{\theta}} = \frac{1+\underline{\theta}}{2}$ for $\mu \leq \tilde{\mu}_1$. Thus, $\underline{\theta} < \overline{\theta} < 1$ implies $\underline{x}^* < \frac{1+\overline{\theta}}{2} = \overline{x}^*$, so $\underline{x}^* \neq \overline{x}^*$. Every one-type-acceptance

allocation fixes a policy at a type's ideal point, which likewise differs from z^* at interior beliefs. Therefore all such menus are strictly suboptimal, and the second-period equilibrium menu is separating. □

The interest group's second-period expected payoff is:

$$V_G(\mu) = \begin{cases} \frac{1}{2(1-\mu)}\left((1-\mu)(-(1-\underline{\theta})^2) + \mu(\overline{\theta}-\underline{\theta})^2\right) & \text{if } \mu \leq \tilde{\mu}_1, \\ (1-\mu)\left(-\frac{(1-\overline{\theta})^2}{2} - \frac{(1-\underline{\theta})^2}{2}\right) + \mu\left(-\frac{(1-\overline{\theta})^2}{2}\right) & \text{if } \mu > \tilde{\mu}_1. \end{cases}$$

A.2 Proofs of Remarks 1 and 2

By Lemma 1, the expected payoffs for P and G in equilibrium as a function of belief μ are:

$$V_P(\theta, \mu) = \begin{cases} \frac{(\overline{\theta}-\underline{\theta})(1-\mu-\overline{\theta}+\mu\underline{\theta})}{1-\mu} & \text{if } \theta = \overline{\theta} \text{ and } \mu \leq \tilde{\mu}_1, \\ 0 & \text{if } \theta = \overline{\theta} \text{ and } \mu > \tilde{\mu}_1, \\ 0 & \text{if } \theta = \underline{\theta}; \end{cases}$$

and

$$V_G(\mu) = \begin{cases} \frac{1}{2(1-\mu)}\left((1-\mu)(-(1-\underline{\theta})^2) + \mu(\overline{\theta}-\underline{\theta})^2\right) & \text{if } \mu \leq \tilde{\mu}_1, \\ (1-\mu)\left(-\frac{(1-\overline{\theta})^2}{2} - \frac{(1-\underline{\theta})^2}{2}\right) + \mu\left(-\frac{(1-\overline{\theta})^2}{2}\right) & \text{if } \mu > \tilde{\mu}_1. \end{cases}$$

Equivalently, on $\mu \leq \tilde{\mu}_1$, $V_G(\mu) = \pi_{\underline{\theta}} + \frac{\mu\Delta^2}{2(1-\mu)} = (1-\mu)\pi_{\underline{\theta}} + \mu\left(\pi_{\underline{\theta}} + \frac{\Delta^2}{2(1-\mu)}\right)$, as expressed in Remark 1. Note that both functions are continuous in μ, since they are equal at the cutpoint $\tilde{\mu}_1$ and continuous everywhere else. □

A.3 Proof of Proposition 1

Let $\mu := \mu_0 \in (0,1)$ denote G's prior. We proceed in two steps. Step 1 focuses on separating contracts. Step 2 focuses on pooling contracts. Comparing the solutions of these cases is sufficient. If type θ rejects a menu and then chooses policy y on path, then an equivalent realized allocation and posterior can be replicated by a zero-transfer contract $(y, 0)$ assigned to that type. Thus, a menu in which only one type accepts is payoff-equivalent to a direct menu in which the other type accepts its zero-transfer policy contract. Likewise, if both types reject and choose different policies, then the induced outcome is separating. And if both reject and choose the same policy, then the induced outcome is pooling. Thus, it is

without loss to represent on-path outcomes by menus in which both types accept, and it suffices to solve G's problem within the separating and pooling classes and then compare.

Step 1. Separation. If G offers separating contracts, then P's on-path continuation values are $V_P(\overline{\theta},1) = V_P(\underline{\theta},0) = 0$ in equilibrium. By deviating, type $\underline{\theta}$ would have $V_P(\underline{\theta},\mu) = 0$ for all $\mu \in [0,1]$. By mimicking $\underline{\theta}$, type $\overline{\theta}$ would have $V_P(\overline{\theta},0) = (\overline{\theta}-\underline{\theta})(1-\overline{\theta})$.

Conditional on separation, G's optimization problem is:

$$\max_{\underline{x},\overline{x},\underline{T},\overline{T}\in\mathbb{R}} \mu\left(-(\overline{x}-1)^2 - \overline{T} + \delta_G V_G(1)\right) + (1-\mu)\left(-(\underline{x}-1)^2 - \underline{T} + \delta_G V_G(0)\right)$$
$$s.t.\ (IC_{\overline{\theta}}), (IC_{\underline{\theta}}), (P_{\overline{\theta}}) \text{ and } (P_{\underline{\theta}}),$$

where the constraints are

$$-(\overline{x}-\overline{\theta})^2 + \overline{T} + \delta_P V_P(\overline{\theta},1) \geq -(\underline{x}-\overline{\theta})^2 + \underline{T} + \delta_P V_P(\overline{\theta},0), \qquad (IC_{\overline{\theta}})$$
$$-(\underline{x}-\underline{\theta})^2 + \underline{T} + \delta_P V_P(\underline{\theta},0) \geq -(\overline{x}-\underline{\theta})^2 + \overline{T} + \delta_P V_P(\underline{\theta},1), \qquad (IC_{\underline{\theta}})$$
$$-(\overline{x}-\overline{\theta})^2 + \overline{T} + \delta_P V_P(\overline{\theta},1) \geq 0, \qquad (P_{\overline{\theta}})$$
$$-(\underline{x}-\underline{\theta})^2 + \underline{T} + \delta_P V_P(\underline{\theta},0) \geq 0. \qquad (P_{\underline{\theta}})$$

We discipline beliefs after an off-path *rejection* using the Never-a-Weak-Best-Response (NWBR) refinement (Fudenberg and Tirole, 1991) where it has bite, and otherwise impose $\mu_1 = \tilde{\mu}_1$. In either case, the induced off-path belief is in $[\tilde{\mu}_1, 1]$, so both types of P would anticipate zero continuation surplus following such a rejection (Remark 2).[32] This zero-surplus continuation value is P's outside option rejection payoff that enters the participation constraints below.

We begin our analysis by setting up the Lagrangian:

$$\begin{aligned}\mathcal{L} =& \mu\left(-(\overline{x}-1)^2 - \overline{T} + \delta_G V_G(1)\right) + (1-\mu)\left(-(\underline{x}-1)^2 - \underline{T} + \delta_G V_G(0)\right)\\ &+\lambda_1\left(-(\overline{x}-\overline{\theta})^2 + \overline{T} + \delta_P V_P(\overline{\theta},1) + (\underline{x}-\overline{\theta})^2 - \underline{T} - \delta_P V_P(\overline{\theta},0)\right)\\ &+\lambda_2\left(-(\underline{x}-\underline{\theta})^2 + \underline{T} + \delta_P V_P(\underline{\theta},0) + (\overline{x}-\underline{\theta})^2 - \overline{T} - \delta_P V_P(\underline{\theta},1)\right)\\ &+\lambda_3\left(-(\overline{x}-\overline{\theta})^2 + \overline{T} + \delta_P V_P(\overline{\theta},1)\right)\\ &+\lambda_4\left(-(\underline{x}-\underline{\theta})^2 + \underline{T} + \delta_P V_P(\underline{\theta},0)\right).\end{aligned}$$

[32] There are two cases. If a profitable rejection could come only from type $\overline{\theta}$, NWBR assigns $\mu_1 = 1$. Otherwise, NWBR has no bite, and we set $\mu_1 = \tilde{\mu}_1$. Either way the continuation surplus is zero.

The first-order conditions with respect to $\underline{x}, \overline{x}, \underline{T}, \overline{T}$ are:

$$\frac{\partial \mathcal{L}}{\partial \underline{x}} = -2(1-\mu)(\underline{x}-1) + 2\lambda_1(\underline{x}-\overline{\theta}) - 2\lambda_2(\underline{x}-\underline{\theta}) - 2\lambda_4(\underline{x}-\underline{\theta}) = 0, \tag{A.10}$$

$$\frac{\partial \mathcal{L}}{\partial \overline{x}} = -2\mu(\overline{x}-1) - 2\lambda_1(\overline{x}-\overline{\theta}) + 2\lambda_2(\overline{x}-\underline{\theta}) - 2\lambda_3(\overline{x}-\overline{\theta}) = 0, \tag{A.11}$$

$$\frac{\partial \mathcal{L}}{\partial \underline{T}} = -(1-\mu) - \lambda_1 + \lambda_2 + \lambda_4 = 0, \tag{A.12}$$

$$\frac{\partial \mathcal{L}}{\partial \overline{T}} = -\mu + \lambda_1 - \lambda_2 + \lambda_3 = 0. \tag{A.13}$$

