



^b
**UNIVERSITÄT
BERN**

Faculty of Business, Economics
and Social Sciences

Department of Economics

Strategic Attribute Learning

Jean-Michel Benkert, Ludmila Matyskova,
Egor Starkov

24-11

December, 2024

DISCUSSION PAPERS

Schanzeneckstrasse 1
CH-3012 Bern, Switzerland
<http://www.vwi.unibe.ch>

Strategic Attribute Learning*

Jean-Michel Benkert, Ludmila Matysková, and Egor Starkov†

December 13, 2024

Abstract

A researcher allocates a budget of informative tests across multiple unknown attributes to influence a decision-maker. We derive the researcher's equilibrium learning strategy by solving an auxiliary single-player problem. The attribute weights in this problem depend on how much the researcher and the decision-maker disagree. If the researcher expects an excessive response to new information, she forgoes learning altogether. In an organizational context, we show that a manager favors more diverse analysts as the hierarchical distance grows. In another application, we show how an appropriately opposed advisor can constrain a discriminatory politician, and identify the welfare-inequality Pareto frontier of researchers.

Keywords: Attributes, Information acquisition, Gaussian distribution, Strategic learning
JEL classification: D72, D81, D83

*For valuable comments, we thank: Arjada Bardhi, Nathan Hancart, Toomas Hinnosaar, Alessandro Spano, Igor Letina, Antoine Loeper, Marc Möller, Nick Netzer, Francisco Poggi, Johannes Schneider, Armin Schmutzler, Jakub Steiner, Peter Norman Sørensen, Dezső Szalay, and the audiences at the Universities of Alicante, Bern, Copenhagen, Madrid Carlos III, Mannheim, Zurich, and CERGE-EI, as well as the Nordic Theory Meetings, the Annual Meeting of European Association of Young Economists, the Lisbon Meetings in Game Theory and Applications, the Barcelona Summer Forum, EEA-ESEM, and the Annual Meeting of the VfS. Matysková gratefully acknowledges funding from the Generalitat Valenciana (Prometeo/2021/073) and from the grant PID2022-142356NB-I00 financed by MICIU/AEI/10.13039/501100011033 and FEDER, UE.

†Benkert: Department of Economics, University of Bern. Matysková: Department of Economics, University of Alicante. Starkov: Department of Economics, University of Copenhagen. Email: jean-michel.benkert@unibe.ch, ludmila.matyskova@ua.es, egor.starkov@econ.ku.dk.

1 Introduction

Decision-making under uncertainty often involves multiple dimensions of potentially unequal importance. For example, when designing a new product, a company may consider the product’s uncertain reception across various market segments, some of which may be more important than others. Similarly, a politician’s optimal policy may need to address the diverse and uncertain needs of distinct social groups, with some groups potentially carrying more weight than others do.

In both examples, decision-makers must learn about the various dimensions, or attributes, influencing their decisions but often lack the resources for a thorough investigation. Consequently, a specialized researcher is often in charge of the learning task. However, the researcher may prioritize attributes differently from the decision-maker and strategically procure information about the unknown attributes to influence the final decision.

In this paper, we explore questions pertaining to such strategic concerns in learning about complex decisions. First, what kinds of biases may arise when multiple attributes are at play, and how do they interact? For instance, in an organizational application, we explore two orthogonal biases—hierarchical distance and diversity—and analyze when a manager (decision-maker) prefers diverse analysts (researchers) as a function of hierarchical distance. Second, how can biases mitigate socially undesirable decisions? For instance, in the context of political discrimination, we show how an appropriately opposed advisor (researcher) can constrain a discriminatory politician whose decision affects well-being of various social groups (attributes). More broadly, how does preference misalignment in a multiattribute environment affect learning, when researchers shape not only the *extent* of learning but also its *direction*?

To address these questions, we develop a framework that captures the strategic interaction between a *researcher*, who learns about an unknown state of the world, and a *decision-maker*, who acts based on the resulting information. Crucially, we assume that the state of the world consists of multiple *attributes* and capture the player’s preference misalignment by allowing for differing weights on the attributes. As such, our main contribution is a tractable model for studying how preference misalignment affects multi-attribute learning and decision-making.

In Section 2, we formally introduce our framework. In our model, the attributes are independently and normally distributed and jointly determine the state of the world and thus the players' preferred decisions. The state can be imperfectly learned by allocating a given budget of *tests* across different attributes. The more tests are allocated to an attribute, the more informative a signal about it becomes. Players aim to minimize the quadratic loss between the final decision and their *bliss point*, which is a weighted sum of attributes.¹

We begin by analyzing a helpful benchmark in Section 3: a situation in which a single player controls both the decision and the learning, where the optimal learning strategy is the primary focus of our analysis. Theorem 1 shows that—absent strategic motives—the player chooses the test allocation which achieves the highest reduction in residual uncertainty regarding her optimal decision. In particular, optimal learning reflects both the attribute's relevance to the agent and its prior uncertainty. Put differently, the more important and more uncertain an attribute is, the more tests are allocated to it. Moreover, the budget size determines how many attributes the agent learns about.

In Section 4, we return to the primary framework, in which a researcher controls the learning and a decision-maker makes the decision. Our main result (Theorem 2) shows that the researcher's equilibrium test allocation then coincides with the single-player solution with *auxiliary weights*, which are different from the actual weights of either player. We show that these auxiliary weights decrease as the misalignment between players grows. If the researcher values an attribute less (or more) than the decision-maker does, he views the decision-maker's reaction to any information about that attribute as excessive (or insufficient). When the misalignment results in overreaction, it can lead the researcher to abstain from learning about the attribute entirely. Notably, when all of the researcher's weights are sufficiently low relative to the decision-maker's weights, the researcher entirely forgoes learning—contrasting with the single-player case, where learning generically takes place. The fact that the researcher's equilibrium test allocation coincides with the single-player solution with auxiliary weights clarifies how the strategic motive influences the researcher's learning, reducing its impact to a shift in attribute weights. Thus, by comparing these auxiliary weights to the researcher's original weights, we can pinpoint how misalignment drives adjustments in the researcher's learn-

¹Our model is a static variant of the framework in Liang, Mu, and Syrgkanis (2022), who study a single agent's dynamic learning under attribute correlation.

ing strategy—insight that would be harder to achieve if the strategic solution deviated fundamentally from the single-player case.

We explore two economical applications, each with a tailored parametrization of player misalignment. These applications demonstrate the versatility of our analytical tools in studying misalignment and yield new insights within each specific context. In our first application, we examine how biases across layers of hierarchy influence learning in organizations. Here, an analyst (researcher) gathers information, and a manager (decision-maker) acts on it. While the manager trivially prefers an analyst whose preferences match hers, we examine the preferred analyst when full alignment is unfeasible. We parametrize misalignment by decomposing the analyst’s preferences into *sensitivity* (how closely the players align in absolute terms) and *distortion* (how much the players disagree in relative terms). Proposition 2 shows that the manager may prefer analysts with high distortion when the analysts have low sensitivity. An analyst with low sensitivity views the manager’s reaction to any information as excessive and thus may completely forgo learning. In such cases, the manager then prefers an analyst who is very distorted toward some attribute, as then the manager’s reaction to information about that attribute is no longer viewed as excessive, resulting in some learning. As any learning is better for the manager than no learning, she prefers a strongly distorted analyst over one with little or no distortion. Assuming that the more hierarchical the organization, the bigger the difference in the sensitivity between the managers and the analysts, our results suggest balancing diversity and uniformity (reflected in distortion) in organizations as a function of hierarchical distance.

In our second application, we explore the issue of discrimination. A politician (decision-maker) chooses a policy to meet the uncertain needs of two a priori identical social groups (attributes) but favors one group over the other. With an unchecked politician—who controls both decision-making and learning—such favoritism results in inequality and reduces utilitarian welfare compared to the situation without favoritism. We examine how delegating learning to an advisor (researcher) with different preferences can mitigate these negative effects. Proposition 4 demonstrates that there exists a welfare-inequality Pareto frontier of potential advisors that strictly dominates unchecked discrimination. At one end of this frontier is an impartial advisor, who—given the politician’s favoritism—maximizes welfare, but leaves the groups with unequal outcomes (Lemma 3). At the other end of the frontier is an advisor who is suitably partial towards

the disadvantaged group. Such an advisor learns precisely enough about this group to counteract the politician’s favoritism, and thus eliminates inequality when appointed (Lemma 4). However, it comes at a cost to welfare, as the politician then relies less on the advisor’s information, and information is thus “wasted.” Notably, the level of advisor’s partiality required to eliminate inequality does not increase with the politician’s favoritism but exhibits a *non-monotone* relationship.

In Section 7, we outline two extensions developed in detail in the online appendix. The first examines multiple researchers in the context of media markets, showing that while competition between media outlets (researchers) leads to polarization in equilibrium, this outcome is actually beneficial for a voter (decision-maker). The second extension addresses uncertainty about the decision-maker’s preferences. We study this within a dual-self model, where a single agent may be either sophisticated or naïve about potential changes in her preferences between the learning and the decision-making stages, finding that naïvete can, in fact, be advantageous.

We conclude in Section 8 by discussing our results and prospective avenues for future research.

Related literature. Our work is closely related to Bardhi (2024), who examines a strategic multi-attribute learning problem where a project’s payoff is composed of correlated attributes, weighted differently by a decision-maker and a researcher. The researcher selects which attributes to sample (and perfectly learn), given that he can sample a limited number. Bardhi (2024) primarily focuses on the role of correlation between attributes. In contrast, we aim to provide a tractable model to study preference misalignment. First, we focus on independent attributes to separate misalignment effects from correlation. Second, we consider noisy learning, where the choice involves not only which attributes to learn about but also how much to learn about each. Our analysis demonstrates that with as few as two attributes, noisy learning effectively captures the complexity of bias interactions.²

Kirneva (2023), building on Tamura (2018), considers a related multidimensional setting, where the decision-maker can acquire additional costly information on top of what

²More broadly, our framework is related to the Gaussian-sampling literature such as Bardhi and Bobkova (2023), Callander (2011), and Carnehl and Schneider (2024).

is provided by the researcher. The need to influence the decision-maker’s learning strategy (and not only the decision) then shapes the researcher’s learning strategy. The researcher may then provide partial information about the attribute on which the players’ preferences are misaligned in order to divert the decision-maker’s learning away from it.³ In contrast, we do not allow for independent learning by the decision-maker, creating a distinct strategic environment and shutting down the “attention diversion” channel.

Our framework builds on Liang et al. (2022), adapting their dynamic model of multiple attributes to a static form. Furthermore, while they study a single agent’s dynamic learning under attribute correlation, we focus on strategic learning and abstract from correlation to highlight the role of preference misalignment between players. Notably, there is a key parallel in our findings. Both papers show that a “greedy” learning strategy, which achieves the highest reduction in residual uncertainty and is optimal in a single-player model without correlation, remains optimal in the presence of correlation (Liang et al., 2022) and in the presence of strategic motives (our paper), respectively.

Our first application, on diversity in organization, connects to the literature on delegated expertise, where a decision-maker delegates learning about the state of interest to a biased researcher.⁴ We explore the effects of different biases between the researcher and the decision-maker, which is related to studies by Ball and Gao (2024), Che and Kartik (2009), and Ilinov et al. (2022). These papers show that delegating learning to a researcher with some preference misalignment can be optimal in single-dimensional settings, as misalignment encourages the acquisition of more costly information. In contrast, we explore a multidimensional setting, where the researcher decides *which attributes* to learn about and to what extent. Our findings in Section 6.1 reveal that higher misalignment on a certain type of bias can be beneficial, as it may prevent the breakdown of learning caused by misalignment on another, orthogonal, type of bias.

Our second application, on discrimination in policymaking, relates to Fosgerau et al. (2023) and Echenique and Li (2023), who explore how decision-makers’ strategic learn-

³Similar insights—namely, that when the receiver can acquire information, the sender’s choices are primarily driven by a desire to influence the learning process rather than the final decision—have also been obtained by Matveenko and Starkov (2023).

⁴Pioneered by Demski and Sappington (1987), this literature has seen renewed interest, for instance, see Deimen and Szalay (2019) and Lindbeck and Weibull (2020).

ing choices reinforce discrimination.⁵ They show that employers’ discriminatory beliefs shape job candidates’ incentives to invest in their skills, potentially leading to a self-sustaining discriminatory equilibrium. In contrast, we focus on countering discriminatory tendencies by separating learning from decision-making and delegating it to independent advisors. Further, our results in Section 6.2 complement Liang et al. (2024). They study the fairness-accuracy Pareto frontier in the context of an interaction between an egalitarian researcher and a utilitarian decision-maker, where the researcher may coarsen or ban information about certain attributes. In contrast, we examine the welfare-inequality frontier in a setting where the decision-maker is inherently unfair but can be paired with a researcher holding different preferences, who can flexibly allocate informative tests across different attributes.

2 Model

2.1 Setup

Players. There are two players, a decision-maker (D , she) and a researcher (R , he). First, the researcher chooses the learning strategy to examine payoff-relevant attributes of the state of the world by choosing how to allocate a budget of tests T . Second, the decision-maker observes the results of these tests and makes a decision.

Attributes. There is an unknown state of the world $\tilde{\theta} = (\tilde{\theta}_1, \dots, \tilde{\theta}_K)$ consisting of $K \geq 2$ attributes.⁶ All attributes $\tilde{\theta}_k$ for $k = 1, \dots, K$ are jointly normally distributed with commonly known prior means $\mu_k^0 \in \mathbb{R}$ and prior variances $\Sigma_k^0 > 0$; that is, $\tilde{\theta}_k \sim \mathcal{N}(\mu_k^0, \Sigma_k^0)$. Moreover, the attributes are independent: $\tilde{\theta}_l \perp \tilde{\theta}_j$ for $l \neq j$.

Actions. The decision-maker chooses a decision $d \in \mathbb{R}$. The researcher, given an exogenous budget of tests $T > 0$, chooses a test allocation $\tau \in \mathcal{T} := \{\tau \in \mathbb{R}_+^K : \tau_1 + \dots + \tau_K \leq T\}$, where τ_k is the amount of tests allocated to learn about attribute

⁵See Onuchic (2024) for a recent overview of theories of discrimination.

⁶We denote a random variable and its realization by \tilde{x} and x , respectively.

$\tilde{\theta}_k$. Allocating τ_k to attribute $\tilde{\theta}_k$ yields a single observation of a signal,

$$\tilde{s}_k(\tau) = \tilde{\theta}_k + \tilde{\varepsilon}_k, \quad (1)$$

where $\tilde{\varepsilon}_k \sim \mathcal{N}(0, 1/\tau_k)$, and where $\tilde{\varepsilon}_k$ is independent of all other random variables.⁷ When the researcher chooses $\tau = (0, \dots, 0)$, we say *the researcher abstains from learning*. The test allocation $\tau = (\tau_1, \tau_2, \dots, \tau_K)$ and the subsequent signal realizations $s = (s_1, s_2, \dots, s_K)$ are publicly observed, and the players are symmetrically informed throughout the game.

Payoffs. Each player $i = D, R$ has a bliss point, which is a linear combination of the realized attribute values:

$$b^i(\theta) := \alpha_1^i \theta_1 + \dots + \alpha_K^i \theta_K$$

for commonly known weights $\alpha^i = (\alpha_1^i, \dots, \alpha_K^i) \in \mathbb{R}_+^K$.⁸ The players may weigh attributes differently, which is the source of their conflict. When the decision-maker takes a decision $d \in \mathbb{R}$ and the realized state is θ , player i obtains ex post utility

$$u^i(d, \theta) = -(d - b^i(\theta))^2.$$

From the ex ante perspective, i 's bliss point $b^i(\theta)$ is distributed according to

$$\tilde{b}^i := b^i(\tilde{\theta}) \sim \mathcal{N}\left(b_0^i, \sigma_0^{2,i}\right),$$

where $b_0^i = \alpha_1^i \mu_1^0 + \dots + \alpha_K^i \mu_K^0$ is player i 's ex ante expectation of their bliss point, and $\sigma_0^{2,i} = (\alpha_1^i)^2 \Sigma_1^0 + \dots + (\alpha_K^i)^2 \Sigma_K^0$ is the prior variance of player i 's bliss point.

⁷When $\tau_k = 0$, the signal $\tilde{s}_k(\tau)$ is considered completely uninformative about $\tilde{\theta}_k$.

⁸The non-negativity assumption $\alpha_k^i \geq 0$ simplifies the exposition and is without loss of generality. If the players' weights α_k^D and α_k^R had opposite signs, the researcher would never learn about attribute $\tilde{\theta}_k$, effectively equivalent to setting $\alpha_k^R = 0$. Additionally, since any attribute can be replaced by its negative, if $\alpha_k^D, \alpha_k^R < 0$ for some k , we can instead consider $-\tilde{\theta}_k$ with positive weights $-\alpha_k^D, -\alpha_k^R > 0$.

Timing. First, at the ex ante stage, the researcher chooses a publicly observable test allocation τ to maximize his ex ante expected utility

$$V^R(\tau) := \mathbb{E} \left[u^R \left(d(\tilde{s}, \tau), \tilde{\theta} \right) \right], \quad (2)$$

where, given some anticipated decision strategy $d(s, \tau)$, the expectation is taken over the induced distribution of the decisions and the state of the world. Next, nature draws the state realization θ and signal realizations s . The players observe s and update their beliefs about the state (and bliss points). Then, at the interim stage, the decision-maker chooses decision d to maximize her interim expected utility

$$U^D(d, s, \tau) := \mathbb{E} \left[u^D \left(d, \tilde{\theta} \right) \right], \quad (3)$$

where the expectation is taken over the state. Finally, at the ex post stage, payoffs are realized.

Equilibrium. A weak Perfect Bayesian Equilibrium of the game (henceforth, “equilibrium”) is a triple $(\tau^*, d^*(s, \tau), \tilde{\theta}(s, \tau))$ such that

1. the researcher’s test allocation strategy $\tau^* \in \mathcal{T}$ maximizes his expected utility (2) given the decision-maker’s strategy $d^*(s, \tau)$ and her prior belief; whenever the researcher is indifferent between some τ and abstaining from learning, he chooses the latter;⁹
2. the decision-maker’s decision strategy $d^*(s, \tau) : \mathbb{R}^K \times \mathcal{T} \rightarrow \mathbb{R}$ maximizes her expected utility (3) given her posterior beliefs $\tilde{\theta}(s, \tau)$; and
3. the decision-maker’s posterior beliefs $\tilde{\theta}(s, \tau) : \mathbb{R}^K \times \mathcal{T} \rightarrow \Delta(\mathbb{R}^K)$ are obtained via Bayes’ rule for all possible signal realizations s , given the researcher’s choice τ .

2.2 Preliminary analysis: belief updating

Given test allocation $\tau = (\tau_1, \dots, \tau_K)$ and signal realizations $s = (s_1, \dots, s_K)$, let $\tilde{\theta}_k(s_k, \tau_k)$ and $\tilde{b}^i(s, \tau)$ denote the posterior beliefs about attribute $\tilde{\theta}_k$ and about player

⁹The tie-breaking rule resolves a possible multiplicity of equilibria, which only arises in non-generic cases.

i 's bliss point, respectively. For every k , we get

$$\tilde{\theta}_k(s_k, \tau_k) \sim \mathcal{N}\left(\hat{\mu}(s_k, \tau_k), \hat{\Sigma}_k(\tau_k)\right), \quad (4)$$

where

$$\hat{\mu}_k(s_k, \tau_k) = \frac{1}{1 + \tau_k \Sigma_k^0} \mu_k^0 + \frac{\tau_k \Sigma_k^0}{1 + \tau_k \Sigma_k^0} s_k, \quad (5)$$

$$\hat{\Sigma}_k(\tau_k) = \frac{\Sigma_k^0}{1 + \tau_k \Sigma_k^0}. \quad (6)$$

Since $b^i(\theta) = \sum_k \alpha_k^i \theta_k$ and attributes are independent, we get

$$\tilde{b}^i(s, \tau) \sim \mathcal{N}\left(\sum_k \alpha_k^i \hat{\mu}_k(s_k, \tau_k), \hat{\sigma}^{2,i}(\tau)\right), \quad (7)$$

where

$$\hat{\sigma}^{2,i}(\tau) = (\alpha_1^i)^2 \hat{\Sigma}_1(\tau_1) + \dots + (\alpha_K^i)^2 \hat{\Sigma}_K(\tau_K). \quad (8)$$

Due to the properties of normal distribution, the variances $\hat{\Sigma}_k(\tau_k)$, and hence also the variance of $\tilde{b}^i(s, \tau)$, do not depend on signal realizations s .

3 The single-player case

We begin by analyzing the scenario in which the decision-maker and the researcher share identical preferences, referred to as the single-player case. Here, we omit the players' superscripts $i = D, R$ and set $\alpha := \alpha^R = \alpha^D$, $b(\theta) := b^R(\theta) = b^D(\theta)$, and so forth. For simplicity, we refer to this single player as the *agent* in this section. In the learning stage, the agent allocates tests to learn about the attributes, and in the decision stage, after observing the signal realizations, the agent takes a decision.

The optimal decision strategy is straightforward.¹⁰ Given a test allocation τ and signal realizations s , the expected utility (3) is maximized by

$$d^*(s, \tau) = \mathbb{E}[\tilde{b}(s, \tau)] = \alpha_1 \hat{\mu}_1(s_1, \tau_1) + \dots + \alpha_K \hat{\mu}_K(s_K, \tau_K), \quad (9)$$

¹⁰We use "optimal" strategy to denote the equilibrium solution defined in Section 2 when the researcher's and decision-maker's weights coincide.

where $\hat{\mu}_k(s_k, \tau_k)$ for every $k = 1, \dots, K$ is given by equation (5). Conditional on τ , the optimal decision is ex ante distributed as

$$\tilde{d}^*(\tau) := d^*(\tilde{s}, \tau) \sim \mathcal{N}(b_0, \psi(\tau)), \quad (10)$$

where the variance is given by $\psi(\tau) := \sigma_0^2 - \hat{\sigma}^2(\tau)$. Anticipating the optimal decision strategy, the problem of choosing a test allocation can be reformulated as follows.

Lemma 1. In the single-player case, the optimal test allocation is given by

$$\arg \max_{\tau \in \mathcal{T}} V(\tau) = \arg \min_{\tau \in \mathcal{T}} \hat{\sigma}^2(\tau). \quad (11)$$

Lemma 1 establishes that the agent chooses τ to minimize the residual uncertainty $\hat{\sigma}^2$ regarding the bliss point \tilde{b} (i.e., the residual uncertainty concerning the optimal decision) that remains at the posterior beliefs. Our first main result below characterizes the test allocation that accomplishes this objective. For clarity, we present the characterization in the case of two attributes, with the general form for $K \geq 2$ available in the appendix.

Theorem 1. Consider the single-player case and let $K = 2$. Without loss of generality, let $\alpha_1 \Sigma_1^0 \geq \alpha_2 \Sigma_2^0$. Then, the unique optimal test allocation $\tau^* = (\tau_1^*, \tau_2^*)$ is as follows: if $\alpha_1 = \alpha_2 = 0$, then $\tau^* = (0, 0)$; otherwise,

$$\tau_1^* = \begin{cases} T & \text{if } T \leq \bar{T} \\ \bar{T} + \frac{\alpha_1}{\alpha_1 + \alpha_2}(T - \bar{T}) & \text{if } T > \bar{T} \end{cases}$$

and $\tau_2^* = T - \tau_1^*$, where $\bar{T} = \frac{\alpha_1}{\alpha_2 \Sigma_2^0} - \frac{1}{\Sigma_1^0}$.

