



WARWICK

ECONOMICS

QAPEC

Quantitative and Analytical Political Economy Research Centre

Discussion Paper Series

The Choice of Political Advisors

Dimitri Mogrow, Hyungmin Park & Francesco Squintani

[This paper also appears as Department working paper 1582](#)

and also

[CRETA Discussion paper 94](#)

Original version October 2025

Revised September 2026

No: 29

QAPEC

Quantitative and Analytical Political Economy Research Centre

Department of Economics
University of Warwick, Coventry,
CV4 7AL, United Kingdom

<https://warwick.ac.uk/fac/soc/economics/research/centres/qapec/discussionpapers>

The Choice of Political Advisors^{*}

Dimitri Migrow
University of Edinburgh[†]

Francesco Squintani
University of Warwick and CEPR[‡]

Hyungmin Park
University of Warwick[§]

September 2026

Abstract

Political leaders choose advisors by balancing political alignment, competence, and diversity of views. We study a leader who can consult one or two strategically communicating experts at a cost. One is more aligned but has less precise or more redundant information; the other is better informed but less aligned. The leader's optimal consultation choice exhibits a reversal. When both are consulted, increasing the better-informed expert's ideological distance first causes the leader to dismiss the other advisor and rely on that expert alone; as the distance grows further, she switches to the more aligned advisor. The same reversal persists with continuous signals, and the dismissal of the more aligned consultant can occur even when informative transmission by both experts remains feasible. When the better-informed expert's alignment is uncertain, a greater probability of alignment can induce consultation of both. We relate these results to advisor turnover across political administrations.

^{*}For helpful comments and discussions we thank Kim Sau Chung, Wouter Dessein, Torun Dewan, Marina Halac, Navin Kartik, Adam Meirowitz, Meg Meyer, Giacomo Ponzetto, Jesse Shapiro, Mike Ting and Leonard Wantchekon, as well as audiences at seminars/conferences, including Barcelona GSE Summer Forum, Columbia University, Harvard University, Kellogg School of Management, University of Michigan, NYU Abu Dhabi, Princeton University, and Yale University. This paper also benefited from conversations with Dr. Davide Taliente, Advisor at the *Ministero per lo Sviluppo Economico*, and then *Ministero delle Imprese e del Made in Italy*, Italian Government. (The views expressed in this paper do not necessarily reflect the views of the Italian Government, or of any of its bodies).

[†]Department of Economics, 30-31 Buccleuch Place, Edinburgh EH8 9JT, UK;
dimitri.migrow@gmail.com

[‡]Department of Economics, Gibbet Hill Road, Coventry CV4 7AL, UK; f.squintani@warwick.ac.uk

[§]Department of Economics, Gibbet Hill Road, Coventry CV4 7AL, UK; hyungmin.park@warwick.ac.uk

1 Introduction

Political leaders rely on a variety of advisors, and the quality of this advice often shapes the success of their decisions. As Niccolò Machiavelli noted in *The Prince* (Ch. 22):

The first opinion that one forms of a prince, and of his understanding, is by observing the men he has around him.

Advisor selection involves a fundamental trade-off. The pool of potential appointees is limited and heterogeneous in political alignment, competence and diversity of views.¹ Should the leader favor aligned advisors or more competent experts? Should she include ideologically distant consultants to gain from diverse perspectives?

These dilemmas have long been recognized by political scholars, and pervade all political appointments, from top executive nominations to the assignment of technical policy experts. In the context of Presidential appointments for example²:

The President is interested in staffing his administration with people personally and politically loyal to him, who will at the same time bring competence and support to the programs he espouses. (Mann, 1964)

The tension between the White House’s need for loyal and committed appointees and the country’s requirement for men and women of exceptional ability to run the executive branch appears to require a trade-off between loyalty and competence. (Edwards, 2001)

More fundamentally, Max Weber in *Economy and Society* (Vol. 1, Part III) identified the same tension as a general leadership problem under informational asymmetry:

Since the specialized knowledge of the expert became more and more the foundation for the power of the officeholder, an early concern of the ruler was how to exploit the special knowledge of experts without having to abdicate in their favor.

¹For example, Lewis (2012) notes that “the pool of loyal and competent persons who are also politically connected in useful ways is limited.”

²Among others, Cohen (1998) argues that appointing loyal politicians may result in an “amateur government,” as such appointees often lack administrative expertise, while Moe (1985) emphasizes the importance of “responsive competence.”

Hence, a natural trade-off between political alignment and competence arises as a systematic constraint on political selection.^{3,4}

Further, democratic leaders must decide how much ideological diversity to seek among the experts they consult. Overreliance on like-minded consultants can produce “group-think” and impair deliberation, while including ideologically distant voices can improve decision-making but may increase the risk of biased or strategically manipulated advice. Empirical studies document the strategic inclusion of diverse appointees, and leaders often consult multiple experts to improve information aggregation and reduce reliance on any single perspective (Ingraham et al., 1995; Bertelli and Feldmann, 2007; Lewis, 2008).⁵

These strategic tensions – between loyalty, competence, and ideological diversity – are present across political systems. Rooted in the Weberian tradition that separates administration from politics, many European countries and the UK expect senior civil servants to provide impartial advice (Putnam, 1973). Yet, the political context in which advice is given often blurs this boundary. Civil servants may be formally neutral but still perceived as ideologically misaligned, prompting political leaders to question their objectivity. Simply mandating neutrality may in fact reduce the effectiveness of advice by suppressing independence rather than ensuring it.

