WORKING PAPERS

N° 1706

JANUARY 2026

"A Demand-Side Driven Explanation of Niche Lobbying: A Theory and Some Application to Climate-Biodiversity Policy"

Perrin Lefebvre and David Martimort

Toulouse School of Economics

A Demand-Side Driven Explanation of Niche Lobbying: A Theory and Some Application to Climate-Biodiversity Policy

Perrin Lefebvre
UNamur, DeFiPP - perrin.lefebvre@unamur.be

David Martimort
Toulouse School of Economics, david.martimort@tse-fr.eu

Abstract. This paper develops a model of niche lobbying in which interest groups endogenously specialize in the acquisition of distinct types of policy-relevant information. Contrary to the view that niche strategies are chosen to soften competition and secure autonomy, we show that specialization arises as a self-enforcing equilibrium even though groups would prefer to compete over the same informational dimensions. The mechanism is demand-driven: when information acquisition is private and non-verifiable, the decision-maker's inference from silence intensifies informational pressure on specialized groups, increasing the burden of information acquisition. We discuss the implications of these results for interest groups influence in climate and biodiversity policy.

Keywords. Lobbying, Information Acquisition, Niche Expertise, Hard Information Communication, Specialization.
JEL classification. D72, D82, D83.

1. Introduction

Motivation. A salient feature of lobbying organizations is the widespread development of *niche strategies*. Many interest groups concentrate their activity on narrowly defined policy issues or develop highly specialized expertise. This pattern has been extensively documented and is commonly interpreted as the outcome of deliberate organizational choices. In environments where political attention, financial resources, or memberships are scarce, groups are argued to differentiate so as to secure a protected niche and reduce direct competition with rival organizations (Clark and Wilson, 1961; Wilson, 1995; Gray and Lowery, 1996, 1997; Lowery, 2007). As Wilson (1995, p. 263) puts it, for groups and associations *"seeking to maintain themselves (...) [t]he easiest and most prudent strategy is to develop*

We thank Morten Bennedsen, Vianney Dequiedt, Sidartha Gordon and Michel Le Breton for helpful discussions and comments. This paper also benefited from remarks by participants to the TOM seminar (PSE), to the LEI seminar (CREST) and to the Journées de Microéconomie Appliquées in Clermont-Ferrand. This paper has been funded by the Agence Nationale de la Recherche under grants ANR-17-EURE-0010 (Investissements d'Avenir program) and ANR-24-CE26-7833. The usual disclaimer applies.

autonomy – that is, a distinctive area of competence (...) and undisputed juridiction over a function, service, goal, or cause."

These explanations emphasize the *supply side* of lobbying. Niche strategies are portrayed as the outcome of groups' incentives to soften competition, preserve autonomy, and secure organizational survival. In this paper, we propose a complementary and fundamentally different explanation. We show that niche lobbying can arise as a robust *equilibrium outcome* of the informational environment, even though it is not the outcome that interest groups themselves would choose if they could coordinate or commit to an organization of lobbying.

Mechanism. The key mechanism is driven by the *demand side*, i.e., the beliefs and reactions of the decision-maker (thereafter DM) when no information is transmitted. When information acquisition is endogenous, privately chosen by groups, and unobservable to the DM, the absence of disclosure is itself informative. In particular, when a group is expected to specialize in a given dimension of information, silence may be interpreted as evidence that unfavorable information has been deliberately withheld. This inference strengthens the group's incentive to acquire information in order to avoid inducing unfavorable beliefs. The pressure to gather information is therefore especially strong when a group is perceived as a monopolist over a given dimension.

By contrast, when several groups with conflicting objectives compete over the same informational dimension, the absence of disclosure becomes ambiguous. Suspicion that it reflects the strategic withholding of unfavorable information now applies to both groups, thereby mitigating the DM's reaction to the lack of information. As a result, the informational cost of remaining silent is reduced, weakening incentives to invest in information acquisition. Paradoxically, competition over the same informational dimension relaxes pressure on interest groups, while specialization into distinct niches intensifies it.

This logic leads to a striking reversal of the usual interpretation of niche strategies. In our model, niche lobbying does not provide groups with a *"quiet life"* (ła Bertrand and Mullainathan, 2003). On the contrary, specialization amplifies informational pressure and forces groups to invest more heavily in costly information acquisition. If groups could jointly decide how lobbying should be organized, they would prefer to avoid such differentiation. Instead, niche strategies emerge endogenously because they are self-enforcing equilibrium outcomes.

This insight is part of a broader paradox. In our environment, interest groups would strictly prefer to commit ex ante to remaining inactive in the lobbying process. Because the DM's payoff is linear in beliefs, groups derive no informational rents from persuasion: on average, information acquisition does not improve outcomes from their perspective. Consequently, groups favor any arrangement that minimizes the need to gather information. Yet such commitments are not credible. Once a group is expected to be inactive, it has a strict incentive to deviate and acquire information in order to influence the DM's action. The lobbying process thus traps groups in equilibria that are, from the groups' viewpoint, collectively undesirable but individually unavoidable.

The model we develop formalizes these forces in a multidimensional framework with endogenous information acquisition, verifiable disclosure, and limited commitment. We show that specialization on single dimensions of information arises generically, that differentiation across groups is robust, and that these outcomes persist with many groups and many dimensions. Our analysis highlights how the informational environment faced by the DM can shape the organization of lobbying in ways that run counter to the preferences of interest groups themselves.

Organization. Section 2 reviews the related literature, and Section 3 presents the model. Section 4 analyzes the case of two interest groups, whose objectives may be either conflicting (Section 4.1) or congruent (Section 4.2). Section 5 extends the analysis to the case of K groups with heterogeneous objectives. Section 6 highlights an important commitment problem faced by groups, namely their ex ante incentive to remain inactive. Section 7 derives implications for the types of information groups acquire and for lobbying strategies at the sectoral level. Section 8 illustrates the model's predictions in the context of climate and biodiversity policies. Proofs omitted from the main text are relegated to an Appendix.

2. Literature Review

This paper contributes to several strands of the literature on information transmission, strategic disclosure, and lobbying. Its main contribution is to show how endogenous information acquisition and belief formation by a DM jointly shape the equilibrium organization of lobbying activity.

Verifiable information and strategic disclosure. A large literature studies environments in which informed parties strategically disclose verifiable information. Starting with Milgrom (1981), Matthews and Postlewaite (1985), and Milgrom and Roberts (1986), this work emphasizes that the absence of disclosure can itself be informative. Our analysis builds on this insight but departs from much of the literature by endogenizing information acquisition and allowing multiple parties to choose privately and strategically whether to become informed in the first place.

Bhattacharya and Mukherjee (2013) study disclosure by experts with exogenously given probabilities of being informed and show that DMs may benefit from interacting with experts holding extreme or conflicting preferences. In contrast, information acquisition is endogenous in our framework, and incentives to acquire information depend critically on how silence is interpreted by the DM.

Endogenous information acquisition and persuasion. Several recent contributions emphasize that information acquisition may be driven less by the prospect of persuasion than by the desire to avoid unfavorable decisions when no information is provided. Kartik, Lee, and Wu (2017) show that when experts can acquire information at a cost, adding more experts may reduce incentives to become informed and may harm the DM. Related insights appear in Henry (2009) and Wong and Yang (2018).

Our model generalizes this logic to a multidimensional environment with multiple interest groups. A central mechanism is that information acquisition is privately chosen and unobservable ex ante. As a result, silence is interpreted as potentially strategic, which creates endogenous pressure to acquire information even when persuasion yields no expected rents. This mechanism plays a key role in shaping equilibrium specialization patterns.

Specialization, selective investigation, and expert panels. A growing literature studies how experts or interest groups specialize across informational dimensions. Bennedsen and Friedman (2002) analyze specialization in the provision of verifiable information to legislators, with an emphasis on institutional features of legislatures. More recently, Gong and Yang (2022) introduce the notion of *selective investigation*, where experts are assumed to focus on a single dimension of information. In contrast to this work, specialization is not imposed in our model. Instead, it emerges endogenously as an equilibrium outcome of resource constraints, belief updating, and strategic interaction among groups. Moreover, while much of the literature focuses on how a DM should optimally select or design a panel of experts, we study how the organization of lobbying arises from the incentives faced by interest groups themselves.