Theorem 1 shows that the agent's optimal test allocation depends on the weights α_k solely through their *ratios*. The ex ante marginal value of learning about attribute $\tilde{\theta}_k$ is proportional to the term $\alpha_k \Sigma_k^0$, which reflects both the attribute's relevance to the agent (weight α_k) and its prior uncertainty (prior variance Σ_k^0). The budget size T determines whether the agent learns exclusively about the attribute with the higher ex ante marginal value of learning ($T \leq \bar{T}$) or if he learns about both attributes ($T > \bar{T}$). Note that allocating τ_k tests to learn about attribute $\tilde{\theta}_k$ decreases its interim marginal value of learning, which is given by $\left(\frac{\alpha_k \Sigma_k^0}{1 + \tau_k \Sigma_k^0}\right)^2$. Evaluating the interim marginal value of learning for attribute $\tilde{\theta}_1$ at $\tau_1 = \bar{T}$, we obtain $(\alpha_2 \Sigma_2^0)^2$, which matches the ex ante

marginal value of learning about attribute $\tilde{\theta}_2$. Hence, with \bar{T} tests allocated to attribute $\tilde{\theta}_1$ and none to $\tilde{\theta}_2$, the marginal values of learning about the two attributes are equalized. Therefore, any test budget $T \leq \bar{T}$ is allocated exclusively to attribute $\tilde{\theta}_1$, and the agent learns about attribute $\tilde{\theta}_2$ only when more tests are available. These additional tests are allocated between $\tilde{\theta}_1$ and $\tilde{\theta}_2$ in proportion to weights α_1 and α_2 , keeping the interim marginal values of learning about the two attributes equal.¹¹

Note that the agent exhausts the entire budget T , except in the degenerate case where neither attribute holds any importance ($\alpha_1 = \alpha_2 = 0$). In this scenario, consistent with our equilibrium selection rule, the agent abstains from learning.

4 Strategic players

We now turn to the general case where the players have different preferences, leading to a conflict of interest. We proceed by backward induction, first solving the decision-maker's problem and then addressing the researcher's learning problem.

The decision-maker's problem is analogous to the decision stage in the single-player problem with $\alpha = \alpha^D$. In particular, maximizing the expected utility (3) over d given some test allocation τ yields the equilibrium decision strategy $d^*(s, \tau) = \mathbb{E}[\tilde{b}^D(s, \tau)] = \alpha_1^D \hat{\mu}(s_1, \tau_1) + \dots + \alpha_K^D \hat{\mu}(s_K, \tau_K)$. Thus, for a given τ , the equilibrium decision is ex ante distributed as $\tilde{d}^*(\tau) := d^*(\tilde{s}, \tau) \sim N(b_0^D, \psi^D(\tau))$, where $\psi^D(\tau) := \sigma_0^{2,D} - \hat{\sigma}^{2,D}(\tau)$ and $\hat{\sigma}^{2,D}(\tau)$ is given by (8).

Turning to the researcher's learning problem, Lemma 1 shows that if the researcher's optimal decisions were implemented, he would choose τ to minimize $\hat{\sigma}^{2,R}(\tau)$, the posterior uncertainty about his optimal decision. In a strategic setting, however, the researcher prioritizes attributes where his weights align closely with those of the decision-maker. The following lemma captures the essence of the researcher's maximization problem.

¹¹As noted in the literature review, Liang et al. (2022) investigate an agent's optimal dynamic information acquisition problem in a framework similar to ours but which allows for correlated attributes. Our single-agent solution can also be obtained using their results with suitable modifications.

Lemma 2. The researcher’s equilibrium test allocation solves

$$\max_{\tau \in \mathcal{T}} V^R(\tau) = \max_{\tau \in \mathcal{T}} \left\{ \underbrace{2 \sum_k \alpha_k^D \alpha_k^R \Sigma_k^0 \hat{\Sigma}_k(\tau_k) \tau_k}_{2 \operatorname{cov}(\tilde{d}^*(\tau), \tilde{b}^R)} - \underbrace{\left(\sigma_0^{2,D} - \hat{\sigma}^{2,D}(\tau) \right)}_{\psi^D(\tau)} + V^R(\emptyset) \right\}, \quad (12)$$

where $\hat{\Sigma}_k(\tau_k)$ for every $k = 1, \dots, K$ are given by equation (6), $\hat{\sigma}^{2,D}(\tau)$ by equation (8), and $V^R(\emptyset)$ is the researcher’s expected utility under no learning (constant in τ).

Lemma 2 shows that the researcher pursues two interdependent objectives. First, he seeks to align the final decision with his own bliss point: $\uparrow \operatorname{cov}(\tilde{d}^*(\tau), \tilde{b}^R)$. Second, he aims to prevent the decision-maker from overreacting to new information relative to the researcher’s bliss point: $\downarrow (\psi^D(\tau) - \operatorname{cov}(\tilde{d}^*(\tau), \tilde{b}^R))$. Together, these effects constitute the added value of learning.¹²

To concisely formulate our main result—the solution to the researcher’s problem—we introduce a final piece of notation. For all $k = 1, \dots, K$, let

$$\begin{aligned} \Delta_k &:= |\alpha_k^R - \alpha_k^D|, \\ \lambda_k &:= 2\alpha_k^R \alpha_k^D - (\alpha_k^D)^2 = (\alpha_k^R)^2 - (\Delta_k)^2. \end{aligned} \quad (13)$$

Hence, Δ_k is the distance between the players’ weights on that attribute, and λ_k is a parameter that reflects the aforementioned trade-off.


Theorem 2. The equilibrium test allocation τ^* coincides with the solution of the single-player problem with *auxiliary weights*

$$\hat{\alpha}_k := \sqrt{\max\{\lambda_k, 0\}} \quad \text{for all } k = 1, \dots, K. \quad (14)$$

Theorem 2 reveals a key insight: even when the researcher and the decision-maker have different weights, the researcher’s equilibrium test allocation coincides with a solution to the single-player problem with some auxiliary weights. Thus, the effect of the strategic motive on the researcher’s learning reduces to a shift in attribute weights. Therefore, by comparing these auxiliary weights to the researcher’s original weights, we can pinpoint

¹²The trade-off illustrated by Lemma 2 also arises in Theorem 2 in Bardhi (2024). However, the specifics of ψ^D and $\operatorname{cov}(\tilde{d}^*(\tau), \tilde{b}^R)$ differ between the two models due to the distinct ways in which uncertainty and player payoffs are modeled.

FIGURE 1. THE RESEARCHER’S SOLUTION: AUXILIARY SINGLE-PLAYER WEIGHTS



Notes: The figures depict $\hat{\alpha}_k$ as a function of the decision-maker’s weight, α_k^D , in panel (a), and as a function of the researcher’s weight, α_k^R , in panel (b), keeping the weight of the other player constant.

how misalignment drives adjustments in the researcher’s learning strategy—insight that would be harder to achieve if the strategic solution deviated fundamentally from the single-player case.

From the researcher’s perspective, the ex ante marginal value of learning about attribute $\tilde{\theta}_k$ is proportional to the parameter λ_k , which—unlike in the single-player case—can be negative. Theorem 2 explains the link between the auxiliary weights $\hat{\alpha}_k$ and the parameters λ_k . Note that λ_k (and thus $\hat{\alpha}_k$) is decreasing in the misalignment between the players’ weights Δ_k . When $\Delta_k > 0$, the researcher views the decision-maker’s reaction to information about $\tilde{\theta}_k$ as either excessive ($\alpha_k^D > \alpha_k^R$) or insufficient ($\alpha_k^D < \alpha_k^R$), thereby reducing the researcher’s incentives to learn about that attribute. If the reaction is too strong ($\alpha_k^D \geq 2\alpha_k^R$) or too weak ($\alpha_k^D = 0$), the researcher avoids learning about that attribute entirely. If $\lambda_k \leq 0$ for all k , the researcher abstains from learning, as any information would lead to undesirable overreaction (or no reaction) by the decision-maker, making the status quo decision the researcher’s preferred outcome.

When players disagree on the weight of an attribute, the auxiliary weight differs from the weight of either player. Figure 1 illustrates how $\hat{\alpha}_k$ depends on α_k^R and α_k^D , keeping the weight of the other player fixed. Panel (a) shows that $\hat{\alpha}_k$ always satisfies $\hat{\alpha}_k \leq \alpha_k^R$: the misalignment effectively reduces the importance the researcher assigns to the attribute. Nevertheless, the extent and direction of any distortion in the equilibrium test allocation—compared to the researcher’s non-strategic optimum—then depend on how the *ratios* of the auxiliary weights differ from the ratios of the researcher’s weights.

5 Equivalent payoff specifications

This section introduces alternative payoff structures that result in the same equilibrium test allocation as the baseline model. The objective of presenting these frameworks is to offer alternatives that may be better tailored to particular economic applications, thus demonstrating the adaptability of our model and the analytical tools detailed in Sections 3 and 4.

Baseline model. Introduced in Section 2, the baseline model features a decision-maker taking a single decision, $d \in \mathbb{R}$. Given a decision d and a realized state $\theta = (\theta_1, \dots, \theta_K)$, the utility of player $i = R, D$ is

$$u^i(d, \theta) = -(d - b^i(\theta))^2 = -\left(d - \sum_k \alpha_k^i \theta_k\right)^2. \quad (15)$$

Framework A. In this framework, the decision-maker also takes a single decision, $d \in \mathbb{R}$. Given a decision d and a realized state $\theta = (\theta_1, \dots, \theta_K)$, the utility of player $i = R, D$ is

$$u_A^i(d, \theta) = -\sum_k \alpha_k^i (d - \theta_k)^2, \quad (16)$$

where $\sum_k \alpha_k^i = 1$ for both players.

Framework B. This framework features the decision-maker simultaneously taking K distinct decisions $d_1, \dots, d_K \in \mathbb{R}$. Given decisions $d = (d_1, \dots, d_K)$ and a realized state $\theta = (\theta_1, \dots, \theta_K)$, the utility of player $i = R, D$ is

$$u_B^i(d, \theta) = -\sum_k (d_k - \alpha_k^i \theta_k)^2. \quad (17)$$

In all three models: (i) players aim to minimize a loss given by a quadratic distance, and (ii) the weights α^i determine the relative importance of different attributes. Beyond these similarities, the frameworks differ in how the players aggregate losses across attributes and decisions. Despite these differences, the following proposition (proved in

Online Appendix C.1) shows that Frameworks A and B are equivalent to the baseline model in terms of the equilibrium test allocation.

Proposition 1. Given the decision-maker’s weights α^D , the researcher’s weights α^R , the test budget T , and the prior distribution of the state $\tilde{\theta}$, the researcher’s equilibrium test allocations in the baseline model, Framework A, and Framework B are identical.

The displayed flexibility allows us to study a rich set of economic problems. For instance, we use Framework A to study discrimination (Section 6.2) and media polarization (Online Appendix B.1). To illustrate an application of Framework B, consider a portfolio choice problem, where an investor makes decisions d_k on how much to invest in each asset $k = 1, \dots, K$. Here, $\tilde{\theta}_k$ represents the uncertain future return on asset k . The investor evaluates the future returns objectively ($\alpha_k^D = 1$ for all k), while an advisor may have an incentive to steer the investor towards some assets and away from others: $\alpha_k^R \neq 1$ for some k .

6 Applications


In this section, we illustrate how our framework can be used to study preference misalignment in two applications, each with a tailored parametrization of the misalignment. First, we adopt an organizational setting to showcase how biases across hierarchical layers impact learning in organizations. Next, we analyze a policy-making model with a discriminatory politician and explore how introducing independent researchers to inform policy decisions can mitigate negative effects of discrimination on welfare and inequality.

In the applications, we adopt the simplest specification, a two-attribute case, which is already sufficiently rich to demonstrate the tension between the players arising in the multi-attribute context. Throughout, we assume that both weights of the decision-maker are strictly positive, $\alpha_k^D > 0$ for $k = 1, 2$.

6.1 Diversity in organizations

In this application, we consider a stylized model of an organization, comprising a manager (decision-maker) and an analyst (researcher). This setup reflects an organizational

FIGURE 2. THE BIAS DECOMPOSITION OF AN ANALYST



Notes: The dashed line represents different α^R values where γ is fixed at a positive value and various values of β are considered.

hierarchy where the manager holds authority over the analyst and takes decisions based on the analyst's information. Given a diverse pool of analysts with varying weights, the manager selects one to provide information. The key question is which analyst the manager will choose.

First, it is immediate that the manager's ideal choice would be an analyst whose weights match hers. However, our focus is on identifying the type of analyst the manager would prefer when, due to an inherent hierarchical structure, perfect alignment is unachievable. To analyze this question, we begin by establishing a suitable parametrization of the misalignment between the players' weights.

We fix the manager's weight vector, $\alpha^D = (\alpha_1^D, \alpha_2^D)$, and define an orthogonal vector $\bar{\alpha}^D := (-\alpha_2^D, \alpha_1^D)$. Next, we express the analyst's weight vector α^R as a linear combination of the manager's weight vector and its orthogonal counterpart as


$$\alpha^R(\beta, \gamma) := \beta\alpha^D + \gamma\bar{\alpha}^D \quad (18)$$

for some $\beta \in \mathbb{R}_+$ and $\gamma \in \Gamma(\beta) := \left[-\beta\frac{\alpha_2^D}{\alpha_1^D}, \beta\frac{\alpha_1^D}{\alpha_2^D}\right]$.¹³ Here, β represents the absolute bias, referred to as *sensitivity* (to new information), and γ represents the relative bias, referred to as *distortion*. We say that *the analyst becomes more sensitive* when β increases.¹⁴

¹³The constraint $\gamma \in \Gamma(\beta)$ follows from the requirement that $\alpha_k^R \geq 0$ for both $k = 1, 2$. As noted in Footnote 8, this restriction is without loss and only for ease of exposition.

¹⁴Note that when $\beta = 1$, the analyst has the same sensitivity as the manager.

FIGURE 3. EQUILIBRIUM TEST ALLOCATION



Notes: Panel (a) depicts the undistorted analyst's equilibrium test allocation: no testing if $\beta \leq 1/2$, and manager's first best if $\beta > 1/2$. Panel (b) details the equilibrium test allocation for different (β, γ) pairs: the analyst learns (i) only about attribute $\tilde{\theta}_1$ if $\alpha^R \in L_1$; (ii) only about attribute $\tilde{\theta}_2$ if $\alpha^R \in L_2$; (iii) about both attributes if $\alpha^R \in L_{12}$; (iv) about no attribute if $\alpha^R \in L_0$. The black thick lines captures analysts with increasing distortion towards attribute $\tilde{\theta}_2$ from an undistorted analyst (dashed line) for fixed sensitivity at $\beta < 1/2$ (line closer to the origin) and $\beta > 1/2$ (line further away from the origin).

We further say *the analyst becomes more distorted* when $\gamma \geq 0$ increases or $\gamma \leq 0$ decreases. Specifically, when $\gamma > 0$ ($\gamma < 0$), the analyst exhibits bias towards attribute $\tilde{\theta}_2$ (attribute $\tilde{\theta}_1$). When $\gamma = 0$, we call the analyst *undistorted*. Figure 2 visualizes this bias decomposition.¹⁵ In the context of our organizational application, we posit that analysts who are separated by more layers of hierarchy from the manager are less sensitive.

First, we examine how the manager's payoff changes with the analyst's distortion γ . Proposition 2 below shows that the manager is worse off as the analyst becomes more distorted, provided the analyst is sufficiently sensitive (high β). However, the converse is true when the analyst is too insensitive (low β).

Proposition 2. Fix the analyst's level of sensitivity $\beta \in \mathbb{R}_+$.

- (i) If $\beta \leq 1/2$ (the analyst is too insensitive), then the manager's expected equilibrium

¹⁵Alternative decompositions of the analyst's weights in terms of "absolute" and "relative" bias exist. For instance, one could present both α^R and α^D in polar coordinates and let β be the difference in distances from the origin and γ be the difference in angles. The insights from our Propositions 2 and 3 would extend to such alternative representations.

payoff weakly *increases* as the analyst becomes more distorted.

- (ii) If $\beta > 1/2$ (the analyst is sufficiently sensitive), then the manager's expected equilibrium payoff weakly *decreases* as the analyst becomes more distorted.

The manager and an undistorted analyst share the same weight ratios, so the analyst's equilibrium learning strategy aligns with the manager's first-best solution, provided the analyst's sensitivity is high enough to motivate learning ($\beta > 1/2$). In this case, if the analyst's distortion increases, his equilibrium learning strategy deviates further from the manager's first best, making the manager worse off. Conversely, when $\beta \leq 1/2$, an undistorted analyst abstains from learning, perceiving the manager's reaction to any information as excessive. Then, if the analyst's distortion increases, his preferences align more closely with the manager's on one particular attribute, potentially prompting him to learn about it. While the added distortion does increase misalignment on the other attribute, it does not negatively impact learning since the initial misalignment was already sufficient to deter the analyst from learning about it. Overall, since any acquired information is preferable to none, the manager ultimately benefits.

The results, which are illustrated in Figure 3, highlight the importance of both diversity and uniformity within organizations. Our findings suggest a strategic approach to organizational structure if employees lower down the hierarchy are less engaged and thus less sensitive to new information than the managers higher up are. In organizations with many layers of hierarchy and thus potentially significant differences in engagement, it is beneficial to have substantial diversity at the lower levels, as indicated by relatively high degrees of distortion $|\gamma|$ compared to the managers. Conversely, in smaller, less hierarchical organizations, a more uniform workforce with minimal distortion is preferable as the variation in engagement levels among managers and analysts is less problematic in these settings.

The above underscores the nuanced role of distortion in shaping the analyst's alignment with the manager's objectives. However, when we shift focus from distortion to sensitivity, the effects become more straightforward.

Proposition 3. Fix the analyst's distortion $\gamma \in \mathbb{R}$. Then, the manager's expected equilibrium payoff is weakly increasing in the analyst's sensitivity β .

The manager always prefers a more sensitive analyst, as this leads to an equilibrium

test allocation that aligns more closely with her first-best solution. In other words, the manager seeks highly engaged employees who are sensitive to new information relevant to the organization. To prove this result, we show that for any distortion γ , the equilibrium test allocation τ^* increasingly aligns with the manager's first best as β increases, effectively neutralizing the impact of distortion.

6.2 Discrimination, welfare and inequality

In this application, we explore a scenario where a politician decides on a policy affecting two social groups and favors one group. We investigate how appointing an advisor with different preferences, who strategically curates the information provided to the politician, can mitigate the negative impact of this type of discrimination on utilitarian welfare and inequality.

Suppose there are two social groups $k = 1, 2$. Each group k has an unknown optimal policy $\tilde{\theta}_k \sim \mathcal{N}(0, 1)$, and so the two groups are a priori identical. A politician (decision-maker) decides on a common policy, $d \in \mathbb{R}$. The utility of group k is given by $u_k(d, \theta_k) = -(d - \theta_k)^2$. A budget $T = 1$ of tests is available to inquire about the optimal policies of both groups. An advisor (researcher) chooses a test allocation $\tau = (\tau_1, \tau_2) \in \mathcal{T}$. The learning process and notation are the same as in the baseline model.

The politician and the advisor are utilitarian with particular weights. The politician cares about both groups but favors group $k = 1$, her payoff being

$$u^D(d, \theta_1, \theta_2; \delta) = \underbrace{\frac{1 + \delta}{2}}_{\alpha_1^D(\delta):=} u_1(d, \theta_1) + \underbrace{\frac{1 - \delta}{2}}_{\alpha_2^D(\delta):=} u_2(d, \theta_2)$$

for a given level of discrimination $\delta \in (0, 1)$. On the other hand, the advisor (possibly) favors group $k = 2$, his payoff being

$$u^R(d, \theta_1, \theta_2; p) = \underbrace{\frac{1 - p}{2}}_{\alpha_1^R(p):=} u_1(d, \theta_1) + \underbrace{\frac{1 + p}{2}}_{\alpha_2^R(p):=} u_2(d, \theta_2)$$

for a given level of partiality $p \in [0, 1)$. When $p = 0$, indicating the advisor cares equally about both groups, we call the advisor *impartial*.

Our focus in this application is on *welfare* $W(p, \delta)$, defined as the sum of ex ante equilibrium expected payoffs of both groups (and is hence utilitarian with equal weights), and *inequality* $I(p, \delta)$, defined as their difference (and is hence an egalitarian measure):

$$W(p, \delta) := \mathbb{E} \left[u_1 \left(\tilde{d}^*(\tau^*), \tilde{\theta}_1 \right) \right] + \mathbb{E} \left[u_2 \left(\tilde{d}^*(\tau^*), \tilde{\theta}_2 \right) \right], \quad (19)$$

$$I(p, \delta) := \left| \mathbb{E} \left[u_1 \left(\tilde{d}^*(\tau^*), \tilde{\theta}_1 \right) \right] - \mathbb{E} \left[u_2 \left(\tilde{d}^*(\tau^*), \tilde{\theta}_2 \right) \right] \right|. \quad (20)$$

In the expressions above, the right-hand side depends on p and δ via the advisor's choice of test allocation τ^* and the politician's policy choice $\tilde{d}^*(\tau^*)$, which can be more fully described as $\tilde{d}^*(\tau^*(p, \delta), \delta)$. When $I(p, \delta) = 0$, we say there is *equality* in equilibrium.¹⁶

As a benchmark, we consider the case of *unchecked discrimination*, when the politician controls both learning and decision-making. Here, welfare and inequality are defined analogously to equations (19) and (20), with the politician selecting the optimal test allocation instead of an advisor doing so. In this context, increased discrimination negatively impacts both welfare and inequality: welfare declines while inequality rises. We aim to understand how appointing an advisor can mitigate these adverse effects. We analyze a scenario where welfare and inequality are influenced solely through the advisor's role in the learning stage, while the politician retains full control over decision-making.

First, we ask which advisor maximizes welfare. Since both welfare and the preferences of the impartial advisor are utilitarian with equal weights, it immediately follows that appointing an impartial advisor maximizes welfare, regardless of the politician's level of discrimination. In contrast, appointing a partial advisor or allowing unchecked discrimination yields lower welfare. The next lemma describes the (welfare-maximizing) learning strategy of the impartial advisor.

Lemma 3. For every level of discrimination δ , welfare is maximized by appointing the impartial advisor ($p = 0$). Moreover, the impartial advisor's equilibrium test allocation $\tau^*(0, \delta) = (1/2, 1/2)$ is *independent* of δ .

As the level of discrimination δ increases, two effects emerge. From the advisor's perspective, information about group 1 is increasingly overemphasized in the politician's

¹⁶This model fits Framework A in Section 5. By Proposition 1, Theorem 2 can be used to find the equilibrium test allocation.

decision, while information about group 2 is increasingly underused. Each effect independently *reduces* the advisor’s incentive to learn about the respective group. However, for the impartial advisor, these effects are equal in magnitude and cancel each other out. This leads him to be unresponsive to changes in δ and always choose an equal test allocation. Consequently, as δ increases, inequality increases under an impartial advisor.