This paper provides a first formal analysis of the trade-offs a leader faces when advisors differ in political alignment, competence and informational diversity. A leader must decide under uncertainty and may consult one or two experts, at a cost for each. Each observes a private binary signal informative about the state of the world and, if consulted, sends an unverifiable report. One expert is less biased relative to the leader, but either observes a noisier signal (Section 3) or information correlated with the leader’s (Section 7); the other is better informed but more biased.

Section 4 characterizes the pure and mixed strategy equilibria. In a pure strategy, each consulted advisor either reports his signal truthfully or sends a message unrelated to what he knows. Truthful reporting requires his bias not to exceed a threshold determined in equilibrium; randomization is possible only at that threshold. The threshold rises when

³Similar trade-offs are uncovered by Bellodi et al. (2026) in the context of regulator assignments. They show that, for the average rule, about 1/3 of potential assignees have relevant prior experience, where the pool is limited to regulators in the originating bureau who are observed writing rules both before and after that rule. Restricting assignments to co-partisans would forgo about 33% of maximally available expertise.

⁴A structurally similar trade-off has been formally analyzed in the context of candidate selection under list proportional representation. Buisseret and Prato (2022) examine how parties allocate list ranks to candidates by balancing competence with electability.

⁵For example, the Bush administration appointed both Colin Powell, who opposed the Iraq War, and Donald Rumsfeld, who supported it (Saunders, 2018). President Obama appointed Republican Ray LaHood as Secretary of Transportation to foster bipartisan collaboration (Newton-Small, 2008).

the other expert’s message is uninformative or he is not consulted, making truth-telling less demanding. The leader does not consult an expert whose message she expects to be uninformative.

The leader’s optimal consultation strategy relates the experts’ biases to consultation in a subtle way. When both biases are sufficiently small, consulting both maximizes her welfare, unless consultation costs are too high. As the bias of the better-informed expert increases, the leader first dismisses his colleague, despite the latter’s closer alignment: both can no longer be truthful with positive probability, but the more informed expert can still be truthful when the other sends an uninformative message or is not consulted. The more aligned consultant therefore contributes no information, and dismissing him saves his consultation cost.⁶ As the better-informed expert’s bias grows further, he too can no longer be truthful, and the leader switches to his less competent but more aligned colleague.

These consultation reversals persist with continuous, uniformly distributed signals (Section 5), where communication is partitional and never fully revealing. The result is stronger in one respect. In the binary model, the leader turns to the better-informed expert alone when informative communication by both ceases to be feasible. With continuous signals, she can make this change while both can still communicate informatively. The reason is that adding the more aligned consultant changes the better-informed expert’s communication incentives and can induce him to communicate more coarsely. The resulting loss of information may exceed the information contributed by the additional consultant himself. As the better-informed expert’s bias grows further, the leader again switches to the more aligned one.

For tractability, the continuous-signal analysis focuses on equilibria in which each informative advisor divides his signal space into two intervals. We restrict the comparison to a parameter region in which these are the most informative single-advisor equilibria, and verify numerically that consulting both advisors does not generate richer communication. The formal characterization and numerical checks are provided in Section 5 and Part C of the Supplementary Appendix.

We relate these consultation patterns to advisor turnover during leadership transitions. A new leader whose views break with her predecessors’ is ideologically distant from the experts appointed before her: appointees who share her views lack the incumbents’ specialized expertise, and our analysis shows that introducing her own appointees can hinder

⁶This result would hold unchanged if the advisors could also send public messages without being consulted by the leader, for example, by speaking to the press. In the best equilibrium, the leader does not consult the more aligned expert, ignores his “free” public advice and bases her decision only on the better-informed one’s consultation.

her ability to extract reliable information from the experienced experts. If her views are so divergent that she cannot rely on incumbent experts for trustworthy advice, replacing them with her own appointees is optimal; if instead ideological distance does not preclude effective communication and collaboration, she is better off retaining and relying exclusively on the existing experts. These considerations help explain observed patterns of advisor turnover following leadership transitions, as discussed in Section 4.

We extend the baseline model by allowing for uncertainty over the political alignment of the better-informed expert: it is uncertain whether he is more biased than his less informed colleague, whose preferences are known, or just as closely aligned with the leader. The motivation is that elected leaders must make their political views manifest to gain electoral support, whereas unelected consultants often keep their leanings confidential— withholding them being a crucial aspect of the professional conduct by which an expert establishes the credibility of his advice.