The complementary slackness conditions are:

$$\lambda_1 \left(-(\overline{x}-\overline{\theta})^2 + \overline{T} + \delta_P V_P(\overline{\theta},1) + (\underline{x}-\overline{\theta})^2 - \underline{T} - \delta_P V_P(\overline{\theta},0)\right) = 0, \tag{A.14}$$

$$\lambda_2 \left(-(\underline{x}-\underline{\theta})^2 + \underline{T} + \delta_P V_P(\underline{\theta},0) + (\overline{x}-\underline{\theta})^2 - \overline{T} - \delta_P V_P(\underline{\theta},1)\right) = 0, \tag{A.15}$$

$$\lambda_3 \left(-(\overline{x}-\overline{\theta})^2 + \overline{T} + \delta_P V_P(\overline{\theta},1)\right) = 0, \tag{A.16}$$

$$\lambda_4 \left(-(\underline{x}-\underline{\theta})^2 + \underline{T} + \delta_P V_P(\underline{\theta},0)\right) = 0. \tag{A.17}$$

Thus,

$$\lambda_1 \left(-(\overline{x}-\overline{\theta})^2 + \overline{T} + (\underline{x}-\overline{\theta})^2 - \underline{T} - \delta_P(\overline{\theta}-\underline{\theta})(1-\overline{\theta})\right) = 0, \tag{A.18}$$

$$\lambda_2 \left(-(\underline{x}-\underline{\theta})^2 + \underline{T} + (\overline{x}-\underline{\theta})^2 - \overline{T}\right) = 0, \tag{A.19}$$

$$\lambda_3 \left(-(\overline{x}-\overline{\theta})^2 + \overline{T}\right) = 0, \tag{A.20}$$

$$\lambda_4 \left(-(\underline{x}-\underline{\theta})^2 + \underline{T}\right) = 0. \tag{A.21}$$

Suppose first that $\lambda_2 = 0$. By (A.12), we know that $\lambda_4 = \lambda_1 + (1-\mu) > 0$. This implies that $\left(-(\underline{x}-\underline{\theta})^2 + \underline{T}\right) = 0 \iff \underline{T} = (\underline{x}-\underline{\theta})^2$. Then, by (A.13), we know that $\mu = \lambda_1 + \lambda_3$. Given that λ_1, λ_3 are non-negative, we know that there are three cases: (i) both $\lambda_1 > 0$ and $\lambda_3 > 0$, (ii) $\lambda_1 = \mu$ and $\lambda_3 = 0$, and (iii) $\lambda_1 = 0$ and $\lambda_3 = \mu$.

Case (i). The fact that $\lambda_1 > 0$ and $\lambda_3 > 0$ implies that constraints $(IC_{\overline{\theta}}), (P_{\overline{\theta}})$ and $(P_{\underline{\theta}})$ are all binding. These binding constraints plus the first-order conditions create a system of

7 equations and 7 unknowns with the following solution:

$$\underline{x}^* = \frac{1}{2}(\overline{\theta} + \underline{\theta} - \delta_P(1-\overline{\theta})),\ \overline{x}^* = \frac{1+\overline{\theta}}{2},\ \underline{T}^* = \frac{1}{4}(\Delta - \delta_P(1-\overline{\theta}))^2,\ \overline{T}^* = \frac{(1-\overline{\theta})^2}{4},$$
$$\lambda_1^* = \frac{(1+\delta_P)(1-\mu)(1-\overline{\theta})}{\Delta},\ \lambda_3^* = \frac{\delta_P(\mu(1-\overline{\theta}) + \overline{\theta} - 1) + \mu(-\underline{\theta}) + \mu + \overline{\theta} - 1}{\Delta},$$
$$\lambda_4^* = \frac{(1-\mu)(\delta_P(1-\overline{\theta}) + 1 - \underline{\theta})}{\Delta}.$$

Under the condition $\mu > \tilde{\mu}_0$, the multiplier $\lambda_3^* > 0$ and $(IC_{\underline{\theta}})$ is satisfied, which implies that $\left(\underline{x}^*, \overline{x}^*, \underline{T}^*, \overline{T}^*, \lambda_1^*, \lambda_2^*, \lambda_3^*, \lambda_4^*\right)$ with $\lambda_2^* = 0$ satisfies the Kuhn-Tucker first-order necessary conditions of the maximization problem.

Case (ii). Now $(IC_{\overline{\theta}})$ and $(P_{\underline{\theta}})$ are binding and, together with the first-order conditions, create a system of 6 equations and 6 unknowns with the following solution:

$$\underline{x}' = \frac{1 + \underline{\theta} - \mu(1+\overline{\theta})}{2(1-\mu)},\ \overline{x}' = \frac{1+\overline{\theta}}{2},$$
$$\underline{T}' = \frac{(1 - \underline{\theta} - \mu(1 + \overline{\theta} - 2\underline{\theta}))^2}{4(1-\mu)^2},$$
$$\overline{T}' = \frac{\mu\left((1-4\delta_P)\overline{\theta}^2 + \overline{\theta}(4\delta_P(\underline{\theta}+1) - 4\underline{\theta} + 2) - 4(\delta_P+1)\underline{\theta} + 4\underline{\theta}^2 + 1\right)}{4(\mu-1)}$$
$$+ \frac{(\overline{\theta}-1)((4\delta_P+3)\overline{\theta} - 4(\delta_P+1)\underline{\theta} + 1)}{4(\mu-1)},$$
$$\lambda_1' = \mu,\ \lambda_3' = 0,\ \lambda_4' = 1.$$

Replacing these values on the constraint $(P_{\overline{\theta}})$ we obtain that it is satisfied if $\mu \leq \tilde{\mu}_0$.

Also, the constraint $(IC_{\underline{\theta}})$ is equivalent to $\delta_P(1-\mu)(1-\overline{\theta}) < \Delta$, which covers the case with $\delta_P = 0$. In that no-reputation boundary case, $(IC_{\underline{\theta}})$ is $0 < \Delta$, which holds for all μ, so $(IC_{\underline{\theta}})$ is slack and $\hat{\mu}_0 = 0$. For $\delta_P > 0$, dividing by $\delta_P(1-\overline{\theta}) > 0$ gives the equivalent condition $\mu > \frac{\delta_P(1-\overline{\theta}) - \Delta}{\delta_P(1-\overline{\theta})}$. This cutpoint is below $\tilde{\mu}_0$ and positive if and only if $\delta_P > \delta_P^* = \frac{\Delta}{1-\overline{\theta}}$. Thus, if $\max\{0, \frac{\delta_P(1-\overline{\theta}) - \Delta}{\delta_P(1-\overline{\theta})}\} < \mu \leq \tilde{\mu}_0$, the vector $\left(\underline{x}', \overline{x}', \underline{T}', \overline{T}', \lambda_1', \lambda_2', \lambda_3', \lambda_4'\right)$ with $\lambda_2' = 0$ satisfies the Kuhn-Tucker first-order necessary conditions of the maximization problem.

Case (iii). Now $(P_{\overline{\theta}})$ and $(P_{\underline{\theta}})$ are binding and, together with the first-order conditions, create a system of 6 equations and 6 unknowns with the following solution:

$$\underline{x}'' = \frac{1+\underline{\theta}}{2},\ \overline{x}'' = \frac{1+\overline{\theta}}{2},\ \underline{T}'' = \frac{1}{4}(1-\underline{\theta})^2,\ \overline{T}'' = \frac{1}{4}(1-\overline{\theta})^2,$$
$$\lambda_1'' = 0,\ \lambda_3'' = \mu,\ \lambda_4'' = (1-\mu).$$

Replacing these values on the constraint $(IC_{\overline{\theta}})$ yields $-(1+\delta_P)(1-\overline{\theta})\Delta \geq 0$, which is false. Thus, there are no solutions that satisfy the conditions in case (iii).

Now, suppose $\lambda_2 > 0$. Suppose also that $\lambda_4 > 0$. We have three other cases: (iv) both $\lambda_1 > 0$ and $\lambda_3 > 0$, (v) $\lambda_1 > 0$ and $\lambda_3 = 0$, and (vi) $\lambda_1 = 0$ and $\lambda_3 > 0$.

Case (iv). The fact that $\lambda_1 > 0$, $\lambda_2 > 0$ and $\lambda_3 > 0$ implies that all the constraints $(IC_{\overline{\theta}}), (IC_{\underline{\theta}}), (P_{\overline{\theta}})$ and $(P_{\underline{\theta}})$ are binding. These binding constraints plus the first-order conditions create a system of 8 equations and 8 unknowns with the following solution:

$$\underline{x}^* = \frac{1}{2}(\overline{\theta} + \underline{\theta} - \delta_P(1-\overline{\theta})),\ \overline{x}^* = \frac{\overline{\theta}+\underline{\theta}}{2}, \underline{T}^* = \frac{1}{4}(\Delta - \delta_P(1-\overline{\theta}))^2,\ \overline{T}^* = \frac{1}{4}\Delta^2,$$
$$\lambda_1^* = \frac{(1+\delta_P)(1-\overline{\theta})(1-\mu)}{\Delta},\ \lambda_2^* = -\frac{\mu(1-\underline{\theta})}{\Delta},\ \lambda_3^* = \frac{-\delta_P(1-\mu)(1-\overline{\theta}) - (1-\overline{\theta})}{\Delta},$$
$$\lambda_4^* = \frac{\delta_P(1-\mu)(1-\overline{\theta}) + (1-\underline{\theta})}{\Delta}.$$

In this case we have that $\lambda_2^* < 0$. Thus, there are no solutions that satisfy the conditions in case (iv).