The above raises the question of whether a partial advisor can counterbalance the politician’s discrimination and restore equality. If so, does the required level of partiality p increase with the level of discrimination δ ? The next lemma addresses these questions.

Lemma 4. For every politician’s level of discrimination δ , there exists a unique advisor’s level of partiality $\hat{p}(\delta) > 0$ that ensures equality in equilibrium: $I(\hat{p}(\delta), \delta) = 0$. The function $\hat{p}(\delta)$ is continuous and *non-monotone*: there exists a unique $\bar{\delta} \in (0, 1)$ such that $\hat{p}(\delta)$ is strictly increasing for $\delta < \bar{\delta}$ and strictly decreasing for $\delta > \bar{\delta}$.

As noted earlier, increases in δ reduce the advisor’s incentive to learn about both groups, as information about group 1 is increasingly overemphasized in the politician’s decision, while information about group 2 is increasingly underused. For an advisor with partiality $p > 0$, the former effect dominates the latter, prompting him to learn more about group 2 as δ increases. Moreover, the disparity between the two effects is greater for advisors with higher partiality p , as they prioritize group 2 over group 1 more. As such, advisors with higher partiality adjust their learning strategies more sharply in response to changes in δ than do those with lower partiality.

On the other hand, as the level of discrimination δ increases, restoring equality requires more wasted information: the advisor must learn less about group 1 and more about group 2, even though the politician becomes less responsive to information about group 2. Hence, the politician’s decision increasingly relies on her prior beliefs—where there is no disagreement between groups—rather than the advisor’s information. As the amount of “effectively” utilized information diminishes with higher δ , the adjustments required in the learning strategy to restore equality become progressively smaller.

In summary, two effects are at play: advisors with higher partiality p respond more strongly to changes in δ , while the degree of adjustment in tests to achieve equality diminishes with higher δ . These dynamics then lead to the non-monotonicity result in Lemma 4. Additionally, since restoring equality involves wasting more information as δ

increases, it results in a welfare loss. Indeed, welfare $W(\hat{p}(\delta), \delta)$ decreases with δ .

As we have seen, welfare is maximized with an impartial advisor, but then inequality increases with δ . Conversely, with partial advisors who restore equality, welfare then decreases with δ . Hence, welfare maximization and inequality minimization are misaligned objectives, as captured in the following proposition.

Proposition 4. For any level of discrimination δ , a welfare-inequality Pareto frontier is formed by advisors having partiality levels $p \in [0, \hat{p}(\delta)]$, where both welfare and inequality strictly decrease in p . Moreover, welfare with the equality-restoring advisor $\hat{p}(\delta)$ is strictly higher than under unchecked discrimination.

Proposition 4 illustrates a trade-off between welfare and inequality. Reducing inequality requires wasting more information, achieved by advisors with higher levels of partiality, while increasing welfare requires using information more efficiently, achieved by less partial advisors. This dynamic gives rise to the Pareto frontier formed by advisors ranging from impartial to those restoring equality. Moreover, any advisor along this frontier unambiguously improves both welfare and inequality compared to unchecked discrimination. Notably, even an advisor who restores equality results in higher welfare than does unchecked discrimination. This is because welfare is concave in test allocation, and the learning strategy of an unchecked politician is heavily skewed towards learning about the needs of group 1, whereas an “equalizing” advisor learns more evenly about both groups (even though his learning is skewed towards group 2).

7 Extensions

In this section, we outline two extensions, presented in more detail in Online Appendix B. The first examines media polarization by introducing two researchers into the model, representing media outlets competing to influence a voter (decision-maker). Each outlet and the voter seek implementation of their preferred policy mix on two policy issues. In the model, first, the voter allocates her attention between the two outlets. Second, the outlets simultaneously decide how much coverage to devote to the two policy issues. The combined attention and coverage provide the voter with signals about the two policy issues. Finally, the voter casts her ballot. We provide conditions under which media

outlets become fully polarized, with each outlet covering only a single issue.¹⁷ This polarization benefits the voter, however, as she can acquire her first-best information by splitting her attention appropriately between the two outlets. Hence, when media outlets care about the implemented policy mix, polarization occurs in equilibrium but helps, rather than harms, the voter. The voter thus prefers a polarized media duopoly to a monopoly of either media outlet (which is then more moderate in its coverage).

In the second extension, we introduce uncertainty about the decision-maker’s preferences within a dual-self model. A single agent controls both learning and decision-making, but her preferences might change between the two stages due to factors like job loss, illness, or self-control issues. We study how this uncertainty shapes the learning strategy and compare welfare outcomes between naïve agents (unaware of potential preference changes) and sophisticated agents (aware of potential changes) as in O’Donoghue and Rabin (1999). We identify conditions under which the sophisticate may opt for *strategic ignorance*, completely avoiding learning about an attribute to prevent excessive overreaction by the future self. Further, we delineate when the expected utilities of the naïf and the sophisticate differ and in which direction. Notably, if the welfare criterion is measured by the changed utility, the naïf can outperform the sophisticate. This happens when the preference change increases the weight on the affected attribute, making the sophisticate’s “hedging” strategy—effectively assigning less weight to that attribute due to possible overreaction by future self—counterproductive, regardless of whether preferences ultimately change.

8 Conclusion

We present a model of delegated learning in a multi-attribute context. We explore the strategic interaction between a decision-maker, who makes a decision influenced by the state of the world, and a researcher, who allocates a budget of tests to procure information about various attributes of the state. The players have different preferences for the final decision, as both seek to align the decision with the “weighted state” but assign

¹⁷Research shows that individuals tend to seek information that confirms their prior beliefs, interpreting ambiguous information in a way that aligns with these beliefs (see Fryer et al., 2019; Nimark and Sundaresan, 2019; Olszewski, 2021; D’Amico and Tabellini, 2022 for some examples). Additionally, it has been suggested that information providers, be it media or peers on social media, may contribute to polarization by strategically selecting the information they supply (Pogorelskiy and Shum, 2019; Germano et al., 2022; Aina, 2024).

different weights to the attributes. Our insights shed light on how the researcher's preference misalignment affects the equilibrium learning strategy and the players' resulting payoffs.

We characterize the equilibrium test allocation in this strategic interaction. We show it coincides with the solution to the single-player problem for certain attribute weights, which generally differ from both the decision-maker's and the researcher's weights. A key insight is that the researcher may forgo learning altogether if he anticipates an excessive overreaction of the decision-maker to new information. More broadly, preference misalignment leads to reduced emphasis on attributes that spark greater disagreement between the two players. We also show that our characterization applies beyond our baseline model and extends to other models and preference formulations.

The model provides a fruitful foundation for applied work, as demonstrated by the applications and extensions explored herein. Our framework is both rich and tractable, making it adaptable to a wide variety of settings. In addition to the problems considered in our applications and extensions, our model can be applied to many other problems. These include but are not limited to strategic management under uncertainty, biased portfolio advising, and optimal delegation (of learning and decision rights).

A Main Proofs

A.1 Proofs for Section 3: The single-player case

A.1.1 Proof of Lemma 1

Given the optimal decision strategy $d^*(s, \tau)$ given by equation (9), the agent chooses $\tau \in \mathbb{R}_+^K$ with $\sum_k \tau_k \leq T$ to maximize

$$\begin{aligned}
 V(\tau) &= \mathbb{E} \left[- \left(\tilde{d}^*(\tau) - \tilde{b} \right)^2 \right] = -\mathbb{E} \left[\left(\tilde{d}^*(\tau) \right)^2 \right] + 2\mathbb{E} \left[\tilde{d}^*(\tau) \tilde{b} \right] - \mathbb{E} \left[\tilde{b}^2 \right] \\
 &= - \left(\text{var} \left(\tilde{d}^*(\tau) \right) + \left(\mathbb{E} \left[\tilde{d}^*(\tau) \right] \right)^2 \right) + 2 \left(\text{cov} \left(\tilde{d}^*(\tau), \tilde{b} \right) + \mathbb{E} \left[\tilde{d}^*(\tau) \right] \mathbb{E} \left[\tilde{b} \right] \right) \\
 &\quad - \left(\text{var} \left(\tilde{b} \right) + \left(\mathbb{E} \left[\tilde{b} \right] \right)^2 \right) \\
 &= -(\psi(\tau) + (b_0)^2) + 2(\psi(\tau) + b_0 b_0) - (\sigma_0^2 + (b_0)^2) = \psi(\tau) - \sigma_0^2 = -\hat{\sigma}^2(\tau),
 \end{aligned}$$

where $\hat{\sigma}^2(\tau)$ is given by equation (8).¹⁸ □

A.1.2 Proof of Theorem 1

We prove a general characterization of the optimal test allocation for $K \geq 2$, from which Theorem 1 immediately follows.

Theorem 3. In the single-player case, without loss of generality, let $\alpha_1 \Sigma_1^0 \geq \dots \geq \alpha_K \Sigma_K^0$. Then, the unique optimal test allocation τ^* is as follows: if $\alpha_k = 0$ for all k , then $\tau^* = (0, \dots, 0)$; otherwise,¹⁹

$$\tau_k^* = (T - \bar{T}_J) \frac{\alpha_k}{\sum_{l=1}^J \alpha_l} + \sum_{j=k}^{J-1} \left((\bar{T}_{j+1} - \bar{T}_j) \frac{\alpha_k}{\sum_{l=1}^j \alpha_l} \right)$$

for all $k = 1, \dots, J$ and $\tau_k^* = 0$ for all $k = J+1, \dots, K$, where the budget thresholds $\bar{T}_1 = 0 \leq \bar{T}_2 \leq \dots \leq \bar{T}_{K+1}$ are given by

$$\bar{T}_k := \frac{\sum_{l=1}^{k-1} \alpha_l}{\alpha_k \Sigma_k^0} - \sum_{l=1}^{k-1} \frac{1}{\Sigma_l^0}$$

for all $k = 1, \dots, K$, $\bar{T}_{K+1} := +\infty$, and $J = 1, \dots, K+1$ is such that $T \in [\bar{T}_J, \bar{T}_{J+1})$.

Proof. By Lemma 1, the agent's objective function when choosing a test allocation $\tau \in \mathcal{T}$ is given by

$$V(\tau) = -\hat{\sigma}^2(\tau) = -\sum_{k=1}^K \left[\alpha_k^2 \left(\frac{1}{\Sigma_k^0} + \tau_k \right)^{-1} \right]. \quad (21)$$

If $\alpha = (0, \dots, 0)$, then $V(\tau) = 0$ for all τ . Hence, any test allocation $\tau \in \mathcal{T}$ is optimal. By the assumption that in case of indifference, the agent abstains from learning, we get the first part of the statement.

For the remainder of the proof, assume $\alpha_1 > 0$. It is immediate that (21) is concave in τ (and all constraints are convex in τ), hence we can use the Karush-Kuhn-Tucker

¹⁸In the second equality, we used $\text{var}(\tilde{X}) = \mathbb{E}[\tilde{X}^2] - (\mathbb{E}[\tilde{X}])^2$ and $\text{cov}(\tilde{X}, \tilde{Y}) = \mathbb{E}[\tilde{X}\tilde{Y}] - \mathbb{E}[\tilde{X}]\mathbb{E}[\tilde{Y}]$. In the third equality, we used the fact that \tilde{b} is a mean-preserving spread of \tilde{d} . Hence, we can write $\tilde{b} = \tilde{d} + \tilde{\nu}$ where $\tilde{\nu}$ is a normally distributed noise independent of \tilde{d} . Thus, $\text{cov}(\tilde{d}, \tilde{b}) = \text{cov}(\tilde{d}, \tilde{d} + \tilde{\nu}) = \text{cov}(\tilde{d}, \tilde{d}) = \text{var}(\tilde{d}) = \psi(\tau)$. The last equality follows from $\psi^i(\tau) = \sigma_0^{2,i} - \hat{\sigma}^{2,i}(\tau)$.

¹⁹We apply the convention $\sum_{j=k}^l x_j = 0$ when $l < k$.

technique to find the solution. Let ξ denote the Lagrange multiplier on the budget constraint and η_k denote the Lagrange multiplier on the non-negativity constraint for τ_k , with $\eta := (\eta_1, \dots, \eta_K)$. Then the KKT first order conditions can be written as follows:

$$\alpha_k^2 \left(\frac{1}{\Sigma_k^0} + \tau_k \right)^{-2} = \xi - \eta_k \quad \text{for all } k = 1, \dots, K, \quad (22)$$

$$0 = \eta_k \tau_k \quad \text{for all } k = 1, \dots, K, \quad (23)$$

$$\eta_k \geq 0, \tau_k \geq 0 \quad \text{for all } k = 1, \dots, K, \quad (24)$$

$$0 = \xi \left(T - \sum_{k=1}^K \tau_k \right), \quad (25)$$

$$T \geq \sum_{k=1}^K \tau_k, \xi \geq 0. \quad (26)$$

Observe that the marginal benefit of increasing τ_1 is $\alpha_1^2 \left(\frac{1}{\Sigma_1^0} + \tau_1 \right)^{-2} > 0$ since $\alpha_1 > 0$. Hence, for $k = 1$, the LHS of condition (22) is strictly positive, and so must be the RHS, implying $\xi > 0$ (since $\eta_1 \geq 0$). Therefore, from conditions (25) and (26), the budget constraint must be binding in the optimum: $T = \sum_{k=1}^K \tau_k$.

Consider an arbitrary $\xi > 0$. Then for every k there exists a unique pair $(\tau_k(\xi), \eta_k(\xi))$ satisfying conditions (22)-(24). This pair is given by

$$\tau_k(\xi) = \begin{cases} \frac{\alpha_k}{\sqrt{\xi}} - \frac{1}{\Sigma_k^0} & \text{if } \xi \leq (\alpha_k \Sigma_k^0)^2 \\ 0 & \text{otherwise} \end{cases}, \quad \eta_k(\xi) = \begin{cases} 0 & \text{if } \xi \leq (\alpha_k \Sigma_k^0)^2 \\ \xi - (\alpha_k \Sigma_k^0)^2 & \text{otherwise} \end{cases} \quad (27)$$

To find a solution, next we need to find value $\xi^* > 0$ under which also the budget constraint $T = \sum_k \tau_k(\xi^*)$ is satisfied. Note that $\tau_k(\xi)$ are continuous and weakly decreasing in ξ for all k , and strictly decreasing whenever $\tau_k(\xi) > 0$.

Further, for all k we have $\lim_{\xi \rightarrow +\infty} \tau_k(\xi) = 0$ and $\lim_{\xi \rightarrow 0} \tau_k(\xi) = +\infty$. The above imply that function $B(\xi) := \sum_{k=1}^K \tau_k(\xi)$ is continuous, strictly decreasing in ξ whenever $B(\xi) > 0$ (i.e. whenever there exist k with $\tau_k(\xi) > 0$), and has the same limit properties. By the intermediate value theorem, there exists unique $\xi^* > 0$ such that $B(\xi^*) = T > 0$ (the budget constraint binds). This value ξ^* together with the corresponding pairs $(\tau_k(\xi^*), \eta_k(\xi^*))$ for every k then solve the entire FOC system (22)–(26), and thus uniquely

solve the player's problem from Lemma 1.

It remains to show that the solution $\tau^* := (\tau_1(\xi^*), \dots, \tau_K(\xi^*))$ satisfies the representation of the statement. Let us use the substitution $\zeta := \frac{1}{\sqrt{\xi}}$, and let us redefine the functions $\tau_k(\xi)$ and $B(\xi)$ to be functions of ζ . Hence, let $\hat{\tau}_k(\zeta) = \max \left\{ 0, \alpha_k \zeta - \frac{1}{\Sigma_k^0} \right\}$ and $\hat{B}(\zeta) = \sum_{k=1}^K \hat{\tau}_k(\zeta)$. Like before, let $\zeta^* = \frac{1}{\sqrt{\xi^*}}$ be the unique solution to $\hat{B}(\zeta^*) = T$. Let further $\zeta_k := \frac{1}{\alpha_k \Sigma_k^0}$ for all k , and note that $\hat{\tau}_k(\zeta) > 0$ if and only if $\zeta > \zeta_k$, as well as the fact that by assumption, $\zeta_1 \leq \dots \leq \zeta_K$.

To calculate τ^* , we can use the fundamental theorem of calculus:

$$\tau_k^* = \hat{\tau}_k(\zeta^*) = \int_{\min\{\zeta_k, \zeta^*\}}^{\zeta^*} \frac{d\hat{\tau}_k(\zeta)}{d\zeta} d\zeta = \int_{\min\{\zeta_k, \zeta^*\}}^{\zeta^*} \alpha_k d\zeta = \alpha_k (\zeta^* - \min\{\zeta_k, \zeta^*\}).$$

Let $J = 1, \dots, K + 1$ be such that $\zeta^* \in [\zeta_J, \zeta_{J+1})$, with $\zeta_{K+1} := +\infty$.²⁰ Then for all $k = J + 1, \dots, K$ we have $\tau_k^* = 0$, and for all $k = 1, \dots, J$ we have:

$$\begin{aligned} \zeta^* - \zeta_k &= (\zeta^* - \zeta_J) + (\zeta_J - \zeta_{J-1}) + \dots + (\zeta_{k+1} - \zeta_k) \\ \Rightarrow \tau_k^* &= (\zeta^* - \zeta_J) \alpha_k + \sum_{j=k}^{J-1} (\zeta_{j+1} - \zeta_j) \alpha_k. \end{aligned} \quad (28)$$

To find $(\zeta^* - \zeta_J)$ and $(\zeta_{j+1} - \zeta_j)$ in this case, we again invoke the fundamental theorem of calculus:

$$\begin{aligned} T = \hat{B}(\zeta^*) &= \int_{\zeta_1}^{\zeta^*} \frac{d\hat{B}(\zeta)}{d\zeta} d\zeta = \int_{\zeta_1}^{\zeta^*} \sum_{k=1}^K \frac{d\hat{\tau}_k(\zeta)}{d\zeta} d\zeta \\ &= \int_{\zeta_1}^{\zeta_2} \alpha_1 d\zeta + \int_{\zeta_2}^{\zeta_3} (\alpha_1 + \alpha_2) d\zeta + \dots + \int_{\zeta_J}^{\zeta^*} \left(\sum_{l=1}^J \alpha_l \right) d\zeta = \sum_{l=1}^J (\zeta^* - \zeta_l) \alpha_l \end{aligned} \quad (29)$$

Specifically, then, if we let

$$\bar{T}_k := \hat{B}(\zeta_k) = \sum_{l=1}^{k-1} (\zeta_k - \zeta_l) \alpha_l = \frac{\sum_{l=1}^{k-1} \alpha_l}{\alpha_k \Sigma_k^0} - \sum_{l=1}^{k-1} \frac{1}{\Sigma_l^0},$$

²⁰Note that it must hold that $\zeta^* > \zeta_1$, which implies that there exists $J = 1, \dots, K$ with $\zeta^* \in [\zeta_J, \zeta_{J+1})$. Otherwise, if $\zeta^* \leq \zeta_1$, it would imply that $\hat{\tau}_k(\zeta^*) = 0$ for all k , hence $\hat{B}(\zeta^*) = 0$, which is a contradiction to $\hat{B}(\zeta^*) = T > 0$.

expression (29) implies that $T - \bar{T}_J = (\zeta^* - \zeta_J) \sum_{l=1}^J \alpha_l \iff \zeta^* - \zeta_J = \frac{T - \bar{T}_J}{\sum_{l=1}^J \alpha_l}$ and, similarly, $\bar{T}_{k+1} - \bar{T}_k = (\zeta_{k+1} - \zeta_k) \sum_{l=1}^k \alpha_l \iff \zeta_{k+1} - \zeta_k = \frac{\bar{T}_{k+1} - \bar{T}_k}{\sum_{l=1}^k \alpha_l}$ for any $k = 1, \dots, K-1$. Plugging these into (28) yields the solution presented in the statement. \square

A.2 Proofs for Section 4: Strategic players

A.2.1 Proof of Lemma 2

By following the same steps as in the proof of Theorem 4.1 in Bardhi (2024), given the test allocation τ and the decision-maker's equilibrium decision strategy $d^*(s, \tau) = \mathbb{E} \left[\tilde{b}^D(s, \tau) \right]$, the researcher's expected payoff is

$$V^R(\tau) = \mathbb{E} \left[- \left(\tilde{d}^*(\tau) - \tilde{b}^R \right)^2 \right] = 2 \operatorname{cov} \left(\tilde{d}^*(\tau), \tilde{b}^R \right) - \psi^D(\tau) + V^R(\emptyset)$$

where $V^R(\emptyset) = - (b_0^D - b_0^R)^2 - \sigma_0^{2,R}$ is the researcher's expected utility under no learning. To complete the proof, let us show that

$$\operatorname{cov} \left(\tilde{d}^*(\tau), \tilde{b}^R \right) = \sum_k \alpha_k^D \alpha_k^R \Sigma_k^0 \hat{\Sigma}_k(\tau_k) \tau_k.$$

The researcher's bliss point satisfies $\tilde{b}^R = \sum_k \alpha_k^R \tilde{\theta}_k$. The random variable capturing the equilibrium decision-maker's decision satisfies

$$\begin{aligned} \tilde{d}^*(\tau) &= \sum_k \alpha_k^D \hat{\mu}_k(\tilde{s}_k, \tau_k) = \sum_k \alpha_k^D \left(\frac{1}{1 + \tau_k \Sigma_k^0} \mu_k^0 + \frac{\tau_k \Sigma_k^0}{1 + \tau_k \Sigma_k^0} \tilde{s}_k \right) \\ &= \sum_k \alpha_k^D \left(\frac{1}{1 + \tau_k \Sigma_k^0} \mu_k^0 + \frac{\tau_k \Sigma_k^0}{1 + \tau_k \Sigma_k^0} (\tilde{\theta}_k + \tilde{\varepsilon}_k) \right) \end{aligned}$$

Hence,

$$\begin{aligned} \operatorname{cov}(\tilde{d}^*(\tau), \tilde{b}^R) &= \operatorname{cov} \left(\sum_k \alpha_k^D \frac{\tau_k \Sigma_k^0}{1 + \tau_k \Sigma_k^0} \tilde{\theta}_k, \sum_k \alpha_k^R \tilde{\theta}_k \right) = \sum_k \alpha_k^D \alpha_k^R \tau_k \frac{\Sigma_k^0}{1 + \tau_k \Sigma_k^0} \operatorname{var}(\tilde{\theta}_k) \\ &= \sum_k \alpha_k^D \alpha_k^R \tau_k \hat{\Sigma}_k(\tau_k) \Sigma_k^0, \end{aligned}$$

where we used (i) $\operatorname{cov}(\tilde{\theta}_k, \tilde{\varepsilon}_j) = 0$ for all k and j , since the noise terms are independent of all other random variables; (ii) $\operatorname{cov}(\tilde{\theta}_k, \mu_j^0) = 0$ for all k and j since μ_j^0 is a constant;

and (iii) $\text{cov}(\tilde{\theta}_k, \tilde{\theta}_j) = 0$ for all $k \neq j$ since the attributes are independent. \square

A.2.2 Proof of Theorem 2

Recall that $\psi^D(\tau) = \sigma_0^{2,D} - \hat{\sigma}^{2,D}(\tau)$ is the variance of the decision-maker's decision, where $\hat{\sigma}^{2,D}(\tau)$ is given by (8). Then, from Lemma 2 and dropping terms that are constant in τ , the researcher's value is given by

$$\begin{aligned} V^{st}(\tau) &= 2 \text{cov}(\tilde{d}^*(\tau), \tilde{b}^R) + \tilde{\sigma}^{2,D}(\tau) = 2 \sum_k \alpha_k^D \alpha_k^R \Sigma_k^0 \tau_k \hat{\Sigma}_k(\tau_k) + \sum_k (\alpha_k^D)^2 \hat{\Sigma}_k(\tau_k) \\ &= \sum_k \left(2\alpha_k^D \alpha_k^R \Sigma_k^0 \tau_k + (\alpha_k^D)^2 \right) \frac{\Sigma_k^0}{1 + \tau_k \Sigma_k^0} \end{aligned}$$

subject to the non-negativity and the budget constraints. The partial derivatives of $V^{st}(\tau)$ with respect to $\tau_k \geq 0$ are

$$\frac{\partial V^{st}(\tau)}{\partial \tau_k} = \lambda_k \left(\frac{\Sigma_k^0}{1 + \tau_k \Sigma_k^0} \right)^2 = \lambda_k \left(\frac{1}{\Sigma_k^0} + \tau_k \right)^{-2} \quad (30)$$

where we denote $\lambda_k := \alpha_k^D (2\alpha_k^R - \alpha_k^D)$ for every k .