Many of our earlier findings carry over to this model of “uncertain trade-off,” but the comparative statics become richer. Increasing the bias of the less informed expert, together with his better-informed colleague’s bias whenever their views coincide, can now lead to the latter’s dismissal. This cannot occur in the baseline model, where only the less informed expert’s preferences change. And although the better-informed consultant becomes more attractive *ex ante*, his less informed colleague is consulted more frequently.

Furthermore, the comparative statics are no longer limited to the experts’ biases: increasing the likelihood that the better-informed expert is no more biased than his colleague can lead the leader to shift from consulting only the former to consulting both. This complements our earlier result that a newly elected leader should rely primarily on pre-existing, experienced consultants: if she discovers over time that their views are more aligned with hers than expected, she can improve decision-making by adding loyal, if less experienced, consultants—developments that, as we detail in Section 6, are not uncommon in government appointments.

2 Literature Review

Our work contributes to the political economy literature on advice. Among others, Battaglini (2002) shows how full information revelation arises with biased, perfectly informed experts, while Dewatripont and Tirole (1999) analyze decision-making with competing advocates who conceal unfavorable information but cannot fabricate evidence. Che and Kartik (2009) show that disagreement can induce greater effort to acquire verifi-

able information, and Morris (2001) studies communication by an expert concerned with appearing impartial. Ottaviani and Sørensen (2001) analyze speech order among experts signaling competence, and Dewan and Squintani (2018) model leadership as emerging from communication among ideologically diverse agents. We add to this literature by examining the trade-off between advisor alignment and competence.⁷

Methodologically, we operate within the cheap talk framework pioneered by Crawford and Sobel (1982). Building upon the work of Hagenbach and Koessler (2010) and Galeotti et al. (2013), we formulate a rich yet tractable model for a leader’s choice of advisors. We introduce several features that have not been analyzed jointly before, along with others that are entirely novel to the literature. The decision maker may be imperfectly informed and may consult multiple experts who are also imperfectly informed. We consider players with heterogeneous information and formalize the notions of informativeness, precision, and correlation.⁸ We also allow for the possibility that the preference of one of the advisors is private information.⁹ Relatedly, Ambrus and Lu (2014) study communication by multiple imperfectly informed senders and establish conditions for almost full revelation as observation noise becomes small. Goltsman and Pavlov (2011) instead study one sender communicating with multiple receivers, comparing public and private communication.

The trade-off between selecting aligned or competent subordinates is not unique to democracies, though the selection of consultants in autocratic settings differs in important ways. Our model abstracts from internal power challenges, whereas Egorov and Sonin (2011) show that dictators facing threats of treason may favor loyal but less competent subordinates, since competence can both strengthen regime stability and increase the risk of rebellion.

3 Baseline Model: Alignment vs. Competence

A leader (player 0, “she”) must make a decision $\hat{y} \in \mathbb{R}$ to maximize her utility $u_0(\hat{y}, x) = -(\hat{y} - x)^2$, where x is an unknown state, uniformly distributed on $[0, 1]$. Before making her decision, the leader has the option to consult one or two advisors (both “he”), indexed by

⁷More distantly related to our work, Foerster and Habermacher (2025) study a model of competition between informed agents who differ in preferences and competence that encompasses both advice and lobbying. Each agent posts a menu of transfers. The leader chooses an agent and an action, and the corresponding transfer is established.

⁸Ottaviani and Sørensen (2006) and Denisenko et al. (2026) analyze single-expert models where signal precision reflects competence. The former emphasizes concerns for competence; the latter focuses on strategic disclosure under verifiable communication.

⁹The literature has focused on the case in which there is a single expert, and only his bias is private information; see, for example, Morgan and Stocken (2003) and Li and Madarász (2008).

$i = 1, 2$, at a cost $c > 0$ each. Each advisor i observes a private binary signal $s_i \in \{0, 1\}$ informative about x , and, if consulted, sends a message $\hat{m}_i \in \{0, 1\}$ to the leader.¹⁰ We let $A \subseteq \{1, 2\}$ be the set of consulted advisors. If both advisors are consulted, they send their messages simultaneously.

Each advisor i has a payoff function depending on the known bias b_i , $u_i(\hat{y}, x) = -(\hat{y} - x - b_i)^2$. Advisor 1 is more closely aligned with the leader than advisor 2, in the sense that $|b_1| < |b_2|$ and $b_1 \neq 0$. Advisor 2, however, possesses higher-quality information than advisor 1. We here focus on the case in which s_1 is less precise than s_2 . Specifically, there are two binary signals s'_1 and s_2 , conditionally independent given x . Each signal equals 1 with probability x and 0 otherwise. But while advisor 2 directly observes signal s_2 , the signal s_1 observed by advisor 1 is a noisy version of s'_1 : with probability $p \in (1/2, 1)$, $s_1 = s'_1$, and with probability $1 - p$, $s_1 = 1 - s'_1$.¹¹

A pure strategy of an advisor $i \in A$ is a mapping $m_i : s_i \mapsto \hat{m}_i$. The leader's decision strategy y maps the received advisors' messages $\hat{\mathbf{m}}_A$ to an action $\hat{y} \in \mathbb{R}$. Mixed strategies μ_i are defined as usual. The players play a possibly mixed perfect Bayesian equilibrium (A, μ, y) . By strict concavity of $u_0(\hat{y}, x)$, the leader plays a pure strategy at the decision stage.