Bayesian Persuasion and Commitment. This paper is also related to the literature on Bayesian persuasion following Kamenica and Gentzkow (2011). A key distinction is that we assume substantially weaker commitment power. Interest groups cannot commit to disclosure strategies or to the type or precision of signals they acquire. Recent work has explored persuasion with endogenous information acquisition and partial commitment, but typically assumes that information structures are observable or controlled by the sender (e.g. Gentzkow and Kamenica, 2016, 2017a, 2017b; Wu, 2020), or abstracts from competition and specialization across informational dimensions (Kartik et al. 2017; Wong and Yang, 2018).

Our framework instead focuses on privately chosen, unobservable search effort by multiple groups, which fundamentally alters incentives and drives the emergence of niche lobbying. By contrast, limited commitment is central to our results. Because groups cannot commit to remain uninformed or silent, they become trapped in equilibria that require costly information acquisition, even though such activity yields no expected rents. This mechanism distinguishes our analysis from standard persuasion models.[1]

Lobbying and competition. Finally, our paper relates to the broader literature on lobbying and political influence. Classic contributions emphasize how groups differentiate to reduce competition and secure autonomy (Clark and Wilson, 1961; Wilson, 1995; Gray and Lowery, 1996, 1997; Lowery, 2007). More recent theoretical and empirical work continues to investigate how access, competition, and information shape political influence.

On the theoretical side, recent models analyze how competition among lobbyists or experts affects incentives to acquire and disclose information (Kartik et al. 2017; Bhattacharya et al. 2018; Ekmekci and Kos, 2023). Related contributions emphasize the role of strategic silence and limited commitment. Henry (2009) and Wong and Yang (2018) show that information acquisition may be driven by defensive motives, namely the desire to avoid unfavorable decisions when no information is provided. More recently, Kartik, Squintani, and Tinn (2024) study information revelation in political environments and highlight how silence and disclosure interact with voters' or DMs' inferences. On the empirical side, Lowery (2007) documents how competition among interest groups affects lobbying activity,

[1] Given the linearity of groups' payoffs in the DM's beliefs, groups' ability to commit to a message strategy would not lead to more favorable decisions by the DM. The only benefit groups could obtain from commitment would be a commitment either to refrain from communication or to transmit uninformative messages, thereby saving on information-gathering costs; such a commitment would be equivalent to remaining inactive.

while more recent work such as Awad (2025) studies how access and network position shape the effectiveness of persuasive lobbying.

From a technical viewpoint, our model builds on Kartik et al. (2017). As in their model, linearity of the DM's payoff in beliefs rules out concavification and eliminates any scope for Bayesian persuasion (Kamenica and Gentzkow, 2011). On average, groups derive no informational rents from persuasion. This conclusion is robust to richer message spaces and to allowing groups to commit to communication strategies. However, we extend their framework to a multidimensional environment in which uncertainty stems from incomplete coverage of policy-relevant dimensions rather than from noisy signals.[2] This extension allows us to analyze how the multidimensional structure, combined with a hard resource constraint, creates incentives for specialization and differentiation, despite the fact that groups collectively would favor concentration in information acquisition.[3]

More broadly, our contribution differs from the previously cited literature by highlighting a demand-side mechanism. In our model, niche strategies arise not because they soften competition among groups, but because they intensify the informational pressure generated by the DM's inference from silence. When multiple groups with conflicting preferences compete over the same informational dimension, non-disclosure becomes ambiguous and attenuates adverse inference, thereby reducing incentives to acquire information. Similarly, when congruent groups cooperate on the same dimension, free-riding emerges and more than offsets the DM's inference from silence, again weakening incentives to acquire information. In both cases, it is the intensification of informational effort under differentiation-driven by the DM's belief formation-that makes differentiation an equilibrium outcome.

This reversal of the usual logic offers a new perspective on the organization of lobbying activity and complements recent work on strategic information acquisition and disclosure by showing how belief formation alone can sustain specialization as an equilibrium outcome, even when it is not preferred by interest groups themselves.

[2]Whereas Kartik et al. (2017) consider signals that are imperfectly correlated with the state of the world, we assume that whenever a group acquires information about a given dimension of the decision problem, that signal is fully informative about that dimension. Because signals concern only a subset of dimensions, they remain imperfectly informative about the aggregate state.

[3]This result echoes Ward (2004)'s insight that competition between groups may generate excessive dissipation of effort from the groupsâĂŹ own perspective. The underlying mechanism, however, is fundamentally different. Our model hinges on the endogenous acquisition of information, whereas WardâĂŹs result arises from the aggregation of the direction of group efforts.

3. The Model

Decision and state of the world. DM chooses an action $x \in \mathbb{R}_+$ as a function of the state of the world θ. This state of the world aggregates multiple independent dimensions, with cardinality N, and we write

$$\theta = \sum_{i=1}^{N} \alpha_i \theta_i.$$

Each component θ_i is independently distributed on $\{0,1\}$, with $\mathbb{P}(\theta_i = 1) = p$. Dimensions differ in their relevance for the final decision, with weights ordered as $1 = \alpha_N \leq \cdots \leq \alpha_1 \equiv \alpha$.

DM's payoff is

$$U_{DM}(x,\theta) = \theta x - \frac{x^2}{2}.$$

Conditional on beliefs, the optimal decision is therefore $x = \mathbb{E}(\theta)$.

The interpretation is that policy choices depend on several distinct aspects, each requiring targeted investigation. Examples include economic versus environmental impacts of regulation, revenue versus expenditure effects of fiscal policy, or heterogeneous effects across population groups. The weights α_i capture the relative policy importance of each dimension.

Interest groups and information acquisition. There is a finite set $\mathcal{K}$ of K interest groups that may influence DM by providing information. Groups are of two types. Progressive groups (P) prefer higher actions, while Conservative groups (C) prefer lower actions, independently of the true state of the world. This assumption abstracts from any alignment of objectives between the groups and DM regarding accurate decision-making. We assume $K \leq N$.

Let $\mathcal{K}_P$ denote the set of Progressive groups, with cardinality K_P, and $\mathcal{K}_C$ the set of Conservative groups, with $K_C = K - K_P$.

Information acquisition is constrained by a limited informational resource. Each group k is endowed with a fixed information-acquisition capacity $E^k \in (0,1)$ that is specific to that group.[4] Groups allocate this capacity across dimensions. Specifically, group k chooses

[4] Assuming $E^k < 1$ simplifies the analysis by ruling out situations where a group would get a fully informative message about one dimension ($e_i^k = 1$) before investing effort in another information. But the main intuitions would not be modified if this constraint were removed, provided E^k remains small enough compared to the number of groups and dimensions of search.

effort levels

$$e_i^k \in [0,1], \qquad i = 1, \dots, N,$$

subject to the resource constraint

$$\sum_{i=1}^{N} e_i^k \leq E^k. \tag{3.1}$$

Effort e_i^k represents the probability with which group k becomes informed about dimension i. If informed, the group observes a perfectly informative signal $\sigma_i^k = \theta_i$; otherwise, it observes no information ($\sigma_i^k = \emptyset$). Let $e^k = (e_1^k, \dots, e_N^k)$ denote group k's effort allocation.

Search across dimensions is thus fully substitutable: allocating more effort to one dimension necessarily crowds out effort on others. This captures limited attention, investigative capacity, or organizational resources.

Resources devoted to information gathering are distracted from other potential uses. Group k thus derives a gain $U(E^k - \sum_{i=1}^N e_i^k)$ from resources non-invested in this process. The gain function U is twice-differentiable on $x \in \mathbb{R}_+$, non-decreasing and concave with the Inada condition $U'(0) = +\infty$ being satisfied.[5] For simplicity, we assume an extreme form of decreasing marginal returns for the groups' gain function U; namely there exists $\bar{x} > 0$ such that $U'(x) = 0$ for $x \geq \bar{x}$. On top, groups have enough initial resources so that

$$E_k > \bar{x} \quad \forall k \in \mathcal{K}. \tag{3.2}$$

Together those conditions ensure that, whatever the equilibrium, no group remains inactive. Our results hold for a broader set of utility functions if we assume a sufficiently small marginal utility at E^k. Yet, imposing this condition would make the exposition of proofs with multiple active congruent groups excessively cumbersome, without any insight gain.

Groups' preferences depend only on the induced policy outcome. Progressive and Conservative groups have utilities

$$V_P^k(x, e^k) = x + U(E^k - \sum_{i=1}^{N} e_i^k), \qquad V_C^k(x) = -x + U(E^k - \sum_{i=1}^{N} e_i^k).$$

Information transmission. Following Milgrom (1981) and Kartik et al. (2017), information is verifiable. Groups cannot misreport signals but may choose whether or not to

[5]The analysis would remain valid under heterogeneous utility functions across groups. However, such heterogeneity is already captured by variation in resource endowments E^k, rendering the introduction of group-specific utility functions redundant.

disclose them. The message space of group k is

$$M^k = \prod_{i=1}^{N} m_i^k,$$

where $m_i^k \in \{\emptyset, \theta_i\}$ if group k is informed about dimension i, and $m_i^k = \emptyset$ otherwise.