Case (v). Now $(IC_{\overline{\theta}}), (IC_{\underline{\theta}})$ and $(P_{\underline{\theta}})$ are binding and, together with the first-order conditions, create a system of 7 equations and 7 unknowns with the following solution:

$$\hat{\underline{x}} = \frac{1}{2}(1 + \underline{\theta} - \delta_P\mu(1-\overline{\theta})),\ \hat{\overline{x}} = \frac{1}{2}(1 + \underline{\theta} + \delta_P(1-\mu)(1-\overline{\theta})),$$
$$\hat{\underline{T}} = \frac{1}{4}(\delta_P\mu(1-\overline{\theta}) - (1-\underline{\theta}))^2,$$
$$\hat{\overline{T}} = \frac{1}{4}(\delta_P(1-\mu)(1-\overline{\theta}) + (1-\underline{\theta}))^2,$$
$$\hat{\lambda}_1 = \frac{\delta_P(1-\overline{\theta})\mu\,(1-\mu)}{\Delta},\ \hat{\lambda}_2 = \mu\frac{(\delta_P(1-\mu)(1-\overline{\theta}) - \Delta)}{\Delta},\ \hat{\lambda}_4 = 1.$$

We have that $\hat{\lambda}_1 \geq 0$, but $\hat{\lambda}_2 \geq 0$ if and only if $\mu \leq \frac{\delta_P(1-\overline{\theta})-\Delta}{\delta_P(1-\overline{\theta})}$. Replacing the solution on the constraint $(P_{\overline{\theta}})$ we obtain that it is satisfied. Thus, when $\frac{\delta_P(1-\overline{\theta})-\Delta}{\delta_P(1-\overline{\theta})} > 0$, if $\mu \leq \frac{\delta_P(1-\overline{\theta})-\Delta}{\delta_P(1-\overline{\theta})}$, the vector $\left(\hat{\underline{x}}, \hat{\overline{x}}, \hat{\underline{T}}, \hat{\overline{T}}, \hat{\lambda}_1, \hat{\lambda}_2, \hat{\lambda}_3, \hat{\lambda}_4\right)$ with $\hat{\lambda}_3 = 0$ satisfies the Kuhn-Tucker first-order necessary conditions of the maximization problem.

Case (vi). Now $(IC_{\underline{\theta}}), (P_{\overline{\theta}})$ and $(P_{\underline{\theta}})$ are binding and, together with the first-order conditions, create a system of 7 equations and 7 unknowns with the following solution:

$$\underline{x}'' = \frac{1+\underline{\theta}}{2},\ \overline{x}'' = \frac{\overline{\theta}+\underline{\theta}}{2},\ \underline{T}'' = \frac{1}{4}(1-\underline{\theta})^2,\ \overline{T}'' = \frac{\Delta^2}{4},$$
$$\lambda_2'' = -\frac{\mu(1-\underline{\theta})}{\Delta},\ \lambda_3'' = -\frac{(1-\overline{\theta})\mu}{\Delta},\ \lambda_4'' = \frac{\Delta + \mu(1-\overline{\theta})}{\Delta}.$$

We have that $\lambda_2'' < 0$. Thus, there are no solutions that satisfy the conditions in case (vi).

It remains to rule out $\lambda_4 = 0$. Setting $\lambda_4 = 0$ in (A.12) gives $\lambda_2 = (1-\mu) + \lambda_1 > 0$, so this is the remaining $\lambda_2 > 0$ subcase; and by (A.13), $\lambda_3 = \mu - \lambda_1 + \lambda_2 = 1 > 0$. Therefore $(IC_{\underline{\theta}})$ and $(P_{\overline{\theta}})$ bind for every admissible $\lambda_1 \geq 0$. Binding $(P_{\overline{\theta}})$ gives $\overline{T} = (\overline{x} - \overline{\theta})^2$, and the first-order condition for $\overline{x}$ yields:

$$\overline{x} = \frac{\mu - (1-\mu)\underline{\theta} + \overline{\theta} + \lambda_1 \Delta}{2\mu}, \qquad \text{hence} \qquad 2\overline{x} - \underline{\theta} - \overline{\theta} = \frac{\mu(1-\overline{\theta}) + \Delta(1+\lambda_1)}{\mu} > 0.$$

Because $(IC_{\underline{\theta}})$ binds, an adversary's equilibrium payoff equals her payoff from the ally-contract,

$$-(\underline{x} - \underline{\theta})^2 + \underline{T} = -(\overline{x} - \underline{\theta})^2 + \overline{T} = (\overline{x} - \overline{\theta})^2 - (\overline{x} - \underline{\theta})^2 = -\Delta\big(2\overline{x} - \underline{\theta} - \overline{\theta}\big) < 0,$$

which violates $(P_{\underline{\theta}})$. Note that having no solutions with $\lambda_4 = 0$ is consistent with the adversary's participation constraint binding in every equilibrium, which we use in the proofs of Propositions 2–5.

Note that $\frac{\partial^2 \mathcal{L}}{\partial \underline{x}^2} = -2(1-\mu) + 2(\lambda_1 - \lambda_2 - \lambda_4) = -4(1-\mu) < 0$ by (A.12), and $\frac{\partial^2 \mathcal{L}}{\partial \overline{x}^2} = -2\mu - 2(\lambda_1 - \lambda_2 + \lambda_3) = -4\mu < 0$ by (A.13). Since $\mathcal{L}$ is linear in $(\underline{T}, \overline{T})$, as well as additively separable and strictly concave in $\underline{x}$ and $\overline{x}$, $\mathcal{L}(\cdot; \lambda)$ is concave in $(\underline{x}, \overline{x}, \underline{T}, \overline{T})$. Thus, the Kuhn-Tucker conditions are sufficient for a global maximum.

In sum, the solution of the maximization problem is the following:

$$(\underline{x}, \overline{x}, \underline{T}, \overline{T}) = \begin{cases} \left(\underline{x}^*, \overline{x}^*, \underline{T}^*, \overline{T}^*\right) & \text{if } \mu > \tilde{\mu}_0, \\ \left(\underline{x}', \overline{x}', \underline{T}', \overline{T}'\right) & \text{if } \hat{\mu}_0 < \mu \leq \tilde{\mu}_0, \\ \left(\hat{\underline{x}}, \hat{\overline{x}}, \hat{\underline{T}}, \hat{\overline{T}}\right) & \text{if } 0 \leq \mu \leq \hat{\mu}_0. \end{cases} \tag{A.22}$$

Thus, G's expected payoff from separation equals

$$V^{SEP} = (1-\mu)\left(-(\underline{x} - 1)^2 - \underline{T} + \delta_G V_G(0)\right) + \mu\left(-(\overline{x} - 1)^2 - \overline{T} + \delta_G V_G(1)\right),$$

where $(\underline{x}, \overline{x}, \underline{T}, \overline{T})$ is as in equation (A.22) and $V_G(0)$ and $V_G(1)$ are G's continuation values as in Remark 1.

Step 2. Pooling. Conditional on offering a pooling menu, G's constrained maximization

problem is:

$$\max_{x,T\in\mathbb{R}} \mu\left(-(x-1)^2 - T + \delta_G V_G(\mu)\right) + (1-\mu)\left(-(x-1)^2 - T + \delta_G V_G(\mu)\right)$$
$$s.t.\ (P_{\overline{\theta}}) \text{ and } (P_{\underline{\theta}}),$$

where the constraints are:

$$-(x-\overline{\theta})^2 + T + \delta_P V_P(\overline{\theta},\mu) \geq 0, \qquad (P_{\overline{\theta}})$$
$$-(x-\underline{\theta})^2 + T + \delta_P V_P(\underline{\theta},\mu) \geq 0. \qquad (P_{\underline{\theta}})$$

The left-hand side denotes the payoff of accepting and the right-hand side each type's payoff of rejecting. The solution to this problem is the contract $(x,\, T) = \left(\frac{1+\underline{\theta}}{2}, \frac{(1-\underline{\theta})^2}{4}\right)$. Thus, G's expected payoff of the pooling strategy is equal to

$$V^{POOL} = -\left(\frac{1+\underline{\theta}}{2} - 1\right)^2 - \frac{(1-\underline{\theta})^2}{4} + \delta_G V_G(\mu).$$

Finally, a direct comparison of V^{SEP} and V^{POOL} shows that there exists a cutpoint $\delta_G^\dagger$ such that $V^{SEP} \geq V^{POOL}$ if and only if $\delta_G \geq \delta_G^\dagger$. Specifically,

$$\delta_G^\dagger = \begin{cases} \frac{\delta_P^2(1-\mu)^2(1-\overline{\theta})^2}{\Delta(\mu(\overline{\theta}+\underline{\theta}-2)-2\overline{\theta}+2)} & \text{if } 0 \leq \mu \leq \hat{\mu}_0, \\ \frac{2\delta_P(1-\mu)(1-\overline{\theta})-\Delta}{2+\mu(\overline{\theta}+\underline{\theta}-2)-2\overline{\theta}} & \text{if } \hat{\mu}_0 < \mu \leq \tilde{\mu}_1, \\ \frac{\mu\Delta\left(2\delta_P(1-\mu)(1-\overline{\theta})-\Delta\right)}{(1-\mu)^2(1-\overline{\theta})^2} & \text{if } \tilde{\mu}_1 < \mu \leq \tilde{\mu}_0, \\ \delta_P^2 + 2\delta_P - \frac{\mu(1-\underline{\theta})^2}{(1-\mu)(1-\overline{\theta})^2} + \frac{1}{1-\mu} & \text{if } \tilde{\mu}_0 < \mu. \end{cases}$$

The cutpoint $\delta_G^\dagger$ is continuous in μ. Also, if $\mu > \tilde{\mu}_0$, then $\delta_G^\dagger < 0$ so in this region screening is always optimal. Furthermore, $\delta_G^\dagger$ is increasing in δ_P and decreasing in Δ. □

A.4 Proofs of Propositions 2–5

Throughout this subsection, cases 1–3 refer to the three belief regions of the first-period separating menu in Proposition 1: $\mu_0 < \hat{\mu}_0$ (case 1), $\mu_0 \in (\hat{\mu}_0, \tilde{\mu}_0)$ (case 2), and $\mu_0 > \tilde{\mu}_0$ (case 3).