As is standard in the literature on cheap talk games, there may be multiple equilibria. For any given consultation set A , there always exists a babbling equilibrium (μ_A, y) in the ensuing communication subgame. Let $\mathcal{E}(A)$ denote the set of such equilibria. We assume that the maximum below exists and that, whenever more than one communication equilibrium exists, a most informative one is selected, where "most informative" means the equilibrium maximizing the leader's ex-ante payoff:

$$V_A = \max_{(\mu_A, y) \in \mathcal{E}(A)} U_0(\mu_A, y) = - \min_{(\mu_A, y) \in \mathcal{E}(A)} \mathbb{E}[\text{Var}(x | s_0, \hat{\mathbf{m}}_A)].$$

Here s_0 denotes any private signal observed by the leader and is omitted in sections in which she has no private signal. We call V_A the leader's *consultation value*. With fixed, commonly known biases, all players rank communication equilibria in the same order because advisor i 's payoff is the leader's payoff minus b_i^2 . With private bias types, this Pareto-ranking conclusion need not hold, but we maintain the leader's ex-ante payoff as the selection criterion.

¹⁰In our binary-type framework, by standard arguments it is sufficient to focus on a binary message space.

¹¹Later in Section 7, we consider the case in which s_1 is less informative to the leader than s_2 , not because s_1 is noisier, but because it is correlated with the leader's information while s_2 is conditionally independent of the leader's signal given the state. This allows us to account for the value of advisors' diversity and independence, relative to the leader. We show that the results are qualitatively similar to the case of noisy s_1 considered here.

If $B \subseteq A$, every equilibrium following B can be embedded following A by having the added advisors babble and the leader ignore their reports; hence $V_A \geq V_B$. The leader chooses A to maximize $W_A = V_A - c|A|$. If consultation sets tie, including when $c = 0$ and an added advisor only babbles, we select the set with the fewest advisors. Throughout, V_T for a labeled communication class $\mathcal{E}_T$ denotes that class's payoff and need not equal the unrestricted consultation-set value V_A .

4 Results

For clarity of exposition, we begin by characterizing the pure strategy equilibria $(A, \mathbf{m}, y)$. The mixed strategy analysis is at the end of this section. We proceed backwards, by fixing any arbitrary consultation set A and calculating the equilibria (μ_A, y) of the ensuing communication subgame, which we call “communication equilibria”.

Pure strategy communication equilibrium characterization Given any consultation set A , and consulted advisor messages $\hat{\mathbf{m}}_A \in \{0, 1\}^A$, it is immediate from the leader's objective function that the equilibrium decision rule y is $y(\hat{\mathbf{m}}_A) = E[x|\hat{\mathbf{m}}_A]$ for any message profile $\hat{\mathbf{m}}_A \in \{0, 1\}^A$. The leader's decision matches her posterior expectation of the state x .

Next, consider the communication stage. Up to relabeling the messages $\hat{m}_i$, each consulted advisor i either plays truthfully in equilibrium, $m_i(s_i) = s_i$ for both $s_i = 0, 1$, or babbles: $m_i(s_i = 0) = m_i(s_i = 1)$. Hence, for any non-empty choice of consulted advisors A , we summarize the equilibrium communication strategies $\mathbf{m}_A$ by the set of truthful advisors $T \subseteq A$. The following result characterizes the pure strategy communication equilibria with non-empty A . We henceforth say that an advisor's signal s_i is contrary to his bias b_i if either $s_i = 0$ and $b_i > 0$ or $s_i = 1$ and $b_i < 0$.

Proposition 1 *Fix any non-empty consultation set A . Any equilibrium $(\mathbf{m}_A, y)$ of the ensuing communication subgame is such that for each truthful advisor $i \in T$, given the strategy m_j of the other advisor j (if consulted), for the signal s_i contrary to his bias b_i :*

$$|b_i| \leq \left| \frac{\sum_{\hat{m}_j=0,1} \Delta_i(s_i, \hat{m}_j)^2 \Pr(\hat{m}_j|s_i)}{2 \sum_{\hat{m}_j=0,1} \Delta_i(s_i, \hat{m}_j) \Pr(\hat{m}_j|s_i)} \right|, \quad (1)$$

where $\Delta_i(s_i, \hat{m}_j) = E[x|\hat{m}_i = 1 - s_i, \hat{m}_j] - E[x|\hat{m}_i = s_i, \hat{m}_j]$. Conversely, these conditions are sufficient, as no advisor gains by lying at the signal conforming to his bias. The leader's equilibrium decision rule is $y(\hat{\mathbf{m}}_A) = E[x|\hat{\mathbf{m}}_A]$ for any message profile $\hat{\mathbf{m}}_A \in$




Figure 1: Heuristic argument for truth-telling.