Whenever a group observes $\sigma_i^k = \theta_i$, it has a dominant disclosure strategy. Progressive groups reveal the signal if and only if $\theta_i = 1$, while Conservative groups reveal it if and only if $\theta_i = 0$.

DM's Beliefs. If at least one group transmits an informative message $m_i = \theta_i$, *DM* perfectly infers the value of θ_i. If no informative message is sent about dimension i, *DM* updates beliefs using Bayes' rule, given the anticipated effort allocations of the groups:

$$\theta_i^* = \mathbb{E}\Big(\theta_i \mid \forall k \in \mathcal{K},\ m_i^k = \emptyset\Big).$$

Off-path beliefs matter only if a group unexpectedly reveals information when *DM* does not anticipate disclosure. We impose *passive beliefs*: upon such a deviation, *DM* does not revise her beliefs about the group's effort allocation across other dimensions.[6]

After information transmission occurred, *DM* updates her beliefs and takes the decision $x = \beta$, with

$$\beta \equiv \left(\sum_{i:\,\exists k \in \mathcal{K}:\, m_i^k = \theta_i} \alpha_i \theta_i \right) + \left(\sum_{i:\,\forall k \in \mathcal{K}:\, m_i^k = \emptyset} \alpha_i \theta_i^* \right). \tag{3.3}$$

Timing and Observability. The game unfolds as follows:

1. Groups simultaneously choose effort allocations e^k subject to the resource constraint $\sum_i e_i^k \leq E^k$. Effort choices and research outcomes are unobserved by *DM*.

2. Each group privately observes its signals $\sigma^k = (\sigma_1^k, \ldots, \sigma_N^k)$.

3. Groups simultaneously send messages m^k to *DM*.

4. *DM* updates beliefs and chooses her action taking conditional expectations on the state of nature.

[6] Passive beliefs play the same equilibrium selection role as in Kartik et al. (2017) and ensure stability under sequential information revelation.

4. The Case of Two Groups

Before turning to the general case with K active groups, it is useful to analyze simpler benchmark environments. We begin with two interest groups: one being Progressive (P) while the other is Conservative (C) (Section 4.1) before addressing the scenario of congruent and Progressive groups (Section 4.2).

4.1 *Conflicting Interests*

Throughout this section, each group $g \in \{P, C\}$ chooses an effort vector $e^g = (e_1^g, \dots, e_N^g)$ with $e_i^g \in [0,1]$, subject to the hard resource constraint (3.1). Since $U'(0) = +\infty$, this constraint will never be binding. An interest group always finds optimal to save some resource away from information gathering.

Beliefs when no information is transmitted. Fix a dimension i. If at least one group transmits $m_i = \theta_i$, DM learns θ_i. If no group transmits information about i, then DM updates her beliefs using the effort levels she expects.

With two competing groups and verifiable disclosure, P reveals only if informed and $\theta_i = 1$, whereas C reveals only if informed and $\theta_i = 0$. Hence, conditional on (e_i^P, e_i^C),

$$\Pr(\text{no message on } i \mid \theta_i = 1) = 1 - e_i^P, \qquad \Pr(\text{no message on } i \mid \theta_i = 0) = 1 - e_i^C.$$

Bayes' rule implies that DM's posterior belief in the absence of disclosure is:

$$\theta_i^* = \mathbb{E}\Big(\theta_i \,\Big|\, m_i^P = m_i^C = \emptyset\Big) = \frac{p(1-e_i^P)}{p(1-e_i^P) + (1-p)(1-e_i^C)}. \tag{4.1}$$

Search incentives. Thanks to our quadratic specification, DM's decision is equal to her posterior mean, as defined in (3.3). Since the groups' payoffs only depend on this chosen action, each group chooses e^g to shift β in its preferred direction, subject to (3.1).

Fix an equilibrium candidate and the corresponding beliefs $(\theta_i^*)_{i=1}^N$. Holding fixed the opponent's effort and the induced belief θ_i^*, the marginal (direct) gain for P from increasing e_i^P is proportional to the probability that $\theta_i = 1$ and that disclosure changes DM's belief from θ_i^* to 1. This yields the marginal benefit index

$$v_i^P \equiv \alpha_i \, p \, (1 - \theta_i^*). \tag{4.2}$$

Similarly, for C, increasing e_i^C raises the probability of learning $\theta_i = 0$ and disclosing it, which shifts beliefs from θ_i^* to 0. The corresponding marginal benefit index is

$$v_i^C \equiv \alpha_i \,(1-p)\,\theta_i^*. \tag{4.3}$$

If group $g \in \{C, P\}$ invests in information gathering about dimension i, the equilibrium condition thus writes

$$v_i^g = U'(E^g - \sum_{j=1}^{N} e_j^g). \tag{4.4}$$

Groups increase their search effort as long as the marginal benefit of information acquisition exceeds the marginal opportunity cost of allocating resources to alternative uses.

Existence. Equipped with those conditions, we are ready to provide a first important Lemma.

Lemma 1 (Existence of an equilibrium). *There exists a pure-strategy equilibrium of the game with two competing groups.*

Specialization. The hard resource constraint (3.1) makes search substitutable across dimensions. Generically (i.e. except for knife-edge ties in the indices v_i^g), each group strictly prefer to put all its budget on a single dimension; an otherwise, putting all budget on one dimension remains an optimum.

Lemma 2. *For each group $g \in \{P, C\}$, there exists $k \in \{1, \ldots, N\}$ such that $e_j^g = 0$ for all $j \neq k$.*

This specialization result holds independently of whether groups have aligned or opposing objectives as we will confirm below. While competition between groups affects DM's beliefs and thus the marginal value of information on a given dimension, it does not create complementarities across dimensions within a group's information acquisition problem. It follows that the marginal benefit of effort must be equalized across all dimensions receiving positive effort. As a result, any interior allocation across multiple dimensions leaves the group indifferent between reallocations of effort, and full concentration on a single dimension is optimal.

It is important to note that Lemma 2 concerns the allocation of effort *within* a group, rather than the interaction between groups. Competition shapes the beliefs formed by the

DM when no information is disclosed, and thus affects which dimension a group chooses to specialize in. However, once the target dimension is fixed, competition does not affect the incentive to concentrate effort, which follows solely from the hard resource constraint and the substitutability of effort across dimensions.

DIFFERENTIATION. Using (4.2)–(4.3), note that for a given dimension i, a higher α_i increases both groups' incentives to target that dimension, but in opposite informational directions. When α is close enough to 1, the endogenous belief terms θ_i^* imply that a configuration in which P targets the most relevant dimension and C targets the second-most relevant dimension remains stable to unilateral deviations, while joint concentration is generically fragile. Generically, equilibria entail differentiation.

PROPOSITION 1. *There exists* $\bar{\alpha} > 1$ *such that for* $\alpha \in [1, \bar{\alpha}]$*, the following properties hold.*

1. *There exists an equilibrium in which the two groups specialize in different dimensions.*

2. *There exists* $\eta > 0$ *such that a necessary condition for an equilibrium in which both groups specialize on* θ_1 *is* $\theta_1^* \in [\theta_k^* - \eta, \theta_k^* + \eta]$ *for all* $k \neq 1$*, and* $\lim_{\bar{\alpha} \to 1} \eta = 0$.

3. *There exists no equilibrium in which both groups specialize on a dimension* θ_k *with* $k \neq 1$.

Proposition 1 characterizes how competition and belief formation jointly determine equilibrium specialization patterns.

While Lemma 2 on specialization establishes that each group concentrates its effort on a single dimension, this proposition explains how dimensions are allocated across groups. The key force is DM's inference from silence: when a group is the (expected) sole provider of a given type of information, this inference generates a high marginal benefit of search, generically driving equilibria toward outcomes in which groups differentiate across dimensions.

Item 1. shows that for policy weights α close enough to one, there always exists an equilibrium in which competing groups specialize in different dimensions. Intuitively, when dimensions have comparable policy relevance, specialization on distinct dimensions increases the informational pressure faced by each group. If each group is expected to be the sole source of information on its own dimension, silence is highly informative. Conversely, adverse inference is weaker on dimensions outside a group's domain. Consequently, the marginal benefit of information acquisition is greater for dimensions for which the group is expected to be the sole provider, thereby making differentiation self-enforcing.