Note that when $\lambda_k \leq 0$, the function $V^{st}(\tau)$ is (weakly) decreasing in τ_k . Then the solution dictates $\tau_k^* = 0$ for such k .²¹ Hence, the maximizers of function $V^{st}(\tau)$ are the same as the maximizers of an adjusted function $V^{st,adj}(\tau) := V^{st}(\tau)$ where we set $\alpha_k^D = \alpha_k^R = 0$ for all k with $\lambda_k \leq 0$ (and hence $V^{st,adj}(\tau)$ is independent of τ_k for such k).

For every k , denote $\hat{\alpha}_k = \sqrt{\max\{0, \lambda_k\}}$ and, w.l.o.g., let us relabel the attributes such that $\hat{\alpha}_1 \Sigma_1^0 \geq \dots \geq \hat{\alpha}_K \Sigma_K^0$. By noting that (i) the function $V^{st,adj}(\tau)$ is concave, (ii) for every k , the partial derivative $\frac{\partial V^{st,adj}(\tau)}{\partial \tau_k}$ is the same as the partial derivative $\frac{\partial V(\tau)}{\partial \tau_k}$ of a single-player objective function (21) in which we set $\alpha_k = \hat{\alpha}_k$ for all k , and (iii) the constraints are the same for both (the adjusted strategic and the single-player) maximization problems, we obtain the statement. \square

²¹When $\lambda_k \leq 0$ for all k and there exists at least one j such that $\lambda_j = 0$, we use the assumption that in case of indifference the researcher abstains from learning, thus yielding $\tau^* = (0, 0, \dots, 0)$. Otherwise, setting $\tau_k^* = 0$ whenever $\lambda_k \leq 0$ is uniquely optimal: (i) either there exists at least one j with $\lambda_j > 0$ (and thus the function $V^{st}(\tau)$ is strictly increasing in τ_j), or (ii) $\lambda_k < 0$ for all k (and thus the function $V^{st}(\tau)$ is *strictly* decreasing in every τ_k).

A.3 Proofs for Section 6.1: Diversity in organizations

We first state four lemmas (proved in Online Appendix C.2) that we use to prove both Propositions 2 and 3. Lemma 5 and Lemma 6 characterize the equilibrium learning strategy $\tau^*(\beta, \gamma)$ as a function of γ when $\beta \leq 1/2$ and $\beta > 1/2$, respectively. Lemma 7 establishes that the decision-maker's interim expected payoff is strictly increasing and concave in (τ_1, τ_2) . Lemma 8 derives the researcher's equilibrium test allocation for $\gamma = 0$.

Lemma 5. Let $\alpha_1^D \geq \alpha_2^D > 0$ and assume the analyst's weight vector α^R is given by decomposition (18). For any $\beta \in [0, 1/2]$ there exist unique $\gamma_1(\beta), \gamma_2(\beta) \in \Gamma(\beta)$ such that $\gamma_1(\beta) \leq 0 \leq \gamma_2(\beta)$ and:

$$\tau^*(\beta, \gamma) = \begin{cases} (0, T) & \text{if } \gamma > \gamma_2(\beta), \\ (0, 0) & \text{if } \gamma \in [\gamma_1(\beta), \gamma_2(\beta)], \\ (T, 0) & \text{if } \gamma < \gamma_1(\beta). \end{cases} \quad (31)$$

Further, there exist $\beta_1, \beta_2 \in \mathbb{R}$ with $0 < \beta_2 \leq \beta_1 < 1/2$ such that $\gamma_1(\beta) \in \text{int}(\Gamma(\beta))$ if and only if $\beta \in (\beta_1, 1/2]$, and $\gamma_2(\beta) \in \text{int}(\Gamma(\beta))$ if and only if $\beta \in (\beta_2, 1/2]$.

Lemma 6. Fix the manager's weight vector $\alpha^D = (\alpha_1^D, \alpha_2^D)$ with $\alpha_k^D > 0$ for both $k = 1, 2$. Let the analyst's weight vector $\alpha^R(\beta, \gamma)$ be given by the decomposition (18). Fix $\beta > 1/2$. Then there exist $\gamma_I(\beta), \gamma_{II}(\beta) \in \text{int}(\Gamma(\beta))$ with $\gamma_I(\beta) < \gamma_{II}(\beta)$ such that the analyst's equilibrium test allocation is

$$\tau^*(\beta, \gamma) = \begin{cases} (T, 0) & \text{if } \gamma \leq \gamma_I(\beta) \\ (\tilde{\tau}_1(\beta, \gamma), \tilde{\tau}_2(\beta, \gamma)) & \text{if } \gamma \in (\gamma_I(\beta), \gamma_{II}(\beta)) \\ (0, T) & \text{if } \gamma \geq \gamma_{II}(\beta) \end{cases} \quad (32)$$

where

$$\begin{aligned} \tilde{\tau}_1(\beta, \gamma) &:= \frac{\hat{\alpha}_1(\beta, \gamma)\Sigma_1^0 - \hat{\alpha}_2(\beta, \gamma)\Sigma_2^0}{\Sigma_1^0\Sigma_2^0(\hat{\alpha}_1(\beta, \gamma) + \hat{\alpha}_2(\beta, \gamma))} + \frac{\hat{\alpha}_1(\beta, \gamma)}{\hat{\alpha}_1(\beta, \gamma) + \hat{\alpha}_2(\beta, \gamma)}T \\ \tilde{\tau}_2(\beta, \gamma) &:= T - \tilde{\tau}_1(\beta, \gamma) \end{aligned} \quad (33)$$

and

$$\begin{aligned}\hat{\alpha}_1(\beta, \gamma) &:= \sqrt{\max\left\{(2\beta - 1)(\alpha_1^D)^2 - 2\gamma\alpha_1^D\alpha_2^D, 0\right\}} \\ \hat{\alpha}_2(\beta, \gamma) &:= \sqrt{\max\left\{(2\beta - 1)(\alpha_2^D)^2 + 2\gamma\alpha_1^D\alpha_2^D, 0\right\}}\end{aligned}\tag{34}$$

Moreover, $\tilde{\tau}_1(\beta, \gamma)$ ($\tilde{\tau}_2(\beta, \gamma)$) is strictly decreasing (strictly increasing) in γ for $\gamma \in (\gamma_I(\beta), \gamma_{II}(\beta))$.

Lemma 7. The decision-maker's interim expected equilibrium payoff is a strictly increasing and strictly concave function of (τ_1, τ_2) .

Lemma 8. Fix the test budget $T > 0$ and the decision-maker's weight vector $\alpha^D \in \mathbb{R}_{++}^2$. Suppose $\gamma = 0$ (the agent is not distorted).

- (i) If $\beta \leq 1/2$ (the agent is too insensitive), then the agent optimally chooses no testing: $\tau^*(\beta, 0) = (0, 0)$.
- (ii) If $\beta > 1/2$ (the agent is sufficiently sensitive), then the agent optimally chooses the decision-maker's most preferred test allocation: $\tau^*(\beta, 0) = \tau^*(1, 0)$.

A.3.1 Proof of Proposition 2

Suppose the premise of Proposition 2 holds. Proposition 2(i) is a direct implication of Lemma 5. The manager's interim expected payoff is strictly increasing in τ_1 and τ_2 by Lemma 7. Hence, the manager always strictly prefers any positive test allocation $\tau \neq (0, 0)$ to abstaining from learning.

Now take $\beta > 1/2$. Lemma 8 shows that (given the constraint $\tau \in \mathcal{T}$) the manager's interim expected payoff $V^D(\tau)$ is maximized when the agent's relative bias is $\gamma = 0$. Furthermore, since $V^D(\tau)$ is strictly increasing in τ_1 and τ_2 , the constraint $\tau_1 + \tau_2 \leq T$ must be binding at the manager's most preferred test allocation. Furthermore, Lemma 6 shows when $\beta > 1/2$, the agent's equilibrium test allocation $\tau^*(\beta, \gamma)$ is such that $\tau_1^*(\beta, \gamma)$ and $\tau_2^*(\beta, \gamma)$ are weakly decreasing and weakly increasing in γ , respectively, and $\tau_1^*(\beta, \gamma) + \tau_2^*(\beta, \gamma) = T$. The claim of Proposition 2(ii) then follows from the strict concavity of $V^D(\tau)$ (shown in Lemma 7). \square

A.3.2 Proof of Proposition 3

Suppose $\gamma = 0$. Then, Lemma 8 shows that the researcher's learning strategy is to not learn for $\beta \leq 1/2$ and then to implement the decision-maker's preferred learning strategy. As the decision-maker's interim expected payoff $V^D(\tau)$ is strictly increasing and concave in τ_1 and τ_2 (Lemma 7), this implies that $V^D(\tau)$ is weakly increasing in β for $\gamma = 0$.

Suppose from this point onwards that $\gamma > 0$ (case $\gamma < 0$ is completely analogous). Then, for the researcher's weights to be weakly positive, we must have $\beta \geq \underline{\beta} := \gamma \frac{\alpha_2^D}{\alpha_1^D}$. The equilibrium test allocation is given by Theorems 1 and 2, where we can express $\hat{\alpha}_k$ for $k = 1, 2$ as (34),²² and so

$$\begin{aligned}\lambda_1 &= (2\beta - 1)(\alpha_1^D)^2 - 2\gamma\alpha_1^D\alpha_2^D, \\ \lambda_2 &= (2\beta - 1)(\alpha_2^D)^2 + 2\gamma\alpha_1^D\alpha_2^D.\end{aligned}$$

Note that λ_k for $k = 1, 2$ is strictly increasing in β . Therefore, there exist β_1 and β_2 such that $\hat{\alpha}_k = 0$ for $\beta \leq \beta_k$, and $\hat{\alpha}_k$ is strictly positive and strictly increasing for $\beta > \beta_k$. These values are given by the respective roots of $\lambda_k = 0$:²³

$$\beta_1 := \frac{1}{2} + \gamma \frac{\alpha_2^D}{\alpha_1^D}, \quad \beta_2 := \frac{1}{2} - \gamma \frac{\alpha_1^D}{\alpha_2^D}.$$

Observe that $\beta_2 < 1/2 < \beta_1$ and $\underline{\beta} < \beta_1$. Theorems 1 and 2 then imply that if $\beta > \beta_2$, then $\exists k : \tau_k^*(\beta, \gamma) > 0$ and $\tau_1^*(\beta, \gamma) + \tau_2^*(\beta, \gamma) = T$ (since at least one of the auxiliary weights $\hat{\alpha}_k$ is then strictly positive). Therefore, we can conclude that if $\underline{\beta} > \beta_2$, then $\tau^*(\beta, \gamma) = (0, T)$ for all $\beta \in [\underline{\beta}, \beta_1]$, and if $\underline{\beta} < \beta_2$, then

$$\tau^*(\beta, \gamma) = \begin{cases} (0, 0) & \text{for } \beta \in [\underline{\beta}, \beta_2], \\ (0, T) & \text{for } \beta \in (\beta_2, \beta_1]. \end{cases}$$

As $V^D(\tau)$ is strictly increasing in τ_2^* , it follows that $V^D(\tau)$ is weakly increasing in β in either of the two cases.

Consider now the case when $\beta > \beta_1$, so we have $\hat{\alpha}_1, \hat{\alpha}_2 > 0$. Theorems 1 and 2 imply

²²While Lemma 6 is only stated for $\beta > 1/2$, expression (34) is well-defined for all $\beta > \underline{\beta}$.

²³Note that these β_k are different from the ones defined in the proof of Lemma 5.

that the equilibrium test allocation is given by

$$\tau_1^* = \max \left\{ 0, \min \left\{ \frac{\hat{\alpha}_1 \Sigma_1^0 - \Sigma_2^0}{\Sigma_1^0 \Sigma_2^0 (\hat{\alpha}_2 + 1)} + \frac{\hat{\alpha}_1}{\hat{\alpha}_2 + 1} T, T \right\} \right\}, \quad (35)$$

$$\tau_2^* = T - \tau_1^*.$$

Note that the expression above (and, hence, $V^D(\tau)$) only depends on β through $\frac{\hat{\alpha}_1}{\hat{\alpha}_2}$. Observe that $\lim_{\beta \rightarrow +\infty} \frac{\hat{\alpha}_1}{\hat{\alpha}_2} = \frac{\alpha_1^D}{\alpha_2^D}$, so that $\lim_{\beta \rightarrow +\infty} \tau^*(\gamma, \beta) = \tau^*(\alpha^D)$, so the equilibrium test allocation converges to the decision-maker's optimal (payoff-maximizing) test allocation as $\beta \rightarrow \infty$. Further, $\frac{\hat{\alpha}_1}{\hat{\alpha}_2}$ is strictly increasing in β in the case considered:

$$\frac{\partial}{\partial \beta} \frac{\hat{\alpha}_1}{\hat{\alpha}_2} = \frac{1}{\hat{\alpha}_2} \frac{\partial \hat{\alpha}_1}{\partial \beta} - \frac{\hat{\alpha}_1}{\hat{\alpha}_2^2} \frac{\partial \hat{\alpha}_2}{\partial \beta} = \frac{\hat{\alpha}_1}{\hat{\alpha}_2} \left(\left(\frac{\alpha_1^D}{\hat{\alpha}_1} \right)^2 - \left(\frac{\alpha_2^D}{\hat{\alpha}_2} \right)^2 \right) > 0,$$

where the inequality follows from $\hat{\alpha}_k = \sqrt{\lambda_k}$ and

$$\left(\frac{\alpha_1^D}{\alpha_2^D} \right)^2 > \left(\frac{\hat{\alpha}_1}{\hat{\alpha}_2} \right)^2 = \frac{(2\beta - 1)(\alpha_1^D)^2 - 2\gamma \alpha_1^D \alpha_2^D}{(2\beta - 1)(\alpha_2^D)^2 + 2\gamma \alpha_1^D \alpha_2^D} = \frac{(\alpha_1^D)^2 - \frac{2\gamma}{2\beta - 1} \alpha_1^D \alpha_2^D}{(\alpha_2^D)^2 + \frac{2\gamma}{2\beta - 1} \alpha_1^D \alpha_2^D}.$$

Together with the limit result above, this implies that $\frac{\hat{\alpha}_1}{\hat{\alpha}_2}$ and $\tau^*(\gamma, \beta)$ *monotonically* converge to $\frac{\alpha_1^D}{\alpha_2^D}$ and $\tau^*(\alpha^D)$, respectively, as $\beta \rightarrow \infty$. As $V^D(\tau)$ is concave in τ (see Lemma 7) and maximized by $\tau^*(\alpha^D)$, it must be weakly increasing in $\beta \in (\beta_1, +\infty)$. \square

A.4 Proofs for Section 6.2: Discrimination, welfare and inequality

A.4.1 Proof of Lemma 3

For the first part of the statement, observe that the maximization problem of an impartial advisor is

$$\max_{\tau \in \mathcal{T}} \frac{1}{2} \mathbb{E} \left[u_1 \left(\tilde{d}^*(\tau), \tilde{\theta}_1 \right) + u_2 \left(\tilde{d}^*(\tau), \tilde{\theta}_2 \right) \right].$$

By the definition of welfare, an impartial advisor thus chooses the welfare-maximizing test allocation in equilibrium, since the objectives are scaled versions of each other.

Let us turn to the second part of lemma. Fix $\delta \in (0, 1)$. Using Theorem 2, we can solve

the discrimination model as a solution to a single-agent model with weights

$$\hat{\alpha}_1(p, \delta) := \begin{cases} \frac{1}{2}\sqrt{(1+\delta)(1-2p-\delta)} & \text{if } p < \frac{1-\delta}{2} \\ 0 & \text{otherwise} \end{cases} \quad (36)$$

$$\hat{\alpha}_2(p, \delta) := \begin{cases} \frac{1}{2}\sqrt{(1-\delta)(1+2p+\delta)} & \text{if } p > -\frac{1+\delta}{2} \\ 0 & \text{otherwise} \end{cases} \quad (37)$$

Let $p = 0$. Then $\hat{\alpha}_1(0, \delta) = \hat{\alpha}_2(0, \delta) = \frac{1}{2}\sqrt{1-\delta^2}$. The equilibrium test allocation depends on the weights only through their ratios $\frac{\hat{\alpha}_1(0, \delta)}{\hat{\alpha}_1(0, \delta) + \hat{\alpha}_2(0, \delta)} = \frac{1}{2}$ and $\frac{\hat{\alpha}_2(0, \delta)}{\hat{\alpha}_1(0, \delta) + \hat{\alpha}_2(0, \delta)} = \frac{1}{2}$. These ratios are independent of the level of discrimination δ , which completes the proof. \square

A.4.2 Proof of Lemma 4

We first define several new notions. Given the politician's discrimination $\delta \in (0, 1)$ and the advisor's partiality $p \geq 0$, let $\tau^*(p, \delta) = (\tau_1^*(p, \delta), \tau_2^*(p, \delta))$ denote the respective equilibrium test allocation. Note it holds $\tau_1^*(p, \delta) + \tau_2^*(p, \delta) = 1$, i.e., the budget is always exhausted in equilibrium.²⁴

Given a particular test allocation τ and equilibrium decision strategy, let $V^k(\tau, \delta)$ denote the ex ante expected utility of group k , i.e.,

$$V^k(\tau, \delta) = \mathbb{E} \left[-(\tilde{d}^{*D}(\tau; \delta) - \tilde{\theta}_k)^2 \right].$$

Under the additional constraint that the budget is exhausted, and so the test allocation satisfies $\tau_1 = 1 - \tau_2$, let us define the difference between the expected utilities of group 1 from group 2 as

$$\Delta(\tau_2, \delta) := V^1((1 - \tau_2, \tau_2), \delta) - V^2((1 - \tau_2, \tau_2), \delta). \quad (38)$$

Note that inequality is then given by $I(p, \delta) = |\Delta(\tau_2^*(p, \delta), \delta)|$.

²⁴This result follows from the assumption that the sums of the weights of each player are equal: $\sum_k \alpha_k^R(p) = \sum_k \alpha_k^D(\delta)$. Thus, it cannot simultaneously hold $\alpha_k^R(p) \leq \alpha_k^D(\delta)$ for both k (which is necessary and sufficient for $\hat{\alpha}_1(p, \delta) = \hat{\alpha}_2(p, \delta) = 0$), and there is thus always learning in equilibrium.

We now state two intermediary results (proved in Online Appendix C.3), used to prove Lemma 4.

Lemma 9. For each $\delta \in (0, 1)$, there exists a threshold partiality level $p^{aux}(\delta) \in (0, \frac{1-\delta}{2})$ such that the equilibrium learning strategy satisfies $\tau_2^*(p, \delta) = 1$ for advisors with $p \geq p^{aux}(\delta)$ and

$$\tau_2^*(p, \delta) = \tau_2^{aux}(p, \delta) := \frac{2\hat{\alpha}_2(p, \delta) - \hat{\alpha}_1(p, \delta)}{\hat{\alpha}_1(p, \delta) + \hat{\alpha}_2(p, \delta)} \quad \text{for advisors with } p \in [0, p^{aux}(\delta)],$$

where $\hat{\alpha}_1(p, \delta)$ and $\hat{\alpha}_2(p, \delta)$ are given by equations (36) and (37).

Lemma 10. For each δ there exists a unique value of $\hat{\tau}_2(\delta) \in (1/2, 1)$ for which it holds $\Delta(\hat{\tau}_2(\delta); \delta) = 0$.

Given the politician's level of discrimination δ and her equilibrium decision strategy, the value $\hat{\tau}_2(\delta)$ is the amount of tests allocated to learn about group 2 that guarantees that the expected utilities of both social groups are the same. From the proof of Lemma 10, we obtain

$$\Delta(\tau_2, \delta) = \frac{(1+\delta)(1-\tau_2)}{1+1-\tau_2} - \frac{(1-\delta)\tau_2}{1+\tau_2}.$$

Solving $\Delta(\hat{\tau}_2(\delta), \delta) = 0$ for $\hat{\tau}_2(\delta) \geq 0$, we obtain

$$\hat{\tau}_2(\delta) = \frac{-(1-\delta) + \sqrt{1+3\delta^2}}{2\delta}. \quad (39)$$

From Lemma 9, the equilibrium test allocation satisfies: $\tau_2^*(p, \delta) = 1/2$ at $p = 0$, $\tau_2^*(p, \delta) = 1$ at $p \geq p^{aux}(\delta)$, $\tau_2^*(p, \delta)$ is continuous on $p \in [0, 1]$ and it is strictly increasing in $p \in [0, p^{aux}(\delta)]$. Since $\hat{\tau}_2(\delta) \in (1/2, 1)$, we thus have that for any level of discrimination $\delta \in (0, 1)$ there exists a unique partiality level $\hat{p}(\delta) \in (0, p^{aux}(\delta))$ such that $\hat{\tau}_2(\delta) = \tau_2^*(\hat{p}(\delta), \delta)$. Solving this equation, we obtain

$$\hat{p}(\delta) = \delta \left(\frac{1}{\sqrt{1+3\delta^2}} - \frac{1}{2} \right),$$

which is continuous on $\delta \in (0, 1)$. Taking the first derivative, we get

$$\hat{p}'(\delta) = \frac{1}{(1 + 3\delta^2)^{3/2}} - \frac{1}{2} \begin{cases} > 0 & \text{if } \delta < \bar{\delta} \\ = 0 & \text{if } \delta = \bar{\delta} \\ < 0 & \text{if } \delta > \bar{\delta} \end{cases}$$

where $\bar{\delta} := \frac{1}{3} (1 - 2^{1/2} + 2^{2/3})$. □

A.4.3 Proof of Proposition 4

The first part of the statement follows directly from the following Lemma 11, which is proved in Online Appendix C.3.