The proofs use two facts established above. First, $(P_{\underline{\theta}})$ binds in every equilibrium and period, as its multiplier is positive in every case of the proofs of Lemma 1 and Proposition 1, so $\underline{T}_t = (\underline{x}_t - \underline{\theta})^2$. Thus, the adversary's transfer is equal to her full-information value

$T_{\underline{\theta}} = \frac{(1-\underline{\theta})^2}{4}$ when $\underline{x}_t = x_{\underline{\theta}}$ and to zero when $\underline{x}_t = \underline{\theta}$. By Lemma 1 and Proposition 1 this policy is distorted weakly downward from $x_{\underline{\theta}}$, by at most $\frac{1-\overline{\theta}}{2}$ in the second period and at most $(1+\delta_P)\frac{1-\overline{\theta}}{2}$ in the first. Both bounds are below $1-\underline{\theta} = 2(x_{\underline{\theta}} - \underline{\theta})$, so the policy stays within distance $\frac{1-\underline{\theta}}{2}$ of $\underline{\theta}$ and $\underline{T}_t \leq T_{\underline{\theta}}$. These are equal only absent distortion. As the policy shifts toward $\underline{\theta}$ the transfer decreases, and once it shifts below $\underline{\theta}$ it increases again. In the second period, the distortion is at most $\frac{1-\overline{\theta}}{2} < \frac{1-\underline{\theta}}{2}$, so $\underline{x}_2 > \underline{\theta}$. Recall $\delta_P^* = \frac{\Delta}{1-\overline{\theta}}$, which is the cutpoint on δ_P at which the case-3 policy $x_{\underline{\theta}} - (1+\delta_P)\frac{1-\overline{\theta}}{2}$ reaches $\underline{\theta}$. Define

$$\mu_0^* := \frac{1-\underline{\theta}}{1+\overline{\theta}-2\underline{\theta}}$$

as the prior at which the case-2 policy $x_{\underline{\theta}} - \frac{\mu_0}{1-\mu_0}\frac{\Delta}{2}$ reaches $\underline{\theta}$.

Second, after a separating first-period menu, μ_1 is degenerate, so by Lemma 1 each realized type's second-period contract is its full-information contract. Thus, conditional on separation, each type's policy and transfer shift over time by exactly its first-period distortion. Throughout, $\hat{\mu}_0$ and $\tilde{\mu}_0$ are increasing in δ_P, and the first-period contracts are continuous in (μ_0, δ_P), which follow from Proposition 1. In Proposition 2 and its proof, primed objects ($\hat{\mu}_0'$, $\tilde{\mu}_0'$, and $\delta_G^{\dagger\prime}$) are evaluated at the higher patience level δ_P'. Since both belief cutpoints increase in δ_P, the condition $\mu_0 \in (\hat{\mu}_0', \tilde{\mu}_0)$ keeps the menu in case 2 at every patience level in $[\delta_P, \delta_P']$. If instead $\mu_0 < \hat{\mu}_0'$ with $\mu_0 > \hat{\mu}_0$, the menu transitions from case 2 to case 1 as patience rises, and the conclusions of case 1 below extend across the transition by combining the within-case monotonicity with continuity of the contracts; the case-3-to-case-2 transition is handled within case 3 of the proof. Additionally, each branch of the cutpoint $\delta_G^\dagger$ in the proof of Proposition 1 is increasing in δ_P, and adjacent branches agree at their boundaries, so $\delta_G^\dagger$ is increasing in δ_P at every μ_0; thus, $\delta_G > \delta_G^{\dagger\prime}$ ensures that G separates at every patience level in $[\delta_P, \delta_P']$.

Proof of Proposition 2. Since $\delta_G > \delta_G^{\dagger\prime}$, G separates at every patience level in $[\delta_P, \delta_P']$. Accordingly, we use the characterization of that case in the proof of Proposition 1. There are three cases.

Case 1. If $\mu_0 < \hat{\mu}_0$, then $\underline{x}_1 = x_{\underline{\theta}} - \delta_P \mu_0 \frac{1-\overline{\theta}}{2}$ is greater than $\underline{\theta}$ and decreases in δ_P, so $\underline{T}_1$ decreases. Additionally, the distortion is less than $\frac{1-\overline{\theta}}{2} < \frac{1-\underline{\theta}}{2}$. Meanwhile, $\overline{x}_1 = x_{\underline{\theta}} + \delta_P(1-\mu_0)\frac{1-\overline{\theta}}{2}$ and $\overline{T}_1 = \frac{1}{4}\big((1-\underline{\theta}) + \delta_P(1-\mu_0)(1-\overline{\theta})\big)^2$ increase in δ_P.

Case 2. If $\mu_0 \in (\hat{\mu}_0, \tilde{\mu}_0)$, then $\overline{x}_1 = x_{\overline{\theta}}$ and $\underline{x}_1 = x_{\underline{\theta}} - \frac{\mu_0}{1-\mu_0}\frac{\Delta}{2}$ are independent of δ_P, so $\overline{x}_1$ and $(\underline{x}_1, \underline{T}_1)$ are constant. The ally premium $\overline{T}_1 - T_{\overline{\theta}} = \frac{\Delta}{1-\mu_0}\big[(1-\overline{\theta})(1+\delta_P(1-\mu_0)) - (1-\underline{\theta})\mu_0\big]$ has $\frac{\partial(\overline{T}_1 - T_{\overline{\theta}})}{\partial \delta_P} = \Delta(1-\overline{\theta}) > 0$, so $\overline{T}_1$ increases.

Case 3. Suppose $\mu_0 > \tilde{\mu}_0$. If $\mu_0 > \tilde{\mu}_0'$, then the menu remains in case 3 at every patience

level in $[\delta_P, \delta'_P]$: $(\overline{x}_1, \overline{T}_1) = (x_{\overline{\theta}}, T_{\overline{\theta}})$ is constant, while $\underline{x}_1 = x_{\underline{\theta}} - (1+\delta_P)\frac{1-\overline{\theta}}{2}$ decreases in δ_P and passes through $\underline{\theta}$ at $\delta_P = \delta_P^*$, so $\underline{T}_1$ decreases for $\delta_P < \delta_P^*$ and increases for $\delta_P > \delta_P^*$. If instead $\mu_0 \in (\tilde{\mu}_0, \tilde{\mu}'_0]$, the menu is in case 3 at patience levels below the value δ''_P solving $\tilde{\mu}_0(\delta''_P) = \mu_0$ and in case 2 above it, with the contracts continuous at the transition. The ally-policy equals $x_{\overline{\theta}}$ throughout. The ally-transfer equals $T_{\overline{\theta}}$ up to the transition and increases thereafter, since the case-2 premium $\overline{T}_1 - T_{\overline{\theta}}$ is zero at $\mu_0 = \tilde{\mu}_0$ and increases in δ_P. The adversary-policy decreases up to the transition and is constant thereafter, and $\underline{T}_1$ moves with it according to whether $\underline{x}_1$ lies above or below $\underline{\theta}$, as above. In both subcases, $\overline{x}_1$ is constant, $\overline{T}_1$ increases, $\underline{x}_1$ decreases, and $\underline{T}_1$ decreases while $\underline{x}_1 \geq \underline{\theta}$ and increases once $\underline{x}_1 < \underline{\theta}$, as stated. □

Proof of Proposition 3. Fix $\delta_G > \delta_G^\dagger$, so G separates. There are three cases.

Case 1. If $\mu_0 < \hat{\mu}_0$, then $\overline{x}_1 = x_{\underline{\theta}} + \delta_P(1-\mu_0)\frac{1-\overline{\theta}}{2}$ and $\overline{T}_1 = \frac{1}{4}((1-\underline{\theta}) + \delta_P(1-\mu_0)(1-\overline{\theta}))^2$ decrease in μ_0. Meanwhile, $\underline{x}_1 = x_{\underline{\theta}} - \delta_P\mu_0\frac{1-\overline{\theta}}{2}$ lies above $\underline{\theta}$ and decreases in μ_0, so $\underline{T}_1$ decreases. Thus, both contracts decrease.

Case 2. If $\mu_0 \in (\hat{\mu}_0, \tilde{\mu}_0)$, then $\overline{x}_1 = x_{\overline{\theta}}$ is constant and the ally premium $\overline{T}_1 - T_{\overline{\theta}}$ has $\frac{\partial(\overline{T}_1 - T_{\overline{\theta}})}{\partial \mu_0} = -\frac{\Delta^2}{(1-\mu_0)^2} < 0$, so $\overline{T}_1$ decreases. The policy $\underline{x}_1 = x_{\underline{\theta}} - \frac{\mu_0}{1-\mu_0}\frac{\Delta}{2}$ decreases in μ_0 and passes through $\underline{\theta}$ at $\mu_0 = \mu_0^*$. Thus, $\underline{T}_1$ decreases for $\mu_0 < \mu_0^*$ and increases for $\mu_0 \in (\mu_0^*, \tilde{\mu}_0)$. Finally, if $\delta_P \leq \delta_P^*$, then $\mu_0^* \geq \tilde{\mu}_0$, so $\underline{x}_1 > \underline{\theta}$ always holds and $\underline{T}_1$ decreases throughout this case.