Item 2. shows that equilibria in which both groups specialize on the same dimension are fragile. Such equilibria require DM's beliefs across dimensions to be nearly identical, so that no group has a strict incentive to deviate toward another dimension. The condition $\theta_1^* \in [\theta_k^* - \eta, \theta_k^* + \eta]$ captures this requirement. As α approaches one, even small belief differences across dimensions are sufficient to trigger profitable deviations, and the admissible range η shrinks to zero. This illustrates that concentration equilibria are non-generic.

Item 3. shows that the same logic prevents both groups from specializing in a dimension other than the politically most relevant one. If they were to do so, at least one group would obtain a higher marginal benefit by diverting its search effort toward the more relevant dimension.

Remark. This result on differentiation is robust when the priors p_i on each dimension are sufficiently close to each other. When priors differ widely, equilibria with concentration may arise because the belief terms θ_i^* can overturn the ranking induced by α_i.[7]

4.2 *Congruent Interests*

We now consider two congruent (progressive) groups, indexed by $g \in \{1, 2\}$, with utilities increasing in β. Each group k chooses $e^g = (e_1^k, \ldots, e_N^k)$ where $E_i^k \in [0, 1]$ subject to the resource condition (3.1).

Beliefs when no information is transmitted. Fix a dimension i. If at least one group is informed and $\theta_i = 1$, at least one group reveals. Let $\pi(e_i^1, e_i^2) = e_i^1 + e_i^2 - e_i^1 e_i^2$ denote the probability that at least one group becomes informed about θ_i. If no disclosure occurs, then either $\theta_i = 0$, or $\theta_i = 1$ but no group became informed. Thus,

$$\Pr(\text{no message on } i \mid \theta_i = 1) = 1 - \pi(e_i^1, e_i^2), \qquad \Pr(\text{no message on } i \mid \theta_i = 0) = 1.$$

Bayes' rule yields

$$\theta_i^*(e_i^1, e_i^2) = \frac{(1 - \pi(e_i^1, e_i^2))p}{1 - \pi(e_i^1, e_i^2)p}. \tag{4.5}$$

[7] Unfortunately, there is no clear-cut relationship between the value of p_i and a group's incentives to acquire information about dimension i. For the sake of simplifying the analysis, we thus focus on the symmetric case in which probabilities are equal across dimensions.

Search incentives and strategic substitutability. Holding fixed the opponent's effort and the induced posterior threshold θ_k^*, the marginal gain for group k from increasing e_i^k is proportional to the probability that $\theta_i = 1$, that group k becomes the (marginal) provider of verifiable evidence, and that disclosure moves beliefs from θ_i^* to 1. This delivers the marginal benefit index

$$v_i^k \equiv \alpha_i \, p \, (1 - e_i^{-k}) \left(1 - \theta_i^*(e_i^k, e_i^{-k})\right), \tag{4.6}$$

where $-k$ denotes the other group.

If group k invests in information gathering about dimension i, the equilibrium condition writes like in (4.4):

$$v_i^k = U'(E^k - \sum_{i=j}^{N} e_j^k).$$

Equation (4.6) makes explicit the free-riding force: for a fixed threshold, increasing e_i^{-k} reduces the factor $(1 - e_i^{-k})$ and thus lowers the marginal return to e_i^k. Equilibrium statics are slightly more complex, for an increase in e_i^{-k} also leads to a decrease in θ_i^*, which in turn increases group k's marginal benefit of information gathering. However, it is easy to show that

$$\frac{\partial}{\partial e_i^{-k}} \left((1 - e_i^{-k})(1 - \theta_i^*(e_i^k, e_i^{-k})) \right) < 0. \tag{4.7}$$

The strategic-substitute effect dominates once the endogenous response of θ_i^* is accounted for.

Existence and specialization. Next Lemma mirrors Lemma 1 in this context with congruent groups and proves the existence of an equilibrium

Lemma 3. *There exists a pure-strategy equilibrium of the game with two congruent groups.*

We can now state our next Lemma, which shows that groups generically specialize in one dimension.

Lemma 4. *For each group $k \in \{1,2\}$, there exists $\ell \in \{1,\dots,N\}$ such that $e_i^k = 0$ for all $i \neq \ell$.*

Lemma 4 shows that each group generically concentrates its information-gathering effort on a single dimension. Under an interior allocation, the group is marginally indifferent across dimensions, so concentrating effort on one dimension leaves informational benefits unchanged. When marginal values differ across dimensions, however, the optimal response is to reallocate all search effort toward the dimension with the highest marginal benefit. Hence, multi-dimensional allocation is weakly dominated by full specialization.

The lemma provides a microfoundation for *selective investigation*. Rather than being imposed exogenously, concentration on a single dimension emerges endogenously from optimal behavior under resource constraints. This distinguishes our approach from Gong and Yang (2022) who instead assume specialization as a primitive and highlights how informational structure alone can generate focused investigation.

DIFFERENTIATION AND CONCENTRATION. Group specialization can yield either differentiation across dimensions or concentration on a common dimension. The two propositions that follow characterize these possibilities.

PROPOSITION 2. *There exists $\bar{\alpha} > 1$ such that for $\alpha \in [1, \bar{\alpha}]$, for any two distinct dimensions $k \neq \ell$, there exists an equilibrium in which group 1 specializes in θ_k and group 2 specializes in θ_ℓ.*

The next proposition highlights conditions for a symmetric equilibrium with concentration.

PROPOSITION 3. *Assume that the two groups are identical (same total resource $\bar{E} = \bar{E}^1 = \bar{E}^2$). For all $\alpha \geq 1$, there exists a cutoff $p(\alpha) \in (0, 1)$ such that a symmetric equilibrium in which both groups concentrate on the same dimension exists if and only if $p \geq p(\alpha)$.*

Proposition 2 shows that horizontal differentiation arises even when groups' objectives are perfectly aligned. The driving force behind this result is not, as one might expect, collective optimization of search effort by the groups. Instead, as equation (4.7) suggests and as Proposition 6 below will make explicit, groups would collectively prefer an outcome in which concentration prevails. In such a case, they share the informational pressure generated by *DM*'s interpretation of silence, which allows each group to free ride on the other's information-gathering effort. By contrast, when a group is expected to be the sole source of information on a given dimension, it alone bears the associated informational pressure. By increasing the marginal benefit of information acquisition along that dimension, this pressure renders any deviation unprofitable, thereby making specialization self-enforcing.

For both groups to concentrate on the same dimension, the informational pressure associated with that dimension must be sufficiently strong that, even when shared across groups, the pressure borne by each group exceeds the value of gathering information about a new dimension. This occurs when there is sufficient scope for the absence of information about the dimension groups inquire about to substantially tilt DM's beliefs in a manner unfavorable to the groups' position (i.e., for a sufficiently high value of p).

The contrast between the generality of Proposition 2 and the specific conditions of Proposition 3 highlights that horizontal differentiation is robust to the nature of preferences and constitutes the most generic equilibrium outcome.

5. The Situation with K Groups and N Dimensions

We now turn to the general environment with K interest groups and N dimensions of search. Each group $k \in \mathcal{K}$ chooses an effort vector $e^k = (e^k_1, \dots, e^k_N)$ with $e^k_i \in [0,1]$ subject to the hard resource constraint (3.1).

As before, Progressive groups seek to maximize DM's posterior mean β, while Conservative groups seek to minimize it. Beliefs θ^*_i are defined as the posterior probability that $\theta_i = 1$ conditional on no group transmitting information about dimension i.

The analysis of the two-group cases established two general features: (i) equilibrium search is concentrated on a single dimension for each group, and (ii) search incentives are summarized by a dimension-specific marginal benefit index that depends on α_i and θ^*_i. These insights extend directly to the general case.

For a Progressive group k, the marginal benefit of increasing effort on dimension i when this group is the sole provider of information about θ_i is

$$v_i^{k,P} = \alpha_i \, p \, (1 - \theta^*_i),$$

while for a Conservative group h it is

$$v_i^{h,C} = \alpha_i \, (1 - p) \, \theta^*_i.$$

In such an equilibrium, the condition for group k to exert positive effort on dimension i writes

$$v_i^{k,g} = U'(E^k - \sum_{j=1}^{N} e^k_i),$$

where $g \in \{P, C\}$ denotes the type of group k.

The next proposition generalizes the differentiation results obtained in the two-group cases to a scenario with K groups.