Lemma 11. Welfare $W(p, \delta)$ and inequality $I(p, \delta)$ are both (i) continuous in p ; and (ii) strictly decreasing on $p \in (0, \hat{p}(\delta))$, where $\hat{p}(\delta)$ restores equality. Furthermore, welfare strictly decreases and inequality strictly increases on $p \in (\hat{p}(\delta), p^{aux}(\delta))$, where $p^{aux}(\delta)$ is characterized in Lemma 9; and they are both constant on $p \geq p^{aux}(\delta)$.

Hence, the advisor with $\hat{p}(\delta)$ strictly improves both welfare and inequality outcomes compared to any advisor with $p > \hat{p}(\delta)$, and hence the advisors with $p > \hat{p}(\delta)$ do not constitute the Pareto frontier.

For the second part of the statement, we build on the following Lemma 12, which is proved in Online Appendix C.3. For a given δ , let $(1 - \bar{\tau}_2(\delta), \bar{\tau}_2(\delta))$ denote the optimal learning strategy of a politician in the unchecked discrimination, and let $(1 - \hat{\tau}_2(\delta), \hat{\tau}_2(\delta))$, where $\hat{\tau}_2(\delta)$ is given by equation (39), denote the equilibrium test allocation chosen by the advisor $\hat{p}(\delta)$.

Lemma 12. Fix δ . Then it holds

$$|\bar{\tau}_2(\delta) - 1/2| > |\hat{\tau}_2(\delta) - 1/2|.$$

Let

$$\begin{aligned}\omega(\tau_2, \delta) &:= V^1((1 - \tau_2, \tau_2), \delta) + V^2((1 - \tau_2, \tau_2), \delta) \\ &= -3 - \delta^2 + \frac{1/2(1 + \delta^2) + (1 + \delta)(1 - \tau_2)}{2 - \tau_2} + \frac{1/2(1 + \delta^2) + (1 - \delta)\tau_2}{1 + \tau_2}\end{aligned}$$

denote the sum of the expected utilities of the two groups as a function of test allocation $\tau = (\tau_1, \tau_2)$ under the constraint that the budget is fully used, $\tau_1 = 1 - \tau_2$, and where $V^k(\tau, \delta)$ is given by equation (C.3.2). Since the budget is exhausted in equilibrium (as shown in the proof of Lemma 4), welfare is thus given by $W(p, \delta) = \omega(\tau_2^*(p, \delta), \delta)$.

Observe the following symmetry property: $\omega(\tau_2, \delta) = \omega(1 - \tau_2, \delta)$. Moreover, for any $\delta \in (0, 1)$, the function $\omega(\tau_2, \delta)$ is strictly concave and maximized at $\tau_2 = 1/2$, when the test budget is split equally between the two groups. Therefore, the function $\omega(\tau_2, \delta)$ is strictly decreasing in $|\tau_2 - 1/2|$. Lemma 12 then implies that welfare with an advisor $\hat{p}(\delta)$ is strictly higher than welfare under unchecked discrimination. \square

References

- C. Aina. Tailored Stories. Working paper., 2024. URL https://chiaraaaina.github.io/files/TailoredStories_Aina.pdf.
- I. Ball and X. Gao. Benefitting from Bias: Delegating to encourage information acquisition. *Journal of Economic Theory*, 217:105816, 2024. doi: 10.1016/j.jet.2024.105816.
- A. Bardhi. Attributes: Selective Learning and Influence. *Econometrica*, 92(2):311–353, 2024. doi: 10.3982/ECTA18355.
- A. Bardhi and N. Bobkova. Local Evidence and Diversity in Minipublics. *Journal of Political Economy*, 131(9):2451–2508, September 2023. doi: 10.1086/724322.
- S. Callander. Searching and Learning by Trial and Error. *The American Economic Review*, 101(6):2277–2308, 2011. doi: 10.1257/aer.101.6.2277.
- C. Carnehl and J. Schneider. A quest for Knowledge. *arXiv working paper*, December 2024. doi: 10.48550/arXiv.2102.13434.
- J. D. Carrillo and T. Mariotti. Strategic ignorance as a self-disciplining device. *The Review of Economic Studies*, 67(3):529–544, 2000.

- Y.-K. Che and N. Kartik. Opinions as Incentives. *Journal of Political Economy*, 117 (5):815–860, October 2009. doi: 10.1086/648432.
- L. D'Amico and G. Tabellini. Disengaging from Reality: Online Behavior and Unpleasant Political News. *SSRN Electronic Journal*, 2022. doi: 10.2139/ssrn.4305223.
- I. Deimen and D. Szalay. Delegated Expertise, Authority, and Communication. *American Economic Review*, 109(4):1349–1374, April 2019. doi: 10.1257/aer.20161109.
- J. S. Demski and D. E. Sappington. Delegated expertise. *Journal of Accounting Research*, 25(1):68–89, 1987. doi: 10.2307/2491259.
- F. Echenique and A. Li. Rationally Inattentive Statistical Discrimination: Arrow Meets Phelps. *arXiv working paper*, November 2023. doi: 10.48550/arXiv.2212.08219.
- Facebook. Working to stop misinformation and false news. <https://www.facebook.com/formedia/blog/working-to-stop-misinformation-and-false-news>, 2017. Accessed: 2024-07-30.
- M. Fosgerau, R. Sethi, and J. Weibull. Equilibrium Screening and Categorical Inequality. *American Economic Journal: Microeconomics*, 15(3):201–242, August 2023. doi: 10.1257/mic.20210391.
- R. G. Fryer, Jr, P. Harms, and M. O. Jackson. Updating Beliefs when Evidence is Open to Interpretation: Implications for Bias and Polarization. *Journal of the European Economic Association*, 17(5):1470–1501, October 2019. doi: 10.1093/jeea/jvy025.
- F. Germano, V. Gómez, and F. Sobbrío. Crowding Out the Truth? A simple Model of Misinformation, Polarization and Meaningful Social Interactions. *arXiv working paper*, 2022. doi: 10.48550/arXiv.2210.02248.
- P. Ilinov, A. Matveenko, M. Senkov, and E. Starkov. Optimally Biased Expertise. *arXiv working paper*, 2022. doi: 10.48550/arXiv.2209.13689.
- M. Kirneva. Informing to divert attention. Working paper, 2023.
- A. Liang, X. Mu, and V. Syrgkanis. Dynamically Aggregating Diverse Information. *Econometrica*, 90(1):47–80, 2022. doi: 10.3982/ECTA18324.
- A. Liang, J. Lu, and X. Mu. Algorithm design: A fairness-accuracy frontier. *accepted at Journal of Political Economy*, 2024.

- A. Lindbeck and J. Weibull. Delegation of investment decisions, and optimal remuneration of agents. *European Economic Review*, page 103559, August 2020. doi: 10.1016/j.euroecorev.2020.103559.
- A. Matveenko and E. Starkov. Sparking curiosity or tipping the scales? Targeted advertising with consumer learning. *Journal of Economic Behavior & Organization*, 213:172–192, September 2023. doi: 10.1016/j.jebo.2023.07.008.
- K. P. Nimark and S. Sundaresan. Inattention and belief polarization. *Journal of Economic Theory*, 180:203–228, March 2019. doi: 10.1016/j.jet.2018.12.007.
- T. O’Donoghue and M. Rabin. Doing it now or later. *The American Economic Review*, 89(1):103–124, 1999. doi: 10.1257/aer.89.1.103.
- W. Olszewski. Preferences and information processing under vague information. *Journal of Mathematical Economics*, 94:102461, May 2021. doi: 10.1016/j.jmateco.2020.102461.
- P. Onuchic. Recent contributions to theories of discrimination. September 2024. Unpublished manuscript.
- K. Pogorelskiy and M. Shum. News We Like to Share: How News Sharing on Social Networks Influences Voting Outcomes. *SSRN Electronic Journal*, August 2019. doi: 10.2139/ssrn.2972231.
- W. Tamura. Bayesian persuasion with quadratic preferences. *SSRN Electronic Journal*, page 51, March 2018. doi: 10.2139/ssrn.1987877.
- United Nations. Countering disinformation. <https://www.un.org/en/countering-disinformation>, 2022. Accessed: 2024-07-30.

Online Appendix to “Strategic Attribute Learning”

Jean-Michel Benkert, Ludmila Matysková, Egor Starkov

B Extensions

B.1 Multiple researchers: media polarization

Our first extension explores a model of a media market, in which media polarization can emerge, but—in contrast to conventional wisdom—in a way that does not inherently reduce welfare. Specifically, we consider a voter with some priorities over policy issues and two media outlets with different priorities, which attempt to influence the voter’s choice via their choice of coverage. We show that in a wide range of scenarios, the competition between media outlets polarizes their coverage in equilibrium. However, this polarization is advantageous for the voter, as it offers access to a broader and more diverse range of information. This diversity allows voter to make better-informed decisions compared to relying on a single, moderate outlet that is not fully aligned with the voter.

We adopt Framework A from Section 5 and extend it by introducing two media outlets, A and B , each aiming to influence a decision of voter V that affects two policy issues $k = 1, 2$.²⁵ For $i = A, B, V$, player i ’s utility is

$$u^i(d, \theta) = - \sum_{k=1}^2 \alpha_k^i (d - \theta_k)^2, \quad (\text{B.1.1})$$

where $d \in \mathbb{R}$ is the voter’s decision (the policy platform of the elected candidate), $\alpha_k^i > 0$ are player i ’s weights on respective policy issues $k = 1, 2$ with $\alpha_1^i + \alpha_2^i = 1$, and θ_k is the optimal decision for issue k (i.e., the platform of the candidate offering the best policy

²⁵Our model features a single voter for simplicity. One can think of her as the median voter in a more general model with multiple voters who have different weights on the policy issues. Since the decision is one-dimensional and voters’ preferences are single-peaked, standard logic dictates that the median voter (one with median weights α^i) is pivotal and would thus be the one targeted by the media outlets aiming to influence the vote outcome.

on issue k given state of the world θ).²⁶ Assume w.l.o.g. that media outlet A assigns a higher weight to issue $k = 1$ than outlet B : $\alpha_1^A > \alpha_1^B$.

Each media outlet $m = A, B$ chooses coverage $q^m = (q_1^m, q_2^m) \in \mathbb{R}_+^2$ such that $q_1^m + q_2^m = 1$, where q_k^m describes the fraction of the coverage of media outlet m devoted to issue k .²⁷ The voter has a total budget of attention $T > 0$ and chooses how to allocate it between the two outlets: $t = (t^A, t^B) \in \mathbb{R}_+^2$ such that $t^A + t^B \leq T$. Given q^m and t^m for $m = A, B$, for each issue k the voter observes a realization of signal

$$\tilde{s}_k^m(q_k^m, t^m) = \tilde{\theta}_k + \tilde{\varepsilon}_k^m, \text{ where } \tilde{\varepsilon}_k^m \sim \mathcal{N}\left(0, \frac{1}{q_k^m t^m}\right),$$

and $\tilde{\varepsilon}_k^m$ is independent of all other random variables. Let $s = (s_1^A, s_2^A, s_1^B, s_2^B)$ denote the collection of signal realizations from both outlets on both issues, and \tilde{s} its respective distribution.

To curb the equilibrium multiplicity (though it cannot be avoided entirely, see the discussion below), assume that the voter chooses her attention allocation t first. Media outlets then observe t and simultaneously choose their coverage q^m . The rest of the game proceeds as in the baseline model: nature draws state realization θ and signal realizations s ; the voter observes q^m and s , updates her beliefs and takes decision d ; afterwards, payoffs are realized. The equilibrium concept is analogous to that in the baseline model, and is formally stated in Section B.1.1 below.²⁸

We say that *the voter achieves her best information in equilibrium* when the attention and coverages chosen in equilibrium jointly maximize the voter's ex ante expected payoff, given the voter's equilibrium decision strategy. We can state our main result of this section as follows.²⁹

Proposition B.1.1. If $\alpha_1^A > \alpha_1^V > \alpha_1^B$ and T is large enough, then in the unique

²⁶In the model, all players agree, conditional on θ , which candidate would enact an optimal policy on a given issue k . Tension arises because the single elected candidate will choose policies relevant to many issues, and preferences of the voter and the media outlets differ in terms of weights they assign to different issues.

²⁷In our setting, the media cannot lie, withhold facts, distort information, or provide fake news. This assumption reflects the recent crackdown on fake news by companies and authorities (see, e.g., Facebook, 2017; United Nations, 2022). Therefore, the strategic choice of which topics to cover remains one of the few options available to media editorial rooms.

²⁸It is straightforward to show that an equilibrium of such a sequential-move game would be outcome-equivalent to an equilibrium of a simultaneous-move game, where the media choose q contemporaneously with the voter's choice of t . However, the latter game may also admit other equilibria.

²⁹The proofs of this extension can be found below.

equilibrium the media outlets are polarized, $q^A = (1, 0)$ and $q^B = (0, 1)$, and the voter achieves her best information in equilibrium.

The proposition states that if media outlets are slanted in the opposite directions relative to the voter (and the voter's attention budget is large enough), even small differences in preferences between the voter and the media cause each outlet to focus exclusively on its more favoured issue (as compared to other players), neglecting the other. This occurs even if both media outlets care about both issues (i.e., $\alpha_k^m \in (0, 1)$ for all m, k). The voter's attention allocation and the coverage of each media jointly determine how much information the voter receives on each policy issue. Once the voter chooses her attention, the media outlets engage in a tug-of-war, each seeking to increase the amount of information the voter receives on the outlet's preferred issue. This competition ultimately causes each outlet to cover exclusively a single issue. This polarization and the resulting diverse coverage allow the voter to achieve her best information in equilibrium.

To set the benchmark for the result above, suppose there was only one media outlet. Then, we would be back in the setting of Section 4. In such a setting, the media outlet can shape the total amount of information the voter receives on both issues according to its own preferences and away from the voter's best information. Thus, the voter always weakly prefers a polarized media duopoly over a media monopoly, however moderate the latter may be. We capture this observation in the following statement, which follows from Theorem 2 and Proposition B.1.1.³⁰

Corollary B.1.1. If $\alpha_1^A > \alpha_1^V > \alpha_1^B$ and T is large enough, then the voter strictly prefers the equilibrium when both media outlets $m = A, B$ are available to the equilibrium when only one outlet $m = A, B$ is present in the market.

Finally, we note that the results in this section rely on the assumption that the media outlets are directly interested in the voter's final decision. This is applicable in scenarios where editors have strong ideological stances, or policy-interested outlet owners can affect editorial policy (e.g., Fox News, Russia Today). We can also think of settings

³⁰Both Proposition B.1.1 and Corollary B.1.1 require that $\alpha_1^A > \alpha_1^V > \alpha_1^B$ and T large enough. In Section B.1.1, we generalize the results to all α and T . While equilibrium multiplicity becomes an issue, we show that all equilibria are payoff-equivalent. In a representative equilibrium, at least one media outlet focuses exclusively on its relatively more preferred issue (while the other outlet either does the same or is ignored by the voter). Corollary B.1.1 also continues to hold in general, though with weak preference instead of strict.

other than media markets in which similar outcomes could also arise, such as think tanks or advisors influencing policymakers. In contrast, independent media outlets may instead prioritize engagement and aim to attract attention from consumers in order to boost ad revenue or subscription fees. It is quite intuitive that if attention-seeking is the primary motivation for the media outlets, then they would give the voter the information she wants to see.³¹ We would thus observe less media polarization with attention-seeking than with partisan media, but Proposition B.1.1 suggests that this difference is likely superficial and irrelevant for the actual voting decisions.

B.1.1 Proofs for the extension on media polarization

We begin by formally stating the equilibrium concept for the media model alongside two supplementary results. Lemma 13 then establishes that given (t, q^A, q^B) , the posterior distribution of $\tilde{\theta}|s$ can be related to that in the baseline model by considering *aggregate attention allocation* τ given by $\tau_k := t^A q_k^A + t^B q_k^B$ for each k . Further, Lemma 14 shows that the ex ante payoffs can be expressed as functions of τ , and so the voter’s attention allocation problem and the media outlets’ coverage choice problems can be seen as a joint choice of τ , with every player having some power over it. Then, we provide a general equilibrium characterization, with Proposition B.1.2 covering the cases excluded in Proposition B.1.1, and Corollary B.1.2 being a more general, albeit slightly weaker analog to Corollary B.1.1. We provide a unified proof for Propositions B.1.1 and B.1.2 before proving Corollaries B.1.1 and B.1.2.

Let $\mathcal{R} := \{t \in \mathbb{R}_+^2 \mid t_1 + t_2 \leq T\}$ and $\mathcal{Q} := \{q \in \mathbb{R}_+^2 \mid q_1 + q_2 \leq 1\}$ be the domains of the players’ strategies.

Equilibrium concept. A weak Perfect Bayesian Equilibrium of the game introduced in Section B.1 (referred to in the relevant proofs as “equilibrium”) is a tuple

$$\left(t^*, q^{A^*}(t), q^{B^*}(t), d^*(s, t, q^A, q^B), \tilde{\theta}(s, t, q^A, q^B) \right)$$

such that:

³¹The classic Hotelling model of spatial competition suggests that both media outlets would likely converge on the median voter’s “position” (offer the median voter’s preferred information mix) even if there are many voters with heterogeneous preferences.

1. the coverage choice $q^{m*}(t) : \mathcal{R} \rightarrow \mathcal{Q}$ of each outlet $m = A, B$ maximizes its value (B.1.4) given equilibrium strategies and belief updating rule $(q^{j*}(t), d^*(s, t, q^A, q^B), \tilde{\theta}(s, t, q^A, q^B))$ for $j \neq m$ and realized attention choice t ;
2. the voter's attention allocation $t^* \in \mathcal{R}$ and decision strategy $d^*(s, t, q^A, q^B) : \mathbb{R}^2 \times \mathcal{R} \times \mathcal{Q}^2 \rightarrow \mathbb{R}$ maximize her value (B.1.4) given equilibrium strategies and belief updating rule $(q^{A*}(t), q^{B*}(t), \tilde{\theta}(s, t, q^A, q^B))$;
3. the voter's posterior belief $\tilde{\theta}(s, t, q^A, q^B) : \mathbb{R}^2 \times \mathcal{R} \times \mathcal{Q}^2 \rightarrow \Delta(\mathbb{R}^2)$ about θ is obtained via Bayes' rule given signal realizations s and realized attention allocation and coverage choices (t, q^A, q^B) .

Next, we show that the induced distribution of posterior beliefs in this extension coincides with that in the main model.

Lemma 13. In the media model, the distribution of the voter's posterior belief $\tilde{\theta}(\tilde{s}, t, q^A, q^B)$ is the same as the one induced by (1) and (4)–(6), with $\tau = \tau(t, q^A, q^B)$ being the aggregate attention allocation defined as $\tau = (\tau_1, \tau_2)$ with $\tau_k := t^A q_k^A + t^B q_k^B$ for each k .

Proof. Throughout the proof, we fix some attention allocation and coverages (t, q^A, q^B) and the corresponding aggregate attention allocation $\tau_k = t^A q_k^A + t^B q_k^B$, $k = 1, 2$. For the sake of brevity, we will omit (t, q^A, q^B) and τ from the set of conditioning variables.

The collection (t, q^A, q^B) induces a distribution of the signal vector \tilde{s} . In turn, each signal realization is associated with a posterior belief $\tilde{\theta}(s)$. Ex ante, therefore, the collection (t, q^A, q^B) induces a distribution over posterior beliefs, $\tilde{\theta}(\tilde{s}) \in \Delta(\Delta(\mathbb{R}^2))$. Our goal is to establish that, given the aggregate attention allocation τ associated with the collection (t, q^A, q^B) , the induced distribution coincides with the one described in Section 2 induced by the same τ , where (4)–(6) describe the posterior distribution conditional on s , and (1) implies that $\tilde{s}_k \sim \mathcal{N}\left(\mu_k^0, \Sigma_k^0 + \frac{1}{\tau_k}\right)$ for all k .

Since $\tilde{\theta}_1 \perp \tilde{\theta}_2$ and $\tilde{\theta}_k \perp \tilde{\varepsilon}_j^m$ for all $j, k = 1, 2$ and $m = A, B$, the posterior belief about each attribute is independent of everything related to the other attribute. Hence, we fix some $k = 1, 2$ and work with the (marginal) distributions of $\tilde{\theta}_k$ and \tilde{s}_k^m , which is sufficient to prove the result.