Case 3. If $\mu_0 > \tilde{\mu}_0$, then $(\overline{x}_1, \overline{T}_1) = (x_{\overline{\theta}}, T_{\overline{\theta}})$ and $\underline{x}_1 = x_{\underline{\theta}} - (1+\delta_P)\frac{1-\overline{\theta}}{2}$ are constant. □

Proof of Proposition 4. By Proposition 1 and Lemma 1, on-path histories fall into six classes. Under separation, μ_1 is degenerate and the realized type's second-period contract is full-information; the ally history differs across cases 1–3, labeled (B)–(D), while the adversary histories of the three cases admit a uniform argument and are grouped as (A). Under pooling, $\mu_1 = \mu_0$ and the second-period menu separates, giving an adversary history (E) and an ally history (F). We consider each class in turn.

First, consider separation that reveals an adversary (A). Then $x_1 = \underline{x}_1 < x_{\underline{\theta}} = x_2$, so the policy strictly increases. By the preliminary facts noted above, the separating distortion keeps $\underline{x}_1$ strictly within distance $\frac{1-\theta}{2}$ of $\underline{\theta}$, so $\underline{T}_1 < T_{\underline{\theta}} = T_2$. Both policy and transfer strictly increase.

Next, consider separation that reveals an ally. (B) In case 1, $x_1 = \overline{x}_1 = x_{\overline{\theta}} + \frac{1}{2}\big(\delta_P(1-\mu_0)(1-\overline{\theta}) - \Delta\big) \geq x_{\overline{\theta}} = x_2$, since this case requires $\delta_P(1-\mu_0)(1-\overline{\theta}) \geq \Delta$, and $T_1 = \overline{T}_1 = \frac{1}{4}\big((1-\underline{\theta}) + \delta_P(1-\mu_0)(1-\overline{\theta})\big)^2 > \frac{(1-\overline{\theta})^2}{4} = T_{\overline{\theta}} = T_2$. The policy weakly decreases and the transfer strictly decreases. (C) In case 2, $x_1 = x_2 = x_{\overline{\theta}}$, while $T_1 = \overline{T}_1 \geq T_{\overline{\theta}} = T_2$, since the ally premium $\overline{T}_1 - T_{\overline{\theta}}$, displayed in the proof of Proposition 2, is positive for $\mu_0 < \tilde{\mu}_0$

and vanishes at $\tilde{\mu}_0$. The policy is constant and the transfer weakly decreases, strictly so for $\mu_0 < \tilde{\mu}_0$. (D) In case 3, $x_1 = x_2 = x_{\overline{\theta}}$ and $T_1 = T_2 = T_{\overline{\theta}}$, so both are constant.

Now consider pooling, after which the second-period separating menu reveals the type. If P chooses the adversary-contract (E), then $x_1 = x_{\underline{\theta}}$ and $T_1 = T_{\underline{\theta}}$, while $x_2 = \underline{x}_2 \in (\underline{\theta}, x_{\underline{\theta}})$ and $T_2 = (\underline{x}_2 - \underline{\theta})^2 < T_{\underline{\theta}}$, by the preliminary facts above. Both policy and transfer strictly decrease.

If P chooses the ally-contract (F), then $x_1 = x_{\underline{\theta}} < x_{\overline{\theta}} = x_2$, so the policy strictly increases, and $T_2 - T_1 = \overline{T}_2 - T_{\underline{\theta}}$. If $\mu_0 > \tilde{\mu}_1$, then $\overline{T}_2 = T_{\overline{\theta}}$ by Lemma 1, so $T_2 - T_1 = T_{\overline{\theta}} - T_{\underline{\theta}} < 0$. This history is on path for some parameters: if $2\delta_P(1-\mu_0)(1-\overline{\theta}) > \Delta$, then $\delta_G^\dagger > 0$ (as shown in the proof of Proposition 1) and pooling occurs for $\delta_G < \delta_G^\dagger$. If instead $\mu_0 \leq \tilde{\mu}_1$, then $\overline{T}_2 = T_{\overline{\theta}} + S(\mu_0)$, so $T_2 - T_1 = T_{\overline{\theta}} + S(\mu_0) - T_{\underline{\theta}}$. Since $S'(\mu_0) = -\Delta^2/(1-\mu_0)^2 < 0$, this difference is strictly decreasing in μ_0, and at $\mu_0 = 0$ it equals $T_{\overline{\theta}} + S(0) - T_{\underline{\theta}} = \frac{\Delta}{4}\big(3(1-\overline{\theta}) - (1-\underline{\theta})\big)$. Thus, if $1 - \underline{\theta} < 3(1-\overline{\theta})$, the transfer increases for small μ_0 and decreases for larger μ_0 on this interval; otherwise, it weakly decreases throughout. In every subcase of (F), the policy strictly increases, while the transfer can move in either direction.

Altogether, the transfer strictly increases only in (A) and in (F) when $T_2 > T_1$, and the policy strictly increases in both, so $T_1 < T_2$ implies $x_1 < x_2$. The policy strictly decreases only in (B) and (E), and the transfer strictly decreases in both, so $x_1 > x_2$ implies $T_1 > T_2$. Finally, if $\mu_0 > \tilde{\mu}_0$, then $\delta_G^\dagger < 0$, so G separates and the ally is in case 3: only (A) and (D) arise, and both have $x_1 \leq x_2$ and $T_1 \leq T_2$. Otherwise, (B), (C), and (E) show that weakly decreasing paths, with $x_1 \geq x_2$ and $T_1 \geq T_2$, are also possible. □

Proof of Proposition 5. Since $\delta_G > \delta_G^\dagger$, G separates, so by the preliminary facts above each type's second-period contract is full-information and the over-time shifts equal the first-period distortions: $|\overline{x}_1 - \overline{x}_2| = \overline{x}_1 - x_{\overline{\theta}}$, $|\overline{T}_1 - \overline{T}_2| = \overline{T}_1 - T_{\overline{\theta}}$, $|\underline{x}_1 - \underline{x}_2| = x_{\underline{\theta}} - \underline{x}_1$, and $|\underline{T}_1 - \underline{T}_2| = T_{\underline{\theta}} - \underline{T}_1$. Also by the preliminary facts, the first-period contracts are continuous in (μ_0, δ_P), so each distortion is continuous across the case boundaries and can be signed case by case. Finally, the adversary-transfer shift $T_{\underline{\theta}} - \underline{T}_1 = T_{\underline{\theta}} - (\underline{x}_1 - \underline{\theta})^2$ is largest when $\underline{x}_1 = \underline{\theta}$ and shrinks as $\underline{x}_1$ moves away from $\underline{\theta}$ in either direction.

Part 1. Consider an increase in δ_P. The ally-policy distortion $\overline{x}_1 - x_{\overline{\theta}}$ equals $\frac{1}{2}\big(\delta_P(1-\mu_0)(1-\overline{\theta}) - \Delta\big)$ in case 1, where it increases, and equals 0 in cases 2 and 3. The adversary-policy distortion $x_{\underline{\theta}} - \underline{x}_1$ increases in cases 1 and 3, where it equals $\delta_P\mu_0\frac{1-\overline{\theta}}{2}$ and $(1+\delta_P)\frac{1-\overline{\theta}}{2}$ respectively, and is constant in case 2. The ally-transfer distortion $\overline{T}_1 - T_{\overline{\theta}}$ increases in case 1; equals the ally premium in case 2, which increases in δ_P by the proof of Proposition 2; and equals 0 in case 3. Thus, $|x_1 - x_2|$ for both types and $|\overline{T}_1 - \overline{T}_2|$ weakly increase in δ_P. For the adversary-transfer shift, suppose first that $\mu_0 \leq \tilde{\mu}_0$. In case 1, the policy $\underline{x}_1$ lies above $\underline{\theta}$ and moves toward it, as shown in the proof of Proposition 2, so the shift increases. In case

2, the adversary-contract does not vary with δ_P, so the shift is constant. Thus, $|\underline{T}_1 - \underline{T}_2|$ weakly increases. Next, suppose $\mu_0 > \tilde{\mu}_0$. Then $\underline{x}_1$ decreases in δ_P and crosses $\underline{\theta}$ at $\delta_P = \delta_P^*$, so $|\underline{T}_1 - \underline{T}_2|$ increases over $\delta_P < \delta_P^*$ and decreases over $\delta_P > \delta_P^*$.