Proposition 4. *There exists $\bar{\alpha} > 1$ such that for all $\alpha \in [1, \bar{\alpha})$, there exists an equilibrium in which each group specializes in a distinct dimension.*

Proposition 4 shows that the differentiation forces identified in the two-group case extend naturally to environments with many interest groups and multiple informational dimensions. Specialization within groups is driven by resource constraints and the opportunity cost of allocating effort, and the same belief-driven and free-riding mechanisms that sustain differentiation with two groups continue to operate in larger populations. As a result, the equilibrium organization of lobbying exhibits a stable one-to-one mapping between groups and informational dimensions. When policy weights are sufficiently balanced, it is self-enforcing for each group to occupy a distinct informational niche. This result highlights that differentiation is not a peculiarity of small-group interactions.

The restriction that α be close to one ensures that no single dimension strictly dominates DM's objective. When dimensions have comparable policy relevance, no group has a strict incentive to deviate toward another group's dimension. If policy weights were sufficiently asymmetric, multiple groups would instead concentrate on the most relevant dimension. Thus, the condition on α isolates the informational forces driving differentiation from trivial dominance effects.

Finally, the differentiation equilibrium does not rely on coordination or communication among groups. Each group's specialization decision is individually optimal given the behavior of others. This reinforces the interpretation of niche lobbying as an equilibrium outcome rather than the product of explicit collusion or institutional design.

6. Interest Groups' Preferred Situations

While differentiation emerges as the most robust equilibrium outcome of the lobbying game, it is not the one preferred by interest groups. This section highlights a paradox: groups are collectively trapped in equilibria that require high informational effort, even though each group would strictly prefer to reduce or eliminate its own search activity.

The first observation, which follows directly from the linearity of groups' payoffs in the *DM*'s beliefs, encapsulates the main intuition of this section:

Lemma 5. *The equilibrium payoff of an interest group depends only on the equilibrium level of its own total search effort, and is decreasing in that effort.*

A direct implication is that groups would prefer to commit ex ante to remaining inactive.

Proposition 5. *Each group would strictly prefer to commit to zero information acquisition, independently of the behavior of other groups.*

Proposition 5 highlights a central paradox of informational lobbying. Although interest groups actively invest in information acquisition in equilibrium, each group would strictly prefer to commit ex ante to zero information acquisition. Because DM's action is linear in beliefs, groups at equilibrium extract no informational rents from acquiring and disclosing information. Information acquisition therefore only serves to avoid unfavorable inferences drawn from silence, making it a purely defensive activity. If groups could credibly commit to remaining uninformed, they would strictly reduce their costs without affecting DM's expected action – and hence the groups' expected policy payoff.

The preference for inactivity holds regardless of the number of groups, their objectives, or their equilibrium specialization patterns. Even if other groups continue to acquire information, a group would strictly benefit from committing to zero search, as this removes the cost of acquisition without worsening its influence on DM's beliefs. This shows that lobbying activity is driven by strategic necessity rather than strategic advantage.

Proposition 5 implies that interest groups are trapped in the lobbying process. They incur positive information-acquisition costs in equilibrium despite strictly preferring an outcome in which no group is active. This highlights a tension between individual optimality and collective inefficiency, and helps explain why lobbying activity may persist even when it yields no net benefits to groups.

The ex ante commitment of groups to stay inactive is typically infeasible as soon as condition (3.2) is satisfied. Once a group is expected to participate, it has a strict incentive to use its available resources, as marginal effort always has a first-order effect on beliefs.

When commitment to inactivity is impossible, groups turn to what can be viewed as a second-best outcome.

PROPOSITION 6. *Compared to differentiation,*

- *each group always prefers the situation where other groups (either competing or congruent) gather information about the same dimension as itself, and*
- *all the groups prefer the situation in which they all concentrate on the same dimension N.*

Proposition 6 shows that, conditional on being active, each group prefers to share its informational domain with other groups rather than to be the sole provider of information on a given dimension. This outcome arises from the combination of two effects. First, when a group shares the investigation of a dimension with one or more competing groups, the absence of disclosure becomes less informative. This weakens the adverse inference drawn by DM from silence and reduces the informational pressure faced by any individual group. Second, when a group is in the presence of other congruent groups, the free-riding that occurs within shared information acquisition softens the intensity of each group's information gathering. In both cases, overlap in information acquisition strictly lowers equilibrium search costs. Moreover, if groups could jointly commit to specializing in the least relevant dimension, each would benefit from the ensuing softening of information acquisition relative to any differentiated configuration.

These results clarify why differentiation, although robust, is not favored by interest groups. Groups would ideally prefer to avoid information acquisition altogether; failing that, they prefer to pool on the same dimensions in order to soften competition in persuasion.

This feature, and it contrasts with the typical emergence of differentiation at equilibrium, highlights the demand-side-driven development of niche specialization by groups, and how it departs from explanations that attribute it to a softening of competition among groups.

7. The Structuration of Lobbying

The analysis conducted so far shows how differentiation arises endogenously in equilibrium, but it is silent on which informational dimensions different groups choose to investigate. While a complete characterization of these choices lies beyond the scope of this article, we provide two results that shed light on this issue. Their relevance will become apparent in the illustrative examples presented in the next section.

Group specialization in strong domains. Proposition 4 does not rely on any assumptions about group types: in equilibrium, each group may specialize in any informational dimension. This indeterminacy stems from the symmetry of the model, in which all dimensions are associated with the same probability p. In practice, however, different types of information are likely to be associated with different prior probabilities. For example, information about climate change may be less favorable to the oil industry than information about the economic costs of ecological transition. This raises the question of whether industrial groups should concentrate on dimensions that are *a priori* more favorable to their position, or instead invest in information acquisition that mitigates unfavorable aspects of their activities.

To address this question, we allow the probabilities $p_i = \mathbb{P}(\theta_i = 1)$ to vary across dimensions. Specifically, consider two dimensions, indexed by 1 and 2, and let $\delta \in (0, \min\{p, 1-p\})$ such that $p_1 = p + \delta$ and $p_2 = p - \delta$. For a Progressive group, dimension 1 is more likely to generate information that supports its position. We refer to this dimension as the group's *strong domain.* Analogous terminology applies to a Conservative group.

Since the equilibrium characterizations derived above are robust to small perturbations in the probabilities associated with each dimension, the set of equilibrium specialization patterns remains unchanged for sufficiently small δ. However, when the *DM* is able to influence equilibrium selection, introducing asymmetries in prior probabilities has substantive implications. The following proposition illustrates this point in a symmetric setting when $p = \frac{1}{2}$:

Proposition 7. *Consider a symmetric situation where $p_{1,2} = \frac{1}{2} \pm \delta$, and two groups P and C, with $E^C = E^P$. There exists $\bar{\alpha} > 1$ and $\bar{\delta} > 0$ such that, for $\alpha \in [1, \bar{\alpha}]$ and $\delta \in (0, \bar{\delta}]$, DM prefers equilibria where each group specializes in its strong domain.*

Note that, once again, the result that groups-*ceteris paribus*-specialize in informational dimensions that are more likely to support their position is driven by the demand side of the equilibrium. Specialization in a group's strong domain induces higher information-gathering effort than specialization in its weak domain, holding all else constant.[8]

Sector concentration in lobbying. Thus far, we have assumed that $K \leq N$, so that each group can, in principle, be the sole provider of a distinct type of information. In many environments, however, this assumption may fail. Continuing with the example above, a large number of groups may engage in information acquisition along only two dimensions-or, more generally, along two sets of signals, with signals within each set being highly correlated.

Consider a setting with two relevant informational dimensions,

$$\theta = \alpha\theta_1 + \theta_2,$$

with $\mathbb{P}(\theta_1 = 1) = \mathbb{P}(\theta_2 = 1) = p$. The following proposition characterizes an equilibrium that arises when the environment exhibits a sufficiently strong symmetry.

Proposition 8. *Assume $N = 2$, $K_C = K_P$, $p = \frac{1}{2}$, and there exists $\bar{E}$ such that for all $k \in \mathcal{K}$, $E^k = \bar{E}$. There exists $\bar{\alpha} > 1$ such that for $\alpha \in [1, \bar{\alpha}]$, there exists an equilibrium in which all groups in $\mathcal{K}_P$ specialize in one dimension, and all groups in $\mathcal{K}_C$ specialize in the other dimension.*

The equilibrium characterized in Proposition 8 exhibits sectoral concentration. Interest groups belonging to the same sector and sharing aligned objectives concentrate their information-gathering efforts on a single dimension, while the opposing sector concentrates on the alternative dimension. Although free riding occurs within sectors, each sector can therefore be treated as a cohesive coalition for the purpose of information acquisition.