We first characterize $\tilde{\theta}_k(s)$, the voter's posterior belief about θ_k given some signal realizations s_k^A, s_k^B . It is immediate to verify directly from normal p.d.f.s that if $\tilde{X} \sim \mathcal{N}(\mu, \sigma_X^2)$, $\tilde{\varepsilon} \sim \mathcal{N}(0, \sigma_\varepsilon^2)$ and $Y = X + \varepsilon$, then $\tilde{X}|Y \sim \mathcal{N}(\hat{\mu}, \hat{\sigma}^2)$, where

$$\hat{\mu} = \frac{\frac{1}{\sigma_X^2}}{\frac{1}{\sigma_X^2} + \frac{1}{\sigma_\varepsilon^2}} \mu + \frac{\frac{1}{\sigma_\varepsilon^2}}{\frac{1}{\sigma_X^2} + \frac{1}{\sigma_\varepsilon^2}} Y, \quad \hat{\sigma}^2 = \frac{1}{\frac{1}{\sigma_X^2} + \frac{1}{\sigma_\varepsilon^2}}.$$

This directly implies characterization (4)–(6) in Section 2.2. In the context of this proof, this gives first that $\tilde{\theta}_k|s_k^A \sim \mathcal{N}(\hat{\mu}_{k,A}, \hat{\Sigma}_{k,A})$ with

$$\hat{\mu}_{k,A} = \frac{\frac{1}{\Sigma_k^0}}{\frac{1}{\Sigma_k^0} + t^A q_k^A} \mu_k^0 + \frac{t^A q_k^A}{\frac{1}{\Sigma_k^0} + t^A q_k^A} s_k^A, \quad \hat{\Sigma}_{k,A} = \frac{1}{\frac{1}{\Sigma_k^0} + t^A q_k^A},$$

and then, subsequently, that $(\tilde{\theta}_k|s_k^A)|s_k^B \sim \mathcal{N}(\hat{\mu}_k, \hat{\Sigma}_k)$, where

$$\begin{aligned} \hat{\mu}_k &= \frac{\frac{1}{\hat{\Sigma}_{k,A}}}{\frac{1}{\hat{\Sigma}_{k,A}} + t^B q_k^B} \hat{\mu}_{k,A} + \frac{t^B q_k^B}{\frac{1}{\hat{\Sigma}_{k,A}} + t^B q_k^B} s_k^B = \frac{\frac{1}{\Sigma_k^0} \mu_k^0 + t^A q_k^A s_k^A + t^B q_k^B s_k^B}{\frac{1}{\Sigma_k^0} + t^A q_k^A + t^B q_k^B} \\ &= \frac{\frac{1}{\Sigma_k^0}}{\frac{1}{\Sigma_k^0} + \tau_k} \mu_k^0 + \frac{1}{\frac{1}{\Sigma_k^0} + \tau_k} (t^A q_k^A s_k^A + t^B q_k^B s_k^B), \end{aligned} \quad (\text{B.1.2})$$

$$\hat{\Sigma}_k = \frac{1}{\frac{1}{\hat{\Sigma}_{k,A}} + t^B q_k^B} = \frac{1}{\frac{1}{\Sigma_k^0} + t^A q_k^A + t^B q_k^B} = \frac{1}{\frac{1}{\Sigma_k^0} + \tau_k}. \quad (\text{B.1.3})$$

Comparing (B.1.3) with its counterpart (6) in the baseline model, we note that they coincide. Next, compare (B.1.2) with its counterpart (5) in the baseline model. In order to conclude that the distribution of posteriors is the same as in the baseline model, we need to show that the random variable defined as

$$\tilde{s}_k^f := \frac{1}{\tau_k} (t^A q_k^A \tilde{s}_k^A + t^B q_k^B \tilde{s}_k^B) = \tilde{\theta}_k + \frac{1}{\tau_k} (t^A q_k^A \tilde{\varepsilon}_k^A + t^B q_k^B \tilde{\varepsilon}_k^B)$$

has the same distribution as \tilde{s}_k from the baseline model. Note that $\tilde{s}_k^f|\theta_k \sim \mathcal{N}\left(\theta_k, \frac{1}{\tau_k}\right)$ since $\mathbb{E}[\tilde{\varepsilon}_k^m] = 0$ and $\text{var}\left(\frac{t^m q_k^m}{\tau_k} \tilde{\varepsilon}_k^m\right) = \frac{1}{\tau_k}$, and then $\tilde{s}_k^f \sim \mathcal{N}\left(\mu_k^0, \Sigma_k^0 + \frac{1}{\tau_k}\right)$. For a given τ , these coincide with the distributions of $\tilde{s}_k|\theta_k$ and \tilde{s}_k from the baseline model, respectively. Therefore, the distribution of $\hat{\mu}_k$ above coincides with that of (4), so the distribution of $\tilde{\theta}_k(s_k^f)$ in the media model coincides with the distribution of $\tilde{\theta}_k(s_k)$ in the baseline model when $s_k^f = s_k$. Together, the two latter facts imply that the distribution of $\tilde{\theta}_k(\tilde{s}_k^f)$ in the media model coincides with the distribution of $\tilde{\theta}_k(\tilde{s}_k)$ in the baseline

model. □

Our next lemma below show that all players' respective ex ante expected utilities only depend on (t, q^A, q^B) through $\tau = \tau(t, q^A, q^B)$. Specifically, define player i 's value function as

$$V^i(\tau) := \mathbb{E} \left[- \sum_{k=1}^2 \alpha_k^i \left(d^*(\tilde{s}^f, \tau) - \tilde{\theta}_k(\tilde{s}^f, \tau) \right)^2 \right] \quad (\text{B.1.4})$$

for $i = A, B, V$, where $d^*(\tilde{s}^f, \tau)$ is the voter's equilibrium decision strategy given $\tau = \tau(t, q^A, q^B)$, and $s^f = (s_1^f, s_2^f)$ is a vector of fictitious "aggregate" signals introduced in the proof of Lemma 13.

Lemma 14. In the media model, given (t, q^A, q^B) , the expectations of all players' utilities (B.1.1) are given by values (B.1.4) and only depend on $\tau(t, q^A, q^B)$:

$$\mathbb{E} \left[u^i \left(d^*(\tilde{s}, t, q^A, q^B), \tilde{\theta}(\tilde{s}, t, q^A, q^B) \right) \right] = V^i(\tau(t, q^A, q^B))$$

Proof. Rewriting the expected utility (the left-hand side of the equality in the lemma) using definition (B.1.1), we get

$$\mathbb{E} \left[- \sum_{k=1}^2 \alpha_k^i \left(d^*(\tilde{s}, t, q^A, q^B) - \tilde{\theta}_k(\tilde{s}, t, q^A, q^B) \right)^2 \right]. \quad (\text{B.1.5})$$

The voter's equilibrium decision for any signal realization s is given by

$$d^*(s, t, q^A, q^B) = \sum_{k=1}^2 \alpha_k^V \mathbb{E} \left[\tilde{\theta}_k(s, t, q^A, q^B) \right].$$

By Lemma 13 we know that $\tilde{\theta}_k(\tilde{s}, t, q^A, q^B) \sim \tilde{\theta}_k(\tilde{s}^f, \tau)$ (the two have the same distributions), where $\tilde{s}^f = (\tilde{s}_1^f, \tilde{s}_2^f)$ with $\tilde{s}_k^f = \frac{1}{\tau_k} (t^A q_k^A \tilde{s}_k^A + t^B q_k^B \tilde{s}_k^B)$. This implies that $d^*(\tilde{s}, t, q^A, q^B) \sim d^*(\tilde{s}^f, \tau)$. Then we can write (B.1.5) as

$$\mathbb{E} \left[- \sum_{k=1}^2 \alpha_k^i \left(d^*(\tilde{s}^f, \tau) - \tilde{\theta}_k(\tilde{s}^f, \tau) \right)^2 \right],$$

which is exactly (B.1.4), the definition of $V^i(\tau)$. □

We now make a statement about the media model equilibria in cases not covered in

Proposition B.1.1, and prove both statements simultaneously.

Proposition B.1.2. In each of the cases listed below, all equilibria are payoff-equivalent to the following:

1. if $\alpha_1^A > \alpha_1^V > \alpha_1^B$ and T is not large enough: both media outlets choose the same extreme coverage, either $q^A = q^B = (1, 0)$, or $q^A = q^B = (0, 1)$, and the voter achieves her best information;
2. if $\alpha_1^V \geq \alpha_1^A$: the voter only follows outlet A , $t = (T, 0)$, and outlet A achieves its best information;³²
3. if $\alpha_1^B \geq \alpha_1^V$: the voter only follows outlet B , $t = (0, T)$, and outlet B achieves its best information.

Proof of Propositions B.1.1 and B.1.2. Lemmas 13 and 14 imply that the media setting with four signals $s = (s_1^A, s_2^A, s_1^B, s_2^B)$ can, without loss, be replaced by the baseline model with fictitious aggregate signals $s^f = (s_1^f, s_2^f)$, which have the same distribution conditional on aggregate attention allocation $\tau(t, q^A, q^B)$ as the signals in our main model of Section 2. For brevity, we often drop the arguments of the aggregate attention allocation and write it as simply τ .

Proceeding by backwards induction, we first note that the voter's equilibrium decision $d^*(s^f, \tau)$ is analogous to the single-player and strategic versions of our model, see the proof of Lemma 14. In the remainder of this proof, we assume the voter follows this equilibrium decision strategy $d^*(s^f, \tau)$.

Proceeding further, Lemma 14 implies that the outlets' problem of choosing coverage q^m and the voter's problem of choosing attention allocation t reduce to the problem of maximizing value function (B.1.4) over $\tau(t, q^A, q^B)$. Fix some t and consider the outlets' problem. Let $T_V(t) := t_1 + t_2$ denote the total amount of attention devoted to the two outlets by the voter. Let $\tau^{i^*}(t) := \arg \max_{\tau \in \mathbb{R}_+^2} V^i(\tau)$ s.t. $\tau_1 + \tau_2 \leq T_V(t)$ for all $i = A, B, V$. Note that player i 's best information, by definition, is given by $\max_{t \in \mathcal{R}} \tau^{i^*}(t)$. For both $m = A, B$, $\tau^{m^*}(t)$ is then described by Lemma 2 and Theorem 2 with $T = T_V(t)$. By Lemma 2, if m 's expected payoff (12) is weakly increasing in

³²We define outlet m 's best information by analogy with the voter's best information in the text, i.e., (t, q^A, q^B) must jointly maximize m 's expected payoff given the voter's equilibrium decision strategy.

τ_k , it is also strictly concave in τ_k .³³ From Theorems 1 and 2, we obtain a closed-form solution for $\tau^{i*}(t)$:

$$\begin{aligned}\tau_1^{i*}(t) &= \max \left\{ 0, \min \left\{ \frac{\hat{\alpha}_1^i \Sigma_1^0 - \hat{\alpha}_2^i \Sigma_2^0}{\Sigma_1^0 \Sigma_2^0 (\hat{\alpha}_1^i + \hat{\alpha}_2^i)} + \frac{\hat{\alpha}_1^i}{\hat{\alpha}_1^i + \hat{\alpha}_2^i} T_V(t), T_V(t) \right\} \right\}, \\ \tau_2^{i*}(t) &= T_V(t) - \tau_1^{i*}(t),\end{aligned}\tag{B.1.6}$$

where $\hat{\alpha}_k^i := \sqrt{\max \{0, (\alpha_k^i)^2 - (\alpha_k^i - \alpha_k^V)^2\}}$ for all $k = 1, 2$ and $i = A, B, V$ (note that $\hat{\alpha}_k^V = \alpha_k^V$, so for $i = V$, (B.1.6) coincides with the solution in Theorem 1). Note that $\tau_1^{i*}(t)$ is weakly increasing in α_1^i (subject to the $\alpha_1^i + \alpha_2^i = 1$ constraint). This characterization implies that since $\alpha_1^i + \alpha_2^i = 1$ for all $i = A, B, V$, the auxiliary weights $\hat{\alpha}_k^m$ cannot simultaneously be zero for both $k = 1, 2$ for any $m = A, B$, so $\tau_1^{m*}(t) + \tau_2^{m*}(t) = T_V(t)$ for all t (the media want no attention wasted). We note that subject to this constraint, value $V^i(\tau_1, T_V(t) - \tau_1)$ is either strictly monotone, or strictly concave in τ_1 for all $i = A, B, V$.

The characterization in Theorem 2 also implies that given t , if $\hat{\alpha}_1^m \Sigma_1^0 > \hat{\alpha}_2^m \Sigma_2^0$ and $T_V(t) \leq \frac{\hat{\alpha}_1^m}{\hat{\alpha}_2^m \Sigma_2^0} - \frac{1}{\Sigma_1^0}$ (or for all $T_V(t)$ if $\hat{\alpha}_2 = 0$), then $\tau^{m*}(t) = (T_V(t), 0)$, and vice versa: if $\hat{\alpha}_2^m \Sigma_2^0 > \hat{\alpha}_1^m \Sigma_1^0$ and $T_V(t) \leq \frac{\hat{\alpha}_2^m}{\hat{\alpha}_1^m \Sigma_1^0} - \frac{1}{\Sigma_2^0}$ (or for all $T_V(t)$ if $\hat{\alpha}_1 = 0$), then $\tau^{m*}(t) = (0, T_V(t))$. If none of these conditions apply, then $\tau^{m*}(t)$ is interior (and strictly monotone in α_1^m subject to the $\alpha_1^m + \alpha_2^m = 1$ constraint). Therefore, if we let

$$\hat{T} := \max \left\{ 0, \min_{m \in \{A, B\}} \left\{ \frac{\hat{\alpha}_1^m}{\hat{\alpha}_2^m \Sigma_2^0} - \frac{1}{\Sigma_1^0} \right\}, \min_{m \in \{A, B\}} \left\{ \frac{\hat{\alpha}_2^m}{\hat{\alpha}_1^m \Sigma_1^0} - \frac{1}{\Sigma_2^0} \right\} \right\},$$

then whenever $T_V(t) \leq \hat{T}$, both media outlets optimally prefer the same extreme coverage: $\tau_1^{A*}(t) = \tau_1^{B*}(t) \in \{0, T_V(t)\}$. Note that if $\hat{\alpha}_1^A \Sigma_1^0 > \hat{\alpha}_2^A \Sigma_2^0$ and $\hat{\alpha}_1^B \Sigma_1^0 < \hat{\alpha}_2^B \Sigma_2^0$, then $\hat{T} = 0$, and such a preference profile ($\tau_1^{A*}(t) = \tau_1^{B*}(t) \in \{0, T_V(t)\}$) cannot arise. Further, if $T_V(t) > \hat{T}$, then $\tau^{m*}(t)$ is interior for at least one outlet m , and since $\alpha_1^A > \alpha_1^B$ (and $\alpha_2^A < \alpha_2^B$), we then have $\tau_1^{A*}(t) > \tau_1^{B*}(t)$.

Next, we claim that if $\tau \neq \tau^{m*}(t)$ for some outlet m in equilibrium, then m is either polarized (i.e., $q^m \in \{(1, 0), (0, 1)\}$), or ignored ($t^m = 0$). To see this, proceed by contradiction and consider a strategy profile (t, q^A, q^B) such that $\tau_k(t, q^A, q^B) < \tau_k^{m*}(t)$, $q_k^m < 1$, and $t^m > 0$ for some $k = 1, 2$ and $m = A, B$. Then increasing q_k^m (while

³³This is most evident from (30) in the proof of Theorem 2: one can readily see by analogy that both monotonicity and concavity in this problem depend on whether $\lambda_k^i := \alpha_k^V (2\alpha_k^m - \alpha_k^V) \geq 0$.

decreasing the other weight q_{-k}^m) brings the aggregate test allocation τ closer to m 's optimum $\tau^{m*}(t)$, which strictly increases m 's expected payoff (12).³⁴ This is a profitable deviation for outlet m , and the original strategy profile thus cannot be an equilibrium. The same is true of the case when $\tau_k > \tau_k^{m*}(t)$ and $q_k^m > 0$ and $t^m > 0$ for some $k = 1, 2$, $m = A, B$.

The logic above also implies that as $\alpha_1^A > \alpha_1^B$, there is no equilibrium with τ such that $\tau_1 > \tau_1^{A*}(t)$ or $\tau_1^{B*}(t) > \tau_1$, since in the former case both outlets m want to decrease q_1^m and increase q_2^m (and at least one has the scope for such a deviation), and in the latter they want to do the opposite. Therefore, in equilibrium, the aggregate attention allocation $\tau = \tau(t^*, q^{A*}, q^{B*})$ must be such that $\tau_1 \in [\tau_1^{B*}(t^*), \tau_1^{A*}(t^*)]$ and $\tau_2 = T_V(t^*) - \tau_1 \in [\tau_2^{A*}(t^*), \tau_2^{B*}(t^*)]$.

It remains to consider the voter's attention allocation problem. Since $\tau^{m*}(t)$ only depend on t through $T_V(t)$, the voter's problem can be split into first choosing the total amount of attention to devote to media, T_V , and then choosing how to split it by way of choosing τ such that $\tau_1 \in [\tau_1^{B*}(T_V), \tau_1^{A*}(T_V)]$ and $\tau_2 = T_V - \tau_1$. Note that in the second step, all aggregate allocations in this interval are available to the voter for a given T_V . By setting $t^m = T_V$ for some m and $t^{-m} = 0$, she can induce $\tau = \tau^{m*}(T_V)$. Any allocation τ in the interior of the interval can be achieved by setting $t^A = \tau_1$ and $t^B = \tau_2$, in which case the unique best response for outlet A is $q^A = (1, 0)$, and the unique best response for outlet B is $q^B = (0, 1)$, as argued above. We shall thus refer to τ that satisfy these requirements as *feasible* given some T_V .

The voter never wants any attention wasted, since her payoff is strictly increasing in τ_k for both k . From (B.1.6) we can see that $\frac{d}{dT_V} \tau_k^{m*}(T_V) \in [0, 1]$ for all m, k , hence for any T_V , any τ^a that is feasible given T_V , and any $\varepsilon > 0$, there exists τ^b that is feasible given $T_V + \varepsilon$ and is such that $\tau_k^b \geq \tau_k^a$ for both k , with at least one inequality being strict.³⁵ Therefore, increasing T_V always offers a strict improvement the voter, so in equilibrium, $T_V(t^*) = T$.

We now move on to the second stage of the voter's problem. If $T \leq \hat{T}$, then $\tau_1^{A*}(T) =$

³⁴ "Closer" here is used in the sense of the new aggregate attention allocation $\bar{\tau}$ being a convex combination of the original τ and the optimal $\tau^{m*}(t)$. The payoff function being concave on the relevant interval, and $\tau^{m*}(t)$ being its maximizer then directly imply that $\bar{\tau}$ yields a higher payoff than the original τ .

³⁵ A way to see this is to notice that an increase in T_V implies an increase in the strong set order of the intervals of feasible in τ_1 and τ_2 .

$\tau_1^{B^*}(T) \in \{0, 1\}$, so the set of available τ is a singleton, and the voter is then indifferent between all attention allocations. Suppose now that $T > \hat{T}$. If $\alpha_1^A > \alpha_1^V > \alpha_1^B$, then $\tau_1^{V^*}(T) \in [\tau_1^{B^*}(T), \tau_1^{A^*}(T)]$, so the voter can attain her optimum $\tau^{V^*}(T)$ by following the strategy described above. It is immediate from (B.1.6) that in this case there exists $\hat{T}_V \geq \hat{T}$ such that $\tau^{V^*}(T)$ is interior if and only if $T > \hat{T}_V$. For $T > \hat{T}_V$, $\tau^{V^*}(T)$ is uniquely optimal for the voter, hence $t = (\tau_1^{V^*}(T), \tau_2^{V^*}(T))$, $q^A = (1, 0)$, and $q^B = (0, 1)$ is the unique equilibrium, proving Proposition B.1.1. For $T \leq \hat{T}_V$, as argued above, the voter can still achieve $\tau^{V^*}(T)$, but the equilibrium may not be unique. This proves case 1 of Proposition B.1.2. In case $\alpha_1^V \notin (\alpha_1^A, \alpha_1^B)$, since the voter's value function (11) is strictly concave in τ , she will optimally choose the available allocation closest (in the sense of Footnote 34) to $\tau^{V^*}(T)$. In particular, if $\alpha_1^B \geq \alpha_1^V$, then $\tau_1^{V^*}(T) \leq \tau_1^{B^*}(T)$, so the voter will choose $t = (0, T)$, resulting in $\tau = \tau_B^*(T)$. And if $\alpha_1^V \geq \alpha_1^A$, then $\tau_1^{V^*}(T) \geq \tau_1^{A^*}(T)$, and the voter will choose $t = (T, 0)$, resulting in $\tau = \tau_A^*(T)$. This proves parts 2 and 3 of Proposition B.1.2 and completes this proof. \square

Finally, we state the following general corollary before proving it along with Corollary B.1.1.

Corollary B.1.2. The voter weakly prefers the equilibrium when both media outlets $m = A, B$ are available to the equilibrium when only one outlet $m = A, B$ is present in the market.

Proof of Corollaries B.1.1 and B.1.2. In the monopoly case, if only outlet m is available, then the voter always lends it her full attention, $t^m = T$, since her value (B.1.4) reduces to (11), which is strictly increasing and strictly concave in τ . Therefore, m can choose aggregate attention allocation τ freely subject to $\tau_1 + \tau_2 = T$. Since $\alpha_1^i + \alpha_2^i = 1$ for all i , Theorem 2 implies that m does not want to waste any of the voter's attention. The unique equilibrium aggregate attention allocation is, therefore, $\tau^{m^*}(T)$, as defined in the proof of Propositions B.1.1 and B.1.2, which is m 's best information.

Propositions B.1.1 and B.1.2 above shows that under media duopoly, the equilibrium aggregate attention allocation τ^* is given by $\tau^{V^*}(T)$ if $\alpha_1^A > \alpha_1^V > \alpha_1^B$, by $\tau^{A^*}(T)$ if $\alpha_1^A \leq \alpha_1^V$, and by $\tau^{B^*}(T)$ if $\alpha_1^V \leq \alpha_1^B$. Note that in the two latter cases, it is the voter's preferred media outlet that has the monopoly power, which proves Corollary B.1.2 for those cases. In the former case, $\tau^{V^*}(T)$ being its maximizer of the voter's value

function imply that the voter prefers duopoly, concluding the proof of Corollary B.1.2. If $\tau^{V^*}(T) \in (\tau^{B^*}(T), \tau^{A^*}(T))$, which is the case if and only if $T > \hat{T}_V$, as defined in the proof of Propositions B.1.1 and B.1.2, then it is also the unique maximizer (since the value function is strictly concave on that interval in that case), which proves Corollary B.1.1. \square

B.2 Changing Preferences

In this second extension, we study how a (potential) change in an agent's preferences between learning and decision-making affects his learning strategy. Following the literature on time-inconsistent preferences (e.g., O'Donoghue and Rabin, 1999), we consider sophisticated and naive agents and investigate the impact of such a change in preferences on the learning strategy and on the agent's welfare.

We model this similarly to our baseline model by considering a single agent facing a two-stage problem. In the first stage, the agent acquires information about unknown attributes. In the second stage, the agent takes a decision based on the acquired information. The agent's preference may change between the two stages. For instance, this change may occur due to a change in the agent's circumstances, such as losing the job or falling ill. Alternatively, the change in preferences may be due to the agent's self-control problems, such as succumbing to temptation.

In the first stage, when acquiring information, the agent has utility function

$$u^R(d, \theta) = - (d - \alpha_1^R \theta_1 - \alpha_2^R \theta_2)^2, \quad (\text{B.2.1})$$

where we assume $\alpha_k^R > 0$ for $k = 1, 2$. Afterwards, the agent's preference may change, which we model as a change in the weight of attribute $\tilde{\theta}_2$ occurring with (ex ante) probability $p \in (0, 1)$. In the second stage, when making the decision, the agent has utility function

$$u^D(d, \theta) = \begin{cases} - (d - \alpha_1^R \theta_1 - \alpha_2^R \theta_2)^2 & \text{with probability } 1 - p, \\ - (d - \alpha_1^R \theta_1 - c \alpha_2^R \theta_2)^2 & \text{with probability } p, \end{cases}$$

where $c > 0$ and $c \neq 1$.³⁶ The naive agent believes her preferences will stay the same across both stages, while the sophisticated agent understands they might change. Throughout this section, we simplify the analysis by assuming $\mu^0 = (0, 0)$, so that, absent any new information, the researcher and both types of decision-makers agree on the optimal decision $d^* = 0$. Additionally, we fix $\Sigma_1^0 = \Sigma_2^0 = T = 1$.

As the naive agent incorrectly assumes no possibility of change, the naive's learning problem coincides with that of a single player in Section 3, and the equilibrium test allocation follows from Theorem 1 with weights α^R . In contrast, for the sophisticate, a strategic game similar to that in Section 4 ensues. However, Theorem 2 does not apply directly, as the decision-maker's preferences are stochastic. By appropriately adapting Lemma 2 and then applying the same logic as in the proof of Theorem 2, we find that the sophisticate's equilibrium test allocation can nevertheless be expressed as the solution to a single-player problem with auxiliary weights³⁷

$$(\hat{\alpha}_1, \hat{\alpha}_2) := \left(\alpha_1^R, \alpha_2^R \sqrt{\max\{1 - p + p(c(2 - c)), 0\}} \right).$$

Thus, the potential change in preferences reduces the weight put on attribute $\tilde{\theta}_2$, as the sophisticate anticipates the misalignment between his current and future self. Notably, the next result shows that this reduction may even induce the sophisticate to engage in “strategic ignorance” (see, e.g., Carrillo and Mariotti, 2000).

Proposition B.2.1. The sophisticate avoids learning about attribute $\tilde{\theta}_2$ for any $T > 0$ if and only if $\sqrt{p}(c - 1) \geq 1$.