Part 2. Consider an increase in μ_0. The ally-policy distortion decreases in case 1, and equals 0 in cases 2 and 3. The ally-transfer distortion decreases in case 1; equals the ally premium in case 2, which decreases in μ_0 by the proof of Proposition 3; and equals 0 in case 3. Thus, $|\overline{x}_1 - \overline{x}_2|$ and $|\overline{T}_1 - \overline{T}_2|$ weakly decrease. The adversary-policy distortion increases in cases 1 and 2 and is constant in case 3, so $|\underline{x}_1 - \underline{x}_2|$ weakly increases. Finally, $\underline{x}_1$ decreases in cases 1 and 2, and is constant in case 3. If $\delta_P > \delta_P^*$, then it crosses $\underline{\theta}$ at μ_0^* within case 2; if instead $\delta_P \leq \delta_P^*$, then it remains above $\underline{\theta}$ throughout cases 1 and 2, since $\mu_0^* \geq \tilde{\mu}_0$. Thus, $|\underline{T}_1 - \underline{T}_2|$ increases over $\mu_0 < \min\{\mu_0^*, \tilde{\mu}_0\}$, decreases over $\mu_0 \in (\mu_0^*, \tilde{\mu}_0)$, and is constant over $\mu_0 > \tilde{\mu}_0$. □

Expected policy shift. The claim in the text follows from the contracts above. Under a separating first-period menu, $\mathbb{E}[x_2 - x_1] = \mu_0(x_{\overline{\theta}} - \overline{x}_1) + (1-\mu_0)(x_{\underline{\theta}} - \underline{x}_1)$, which equals $\mu_0 \frac{\Delta}{2}$ in cases 1 and 2 and $(1-\mu_0)(1+\delta_P)\frac{1-\overline{\theta}}{2}$ in case 3. Each is strictly positive for all $\mu_0 \in (0,1)$ and zero only under complete information. Under a pooling first-period menu both types receive $x_{\underline{\theta}}$, so $\mathbb{E}[x_2 - x_1] = \mu_0 \frac{\Delta}{2} - (1-\mu_0)(x_{\underline{\theta}} - \underline{x}_2)$, which is zero for $\mu_0 \leq \tilde{\mu}_1$, where $x_{\underline{\theta}} - \underline{x}_2 = \frac{\mu_0}{1-\mu_0}\frac{\Delta}{2}$, and positive for $\mu_0 > \tilde{\mu}_1$, where $x_{\underline{\theta}} - \underline{x}_2 = \frac{1-\overline{\theta}}{2}$. □

B Proofs of Extensions

We repeatedly use two facts from the proof of Proposition 1. First, after a separating first-period menu μ_1 is degenerate, so by Lemma 1 the second-period menu is a full-information contract. Second, the optimal pooled contract is the adversary-efficient contract $(x_{\underline{\theta}}, T_{\underline{\theta}})$, which both types accept. Thus, let $\Pi^* := (1-\mu_0)\pi_{\underline{\theta}} + \mu_0 \pi_{\overline{\theta}}$ denote G's per-period full-information payoff and Π_1^{SEP} its expected first-period payoff under the separating menu of Proposition 1, so we can decompose G's payoffs to facilitate the following proofs.

Remark B.1. Under separation and pooling, respectively, G's equilibrium cumulative payoffs are $V^{SEP} = \Pi_1^{SEP} + \delta_G \Pi^*$ and $V^{POOL} = \pi_{\underline{\theta}} + \delta_G V_G(\mu_0)$. The value of learning is $W(\mu_0) = \Pi^* - V_G(\mu_0)$, where:

$$W(\mu_0) = \begin{cases} \dfrac{\mu_0 \Delta}{2}\left[(2 - \underline{\theta} - \overline{\theta}) - \dfrac{\Delta}{1-\mu_0}\right] & \text{if } \mu_0 \leq \tilde{\mu}_1, \\ \dfrac{(1-\mu_0)(1-\overline{\theta})^2}{2} & \text{if } \mu_0 > \tilde{\mu}_1, \end{cases}$$

which is strictly positive for all $\mu_0 \in (0,1)$.

PROOF. If the equilibrium menu is separating, G initially receives its expected first-period payoff Π_1^{SEP} and then the full-information continuation $(1-\mu_0)\,V_G(0)+\mu_0\,V_G(1)=(1-\mu_0)\pi_{\underline{\theta}}+\mu_0\pi_{\overline{\theta}}=\Pi^*$, which yields V^{SEP}. If the equilibrium menu is pooling, G initially receives the adversary-efficient payoff $\pi_{\underline{\theta}}$ and then the continuation $V_G(\mu_0)$ from $\mu_1=\mu_0$, which yields V^{POOL}. By Definition 1, together with $V_G(0)=\pi_{\underline{\theta}}$ and $V_G(1)=\pi_{\overline{\theta}}$, the value of learning is $W(\mu_0)=\Pi^*-V_G(\mu_0)$. Substituting the continuation values of Remark 1 gives the displayed cases: for $\mu_0\le\tilde{\mu}_1$, using $\pi_{\overline{\theta}}-\pi_{\underline{\theta}}=\frac{\Delta(2-\underline{\theta}-\overline{\theta})}{2}$ and $V_G(\mu_0)=\pi_{\underline{\theta}}+\frac{\mu_0\Delta^2}{2(1-\mu_0)}$; and for $\mu_0>\tilde{\mu}_1$, $W(\mu_0)=-(1-\mu_0)\pi_{\overline{\theta}}=\frac{(1-\mu_0)(1-\overline{\theta})^2}{2}$. Positivity for the second case is immediate. For the first case, $\frac{\Delta}{1-\mu_0}$ increases in μ_0 to $1-\underline{\theta}$ at $\mu_0=\tilde{\mu}_1$, below $2-\underline{\theta}-\overline{\theta}=(1-\underline{\theta})+(1-\overline{\theta})$, so the bracket is positive. The comparison $V^{SEP}-V^{POOL}$ and the cutpoint follow directly from the two expressions and $W(\mu_0)>0$. □

Note that $V^{SEP}-V^{POOL}=(\Pi_1^{SEP}-\pi_{\underline{\theta}})+\delta_G\,W(\mu_0)$ is strictly increasing in δ_G. Thus, G separates if $\delta_G>\delta_G^{\dagger}$ and pools if $\delta_G<\delta_G^{\dagger}$, where $\delta_G^{\dagger}=\dfrac{\pi_{\underline{\theta}}-\Pi_1^{SEP}}{W(\mu_0)}$. This difference and the cutpoint follow from the two expressions and $W(\mu_0)>0$. This is equivalent to the characterization in the proof of Proposition 1.

B.1 Proof of Proposition 6

The group's ex ante equilibrium payoff is $\max\{V^{SEP},V^{POOL}\}$ and full information yields $(1+\delta_G)\Pi^*$, so the value of early-career information is

$$(1+\delta_G)\Pi^*-\max\{V^{SEP},V^{POOL}\}=\min\big\{(1+\delta_G)\Pi^*-V^{SEP},\;(1+\delta_G)\Pi^*-V^{POOL}\big\}.$$

By Remark B.1 the first term equals $\Pi^*-\Pi_1^{SEP}=D(\mu_0,\delta_P)$ and the second equals $(\Pi^*-\pi_{\underline{\theta}})+\delta_G(\Pi^*-V_G(\mu_0))=\mu_0(\pi_{\overline{\theta}}-\pi_{\underline{\theta}})+\delta_G\,W(\mu_0)$, which establishes the desired representation. It remains to show that this minimum is positive and weakly increasing in δ_G and δ_P.

First, both terms are positive for all $\mu_0\in(0,1)$. The second is positive because $\pi_{\overline{\theta}}-\pi_{\underline{\theta}}=\frac{(1-\underline{\theta})^2-(1-\overline{\theta})^2}{2}>0$ and $W(\mu_0)>0$ by Remark B.1. For the first, consider the relaxed separating program that drops the incentive constraints $(IC_{\overline{\theta}})$ and $(IC_{\underline{\theta}})$ and retains only participation. Its unique optimum is the efficient menu $(c_{\underline{\theta}},c_{\overline{\theta}})$, with value Π^*. This menu violates $(IC_{\overline{\theta}})$, since an ally who takes $c_{\underline{\theta}}$ obtains $-(x_{\underline{\theta}}-\overline{\theta})^2+T_{\underline{\theta}}+\delta_P S(0)=(1+\delta_P)\Delta(1-\overline{\theta})>0$ while her intended contract yields 0. The incentive-feasible set is therefore a strict subset of the relaxed feasible set, so $\Pi_1^{SEP}<\Pi^*$ and therefore $D(\mu_0,\delta_P)>0$.

Second, the value is weakly increasing in δ_P. The second term is a second-period object and does not depend on δ_P. In the separating program, δ_P enters only the right-hand side of

$(IC_{\overline{\theta}})$, through the continuation $\delta_P S(0) = \delta_P \Delta(1-\overline{\theta})$ that an ally would earn by taking the adversary's contract. By the envelope theorem, $\frac{\partial \Pi_1^{SEP}}{\partial \delta_P} = -\lambda_1 \Delta(1-\overline{\theta}) \leq 0$, where $\lambda_1 \geq 0$ is the multiplier on $(IC_{\overline{\theta}})$ and is strictly positive in every case of the proof of Proposition 1. Thus, D is weakly increasing in δ_P, and so is its minimum with a term that is constant in δ_P.

Finally, the value is also weakly increasing in δ_G. The first term does not depend on δ_G. The second increases in δ_G at rate $W(\mu_0) > 0$ and lies below the first if $V^{POOL} \geq V^{SEP}$, i.e., if $\delta_G \leq \delta_G^\dagger$ (Remark B.1). The value of information therefore equals the increasing second term for $\delta_G < \delta_G^\dagger$ and the constant D for $\delta_G \geq \delta_G^\dagger$, with the two expressions equal at $\delta_G = \delta_G^\dagger$. Moreover, it is continuous and weakly increasing in δ_G, and strictly increasing wherever G would otherwise pool. □

B.2 Proof of Proposition 7

Since $\delta_G > \delta_G^\dagger$, with access in both periods G separates and obtains V^{SEP} (Remark B.1). The value of early-career access is $V^{SEP} - B$, where B denotes G's equilibrium payoff if it can lobby only in the second period. We first characterize B and then compare.