Combining the insights from Propositions 7 and 8, we expect such concentration to arise predominantly on informational dimensions corresponding to a sector's strong domain. This pattern is commonly observed in practice, as illustrated in the next section.

[8]In some cases, the nature of a group's activity may affect the cost of acquiring information along certain dimensions. For instance, economic forecasts may be less costly for industries that already conduct similar analyses for their own purposes. However, this need not always be the case: for example, evaluating ecological damages may not be less costly for environmental groups than for industrial ones. Our analysis abstracts from such cost asymmetries.

8. An Illustrative Example: Climate Policy and Biodiversity

We conclude this article with a section that illustrates the model's logic in the context of climate policy and biodiversity protection. These policy areas are particularly well-suited to our framework, as decisions typically depend on multiple dimensions of information and involve repeated interactions with specialized interest groups holding opposing objectives. For instance, climate policy decisions routinely hinge on assessments of physical and ecological damages (e.g. temperature trajectories, tipping points, or biodiversity loss) as well as on evaluations of economic and distributional costs (e.g. impacts on energy prices, employment, or competitiveness). A large body of literature emphasizes this multidimensional nature of climate policy evaluation (e.g. Stern, 2007; Nordhaus, 2017; IPCC, 2022).

In practice, these dimensions are investigated by distinct and recurring actors. Environmental NGOs, scientific coalitions, and research institutes primarily document climate and ecological risks, while industry associations, trade unions, and sectoral lobbies focus on transition costs and economic adjustment.[9] This division closely mirrors the structure of expertise emphasized in the political economy literature on environmental regulation (see, e.g., Aghion et al. 2016; Grossman and Krueger, 1995).

A similar structure arises in biodiversity policy. Decisions regarding protected areas, land-use regulation, or species conservation depend both on ecological assessments of irreversibility and ecosystem services and on evaluations of local economic impacts on agriculture, forestry, or development (e.g. Dasgupta, 2021).[10]

Finally, climate and biodiversity policies are characterized by repeated regulatory episodes and institutionalized consultation procedures, which reinforce expectations about which actors are responsible for providing information on each dimension.[11] Such repetition strengthens the inference drawn from silence: when a specialized group does not disclose

[9]For example, the assessment of climate risks is largely structured around scientific and advocacy organizations contributing to reports such as those of the Intergovernmental Panel on Climate Change (IPCC), whereas the economic consequences of climate regulation are frequently documented by industry associations and sector-specific studies submitted in regulatory impact assessments or parliamentary hearings.

[10]For instance, debates surrounding the expansion of protected areas under the Convention on Biological Diversity or national conservation laws typically rely on ecological assessments produced by conservation organizations and scientific bodies, alongside economic impact analyses produced by agricultural unions, landowners, and local governments.

[11] Examples include repeated Conference of the Parties (COP) negotiations under the UNFCCC, periodic revisions of nationally determined contributions, environmental impact assessment procedures, and standing expert panels or parliamentary hearings dedicated separately to environmental and economic impacts.

information on the dimension it is expected to cover, this absence naturally affects DM's beliefs. This feature makes climate and biodiversity policy a natural application of our belief-driven mechanism.

DM's problem. DM must choose the stringency x of a climate or biodiversity policy, such as the level of a carbon tax, the scope of protected natural areas, or restrictions on land use. The relevant state of the world is given by

$$\theta = \alpha_1 \theta_1 + \alpha_2 \theta_2,$$

where:

- $\theta_1 \in \{0, 1\}$ captures the severity of climate or biodiversity damages (e.g. irreversible ecosystem loss, long-run climate risks, or tipping points), with $\theta_1 = 1$ corresponding to severe damages;
- $\theta_2 \in \{0, 1\}$ captures the magnitude of economic and social adjustment costs (e.g. impacts on employment, energy prices, or agricultural output), with $\theta_2 = 1$ corresponding to lower adjustment costs.

DM chooses a decision equal to the conditional expectation of the state, given information revealed by groups. Note that higher expected damages or lower expected costs lead to more stringent environmental policies.

Environmental organizations prefer higher policy stringency, while industry groups (e.g. energy producers, agriculture, or land developers) prefer lower stringency. Both types of groups can acquire information about either environmental damages or economic costs, but their information acquisition efforts are private and unobservable to the DM.

Niche lobbying in climate and biodiversity debates. In practice, environmental *NGO*s and scientific advocacy groups tend to specialize in producing evidence on climate risks and biodiversity loss, such as reports on species extinction, ecosystem services, or climate tipping points. By contrast, industry groups typically specialize in studies of economic costs, including competitiveness effects, job losses, or short-run adjustment burdens.

This pattern corresponds to a differentiated equilibrium in our model. Each group becomes the primary source of information on a specific dimension, effectively holding a monopoly over that type of evidence. Such specialization may occur at the individual group level or, more broadly, at the sectoral level. Moreover, groups typically specialize in their *strong domain*, focusing on informational dimensions that are *ex ante* more likely to support their position.[12]

The model clarifies why such specialization is self-enforcing. If an environmental NGO is expected to be the main source of information on ecological damages, the absence of disclosure may be interpreted by the DM as evidence that damages are limited or uncertain. To avoid inducing such unfavorable beliefs, the NGO faces strong incentives to invest in information acquisition. More generally, the stronger the domain in which a group specializes, the greater the associated informational pressure. A symmetric logic applies to industry groups specializing in information about economic costs.

As a result, niche lobbying generates high informational pressure: each group must invest substantial resources simply to avoid being perceived as withholding unfavorable information. Specialization therefore increases, rather than decreases, equilibrium effort in information acquisition.

Why groups would prefer overlap. Now consider a counterfactual organization of lobbying in which both environmental and industry groups invest in information about the same dimension, for instance climate damages. In this case, silence becomes ambiguous. The absence of disclosure may reflect either weak evidence of severe damages or the strategic withholding of information by the opposing group.

In the model, this ambiguity weakens the DM's inference from silence and reduces the marginal benefit of information acquisition. Each group can partially free-ride on the possibility that the other group is withholding information. Consequently, equilibrium information-gathering effort is lower.

Importantly, both environmental and industry groups would strictly prefer such overlap to niche specialization. Competition over the same informational dimension relaxes informational pressure and leads to lower equilibrium costs of participation. Yet, because overlap is not generically self-enforcing, groups end up trapped in a differentiated equilibrium that they do not favor.

Policy implications. The analysis has several implications for the design of institutions governing climate and biodiversity policy.

[12]This feature is not trivial: revealing that dimensions that are *a priori* less favorable do not support the opposing position may carry greater informational value.

First, policies that encourage or impose specialization of expert input (e.g. assigning environmental *NGO*s exclusively to damage assessment and industry groups exclusively to cost assessment) may unintentionally intensify informational arms races. Such institutional designs strengthen the inference drawn from silence and increase incentives for costly information acquisition.

Second, encouraging overlap in information provision – for instance by commissioning independent assessments of both environmental damages and economic costs from multiple sources, or by structuring hearings so that competing groups address the same questions—may reduce informational pressure and lower total lobbying effort. While overlap is often viewed as redundant, the model suggests that it can mitigate defensive information acquisition driven by fear of adverse inference.[13]

Third, the results highlight the importance of transparency about information acquisition. When the *DM* can better distinguish between the absence of evidence and the absence of search, the incentive to engage in excessive information gathering is reduced. This suggests a potential role for disclosure requirements regarding research activity itself, not only regarding research outcomes.[14]

Overall, the climate and biodiversity context illustrates the central message of the paper: observed specialization of interest groups does not necessarily reflect efficient or preferred organizational choices. Instead, it may arise from belief-driven incentives and limited commitment, leading to lobbying structures that are costly for groups. However, the impact on society is ambiguous, since specialization combined with niche strategies may also increase the amount of information acquired by policy makers.[15] Further investigation is needed both to investigate the determinants of groups' organization and unveil how this organization affects society.

[13] This insight can be viewed as a generalization of the observation in Corollary 1 of Kartik et al. (2017) that the introduction of a competing expert can reduce the *DM*'s incentives to acquire information. The result that such a softening of information-acquisition effort can be detrimental to the *DM* extends to our framework as well. However, as in Kartik et al. (2017), his effect does not hold uniformly across environments.

[14] This is in the line of the seminal paper by Henry (2009). It contrasts with Gentzkow and Kamenica (2017b), in which imposing disclosure of the signals does not affect the equilibrium outcomes. The main difference lies in the fact that in their model, experts publicly acquire private information.