In the presence of a potential shift in preferences, represented by c , the sophisticate may choose strategic ignorance regarding attribute $\tilde{\theta}_2$. If the change reduces the weight on $\tilde{\theta}_2$ or does not increase it too much, the agent still learns about it, but with less intensity ($\hat{\alpha}_2 < \alpha_2^R$). In contrast, the sophisticate avoids learning about attribute $\tilde{\theta}_2$ if the change is sufficiently high ($c > 2$) and the probability of change is not too small. In short, when the sophisticate expects the future self to overreact to information about $\tilde{\theta}_2$, she may opt for strategic ignorance.

³⁶This setting extends the baseline model by allowing for uncertainty about the decision-maker's preferences. An alternative interpretation is that there are multiple potential decision-makers, and the researcher is uncertain about who will make the decision. We extend our equilibrium concept to this setting in Section B.2.1.

³⁷The formal steps are contained in the proof of Proposition B.2.1.

Having established the differences between the naif and the sophisticate’s equilibrium learning strategies, we want to understand the implications for the agent’s welfare. In this context, the choice of the appropriate welfare criterion is unclear. For instance, if the change in preferences is triggered by succumbing to temptation, the initial utility function u^R seems most appropriate. However, if the change in preferences is due to altered circumstances, such as illness or parenthood, then the changed utility function u^D seems more appropriate. We remain agnostic about the welfare criterion and consider both alternatives.

Proposition B.2.2. If $\alpha_1^R \notin (\hat{\alpha}_2/2, 2\alpha_2^R)$, the expected payoff of the naif and the sophisticate coincide for both welfare criteria. Otherwise, if the welfare criterion is the initial utility, u^R , the welfare of the sophisticate exceeds that of the naif. If the welfare criterion is the changed utility, u^D , the naif’s welfare exceeds that of the sophisticate for $c > 1$ and vice versa for $c < 1$.

The interesting case arises when a potential change in preferences influences equilibrium learning strategies. As shown in the appendix, the sophisticate allocates (at least weakly) more tests to attribute $\tilde{\theta}_1$ than the naif. Anticipating a shift in preferences, the sophisticate reduces emphasis on attribute $\tilde{\theta}_2$ to avoid misalignment with the future self’s choices. When initial utility is used as the welfare criterion, the sophisticate’s expected payoff is higher because they optimize learning, aware of possible preference changes, while the naif makes suboptimal choices. However, when the welfare criterion is the changed utility, the naif can outperform the sophisticate. If $c > 1$, the naif learns more about $\tilde{\theta}_2$, benefiting whether preferences change or not: If a change occurs, the naif has learned more about the now more important attribute $\tilde{\theta}_2$ than the sophisticate (in extreme case, the sophisticate engages in strategic ignorance and learns nothing about $\tilde{\theta}_2$); if no change occurs, the naif has implemented the optimal test allocation due to her naivete. In contrast, the sophisticate’s “hedging” learning strategy against the potential change, which underweighs $\tilde{\theta}_2$, is less effective. However, this is reversed when $c < 1$. Then, the naif’s learning strategy is still optimal when no change occurs but (substantially) worse when a change occurs.

The result in Proposition B.2.2 has a parallel to O’Donoghue and Rabin’s (1999) findings in their “doing it now or later” framework. They find that in case of immediate costs the sophisticate is always better off than the naif (as in our case with the initial

utility). Conversely, in case of immediate gratification, either type of the agent can be better off (as in our case with changed utility). In both, O’Donoghue and Rabin (1999) and our model, the sophisticate may engage in a form of overcompensation. In our case it happens by putting less weight on attribute $\tilde{\theta}_2$ to hedge against preference changes, which leads to a worse outcome *irrespective* of whether preferences actually change. Summarizing, both models highlight that naive strategies can sometimes lead to better outcomes if future preferences or circumstances validate those strategies. This underscores the complexity of modeling changing preferences and how different assumptions about future behavior or the reason for change can impact welfare.

B.2.1 Proofs for the extension on changing preferences

Equilibrium concept. A weak Perfect Bayesian Equilibrium of the game with potentially changing preferences (hereinafter referred to simply as “equilibrium”) is a tuple $(\tau^*, d_c^*(s, \tau), d_u^*(s, \tau), \tilde{\theta}(s, \tau))$ such that:

1. the researcher’s test allocation strategy $\tau^* \in \mathcal{T}$ maximizes his expected payoff given the decision-maker’s strategies $d_c^*(s, \tau)$ and $d_u^*(s, \tau)$;
2. the decision-maker’s decision strategies $d_c^*(s, \tau) : \mathbb{R}^2 \times \mathcal{T} \rightarrow \mathbb{R}$ and $d_u^*(s, \tau) : \mathbb{R}^2 \times \mathcal{T} \rightarrow \mathbb{R}$ maximize her expected payoff given changed and unchanged preferences, respectively, and her posterior beliefs $\tilde{\theta}(s, \tau)$;
3. the decision-maker’s posterior beliefs $\tilde{\theta}(s, \tau) : \mathbb{R}^2 \times \mathcal{T} \rightarrow \Delta(\mathbb{R}^2)$ are obtained via Bayes’ rule given the signal realizations s and the researcher’s choice τ .

Proof of Proposition B.2.1. Observe that the decision-maker’s equilibrium decision is given by

$$\begin{aligned} d_u^*(s, \tau) &= \alpha_1^R \hat{\mu}_1(s, \tau) + \alpha_2^R \hat{\mu}_2(s, \tau) \quad \text{with probability } 1 - p, \\ d_c^*(s, \tau) &= \alpha_1^R \hat{\mu}_1(s, \tau) + c\alpha_2^R \hat{\mu}_2(s, \tau) \quad \text{with probability } p. \end{aligned}$$

Thus, for any test allocation τ , the uncertainty about the decision-maker’s preferences translates accordingly into uncertainty about the decision-maker’s decision. Hence, we

can write the researcher's value as

$$V^R(\tau) = -\mathbb{E} \left[p(\alpha_1^R \hat{\mu}_1(s, \tau) + c\alpha_2^R \hat{\mu}_2(s, \tau) - \tilde{b}^R)^2 + (1-p)((\alpha_1^R \hat{\mu}_1(s, \tau) + \alpha_2^R \hat{\mu}_2(s, \tau)) - \tilde{b}^R)^2 \right].$$

By appropriately adapting Lemma 2, we obtain that the equilibrium test allocation is found by solving

$$\max_{\tau \in \mathcal{T}} -p \left(\frac{(\alpha_1^R)^2 \Sigma_1^0}{1 + \tau_1 \Sigma_1^0} + \frac{(\alpha_2^R)^2 \Sigma_2^0}{1 + \tau_2 \Sigma_2^0} (c^2 + 2c\Sigma_2^0 \tau_2) \right) + (1-p) \left(\sum_{j=1}^2 \frac{(\alpha_j^R)^2 \Sigma_j^0}{1 + \tau_j \Sigma_j^0} \right).$$

Proceeding as in the benchmark model, we can obtain the FOCs for the learning problem.

These are given by

$$\begin{aligned} \left(\frac{\Sigma_1^0}{1 + \Sigma_1^0 \tau_1^*} \right)^2 (\alpha_1^R)^2 &= \lambda, \\ \left(\frac{\Sigma_2^0}{1 + \Sigma_2^0 \tau_2^*} \right)^2 (\alpha_1^R)^2 (1 - p + p(c(2 - c))) &= \lambda. \end{aligned}$$

By the same logic as in the proof of Theorem 2, the result follows. \square

Proof of Proposition B.2.2. The proof proceeds in three steps.

Step 1. We begin by observing that $\alpha_2^R \geq \hat{\alpha}_2 = \alpha_2^R \sqrt{\max\{1 - p + p(c(2 - c)), 0\}}$, where $1 - p + p(c(2 - c))$ is a concave function of c with a maximum at $c = 1$. Thus, $\alpha_2^R \geq \hat{\alpha}_2^S$. Further, the sophisticate's and the naif's learning strategies τ_1^S and τ_N , respectively, read

$$\tau_1^S = \max \left\{ 0, \min \left\{ \frac{2\alpha_1^R - \hat{\alpha}_2}{\alpha_1^R + \hat{\alpha}_2}, 1 \right\} \right\}, \quad \tau_1^N = \max \left\{ 0, \min \left\{ \frac{2\alpha_1^R - \alpha_2^R}{\alpha_1^R + \alpha_2^R}, 1 \right\} \right\}.$$

Hence, $\tau_1^S \geq \tau_1^N$, as the sophisticate puts less weight on the second attribute than the naif and they both put the same weight on the first. Thus, if naif learns only about attribute $\tilde{\theta}_1$, the sophisticate will do the same. The condition for that is $\frac{2\alpha_1^R - \alpha_2^R}{\alpha_1^R + \alpha_2^R} \geq 1 \Leftrightarrow \alpha_1^R \geq 2\alpha_2^R$. Analogously, when the sophisticate learns only about attribute $\tilde{\theta}_2$, the naif will do the same. The condition for that is $\frac{2\alpha_1^R - \hat{\alpha}_2}{\alpha_1^R + \hat{\alpha}_2} \leq 0 \Leftrightarrow \alpha_1^R \leq \frac{\hat{\alpha}_2}{2}$. Hence, when $\alpha_1^R \notin (\hat{\alpha}_2/2, 2\alpha_2^R)$ the two learning strategies coincide so that the expected payoffs of the naif and the sophisticate coincide. This proves the first part of the proposition.

Step 2. For the rest of the proof, let $\alpha_1^R \in (\hat{\alpha}_2/2, 2\alpha_2^R)$. If the relevant welfare criterion is the initial utility (given the decision-maker's equilibrium decision rule), it coincides with the objective function the sophisticate maximizes. Hence, the second part of the proposition follows trivially as the learning strategies for the naif and the sophisticate differ when $\alpha_1^R \in (\hat{\alpha}_2/2, 2\alpha_2^R)$.

Step 3. If the relevant welfare criterion is the decision-maker's utility, the welfare function (given the decision-maker's equilibrium decision rule) is

$$\begin{aligned} V^D(\tau) &= p \left[-\frac{(\alpha_1^R)^2}{1+\tau_1} - \frac{(c\alpha_2^R)^2}{1+\tau_2} \right] + (1-p) \left[-\frac{(\alpha_1^R)^2}{1+\tau_1} - \frac{(\alpha_2^R)^2}{1+\tau_2} \right] \\ &= -\frac{(\alpha_1^R)^2}{1+\tau_1} - \frac{(\alpha_2^R)^2}{2-\tau_1} (1-p+pc^2). \end{aligned}$$

Let us find the value of $\tau_1 \in [0, 1]$ that maximizes this function. The first and second order conditions show that $V^D(\tau)$ is a strictly concave function for $\tau_1 \in [0, 1]$ reaching a unique maximum at

$$\tau_1^D = \max \left\{ 0, \min \left\{ \frac{2\alpha_1^R - \alpha_2^R \sqrt{1-p+pc^2}}{\alpha_1^R + \alpha_2^R \sqrt{1-p+pc^2}}, 1 \right\} \right\}.$$

Since $\alpha_1^R \in (\hat{\alpha}_2/2, 2\alpha_2^R)$, the learning strategies of the naif and sophisticate differ. Furthermore, from Step 1, we then get that the sophisticate always learns strictly more about attribute $\tilde{\theta}_1$: $\tau_1^S > \tau_1^N$.

Note that a function $f(x) := \frac{2\alpha_1^R - x}{\alpha_1^R + x}$ is strictly decreasing in $x \in \mathbb{R}$. Furthermore, since $\alpha_1^R \in (\hat{\alpha}_2/2, 2\alpha_2^R)$, we have

$$\begin{aligned} \tau_1^N &= \max\{f(\alpha_2^R), 0\}, \\ \tau_1^S &= \min\{f(\hat{\alpha}_2), 1\}, \\ \tau_1^D &= \max\{0, \min\{f(\alpha_2^R \sqrt{p+(1-p)c^2}), 1\}\}. \end{aligned}$$

Let $c > 1$. Then $\tau_1^S > \tau_1^N \geq \tau_1^D$ since $\sqrt{1-p+pc^2} > \sqrt{p+(1-p)} = 1$ and thus $\alpha_2^R \sqrt{1-p+pc^2} > \alpha_2^R$. Hence the naif's learning strategy is strictly closer to the maximizer. Due to the strict concavity of $V^D(\tau)$, the naif is strictly better off compared to the sophisticate.

Let $c < 1$. Then $\tau_1^D \geq \tau_1^S > \tau_1^N$. To see that, we show that $\hat{\alpha}_2 > \alpha_2^R \sqrt{1 - p + pc^2}$. Note that we have $\hat{\alpha}_2 = \alpha_2^R \sqrt{1 - p + pc(2 - c)}$, which follows from $c < 1$ and $p > 0$ and the calculations in Step 1. Hence, the auxiliary weight of the sophisticate on attribute $\tilde{\theta}_2$ is not zero. The inequality $\hat{\alpha}_2 > \alpha_2^R \sqrt{1 - p + pc^2}$ then follows by noting that $c^2 = c \times c < c(2 - c)$ for $c < 1$. Hence the sophisticate's learning strategy is strictly closer to the maximizer. Due to the strict concavity of $V^D(\tau)$, we get that the sophisticate is strictly better off compared to the naif. \square

C Additional Proofs

C.1 Proof of Proposition 1

The Baseline Model and Framework A are equivalent up to realized payoffs, since the utility functions (15) and (16) are equivalent up to an exogenous random term. To see this, expand the squares in both utility functions, using the assumption that $\sum_k \alpha_k^i = 1$:

$$\begin{aligned} u^i(d, \theta) &= - \left(d - \sum_k \alpha_k^i \theta_k \right)^2 = -d^2 + 2 \sum_k d \alpha_k^i \theta_k - \left(\sum_k \alpha_k^i \theta_k \right)^2, \\ u_A^i(d, \theta) &= - \sum_k \alpha_k^i (d - \theta_k)^2 = -d^2 + 2 \sum_k d \alpha_k^i \theta_k - \sum_k \alpha_k^i \theta_k^2. \end{aligned}$$

Observe that the difference $u^i(d, \theta) - u_A^i(d, \theta)$ is independent of the decision d . Therefore, the decision-maker's expected utilities (3) only differ in the two settings by a constant independent of d , and the researcher's values (2) only differ by a constant.

Now consider Framework B. The decision-maker's equilibrium decision $d^*(s, \tau)$ in this setting satisfies

$$d_k^*(s, \tau) = \mathbb{E} \left[\alpha_k^D \tilde{\theta}_k \right] = \mathbb{E} \left[\alpha_k^D \tilde{\theta}_k(s_k, \tau_k) \right] = \alpha_k^D \hat{\mu}_k(s_k, \tau_k)$$

for all k , where $\tilde{\theta}_k(s_k, \tau_k)$ is the decision-maker's posterior belief and $\hat{\mu}_k(s_k, \tau_k)$ is its posterior expectation, given by (4) and (5), respectively. Let $\tilde{d}_k^*(\tau) := d_k^*(\tilde{s}, \tau)$ for all k denote the random variable that captures the distribution of the decision-maker's equilibrium decision given τ . The researcher's value is then given by

$$\begin{aligned} V_B^R(\tau) &= \mathbb{E} \left[u_B^R \left(\tilde{d}^*(\tau), \tilde{\theta} \right) \right] = - \mathbb{E} \left[\sum_k \left(\tilde{d}_k^*(\tau) - \alpha_k^R \tilde{\theta}_k \right)^2 \right] \\ &= \sum_k \left[-Var \left(\tilde{d}_k^*(\tau) \right) + 2 \text{cov} \left(\tilde{d}_k^*(\tau), \alpha_k^R \tilde{\theta}_k \right) \right] + C, \end{aligned}$$

where C is some constant that does not depend on τ (note, specifically, that $\mathbb{E} \left[\tilde{d}_k^*(\tau) \right] = \alpha_k^D \mu_k^0$ does not depend on τ). This is analogous (though not equivalent) to Lemma 2.

From (5) and (1) we have that

$$\begin{aligned} \text{Var} \left(\tilde{d}_k^*(\tau) \right) &= \text{Var} \left(\alpha_k^D \frac{\tau_k \Sigma_k^0}{1 + \tau_k \Sigma_k^0} \tilde{s}_k \right) = \frac{\tau_k (\Sigma_k^0)^2}{1 + \tau_k \Sigma_k^0} (\alpha_k^D)^2, \\ \text{cov} \left(\tilde{d}_k^*(\tau), \alpha_k^R \tilde{\theta}_k \right) &= \text{cov} \left(\alpha_k^D \frac{\tau_k \Sigma_k^0}{1 + \tau_k \Sigma_k^0} \tilde{s}_k, \alpha_k^R \tilde{\theta}_k \right) = \frac{\tau_k (\Sigma_k^0)^2}{1 + \tau_k \Sigma_k^0} \alpha_k^D \alpha_k^R, \end{aligned}$$

hence the objective function above reduces to

$$V_B^R(\tau) = \sum_k \frac{\tau_k (\Sigma_k^0)^2}{1 + \tau_k \Sigma_k^0} (2\alpha_k^R - \alpha_k^D) \alpha_k^D + C$$

and its partial derivative w.r.t. τ_k is given by

$$\frac{\partial V_B^R(\tau)}{\partial \tau_k} = (2\alpha_k^R - \alpha_k^D) \alpha_k^D \left(\frac{1}{\Sigma_k^0} + \tau_k \right)^{-2},$$

which is equivalent to (30). The remainder of the proof of Theorem 2 then applies to Framework B, and the statement of Theorem 2 holds. The equilibrium test allocation is, therefore, the same as in the Baseline Model. \square

C.2 Details of the Proofs in Appendix A.3

C.2.1 Proof of Lemma 5

Fix the manager's weight vector $\alpha^D = (\alpha_1^D, \alpha_2^D)$ with $\alpha_1^D \geq \alpha_2^D > 0$. Let $\alpha^R(\beta, \gamma)$ denote the analyst's weight vector as a function of β and γ for a given α^D as defined by (18).³⁸ We thus have

$$\alpha_1^R(\gamma, \beta) = \beta \alpha_1^D - \gamma \alpha_2^D, \quad \alpha_2^R(\gamma, \beta) = \beta \alpha_2^D + \gamma \alpha_1^D. \quad (\text{C.2.1})$$

For $\beta \in \mathbb{R}_+$, let $\gamma_1(\beta) \in \mathbb{R}$ and $\gamma_2(\beta) \in \mathbb{R}$ denote the values of γ defined by equations $\alpha_1^R(\beta, \gamma_1(\beta)) = 1/2\alpha_1^D$, and $\alpha_2^R(\beta, \gamma_2(\beta)) = 1/2\alpha_2^D$. From equations (C.2.1), we get

$$\gamma_1(\beta) = - \left(\frac{1}{2} - \beta \right) \frac{\alpha_1^D}{\alpha_2^D}, \quad \gamma_2(\beta) = \left(\frac{1}{2} - \beta \right) \frac{\alpha_2^D}{\alpha_1^D} \quad (\text{C.2.2})$$

Since $\alpha_1^R(\beta, \gamma)$ ($\alpha_2^R(\beta, \gamma)$) is strictly decreasing (strictly increasing) in γ , the following results hold: Given $\beta \in \mathbb{R}_+$ (a) $\alpha_1^R(\beta, \gamma) \leq 1/2\alpha_1^D$ if and only if $\gamma \geq \gamma_1(\beta)$; and (b)

³⁸For simplicity, we omit the dependence on α^D in the notation.

$\alpha_2^R(\beta, \gamma) \leq 1/2\alpha_2^D$ if and only if $\gamma \leq \gamma_2(\beta)$.

Note that $\gamma_1(\beta)$ ($\gamma_2(\beta)$) is strictly increasing (strictly decreasing) in β . Furthermore, whenever $\beta < 1/2$ we have $\gamma_1(\beta) < 0 < \gamma_2(\beta)$ and for $\beta = 1/2$ we have $\gamma_1(\beta) = 0 = \gamma_2(\beta)$. Taking $\beta \leq 1/2$, combining results (a) and (b) we can then apply Theorem 2 and obtain the analyst's equilibrium test allocation as

$$\tau^*(\beta, \gamma) = (\tau_1^*(\beta, \gamma), \tau_2^*(\beta, \gamma)) = \begin{cases} (T, 0) & \text{if } \gamma < \gamma_1(\beta) \\ (0, 0) & \text{if } \gamma \in [\gamma_1(\beta), \gamma_2(\beta)] \\ (0, T) & \text{if } \gamma > \gamma_2(\beta) \end{cases} \quad (\text{C.2.3})$$

For each β , however, γ also needs to satisfy the constraint $\gamma \in \Gamma(\beta)$. For $\beta \in \mathbb{R}_+$, let

$$\underline{\gamma}(\beta) := -\beta \frac{\alpha_2^D}{\alpha_1^D}, \quad \bar{\gamma}(\beta) := \beta \frac{\alpha_1^D}{\alpha_2^D} \quad (\text{C.2.4})$$

denote the end points of the constraint $\Gamma(\beta)$. Let $\beta_1, \beta_2 \in \mathbb{R}_+$ denote the values of β defined by equations $\gamma_1(\beta_1) = \underline{\gamma}(\beta_1)$ and $\gamma_2(\beta_2) = \bar{\gamma}(\beta_2)$. Solving these equations, we obtain

$$\beta_1 = \frac{1}{2} \frac{(\alpha_1^D)^2}{(\alpha_1^D)^2 + (\alpha_2^D)^2}, \quad \beta_2 = \frac{1}{2} \frac{(\alpha_2^D)^2}{(\alpha_1^D)^2 + (\alpha_2^D)^2}$$

Since $\alpha_1^D \geq \alpha_2^D > 0$, we get $0 < \beta_2 \leq \beta_1 < 1/2$. Furthermore, since $\gamma_1(\beta)$ ($\gamma_2(\beta)$) is strictly increasing (strictly decreasing) in β and since $\underline{\gamma}(\beta)$ ($\bar{\gamma}(\beta)$) is strictly decreasing (strictly increasing) we get the following results: (i) for all $\gamma \in \Gamma(\beta)$ we have $\gamma \geq \gamma_1(\beta)$ if and only if $\beta \leq \beta_1$; and (ii) for all $\gamma \in \Gamma(\beta)$ we have $\gamma \leq \gamma_2(\beta)$ if and only if $\beta \leq \beta_2$. Combining points (i) and (ii) with the equilibrium test allocation defined in equation (C.2.3), we obtain Lemma 5. \square

C.2.2 Proof of Lemma 6

Fix the manager's weight vector $\alpha^D = (\alpha_1^D, \alpha_2^D)$ with $\alpha_k^D > 0$ for both $k = 1, 2$. Let $\alpha_1^R(\beta, \gamma)$, $\alpha_2^R(\beta, \gamma)$, $\gamma_1(\beta)$, $\gamma_2(\beta)$, $\underline{\gamma}(\beta)$, and $\bar{\gamma}(\beta)$ be given by equations (C.2.1), (C.2.2), and (C.2.4) in the proof of Lemma 5.