Without first-period access, each type chooses its first-period policy freely and without a transfer, and G updates before lobbying in the second period. Because the adversary's continuation value is zero at every belief (Remark 2), she enacts her ideal policy $\underline{\theta}$, so equilibria are indexed by the ally's first-period policy. We assign any off-path policy the belief $\mu_1 = 1$, which is admissible under the NWBR refinement.[33]

Two pure-strategy configurations can arise. In a separating equilibrium, the ally enacts $\overline{\theta}$ and reveals herself, receiving 0. Deviating to $\underline{\theta}$ induces belief 0 and yields $-\Delta^2 + \delta_P S(0)$, so separation is an equilibrium if and only if $0 \geq -\Delta^2 + \delta_P S(0)$, which is equivalent to $\delta_P \leq \delta_P^*$.

In a pooling equilibrium, both types enact $\underline{\theta}$, the belief remains μ_0, and the ally receives $-\Delta^2 + \delta_P S(\mu_0)$. Deviating yields her 0, so pooling is an equilibrium if and only if $\mu_0 < \tilde{\mu}_1$ and $\delta_P S(\mu_0) \geq \Delta^2$.

Since $S(\mu_0) < S(0)$ for $\mu_0 > 0$, the second condition for pooling requires $\delta_P > \delta_P^*$, so the two configurations cannot coexist on the same parameter regions. For intermediate δ_P, and for $\mu_0 \geq \tilde{\mu}_1$ with $\delta_P > \delta_P^*$, only mixed-strategy equilibria exist. Given our focus on pure strategies, we omit these and characterize the value of access on the two pure-strategy

[33] The adversary's payoff from any policy other than $\underline{\theta}$ is strictly negative and her continuation is zero at every belief, so she never gains from such a deviation, and any off-path policy is attributable to the ally.

regions. There, the benchmark payoff is

$$B = \begin{cases} -\big[(1-\mu_0)(1-\underline{\theta})^2 + \mu_0(1-\overline{\theta})^2\big] + \delta_G \Pi^* & \text{(separation)}, \\ -(1-\underline{\theta})^2 + \delta_G\, V_G(\mu_0) & \text{(pooling)}, \end{cases}$$

since under separation G enacts each type's ideal policy in the first period and then learns the type, while under pooling both types enact $\underline{\theta}$ and no learning occurs.

The value of access is positive on both regions. On the separation region, G can offer the menu $\{(\underline{\theta}, 0), c_{\overline{\theta}}\}$. Since this menu separates types, the adversary's continuation is zero under either contract and her payoff from $c_{\overline{\theta}}$ is negative, so she strictly prefers $(\underline{\theta}, 0)$. Meanwhile, the ally earns 0 from $c_{\overline{\theta}}$ and weakly less from $(\underline{\theta}, 0)$ if and only if $\delta_P \leq \delta_P^*$. The menu and the benchmark coincide for the adversary but differ for the ally. In that case, G's first-period payoff increases from $-(1-\overline{\theta})^2$ to $\pi_{\overline{\theta}}$, so $V^{SEP} - B \geq \mu_0\big(\pi_{\overline{\theta}} + (1-\overline{\theta})^2\big) = \mu_0 \frac{(1-\overline{\theta})^2}{2} > 0$.

On the pooling region, G's pooling menu yields $V^{SEP} - B \geq V^{POOL} - B = \pi_{\underline{\theta}} + (1-\underline{\theta})^2 = \frac{(1-\underline{\theta})^2}{2} > 0$.

The value of access weakly decreases in δ_P on each region. The benchmark does not depend on δ_P within a region, while with access we have $\frac{\partial V^{SEP}}{\partial \delta_P} = \frac{\partial \Pi_1^{SEP}}{\partial \delta_P} = -\lambda_1 \Delta(1-\overline{\theta}) \leq 0$, as in the proof of Proposition 6. The inequality is strict if $\lambda_1 > 0$, which holds throughout each region if $\delta_P > 0$ and μ_0 is interior. Thus, the value of access decreases in δ_P on the separation region, where $\delta_P \leq \delta_P^*$, as well as the pooling region, where $\delta_P > \delta_P^*$ and $\mu_0 < \tilde{\mu}_1$.

The value of access weakly increases in δ_G. With access, $\frac{\partial V^{SEP}}{\partial \delta_G} = \Pi^*$. On the separation region, $\frac{\partial B}{\partial \delta_G} = \Pi^*$, so the value of access is constant in δ_G. On the pooling region, $\frac{\partial B}{\partial \delta_G} = V_G(\mu_0)$, so it increases in δ_G at rate $\Pi^* - V_G(\mu_0) = W(\mu_0) > 0$. □

B.3 Proof of Proposition 8

Consider a policy cap $\overline{y} \in [\max\{x_{\underline{\theta}}, \overline{\theta}\}, x_{\overline{\theta}})$. Such a cap binds the ally's policy but not the adversary's. The baseline ally-policy weakly exceeds $x_{\overline{\theta}}$ whenever such a contract is offered (Lemma 1, Proposition 1), so any cap $\overline{y} < x_{\overline{\theta}}$ will bind it. The adversary-contract from the baseline is not constrained because $\underline{x}_t \leq x_{\underline{\theta}} \leq \overline{y}$ in every menu and each period (same results). Furthermore, since $\overline{y} \geq \overline{\theta}$, an ally can enact her ideal point after rejecting, so her reservation payoff is again 0. We show below that her equilibrium continuation value is also unaffected: the cap does not change the adversary's contract, and hence does not change the ally's rent $\max\{S(\mu), 0\}$.

Let $\pi_{\overline{\theta}}^{\overline{y}} := -(\overline{y}-1)^2 - (\overline{y}-\overline{\theta})^2$ denote G's full-information payoff against an ally under the cap. Since $\frac{d}{d\overline{y}} \pi_{\overline{\theta}}^{\overline{y}} = -4(\overline{y} - x_{\overline{\theta}}) > 0$ for $\overline{y} < x_{\overline{\theta}}$, it increases in $\overline{y}$ and equals $\pi_{\overline{\theta}}$ at

$\overline{y} = x_{\overline{\theta}}$.

We now show that $\overline{y}$ does not affect the second-period value of learning, $W(\mu_0)$. In the second-period menu at any belief μ, the adversary's participation constraint binds and fixes her transfer at $(\underline{x}_2 - \underline{\theta})^2$, and G chooses the adversary's policy $\underline{x}_2$. Neither component involves the ally's policy, so the adversary's contract and the ally's equilibrium continuation $V_P(\overline{\theta}, \mu) = \max\{S(\mu), 0\}$ (by Remark 2), are exactly as in Lemma 1. Thus, $\overline{y}$ changes only the ally's policy, from $x_{\overline{\theta}}$ to $\overline{y}$, and her transfer, which adjusts to leave her payoff at $V_P(\overline{\theta}, \mu)$. This adjustment does not induce the adversary to choose the ally-contract. The adversary's payoff from the ally-contract is

$$V_P(\overline{\theta}, \mu) + (\overline{y} - \overline{\theta})^2 - (\overline{y} - \underline{\theta})^2 = V_P(\overline{\theta}, \mu) - \Delta(2\overline{y} - \underline{\theta} - \overline{\theta}).$$

The adversary's intended contract gives payoff zero. Thus, the adversary incentive constraint is satisfied if

$$V_P(\overline{\theta}, \mu) \leq \Delta(2\overline{y} - \underline{\theta} - \overline{\theta}).$$

For $\mu \leq \tilde{\mu}_1$, we have $V_P(\overline{\theta}, \mu) = S(\mu) \leq S(0) = \Delta(1 - \overline{\theta})$, while for $\mu > \tilde{\mu}_1$ this continuation value is zero. Since $\overline{y} \geq x_{\underline{\theta}} = (1 + \underline{\theta})/2$, we have $\Delta(1 - \overline{\theta}) \leq \Delta(2\overline{y} - \underline{\theta} - \overline{\theta})$, so the adversary's incentive constraint is satisfied under the cap $\overline{y}$. With the ally's payoff held at $V_P(\overline{\theta}, \mu)$ and her policy moved from $x_{\overline{\theta}}$ to $\overline{y}$, G's payoff from the ally-contract shifts by the constant $\pi_{\overline{\theta}}^{\overline{y}} - \pi_{\overline{\theta}}$. Therefore the respective continuation values for P and G are $V_P^{\overline{y}}(\theta, \mu) = V_P(\theta, \mu)$ and $V_G^{\overline{y}}(\mu) = V_G(\mu) + \mu\,(\pi_{\overline{\theta}}^{\overline{y}} - \pi_{\overline{\theta}})$. Thus, with $\Pi^{*,\overline{y}} = \Pi^* + \mu_0(\pi_{\overline{\theta}}^{\overline{y}} - \pi_{\overline{\theta}})$, the value of learning is unchanged:

$$W^{\overline{y}}(\mu_0) = \Pi^{*,\overline{y}} - V_G^{\overline{y}}(\mu_0) = \Pi^* - V_G(\mu_0) = W(\mu_0).$$