[15] Beyond the symmetric environment considered in Proposition 7, we are unable to derive a general result regarding the structure preferred by the *DM*. For example, differentiation between two competing groups increases the likelihood that some information is acquired, whereas concentration by both groups on the same dimension provides a stronger guarantee that at least some information will be disclosed. Which of these effects dominates depends in a complex way on parameters of the model.

References

Aghion, P., A. Dechezleprêtre, D. Hémous, R. Martin, and J. Van Reenen (2016). Carbon Taxes, Path Dependency, and Directed Technical Change: Evidence from the Auto Industry. *Journal of Political Economy*, 124: 1-51.

Awad, E., and C. Minaudier (2025). Persuasive Lobbying and the Value of Connections. *American Journal of Political Science*, forthcoming.

Bennedsen, M., and S. E. Friedman (2002). Lobbying Legislatures. *Journal of political Economy*, 110: 919-946.

Bertrand, M., and S. Mullainathan (2003). Enjoying the Quiet Life? *Journal of Political Economy*, 111: 1043-1075.

Bhattacharya, S., and A. Mukherjee (2013). Strategic Information Revelation when Experts Compete to Influence. *The RAND Journal of Economics*, 44: 522-544.

Bhattacharya, S., M. Goltsman, and A. Mukherjee (2018). On the Optimality of Diverse Expert Panels in Persuasion Games. *Games and Economic Behavior*, 107: 345-363.

Brocas, I., and J. Carrillo (2007). Influence Through Ignorance. *The RAND Journal of Economics*, 38: 931-947.

Brocas, I., J. Carrillo, and T. Palfrey (2012). Information Gatekeepers. *Economic Theory*, 51: 307-332.

Chan, J., S. Gupta, F. Li, and Y. Wang (2019). Pivotal Persuasion. *Journal of Economic Theory*, 180: 178-202.

Clark, P. B., and J. Q. Wilson (1961). Incentive Systems: A Theory of Organizations. *Administrative Science Quarterly*, 6: 129-166.

Dasgupta, P. (2021). *The Economics of Biodiversity: The Dasgupta Review.* HM Treasury, London.

Ekmekci, M., and N. Kos (2023). Signaling Covertly Acquired Information. *Journal of Economic Theory*, 214: 105746.

Gentzkow, M., and E. Kamenica (2016). Competition in Persuasion. *he Review of Economic Studies*, 84: 300-322.

Gentzkow, M., and E. Kamenica (2017a). Bayesian Persuasion with Multiple Senders and Rich Signal Spaces. *Games and Economic Behavior*, 104: 41-429

Gentzkow, M., and E. Kamenica (2017b). Disclosure of Endogenous Information. *Economic Theory Bulletin*, 5: 47-56.

Gong, J., and Y. Yang (2022). Expert Panels with Selective Investigation. *The BE Journal of Theoretical Economics*, 22: 369-404.

Gray, V., and D. Lowery (1996). *The Population Ecology of Interest Representation.* University of Michigan Press.

Gray, V., and D. Lowery (1997). Life in a Niche: Mortality Anxiety Among Organized Interests in the American States. *Political Research Quarterly*, 50: 25-47.

Grossman, G., and A. Krueger (1995). Economic Growth and the Environment. *The Quarterly Journal of Economics*, 110: 353-377.

Gul, F., and W. Pesendorfer (2012). The War of Information. *The Review of Economic Studies*, 79: 707-734.

Henry, E. (2009). Strategic Disclosure of Research Results: The Cost of Proving Your Honesty. *The Economic Journal*, 119: 1036-1064.

Henry, E., and M. Ottaviani (2019). Research and the Approval Process: The Organization of Persuasion. *The American Economic Review*, 109: 911-955.

IPCC (2022). *Climate Change 2022: Synthesis Report.* Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press.

Kamenica, E., and M. Gentzkow (2011). Bayesian Persuasion. *The American Economic Review*, 101: 2590-2615.

Kartik, N., F. Lee, and W. Suen (2017). Investment in Concealable Information by Biased Experts. *The RAND Journal of Economics*, 48: 24-43.

Kartik, N., F. Squintani, and K. Tinn (2024). Information Revelation and Pandering in Elections. Working paper.

Lowery, D. (2007). Why do Organized Interests Lobby? A Multi-Goal, Multi-Context Theory of Lobbying. *Polity*, 39: 29-54.

Lukyanov, G., K. Shamruk, T. Su, and A. Wakrim (2022). Public Communication with Externalities. *Games and Economic Behavior*, 136: 177-196.

Matthews, S., and A. Postlewaite (1985). Quality Testing and Disclosure. *The RAND Journal of Economics*, 16: 328-340.

Milgrom, P. (1981). Good News and Bad News: Representation Theorems and Applications. *The Bell Journal of Economics*, 12: 380-391.

Milgrom, P., and J. Roberts (1986). Relying on the Information of Interested Parties. *The RAND Journal of Economics*, 17: 18-32.

Nordhaus, W. (2017). Revisiting the Social Cost of Carbon. *Proceedings of the National Academy of Sciences*, 114: 1518-1523.

Stern, N. (2007). *The Economics of Climate Change: The Stern Review.* Cambridge University Press.

Ward, H. (2004). Pressure Politics: A Game-Theoretical Investigation of Lobbying and the Measurement of Power. *Journal of Theoretical Politics*, 16: 31-52.

Wilson, J. Q. (1995). *Political Organizations.* Princeton University Press.

Wong, T., and L. Yang (2018). When Does Monitoring Hurt? Endogenous Information Acquisition in a Game of Persuasion. *Economics Letters*, 163, 186–189.

Wu, J. (2020). Non-Competing Persuaders. *European Economic Review*, 127: 103454.

Appendix A: Proofs

Proof of Lemma 1. The strategy sets $\{e^g \in [0,1]^N : \sum_i e_i^g \leq E^g\}$ are nonempty, compact, and convex. Given verifiable disclosure, expected payoffs are continuous in (e^P, e^C). A standard fixed-point argument (e.g. Brouwer) applied to the best-response correspondence yields existence. □

Proof of Lemma 2. Fix the opponent's strategy and the implied belief vector $(\theta_i^*)_{i=1}^N$. By (4.4), any dimension with positive effort must satisfy $v_i^P = U'(E^P - \sum_j e_i^P)$. Hence, if $e_j^P > 0$ and $E_P^k > 0$ for $j \neq k$, we must have $v_j^P = v_k^P$. Except for knife-edge parameter values yielding a tie, there is a unique maximizer of v_i^P, so P concentrates all effort on that dimension. The same argument applies to C. □

Proof of Proposition 1. The proof proceeds with several steps;

Step 1. By Lemma 2, each group concentrates on a single dimension. Consider the candidate profile where P targets θ_1 and C targets θ_2. Given (4.1), this implies $\theta_1^* < p < \theta_2^*$, while $\theta_k^* = p$ for $k \notin \{1,2\}$. For $\alpha > 1$ sufficiently close to 1, the induced indices (4.2)–(4.3) satisfy that θ_1 uniquely maximizes v_i^P and θ_2 uniquely maximizes v_i^C, establishing stability.

Step 2. If both groups target θ_1, then for any $k \neq 1$ the no-deviation conditions require that both $v_1^P \geq v_k^P$ and $v_1^C \geq v_k^C$, which jointly force θ_1^* to lie in a shrinking neighborhood of θ_k^* as $\alpha \to 1$.

Step 3. If both groups target θ_k for $k \neq 1$, then $\alpha_k < \alpha_1$ implies that at least one group strictly prefers switching to θ_1, contradicting optimality under (4.4). □

Proof of Lemma 3. As in Lemma 1, the feasible sets are nonempty, compact, and convex, and expected payoffs are continuous. A fixed-point argument yields existence. □

Proof of Lemma 4. Fix the other group's strategy. By (4.7), any two dimensions receiving positive effort must satisfy $v_k^i = v_\ell^i$. Except for knife-edge ties, there is a unique maximizer of v_k^i, so group i concentrates all effort on that dimension. □

Proof of Proposition 2. When $\alpha = 1$, dimensions are symmetric. Consider a profile where group 1 targets k and group 2 targets $\ell \neq k$. Then $\theta_k^* < p$ and $\theta_\ell^* < p$, while $\theta_j^* = p$ for untouched dimensions. Using (4.6) and (4.7), it is the case that for group 1, $v_k^1 > v_j^1$ for all $j \neq k$. Identically for group 2, $v_\ell^2 > v_j^2$ for all $j \neq \ell$. It follows that each group strictly prefers its targeted dimension to any other. Since the relevant inequalities are strict at $\alpha = 1$, the argument extends to $\alpha > 1$ sufficiently close to 1. □

Proof of Proposition 3. We look for an equilibrium in which both groups concentrate on the same dimension (wlog, dimension 1). If such an equilibrium exists, the symmetry of the game also implies that there exists a symmetric equilibrium. We then should have $e_1^1 = e_1^2 = \bar{e}$ and $e_i^k = 0$ for $i \neq 1$.