$\{0, 1\}^A$. Net of the sunk cost c , the leader's expected equilibrium payoff in the subgame is $U_0(\mathbf{m}_A, y) = -E_{\hat{\mathbf{m}}_A} [\text{Var}(x|\hat{\mathbf{m}}_A)]$, and the expected payoff of each advisor i is $U_i(\mathbf{m}_A, y) = -E_{\hat{\mathbf{m}}_A} [\text{Var}(x|\hat{\mathbf{m}}_A)] - b_i^2$.¹²

When $A = \{i\}$, the sum in (1) contains one term and the condition reduces to $|b_i| \leq |\Delta_i(s_i)|/2$, with no conditioning variable for another advisor's message.

For any given consultation set A , omitting the sunk consultation cost c , all players rank the communication equilibria $(\mathbf{m}_A, y)$ in the same order. The term $E_{\hat{\mathbf{m}}_A} [\text{Var}(x|\hat{\mathbf{m}}_A)]$ identifies the residual uncertainty about the state x after observing the advisors' messages $\hat{\mathbf{m}}_A$.

Consider now the truth-telling condition (1) of Proposition 1, and suppose that both advisors are consulted and are believed to be truthful. For each advisor $i = 1, 2$, the term $\Pr(\hat{m}_j|s_i)$ denotes the probability that the leader receives message $\hat{m}_j$ from the other advisor j , given that advisor i receives signal s_i . Because the equilibrium decision of the leader is $y(\hat{\mathbf{m}}_A) = E[x|\hat{m}_i, \hat{m}_j]$ for every $(\hat{m}_i, \hat{m}_j)$, the expression $\Delta_i(s_i, \hat{m}_j) = E[x|\hat{m}_i = 1 - s_i, \hat{m}_j] - E[x|\hat{m}_i = s_i, \hat{m}_j]$ identifies by how much advisor i would move the leader's decision if lying, i.e., if sending message $\hat{m}_i = 1 - s_i$ instead of the truthful message $\hat{m}_i = s_i$.

For further intuition, consider Figure 1 and let $b_i > 0$, so advisor i is biased rightward. Fix any realization of the message $\hat{m}_j$ sent by the other consulted advisor j . For any signal

¹²In the subgame following $A = \emptyset$, the leader plays $y = E[x]$ and receives expected payoff $Eu_0(y) = -\text{Var}(x)$, whereas each advisor i 's payoff is $Eu_i(y) = -\text{Var}(x) - b_i^2$.




Figure 2: Equilibrium bias thresholds and consultation values as functions of p .

$s_i = 0, 1$, advisor i may deviate from equilibrium by sending the false message $\hat{m}_i = 1 - s_i$ instead of the truthful message $\hat{m}_i = s_i$. By doing so, he moves the leader's decision $\hat{y}$ from $E[x|\hat{m}_i = s_i, \hat{m}_j]$ to $E[x|\hat{m}_i = 1 - s_i, \hat{m}_j]$, i.e., by $\Delta_i(s_i, \hat{m}_j)$. Advisor i has no reason to lie when $s_i = 1$, as this can only move $\hat{y}$ to the left and he is biased rightward. When $s_i = 0$, he gains by lying if and only if he moves $\hat{y}$ to the right by not too much: if $\Delta_i(s_i, \hat{m}_j) > 2b_i$, the leader's response to the lie $\hat{m}_i = 1$ "overshoots" i 's bliss point $E[x|s_i, \hat{m}_j] + b_i$ so much that it makes him worse off relative to sending the truthful message $\hat{m}_i = s_i = 0$; if $\Delta_i(s_i, \hat{m}_j) < 2b_i$, he is better off deviating and lying.

The truth-telling condition (1) highlights that equilibrium truthful communication becomes more demanding when the leader consults multiple advisors. Intuitively, each expert is only concerned with how the leader's action compares to his own bliss point. When she receives more reports, each moves her posterior by less, which reduces the risk of "overshooting" one's ideal point. This makes deviations more attractive and truthful reporting harder to sustain. This contrasts with environments where consultants compete for reputation: there, consulting more experts allows the leader to cross-check reports. We conjecture that this disciplines them and facilitates truthful communication. In our pure cheap talk setting, the opposite effect arises.

Communication equilibrium existence and consultation value Suppose that both advisors are consulted, $A = \{1, 2\}$. There are four possible pure strategy communication equilibria $\mathcal{E}_T$, depending on the set $T \subseteq A$ of truthful advisors. Using Proposition 1, we now determine for what parameters p , b_1 and b_2 each exists and its consultation value.

Before proceeding, we recall a useful result from Galeotti et al. (2013) in a setting where each signal $s_i \in \{0, 1\}$ is an i.i.d. Bernoulli trial with $\Pr(s_i = 1|x) = x$. They show that in any communication equilibrium $(\mathbf{m}, y)$ in which the leader has access to k

Bernoulli signals that are independent conditional on x , the bias b_i of each truthful advisor must satisfy $|b_i| \leq 1/(2(k+2))$, and the leader's equilibrium value from consultation is $V(\mathbf{m}, y) = -1/(6(k+2))$.