Next, we show that a tighter cap (lower $\overline{y}$) raises the cutpoint on δ_G that distinguishes pooling menus from separating menus. The pooled contract $(x_{\underline{\theta}}, T_{\underline{\theta}})$ is unaffected by $\overline{y}$, as its policy satisfies $x_{\underline{\theta}} \leq \overline{y}$, so $V^{POOL,\overline{y}} = \pi_{\underline{\theta}} + \delta_G V_G^{\overline{y}}(\mu_0)$. Meanwhile, $V^{SEP,\overline{y}} = \Pi_1^{SEP,\overline{y}} + \delta_G \Pi^{*,\overline{y}}$. Since $W^{\overline{y}} = W$, the argument of Remark B.1 gives a single crossing at $\delta_G^{\dagger,\overline{y}} = \frac{\pi_{\underline{\theta}} - \Pi_1^{SEP,\overline{y}}}{W(\mu_0)}$, with G strictly preferring separation if $\delta_G > \delta_G^{\dagger,\overline{y}}$, strictly preferring pooling if $\delta_G < \delta_G^{\dagger,\overline{y}}$, and being indifferent at equality. Since $V_P^{\overline{y}} = V_P$, $\overline{y}$ does not affect either type's rejection payoff or continuation value. Furthermore, it restricts only the ally's first-period policy. The change in the second-period continuation enters the objective only as the additive constant $\delta_G \Pi^{*,\overline{y}}$. Thus, the problem defining $\Pi_1^{SEP,\overline{y}}$ is the baseline first-period problem of Proposition 1 with the additional constraint $\overline{x}_1 \leq \overline{y}$. Its feasible set shrinks, and the baseline optimum—which is unique by Proposition 1—violates the added constraint,

since $\overline{x}_1 \geq x_{\overline{\theta}} > \overline{y}$. Therefore $\Pi_1^{SEP,\overline{y}} < \Pi_1^{SEP}$. The post-separation continuation $\Pi^{*,\overline{y}}$ also decreases, but this cancels in the cutpoint since $W^{\overline{y}} = W$. Consequently, $\delta_G^{\dagger,\overline{y}} > \delta_G^{\dagger}$. Moreover, since $\Pi_1^{SEP,\overline{y}}$ increases in $\overline{y}$ (shown below), $\delta_G^{\dagger,\overline{y}}$ decreases in $\overline{y}$: a tighter cap raises the screening cutpoint, so G separates under fewer conditions.

Finally, we show that a weaker policy cap increases G's equilibrium payoff. Relaxing the cap enlarges the first-period feasible set, so $\Pi_1^{SEP,\overline{y}}$ increases in $\overline{y}$, and $\Pi^{*,\overline{y}}$ and $V_G^{\overline{y}}(\mu_0)$ increase in $\overline{y}$ through $\pi_{\overline{\theta}}^{\overline{y}}$. Thus, both $V^{SEP,\overline{y}}$ and $V^{POOL,\overline{y}}$ increase in $\overline{y}$, and so does G's equilibrium payoff $\max\{V^{SEP,\overline{y}}, V^{POOL,\overline{y}}\}$. □

B.4 Proofs of Propositions 9 and 10

The politician receives R in the second period if and only if she accepts a first-period contract that induces a posterior above $\tilde{\mu}_1$. Thus, acceptance of such a contract adds $\delta_P R$ to her first-period continuation value. We fix $\delta_P R \leq T_{\overline{\theta}}$, which ensures nonnegative transfers.[34] Because hiring only depends on first-period behavior, the reward is a constant from P's perspective at a given posterior and does not affect the second-period menu. Thus, $V_G^R(\mu) = V_G(\mu)$ at every belief and the full-information continuation Π^* is unchanged.

First, consider a separating first-period menu. The posterior is then $\mu_1 = 1$ after the ally's contract and $\mu_1 = 0$ after the adversary's, so the reward attaches to the ally's contract and not the adversary's. Substituting the continuation values V_P^R, the constraints are

$$-(\overline{x} - \overline{\theta})^2 + \overline{T} + \delta_P R \geq -(\underline{x} - \overline{\theta})^2 + \underline{T} + \delta_P S(0), \tag{$IC_{\overline{\theta}}$}$$

$$-(\underline{x} - \underline{\theta})^2 + \underline{T} \geq -(\overline{x} - \underline{\theta})^2 + \overline{T} + \delta_P R, \tag{$IC_{\underline{\theta}}$}$$

$$-(\overline{x} - \overline{\theta})^2 + \overline{T} + \delta_P R \geq 0, \tag{$P_{\overline{\theta}}$}$$

$$-(\underline{x} - \underline{\theta})^2 + \underline{T} \geq 0. \tag{$P_{\underline{\theta}}$}$$

Setting $\widetilde{T} := \overline{T} + \delta_P R$, these are exactly the separating constraints of Proposition 1 in $(\underline{x}, \overline{x}, \underline{T}, \widetilde{T})$. The reward R is absorbed into $\widetilde{T}$ in ($IC_{\overline{\theta}}$) and ($P_{\overline{\theta}}$), but is absent from ($P_{\underline{\theta}}$). Furthermore, in ($IC_{\underline{\theta}}$) the reward an adversary would collect by taking the ally's contract is exactly offset in the substitution, since $\overline{T} + \delta_P R = \widetilde{T}$. Since G pays $\overline{T} = \widetilde{T} - \delta_P R$ and $V_G^R = V_G$, its objective is the baseline objective in $(\underline{x}, \overline{x}, \underline{T}, \widetilde{T})$ plus the constant $\mu_0 \delta_P R$. The optimal separating menu is therefore analogous to that of Proposition 1 with the ally's transfer reduced by $\delta_P R$, but nonnegative because $\overline{T}_1 \geq T_{\overline{\theta}} \geq \delta_P R$ in every case.

[34] Hiring is conditioned on acceptance, so a rejecting politician is not hired regardless of belief. Therefore rejection yields no rent, and each participation constraint below has a zero right-hand side, as in the baseline.

Consequently,

$$V^{SEP,R} = V^{SEP} + \mu_0\,\delta_P R.$$

Next, consider a pooling menu. The posterior belief is $\mu_1 = \mu_0$, so both types are hired if and only if $\mu_0 > \tilde{\mu}_1$. In that case, the reward relaxes both participation constraints by $\delta_P R$, the group lowers the pooled transfer to $T_{\underline{\theta}} - \delta_P R$, which is nonnegative since $\delta_P R \le T_{\overline{\theta}} \le T_{\underline{\theta}}$, and $V^{POOL,R} = V^{POOL} + \delta_P R$. If instead $\mu_0 \le \tilde{\mu}_1$, then neither type is hired and the pooled contract is unchanged. Thus,

$$V^{POOL,R} = V^{POOL} + \delta_P R \cdot \mathbf{1}\{\mu_0 > \tilde{\mu}_1\}.$$

Setting $V^{SEP,R}$ and $V^{POOL,R}$ equal, and then using $V^{SEP} - V^{POOL} = (\Pi_1^{SEP} - \pi_{\underline{\theta}}) + \delta_G\, W(\mu_0)$ from Remark B.1, G screens if $\delta_G > \delta_G^{\dagger,R}$, where

$$\delta_G^{\dagger,R} = \delta_G^{\dagger} + \frac{\delta_P R\,(\mathbf{1}\{\mu_0 > \tilde{\mu}_1\} - \mu_0)}{W(\mu_0)} = \begin{cases} \delta_G^{\dagger} - \dfrac{\mu_0\,\delta_P R}{W(\mu_0)}, & \mu_0 \le \tilde{\mu}_1, \\ \delta_G^{\dagger} + \dfrac{2\,\delta_P R}{(1-\overline{\theta})^2}, & \mu_0 > \tilde{\mu}_1, \end{cases}$$

and pools if $\delta_G < \delta_G^{\dagger,R}$. The second case uses $W(\mu_0) = \frac{(1-\mu_0)(1-\overline{\theta})^2}{2}$ for $\mu_0 > \tilde{\mu}_1$.

Proof of Proposition 9. Differentiating yields $\frac{\partial \delta_G^{\dagger,R}}{\partial R} = -\frac{\mu_0 \delta_P}{W(\mu_0)} < 0$ for $\mu_0 \le \tilde{\mu}_1$ and $\frac{\partial \delta_G^{\dagger,R}}{\partial R} = \frac{2\delta_P}{(1-\overline{\theta})^2} > 0$ for $\mu_0 > \tilde{\mu}_1$. A larger reward therefore (i) decreases the screening cutpoint, widening the set of δ_G for which G learns, if $\mu_0 \le \tilde{\mu}_1$, and (ii) increases it, narrowing that set, if $\mu_0 > \tilde{\mu}_1$. Moreover, the change is strict for $\delta_G^{\dagger,R} \in (0,1)$. □

Proof of Proposition 10. Fix $\mu_0 \le \tilde{\mu}_1$ and $\delta_G > \delta_G^{\dagger}$. Since $\delta_G^{\dagger,R} = \delta_G^{\dagger} - \frac{\mu_0 \delta_P R}{W(\mu_0)} \le \delta_G^{\dagger}$, the group separates for every admissible R, so its equilibrium payoff is $V^R = V^{SEP,R} = V^{SEP} + \mu_0 \delta_P R$. Thus, $\frac{\partial V^R}{\partial R} = \mu_0 \delta_P$, which is strictly positive for $\mu_0, \delta_P > 0$ and increasing in both δ_P and μ_0. □