Assumption (3.2) guarantees that whatever $p \in (0,1)$, the condition for $\bar{e} > 0$ will not be binding.

A (strict) deviation exists for group k if for some $i \neq 1$, $v_1^k < v_i^k = \alpha_i p(1-p)$. A sufficient condition is that no deviation exists on dimension 2. This condition writes

$$p \geq 1 - \frac{\alpha}{\alpha_2}(1-e)(1-\theta_1^*(e,e)) \equiv p(\alpha).$$

In case different equilibria exist, let us denote E the set of the equilibrium levels of search. We define $p(\alpha)$ as the lowest bound of the set $\{p(\alpha,e) | e \in E\}$. Then, $p \geq p(\alpha)$ is a necessary and sufficient condition for a symmetric equilibrium to exist. □

Proof of Proposition 4. We first consider the benchmark case $\alpha = 1$, so that all dimensions are symmetric from the DM's perspective. Suppose that each group k specializes in a unique dimension d_k, with no two groups targeting the same dimension.

Consider a deviation by group k, assumed without loss of generality to be Progressive. There are two possible types of deviations.

Deviation to an occupied dimension. Suppose group k reallocates effort toward a dimension j already targeted by another group ℓ. Restricting attention to the subgame involving only groups k and ℓ, this deviation would contradict the existence of an equilibrium in which group k concentrates in d_k and group ℓ in d_ℓ. But given the symmetry of the settings when $\alpha = 1$, the existence of this equilibrium is guaranteed by Proposition 1 or Proposition 2, depending on the types of groups involved.

Deviation to an unoccupied dimension. Suppose instead that group k deviates to a dimension j that is not targeted by any group. Then the absence of disclosure on dimension j implies $\theta_j^* = p$. By contrast, on the equilibrium dimension d_k, the absence of disclosure has an informational adverse effect for group k, implying $\theta_{d_k}^* < p$. Since $\alpha = 1$ and the marginal benefit index is strictly higher on d_k than on j, such a deviation cannot be profitable.

Thus, no group has a profitable deviation, and the fully differentiated profile constitutes an equilibrium. Because the relevant inequalities are strict at $\alpha = 1$, continuity implies that the result extends to all $\alpha \in [1, \bar{\alpha})$ for some $\bar{\alpha} > 1$. □

Proof of Lemma 5. At any equilibrium, the law of iterated expectations implies in our case that $\mathbb{E}(x) = \sum_{i=1}^N \alpha_i \mathbb{E}(\theta_i)$. A group's ex ante utility is thus, depending on its type, $\pm\beta + U(E^k - \sum_{i=1}^N e_i^k)$. This payoff only depends on, and is weakly decreasing with, the total effort $\sum_{i=1}^N e_i^k$. □

Proof of Proposition 5. This Proposition unfolds from Lemma 5. If a group commits to $e^k=0$, it never acquires information and never affects beliefs. Since the DM's decision is linear in payoff in beliefs, as well as the policy part of the group's payoff, the group does not lose any informational rents from being inactive, but saves on the opportunity cost associated with $U(E^k - \sum_{i=1}^N e_i^k)$. Thus, inactivity strictly dominates participation. □

Proof of Proposition 6. Consider dimension i, $i \in \{1,\dots,N\}$, and group g. Wlog, $g \in \mathcal{K}_P$. Assume there are in total $K_P \geq 1$ groups of type P and $K_C \geq 0$ groups of type C that specialize in θ_i. In case no information is transmitted, $\theta_i^* = \frac{p\prod_{k\in\mathcal{K}_P}(1-e_i^k)}{p\prod_{k\in\mathcal{K}_P}(1-e_i^k)+(1-p)\prod_{h\in\mathcal{K}_C}(1-e_i^h)}$. The marginal benefit from being informed, for group g, is

$$\frac{v_i^g(e_i)}{\alpha_i p} = \prod_{k\in\mathcal{K}_P\setminus\{g\}}(1-e_i^k)\left(1-\theta_i^*(e)\right) = \prod_{k\in\mathcal{K}_P\setminus\{g\}}(1-e_i^k)\frac{(1-p)\displaystyle\prod_{h\in\mathcal{K}_C}(1-e_i^h)}{p\displaystyle\prod_{k\in\mathcal{K}_P}(1-e_i^k)+(1-p)\displaystyle\prod_{h\in\mathcal{K}_C}(1-e_i^h)}.$$

If $K_P \geq 2$, we can compute, for $k^* \in \mathcal{K}_P\setminus\{g\}$, $\frac{\partial v_i^g(e)}{\partial e_i^{k^*}}$. The sign of this expression is given by $-(1-p)\prod_{h\in\mathcal{K}_C}(1-e_i^h)<0$.

Similarly, if $K_C \geq 1$, we can compute for $h\in\mathcal{K}_C$, $\frac{\partial v_i^g(e)}{\partial e_i^h}$. The sign of this expression is given by $-p\prod_{k\in\mathcal{K}_P}(1-e_i^k)<0$.

It follows that starting with a situation where group g is the sole provider of information about θ_i, adding other groups reduces this group's level of search at equilibrium. This proves Item 1.

Since each group also prefers to be the sole provider on dimension N than on dimension $i<N$, the second Item directly follows: all groups are better off in the situation where all concentrate on dimension N.

□

Proof of Proposition 7. Consider dimension 1, and assume that $\bar{\alpha}=1$. Given the symmetry of the problem, DM's payoff only depends, and positively so, on the effort of each group.

This effort is characterized by equation (4.4), where $v_1^P(e,p)=\frac{p(1-p)}{1-pe}$. When a group is alone on its dimension, $v_1^P(0,p)=(1-p)p>U'(E^P-0)$, where the second inequality comes from assumption (3.2). With $p=\frac{1}{2}$, we have $\frac{\partial v_1^P}{\partial p}(e,p)=\frac{e}{(2-e)^2}>0$. Moreover, $v_1^P(E^P,p)<v_1^P(1,p)=p<U'(E^P-E^P)$.

Let us denote e^* the highest value of $e\in(0,E^P)$ such that (4.4) is satisfied. This effort corresponds to the equilibrium that DM favors. Necessarily, for η small enough $v_1^P(e^*-\eta,p)>U'\left(E^P-(e^*-\eta)\right)$ while $v_1^P(e^*+\eta,p)<U'\left(E^P-(e^*+\eta)\right)$.

Finally, $\frac{\partial v_1^P}{\partial e}(e,p)=\frac{p^2(1-p)}{(1-pe)^2}>0$ for all $p\in(0,1)$. Denoting $e(p)$ the highest equilibrium effort exerted by group P, we thus find $e'(\frac{1}{2})>0$. DM will thus optimally affect each group on its strong domain.

Given that inequalities are strict, they can be extented to $\bar{\alpha}>1$ and $\bar{\delta}$ small enough. □

Proof of Proposition 8. Assume first $\bar{\alpha}=1$. We denote $\tilde{K}\equiv K_C=K_P$. Consider a situation in which all groups in $\mathcal{K}_P$ specialize in dimension 1 and all groups in $\mathcal{K}_C$ specialize in dimension

2. Given the symmetry of the problem, there exists on each dimension an equilibrium in which for all $k,\ell \in \mathcal{K}_P$, $e_1^k = e_1^\ell \equiv e$, and for all $k',\ell' \in \mathcal{K}_P$, $e_2^{k'} = e_2^{\ell'} = e$.

The condition for a group $k \in \mathcal{K}_P$ not to deviate writes

$$(1-e)^{\tilde{K}-1}\left(1-\theta_1^*(e)\right) > \left(1-\theta_2^*(e)\right), \tag{A.1}$$

with $\theta_1^*(e) = \frac{(1-e)^{\tilde{K}}}{(1-e)^{\tilde{K}}+1}$, and $\theta_2^*(e) = \frac{1}{(1-e)^{\tilde{K}}+1}$. Using the values in (A.1), a direct substitution and simplification yields

$$(1-e)^{\tilde{K}-1} > (1-e)^{\tilde{K}}.$$

This condition is always satisfied for $e \in (0,1)$. By symmetry, the same reasoning holds for $k \in \mathcal{K}_C$. Finally, since the inequality is a strict one, the result extends to some $\bar{\alpha} > 1$. □