Fix $\beta > 1/2$. Then we have $\underline{\gamma}(\beta) < \gamma_2(\beta) < 0 < \gamma_1(\beta) < \bar{\gamma}(\beta)$. As shown in the

proof of Lemma 5, it holds that $\alpha_1^R(\beta, \gamma) \leq 1/2\alpha_1^D$ if and only if $\gamma \geq \gamma_1(\beta)$; and $\alpha_2^D(\beta, \gamma) \leq 1/2\alpha_2^D$ if and only if $\gamma \leq \gamma_2(\beta)$. By Theorem 2, the analyst's equilibrium test allocation satisfies

$$\tau^*(\beta, \gamma) = (\tau_1^*(\beta, \gamma), \tau_2^*(\beta, \gamma)) = \begin{cases} (T, 0) & \text{if } \gamma \leq \gamma_2(\beta), \\ (0, T) & \text{if } \gamma \geq \gamma_1(\beta). \end{cases} \quad (\text{C.2.5})$$

Now suppose $\gamma \in [\gamma_2(\beta), \gamma_1(\beta)]$. Let $\tilde{\tau}_k(\beta, \gamma)$ and $\hat{\alpha}_k(\beta, \gamma)$ for $k = 1, 2$ be given by equations (33) and (34). Note that expressions in (34) are well-defined for $\gamma \in [\gamma_2(\beta), \gamma_1(\beta)]$.

Furthermore, we have

$$\tilde{\tau}_1(\beta, \gamma) = \begin{cases} \frac{1}{\Sigma_2^0} + T > T & \text{at } \gamma = \gamma_2(\beta), \\ -\frac{1}{\Sigma_1^0} < 0 & \text{at } \gamma = \gamma_1(\beta), \end{cases}$$

since $\hat{\alpha}_2(\beta, \gamma_2(\beta)) = 0$, $\hat{\alpha}_1(\beta, \gamma_2(\beta)) > 0$, and $\hat{\alpha}_1(\beta, \gamma_1(\beta)) = 0$, $\hat{\alpha}_2(\beta, \gamma_1(\beta)) > 0$. Further, $\tilde{\tau}_1(\beta, \gamma)$ is continuous and strictly decreasing in γ . To see the strict monotonicity, let us take the partial derivative of $\tilde{\tau}_1(\gamma)$ with respect to γ . For ease of notation, we suppress the dependence on β and γ from the right-hand side of the expression. We get

$$\frac{\partial \tilde{\tau}_1(\beta, \gamma)}{\partial \gamma} = \frac{(\Sigma_1^0 + \Sigma_2^0) \left(\hat{\alpha}_2 \frac{\partial \hat{\alpha}_1}{\partial \gamma} - \hat{\alpha}_1 \frac{\partial \hat{\alpha}_2}{\partial \gamma} \right)}{\Sigma_1^0 \Sigma_2^0 (\hat{\alpha}_1 + \hat{\alpha}_2)^2} + \frac{\hat{\alpha}_2 \frac{\partial \hat{\alpha}_1}{\partial \gamma} - \hat{\alpha}_1 \frac{\partial \hat{\alpha}_2}{\partial \gamma}}{(\hat{\alpha}_1 + \hat{\alpha}_2)^2} < 0,$$

where we used $\frac{\partial \hat{\alpha}_1(\beta, \gamma)}{\partial \gamma} = -\frac{\alpha_2^D}{\sqrt{\hat{\alpha}_1(\beta, \gamma)}} < 0$ and $\frac{\partial \hat{\alpha}_2(\beta, \gamma)}{\partial \gamma} = \frac{\alpha_1^D}{\sqrt{\hat{\alpha}_2(\beta, \gamma)}} > 0$ to determine the sign. By the intermediate value theorem, there exist $\gamma_I(\beta), \gamma_{II}(\beta) \in (\gamma_2(\beta), \gamma_1(\beta))$ such that

$$\tilde{\tau}_1(\beta, \gamma) = \begin{cases} T & \text{at } \gamma = \gamma_I(\beta), \\ 0 & \text{at } \gamma = \gamma_{II}(\beta). \end{cases} \quad (\text{C.2.6})$$

From the strict monotonicity of $\tilde{\tau}_1(\beta, \gamma)$, it further holds that

$$\tilde{\tau}_1(\beta, \gamma) \begin{cases} > T & \text{if and only if } \gamma < \gamma_I(\beta), \\ < 0 & \text{if and only if } \gamma > \gamma_{II}(\beta). \end{cases} \quad (\text{C.2.7})$$

Finally, note that $\hat{\alpha}_k(\beta, \gamma)$ as given by (34) are equivalent to (13) and (14) given de-

composition (18). By Theorems 1 and 2, the analyst's equilibrium test allocation for $\gamma \in [\gamma_2(\beta), \gamma_1(\beta)]$ is

$$\tau_1^*(\beta, \gamma) = \max \{0, \min \{\tilde{\tau}_1(\beta, \gamma), T\}\} \quad (\text{C.2.8})$$

$$\tau_2^*(\beta, \gamma) = \max \{0, \min \{\tilde{\tau}_2(\beta, \gamma), T\}\} \quad (\text{C.2.9})$$

Combining the results (C.2.8), (C.2.9) with (C.2.7) and (C.2.5), together with the strict monotonicity result of $\tilde{\tau}_1(\beta, \gamma)$, we obtain the claim of Lemma 6. \square

C.2.3 Proof of Lemma 7

Given a test allocation $\tau = (\tau_1, \tau_2)$, the decision-maker's interim expected payoff is

$$V^D(\tau) = -\frac{(\alpha_1^D)^2 \Sigma_1^0}{1 + \tau_1 \Sigma_1^0} - \frac{(\alpha_2^D)^2 \Sigma_2^0}{1 + \tau_2 \Sigma_2^0}. \quad (\text{C.2.10})$$

The first, the second, and the cross derivatives of the decision-maker's interim expected payoff with respect to $\tau_k \geq 0$ and $\tau_l \neq \tau_k$ are $\frac{\partial V^D(\tau)}{\partial \tau_k} = \left(\frac{\alpha_k^D \Sigma_k^0}{1 + \tau_k \Sigma_k^0}\right)^2 > 0$, $\frac{\partial^2 V^D(\tau)}{\partial \tau_k^2} = -2(\alpha_k^D)^2 \left(\frac{\Sigma_k^0}{1 + \tau_k \Sigma_k^0}\right)^3 < 0$, and $\frac{\partial^2 V^D(\tau)}{\partial \tau_k \partial \tau_l} = 0$. Hence, $V^D(\tau)$ is strictly increasing and strictly concave function. \square

C.2.4 Proof of Lemma 8

Fix the test budget $T > 0$ and the decision-maker's weight vector $\alpha^D \in \mathbb{R}_{++}^2$. Suppose $\gamma = 0$ (the agent is not distorted).

- (i) If $\beta \leq 1/2$ (the agent is too insensitive), then the agent optimally chooses no testing: $\tau^*(\beta, 0) = (0, 0)$.
- (ii) If $\beta > 1/2$ (the agent is sufficiently sensitive), then the agent optimally chooses the decision-maker's most preferred test allocation: $\tau^*(\beta, 0) = \tau^*(1, 0)$.

Suppose the premise of Lemma 8 holds. Part (i) immediately follows from Lemma 5. For part (ii) take $\beta > 1/2$ and let $\tilde{\tau}_1(\beta, \gamma)$ and $\tilde{\tau}_2(\beta, \gamma)$ be given by (33). At $\gamma = 0$, we

get

$$\begin{aligned}\tilde{\tau}_1(\beta, 0) &= \frac{\sqrt{2\beta-1}\alpha_1^D\Sigma_1^0 - \sqrt{2\beta-1}\alpha_2^D\Sigma_2^0}{\Sigma_1^0\Sigma_2^0\sqrt{2\beta-1}(\alpha_1^D + \alpha_2^D)} + \frac{\sqrt{2\beta-1}\alpha_1^D}{\sqrt{2\beta-1}(\alpha_1^D + \alpha_2^D)}T \\ &= \frac{\alpha_1^D\Sigma_1^0 - \alpha_2^D\Sigma_2^0}{\Sigma_1^0\Sigma_2^0(\alpha_1^D + \alpha_2^D)} + \frac{\alpha_1^D}{\alpha_1^D + \alpha_2^D}T\end{aligned}$$

Lemma 6 and Theorem 1 then imply that the agent's equilibrium test allocation coincides with the optimal test allocation of a single player with weight vector $\alpha = \alpha^D$. \square

C.3 Details of the Proofs in Appendix A.4

C.3.1 Proof of Lemma 9

Let us set $T = 1$ and $\Sigma_1^0 = \Sigma_2^0 = 1$. From Theorems 1 and 2, the equilibrium test allocation satisfies $\tau_1^*(p, \delta) = 1 - \tau_2^*(p, \delta)$ and

$$\tau_2^*(p, \delta) = \min \left\{ \frac{\hat{\alpha}_2(p, \delta) - \hat{\alpha}_1(p, \delta)}{\hat{\alpha}_1(p, \delta) + \hat{\alpha}_2(p, \delta)} + \frac{\hat{\alpha}_2(p, \delta)}{\hat{\alpha}_1(p, \delta) + \hat{\alpha}_2(p, \delta)}, 1 \right\} = \min \left\{ \underbrace{\frac{2\hat{\alpha}_2(p, \delta) - \hat{\alpha}_1(p, \delta)}{\hat{\alpha}_1(p, \delta) + \hat{\alpha}_2(p, \delta)}}_{\tau_2^{aux}(p, \delta)}, 1 \right\}$$

where $\hat{\alpha}_1(p, \delta)$ and $\hat{\alpha}_2(p, \delta)$ are given by equations (36) and (37), respectively. The function $\tau_2^{aux}(p, \delta)$ is well-defined, since $\hat{\alpha}_1(p, \delta) + \hat{\alpha}_2(p, \delta) \neq 0$ for all values of p and δ . Note that $\tau_2^{aux}(p, \delta) = 1/2$ when $p = 0$ and $\tau_2^{aux}(p, \delta) = 2$ when $p \geq \frac{1-\delta}{2}$. Furthermore, when $p \in [0, \frac{1-\delta}{2})$ we have

$$\begin{aligned}\frac{\partial \tau_2^{aux}(p, \delta)}{\partial p} &= \frac{\left(2\frac{\partial \hat{\alpha}_2}{\partial p} - \frac{\partial \hat{\alpha}_1}{\partial p}\right)(\hat{\alpha}_1 + \hat{\alpha}_2) - (2\hat{\alpha}_2 - \hat{\alpha}_1)\left(\frac{\partial \hat{\alpha}_1}{\partial p} + \frac{\partial \hat{\alpha}_2}{\partial p}\right)}{(\hat{\alpha}_1 + \hat{\alpha}_2)^2} \\ &= \frac{3}{4} \frac{(1-\delta^2)}{\hat{\alpha}_1\hat{\alpha}_2(\hat{\alpha}_1 + \hat{\alpha}_2)^2} > 0.\end{aligned}\tag{C.3.1}$$

Note that $\tau_2^{aux}(p, \delta)$ is continuous in p for $p \geq 0$.³⁹ Furthermore, $\tau_2^{aux}(p, \delta)$ is strictly increasing in p on $p \in [0, (1-\delta)/2)$. At $p = 0$ we have $\tau_2^{aux}(p, \delta) = 1/2 < 1$ and at $p \geq (1-\delta)/2$ we have $\tau_2^{aux}(p, \delta) = 2 > 1$. Hence, for each value of discrimination $\delta \in (0, 1)$ there exists a threshold partiality level of the advisor $p^{aux}(\delta) \in (0, \frac{1-\delta}{2})$, defined implicitly by the equation $\tau_2^{aux}(p^{aux}(\delta), \delta) = 1$, such that the equilibrium testing strategy

³⁹Since $\hat{\alpha}_1(p, \delta)$ and $\hat{\alpha}_2(p, \delta)$ are both continuous in p , the only possible point of discontinuity of $\tau_2^{aux}(p, \delta)$ is if $\hat{\alpha}_1(p, \delta) + \hat{\alpha}_2(p, \delta) = 0$. Equations (36) and (37) imply this never happens for all feasible values of p and δ . Hence, $\tau_2^{aux}(p, \delta)$ is continuous on $p \geq 0$.

satisfies $\tau_2^*(p, \delta) = \tau_2^{aux}(p, \delta)$ for $p \in [0, p^{aux}(\delta)]$ and $\tau_2^*(p, \delta) = 1$ for $p \geq p^{aux}(\delta)$. \square

C.3.2 Proof of Lemma 10

The ex ante expected utility of group k , $V^k(\tau, \delta)$, can be derived from the expected payoff of a researcher $V^R(\tau)$ in the baseline model in the proof of Lemma 2, where we set the decision-maker's weights to $\alpha_1^D(\delta)$, $\alpha_2^D(\delta)$ and the researcher's weights (now group k weights) to $\alpha_k^R = 1$ and $\alpha_{-k}^R = 0$. In other words, given allocation τ and the politician's equilibrium strategy $d^*(s, \tau; \delta)$, the ex ante expected payoff of group k is the same as that of the researcher in the baseline model who solely cares about attribute $\tilde{\theta}_k$ (with weight $\alpha_k^R = 1$) and not about the other attribute (with weight $\alpha_{-k}^R = 0$). We get

$$\begin{aligned} V^k(\tau, \delta) &= \mathbb{E} \left[-(\tilde{d}^{*D}(\tau; \delta) - \tilde{\theta}_k)^2 \right] \\ &= - (v_0^D(\delta) - \mu_k^0)^2 - \left(\sigma_0^{2,D}(\delta) + \Sigma_k^0 \right) + \hat{\sigma}^{2,D}(\tau; \delta) + 2 \text{cov} \left(\tilde{d}^{*D}(\tau; \delta), \tilde{\theta}_k \right) \\ &= - (\alpha_1^D(\delta) \mu_1^0 + \alpha_2^D(\delta) \mu_2^0 - \mu_k^0)^2 - \left((\alpha_1^D(\delta))^2 \Sigma_1^0 + (\alpha_2^D(\delta))^2 \Sigma_2^0 + \Sigma_k^0 \right) \\ &\quad + \frac{\Sigma_k^0}{1 + \tau_k \Sigma_k^0} \left((\alpha_k^D(\delta))^2 + 2\alpha_k^D(\delta) \Sigma_k^0 \tau_k \right) + \frac{\Sigma_{-k}^0}{1 + \tau_{-k} \Sigma_{-k}^0} (\alpha_{-k}^D(\delta))^2 \end{aligned}$$

Using $\mu_1^0 = \mu_2^0 = 0$, $\Sigma_1^0 = \Sigma_2^0 = 1$, we get

$$\begin{aligned} V^k(\tau, \delta) &= - \left((\alpha_1^D(\delta))^2 + (\alpha_2^D(\delta))^2 + 1 \right) \\ &\quad + \frac{1}{1 + \tau_k} \left((\alpha_k^D(\delta))^2 + 2\alpha_k^D(\delta) \tau_k \right) + \frac{1}{1 + \tau_{-k}} (\alpha_{-k}^D(\delta))^2. \end{aligned} \quad (\text{C.3.2})$$

The difference between the expected utilities of the two groups is

$$V^1(\tau, \delta) - V^2(\tau, \delta) = \frac{2\alpha_1^P(\delta)\tau_1}{1 + \tau_1} - \frac{2\alpha_2^P(\delta)\tau_2}{1 + \tau_2} = \frac{(1 + \delta)\tau_1}{1 + \tau_1} - \frac{(1 - \delta)\tau_2}{1 + \tau_2}$$

where we used $\alpha_1^P(\delta) = \frac{1}{2}(1 + \delta)$ and $\alpha_2^P(\delta) = \frac{1}{2}(1 - \delta)$. The difference between the expected utilities of the two groups under the additional constraint that $\tau_1 + \tau_2 = 1$ is

$$\Delta(\tau_2, \delta) = V^1((1 - \tau_2, \tau_2), \delta) - V^2((1 - \tau_2, \tau_2), \delta) = \frac{(1 + \delta)(1 - \tau_2)}{1 + 1 - \tau_2} - \frac{(1 - \delta)\tau_2}{1 + \tau_2}.$$

The function $\Delta(\tau_2, \delta)$ is continuous and strictly decreasing in τ_2 since $\delta \in (0, 1)$ and thus $\frac{\partial \Delta(\tau_2, \delta)}{\partial \tau_2} = -\frac{(1 + \delta)}{(1 + 1 - \tau_2)^2} - \frac{(1 - \delta)}{(1 + \tau_2)^2} < 0$. Furthermore, we have $\Delta(1/2, \delta) > 0$ and

$\Delta(1, \delta) < 0$. Hence, for each δ there exists a unique value of $\hat{\tau}_2(\delta) \in (1/2, 1)$ for which it holds $\Delta(\hat{\tau}_2(\delta); \delta) = 0$. \square

C.3.3 Proof of Lemma 11

Let

$$\begin{aligned} \omega(\tau_2, \delta) &:= V^1((1 - \tau_2, \tau_2), \delta) + V^2((1 - \tau_2, \tau_2), \delta) \\ &= -3 - \delta^2 + \frac{1/2(1 + \delta^2) + (1 + \delta)(1 - \tau_2)}{2 - \tau_2} + \frac{1/2(1 + \delta^2) + (1 - \delta)\tau_2}{1 + \tau_2} \end{aligned}$$

denote the sum of the expected utilities of the two groups as a function of test allocation $\tau = (\tau_1, \tau_2)$ under the constraint that the budget is fully used, $\tau_1 = T - \tau_2$ with $T = 1$, and where $V^k(\tau, \delta)$ is given by equation (C.3.2). Since the budget is exhausted in equilibrium (as shown in the proof of Lemma 4), welfare is thus given by $W(p, \delta) = \omega(\tau_2^*(p, \delta), \delta)$. The partial derivative of welfare with respect to $p > 0$, whenever differentiable, is then

$$\frac{\partial W(p, \delta)}{\partial p} = \frac{\partial \tau_2^*}{\partial p} \frac{\partial \omega(\tau_2^*, \delta)}{\partial \tau_2} = \frac{\partial \tau_2^*}{\partial p} \left(\underbrace{\frac{1/2(1 + \delta^2) - (1 + \delta)}{(2 - \tau_2^*)^2}}_{< 0, \text{ since } \delta < 1} + \frac{-1/2(1 + \delta^2) + (1 - \delta)}{(1 + \tau_2^*)^2} \right)$$

where we suppressed the dependence of $\tau_2^*(p, \delta)$ on (p, δ) from the RHS for brevity. Note that $|1/2(1 + \delta^2) - (1 + \delta)| > |-1/2(1 + \delta^2) + (1 - \delta)|$ and that $(2 - \tau_2^*)^2 < (1 + \tau_2^*)^2$, since $\tau_2^*(p, \delta) > 1/2$ for $p > 0$. Hence, $\frac{\partial \omega(\tau_2^*, \delta)}{\partial \tau_2} < 0$. We thus have $\text{sign}\left(\frac{\partial W(p, \delta)}{\partial p}\right) = -\text{sign}\left(\frac{\partial \tau_2^*}{\partial p}\right)$.

Let $\tau_2^{aux}(p, \delta)$, $\hat{p}(\delta)$ and $p^{aux}(\delta)$ with $\hat{p}(\delta) < p^{aux}(\delta)$, be the variables defined in the proof of Lemma 4. From equation (C.3.1), it follows

$$\text{sign}\left(\frac{\partial W(p, \delta)}{\partial p}\right) = -\text{sign}\left(\frac{\partial \tau_2^*(p, \delta)}{\partial p}\right) = \begin{cases} -\text{sign}\left(\frac{\tau_2^{aux}(p, \delta)}{\partial p}\right) < 0 & p \in (0, p^{aux}(\delta)) \\ 0 & p > p^{aux}(\delta) \end{cases},$$

and the derivative does not exist at $p = p^{aux}(\delta)$. We thus have that welfare $W(p, \delta)$ is (i) continuous in p (since $V^k(\tau^*(p, \delta), \delta)$ is continuous in p); (ii) strictly decreasing on $p \in (0, p^{aux}(\delta))$; and (iii) constant on $p \geq p^{aux}(\delta)$, where all the advisors use the entire budget to learn exclusively about group 2. Furthermore, the equalizing partiality level

$$\hat{p}(\delta) < p^{aux}(\delta).$$

Next, inequality is given by

$$I(p, \delta) = |\Delta(\tau_2^*(p, \delta), \delta)|$$

where $\Delta(\tau_2^*(p, \delta), \delta)$ is given by equation (38). As we have shown in the proof of Lemma 4, the function $\Delta(\tau_2^*(p, \delta), \delta)$ is continuous in $\tau_2^*(p, \delta)$, strictly decreasing in $\tau_2^*(p, \delta)$, and is zero at $\tau_2^*(\hat{p}(\delta), \delta)$, where $\tau_2^*(\hat{p}(\delta), \delta) < T$ (the advisor who restores equality, $\hat{p}(\delta)$, does not use the entire test budget to learn exclusively about group 2). As shown above, we have: the equilibrium allocation of tests to group 2, $\tau_2^*(p, \delta)$, is continuous in p , strictly increasing on $p < p^{aux}(\delta)$, and constant on $p \geq p^{aux}(\delta)$, where the advisors use the entire test budget to learn exclusively about group 2: $\tau_2^*(p, \delta) = T$. Hence, inequality $I(p, \delta)$ is (i) continuous in p ; (ii) strictly decreasing on $p \in (0, \hat{p}(\delta))$ (iii) zero at $p = \hat{p}(\delta)$; (iv) strictly increasing on $p \in (\hat{p}(\delta), p^{aux}(\delta))$; and (v) constant on $p \geq p^{aux}(\delta)$.

Therefore, welfare $W(p, \delta)$ and inequality $I(p, \delta)$ are both (i) continuous in p ; and (ii) strictly decreasing on $p \in (0, \hat{p}(\delta))$. Furthermore, welfare strictly decreases and inequality strictly increases on $p \in (\hat{p}(\delta), p^{aux}(\delta))$; and they are both constant on $p \geq p^{aux}(\delta)$.

C.3.4 Proof of Lemma 12

By Theorem 1, the politician's optimal learning strategy under unchecked discrimination yields

$$\bar{\tau}_2(\delta) = \begin{cases} \frac{1-3\delta}{2} & \text{if } \delta < \frac{1}{3}, \\ 0 & \text{if } \delta \geq \frac{1}{3}. \end{cases}$$

From the proof of Lemma 4, equation (39), the test allocation restoring equality $\hat{\tau}_2(\delta) = \frac{-(1-\delta) + \sqrt{1+3\delta^2}}{2\delta}$.

Fix $\delta \in [1/3, 1)$. Then $|\bar{\tau}_2(\delta) - 1/2| > |\hat{\tau}_2(\delta) - 1/2|$, since $\hat{\tau}_2(\delta) \in (1/2, 1)$, as shown in

the proof of Lemma 4. Next, fix $\delta \in (0, 1/3)$. Then we have, equivalently,

$$|\bar{\tau}_2(\delta) - 1/2| = \frac{1}{2} - \frac{1 - 3\delta}{2} > \frac{-1 + \delta + \sqrt{1 + 3\delta^2}}{2\delta} - \frac{1}{2} = |\hat{\tau}_2(\delta) - 1/2|$$
$$1 + 3\delta^2 > \sqrt{1 + 3\delta^2}$$

which holds for any $\delta \in (0, 1/3)$.

□