Let $\mathcal{E}_{12}$ denote the equilibrium in which both advisors are truthful. In the appendix, we derive the threshold functions

$$\eta_{12.1}(p) = \frac{1}{4} \frac{2p-1}{(1+p)(2-p)} \quad \text{and} \quad \eta_{12.2}(p) = \frac{1}{4} \frac{1+2p-2p^2}{(1+p)(2-p)},$$

such that the equilibrium exists if and only if $|b_1| \leq \eta_{12.1}(p)$ and $|b_2| \leq \eta_{12.2}(p)$. These thresholds are displayed in Figure 2 (left). As the precision of signal s_1 increases (i.e., as p increases), the threshold $\eta_{12.1}(p)$ becomes less stringent, whereas $\eta_{12.2}(p)$ becomes more demanding. As advisor 1 becomes more informed, he would move the leader's decision more by lying. In line with Proposition 1, the truth-telling condition of advisor 1 is relaxed. Conversely, the effect of a lie by advisor 2 on the leader's decision is reduced, and his truth-telling condition becomes more demanding. Intuitively, the leader's consultation value in this equilibrium, denoted by

$$V_{12}(p) = -\frac{1+2p-2p^2}{12(p+1)(2-p)},$$

increases with p , as displayed in Figure 2 (right). As the signal s_1 becomes more precise, including advisor 1 in the consultations becomes more valuable.¹³

The equilibrium $\mathcal{E}_1$ is such that only advisor 1 is truthful. It exists if and only if $|b_1| \leq \eta_1(p) \equiv (2p-1)/6$. The threshold function $\eta_1(p)$, also plotted in Figure 2 (left), is increasing in p and always lies above $\eta_{12.1}(p)$. The first feature again reflects the fact that a more informative signal induces a larger change in the leader's posterior belief about x . The second one reflects the fact that, in line with Proposition 1, the truth-telling condition for advisor 1 is more demanding when advisor 2 is consulted and is truthful. The leader's consultation value in this equilibrium, $V_1(p) = -(1+2p-2p^2)/18$, increases as his signal becomes more informative.

In equilibrium $\mathcal{E}_2$, only advisor 2 is truthful. Because signal s_2 is a Bernoulli trial with $\Pr(s_2 = 1|x) = x$, this equilibrium exists if and only if $|b_2| \leq \eta_2 = 1/6$, and the corresponding consultation value is $V_2 = -1/18$.

Finally, if both advisors' biases exceed their respective thresholds (i.e., $|b_1| > \eta_1(p)$ and $|b_2| > \eta_2$), neither advisor is truthful. The resulting equilibrium $\mathcal{E}_0$ has consultation value $V_0 = -1/12$.

¹³When $p = 1$, the leader effectively receives two i.i.d. Bernoulli signals from the advisors, which yields $\eta_{12.1} = \eta_{12.2} = 1/8$ and $V_{12}(1) = -1/24$. Similarly, when $p = 1/2$, the leader's only informative signal is s_2 , so that $\eta_{12.1} = 0$, $\eta_{12.2} = 1/6$, and $V_{12}(1/2) = -1/18$.

Each communication equilibrium $\mathcal{E}_T$ exists under the same conditions also in the subgame following a consultation set $A \neq \{1, 2\}$, as long as $T \subseteq A$.

Optimal choice of advisor Having characterized the communication equilibria of any subgame given consultation choice A , the characterization of the equilibrium of the whole game is completed by noting that, because consultation is costly, the leader consults an advisor only if she expects that he will be truthful in equilibrium.

Figure 2 shows that $\eta_{12.1}(p) < \{\eta_{12.2}(p), \eta_1(p)\} < \eta_2$, implying that there exist bias pairs (b_1, b_2) such that only the less biased advisor 1 is truthful in equilibrium, as well as bias pairs for which only the more biased advisor 2 is truthful. Specifically, only advisor 1 is truthful when $|b_1| \leq \eta_1(p)$ and $|b_2| > \eta_2$, while only advisor 2 is truthful when $\eta_1(p) < |b_1| < |b_2| \leq \eta_2$. More generally, when $|b_1| \leq \eta_1(p)$ and $|b_2| \leq \eta_2$, but either $|b_1| > \eta_{12.1}(p)$ or $|b_2| > \eta_{12.2}(p)$ holds, the leader can extract truthful information from one advisor but not both. This highlights a strategic trade-off that emerges in an equilibrium, even when the cost of advice is small. When only one advisor can be truthful, the leader optimally chooses the more informed one, provided that the consultation cost is not too large: since $V_2 > V_1(p)$ for all $p < 1$, she consults advisor 2.

The following proposition accounts for the leader's consultation cost, and characterizes the optimal consultation strategy.

Proposition 2 *There exist thresholds $\bar{c}_{12}(p)$, $\bar{c}_1(p)$ and $\bar{c}$, defined in the appendix, such that for any signal precision $p \in (1/2, 1)$ and $c \leq \bar{c}$, the leader's optimal consultation choices are as follows.*

(i) *When both advisors' biases are sufficiently small, $|b_1| \leq \eta_{12.1}(p)$ and $|b_2| \leq \eta_{12.2}(p)$, the leader consults both advisors if $c \leq \bar{c}_{12}(p)$, and only consults advisor 2 if $c > \bar{c}_{12}(p)$.*

(ii) *When $|b_1| > \eta_{12.1}(p)$ and $|b_2| \leq \eta_{12.2}(p)$, or $|b_2| \in (\eta_{12.2}(p), \eta_2]$, the leader consults advisor 2 who is truthful in this region.*

(iii) *When only advisor 1 can be truthful, $|b_2| > \eta_2$ and $|b_1| \leq \eta_1(p)$, the leader consults advisor 1 if $c \leq \bar{c}_1(p)$ and otherwise consults none.*

If either $c > \bar{c}$ or $|b_1| > \eta_1(p)$, $|b_2| > \eta_2$, then no advisor is consulted.

The cost thresholds $\bar{c}_{12}(p)$ and $\bar{c}_1(p)$ increase in p , reflecting that higher signal precision makes consultation more valuable. (Throughout a bar denotes a cost threshold of this section, a tilde the threshold of $\mathcal{E}_{12}$ against $\mathcal{E}_0$ in Appendix B, and a hat one of Section 6.)

We now present our key comparative statics.

Proposition 3 *Suppose that $c < \bar{c}_{12}(p)$, $|b_1| \leq \eta_{12.1}(p)$, and $|b_2| \leq \eta_{12.2}(p)$, so that both advisors are initially consulted. As the bias b_2 of the more biased advisor increases in absolute value, so that $\eta_{12.2}(p) < |b_2| \leq \eta_2$, the leader drops advisor 1 but retains advisor 2. If $|b_2|$ increases further so that $|b_2| > \eta_2$, the leader dismisses advisor 2 and hires back advisor 1.*




Figure 3: Optimal equilibria when consultation cost is not too high ($c < \bar{c}_{12}(p)$).

Consider the case of positive biases. It may seem counterintuitive that increasing the bias of advisor 2 leads to dropping advisor 1, but the model makes the mechanism clear. When the consultation cost is not too high, $c < \bar{c}_{12}(p)$, and the biases are sufficiently low, $b_1 \leq \eta_{12.1}(p)$ and $b_2 \leq \eta_{12.2}(p)$, the leader benefits from truthful advice from both advisors (region $\mathcal{E}_{12}$ of Figure 3). Once $\eta_{12.2}(p) < b_2 \leq \eta_2$, both advisors can no longer be truthful in equilibrium; the leader must choose and, advisor 2 being more informative, she retains him (region $\mathcal{E}_2$). As b_2 rises beyond η_2 , advisor 2 ceases to be truthful, and the leader instead turns to the less-informative but more-aligned advisor 1 (region $\mathcal{E}_1$).¹⁴

Proposition 3 speaks directly to the organization of political advice during leadership transitions. When a newly elected policymaker departs ideologically from previous administrations, the resulting shift widens the gap between the leader's preferences and those of the most experienced experts appointed under prior governments. In our framework, this

¹⁴This finding also holds under continuous signals, and when allowing for randomization, as shown in later sections.

maps to an increase in the bias b_2 of the more competent expert (advisor 2). The central question becomes whether the leader should retain these experienced but potentially misaligned appointees or replace them with more ideologically aligned, yet less experienced, consultants (advisor 1).

The proposition shows that while previous leaders might have benefited from consulting both kinds of advisors, the new leader is better off relying exclusively on experienced experts, provided that the gap between their preferences is not too large. Were the new leader to combine advice from both an experienced consultant and a less informed loyalist, the former’s information provision incentives would be disrupted. Only if the gap in preferences between the new leader and the experienced expert is sufficiently large would the leader resort to consulting only a loyalist while dismissing the experienced consultant.

This theoretical insight is reflected in how political leaders choose their advisors. Newly elected policymakers often balance competence and continuity with political alignment. For example, in December 2008 President-elect Obama announced that he would retain Robert Gates—Secretary of Defense under President George W. Bush—despite major policy differences, particularly over the Iraq War. While Bush had initiated the conflict, Obama entered office committed to ending it, yet retained Gates to ensure an orderly transition, valuing his experience and leadership during the Iraq surge and ongoing operations in Afghanistan. This decision illustrates a preference for cross-administration continuity and technical expertise (Gates, 2010; Suri, 2018).

If the leader is ideologically too distant from the incumbent bureaucracy, she should instead rely on less experienced but ideologically aligned advisors. This is illustrated by Jimmy Carter’s presidency, where appointments often prioritized personal loyalty over federal governance expertise. As a political outsider with no prior national office, Carter distrusted the Washington establishment, especially in the aftermath of Watergate, and emphasized loyal confidants (Kaufman and Kaufman, 2006). Many senior staff—informally known as the “Georgia Mafia”—were longtime associates from his Georgia years, including Hamilton Jordan, Jody Powell, and Frank Moore, all of whom lacked experience in national politics (Fallows, 1979; Barrow, 2023).

Our analysis has so far focused on pure strategy equilibria. We conclude this section by showing that our main results do not change when we also include mixed strategy equilibria in the analysis.

Mixed strategy equilibria Throughout this paragraph, we assume that the two biases have the same sign, $b_1 b_2 > 0$. Because $|b_1| < |b_2|$, simultaneous reflection of the state,

signals, messages and actions lets us derive the results under $0 < b_1 < b_2$; the case $b_2 < b_1 < 0$ follows from the same reflection. We focus on our main case of interest, presented in Proposition 3. Hence, we suppose that $|b_1| < \eta_{12.1}(p)$ and $c < \bar{c}_{12}(p)$. We earlier proved that, as $|b_2|$ increases continuously, the leader first consults both advisors, then retains only advisor 2, and finally switches to consulting only 1. How does this result change when considering also mixed strategy equilibria?

Consider first the equilibria $\mathcal{M}_i$, in which advisor i randomizes and advisor j babbles. Advisor i can mix only when his signal conflicts with his bias. When they align, sending $\hat{m}_i = s_i$ is a strict best response given the equilibrium decision rule and beliefs. We show in Part B of the Supplementary Appendix that $\mathcal{M}_i$ is dominated by the equilibrium $\mathcal{E}_i$, in which expert i is truthful. First, $\mathcal{M}_i$ is less informative than $\mathcal{E}_i$. Second, he moves the leader's decision less by lying in $\mathcal{M}_i$. The value of Δ_i is therefore smaller, making the no-deviation condition more demanding. This is in line with the discussion after Proposition 1. Among the equilibria $\mathcal{M}_{i,j}$ in which i is truthful and j mixes, $\mathcal{M}_{1,2}$ exists, given $|b_1| < \eta_{12.1}(p)$, only where $\mathcal{E}_{12}$ does and is dominated by it; and $\mathcal{M}_{12}$, in which both advisors randomize, is dominated by $\mathcal{M}_{2,1}$ for the same reason that $\mathcal{M}_i$ is dominated by $\mathcal{E}_i$. The only mixed strategy equilibrium that can be optimal is therefore $\mathcal{M}_{2,1}$, in which advisor 2 is truthful and 1 mixes. It also exists for parameters such that $\mathcal{E}_{12}$ does not and $\mathcal{E}_2$ is the most informative pure strategy equilibrium: there, the leader receives some information from 1 together with 2's truthful message.

Proposition 4 *Suppose $b_1 b_2 > 0$. There exists a continuous threshold function $\eta_{M,2} : (p, |b_1|) \mapsto |b_2|$ such that equilibrium $\mathcal{M}_{2,1}$ exists if and only if $\eta_{12.1}(p)/2 < |b_1| < \eta_{12.1}(p)$ and $|b_2| \leq \eta_{M,2}(p, |b_1|)$. The function $\eta_{M,2}$ is strictly decreasing in $|b_1|$, and such that $\eta_{M,2}(p, |b_1|) = \eta_2$ for $|b_1| = \eta_{12.1}(p)/2$ and $\eta_{M,2}(p, |b_1|) = \eta_{12.2}(p)$ for $|b_1| = \eta_{12.1}(p)$.*

Figure 4 displays the existence region of $\mathcal{M}_{2,1}$, bounded by $\eta_{12.1}(p)/2$, $\eta_{12.1}(p)$ and the curve $\eta_{M,2}$, together with the regions in which $\mathcal{E}_{12}$, $\mathcal{E}_2$ and $\mathcal{E}_1$ are optimal, for $b_2 > b_1 > 0$. The existence region of $\mathcal{M}_{2,1}$ is a strict subset of $\mathcal{E}_2$'s existence region. The reason is that $\mathcal{E}_2$ only requires advisor 2 to be truthful, whereas $\mathcal{M}_{2,1}$ also requires advisor 1 to be indifferent between reporting $\hat{m}_1 = 0$ and $\hat{m}_1 = 1$ when $s_1 = 0$, so that his equilibrium condition holds as an equality. Hence b_1 must be small enough that 1 does not strictly prefer to lie, but also large enough that he does not *strictly* prefer to report $\hat{m}_1 = 0$, which yields $b_1 \geq \eta_{12.1}(p)/2$. For any such b_1 , the indifference condition pins down the probability that 1 is truthful, which we denote by $\mu_{1,s_1}(s_1)$. It increases in b_1 and approaches one as b_1 approaches $\eta_{12.1}(p)$, and with it the upper bound $\eta_{M,2}(p, b_1)$ on b_2 approaches $\eta_{12.2}(p)$, so that it remains (generically) strictly below η_2 .




Figure 4: Mixed strategy equilibrium existence.

For the leader's consultation choice, our case of interest is such that $|b_2| \in (\eta_{12.2}(p), \eta_2]$: equilibrium $\mathcal{E}_{12}$ does not exist, but $\mathcal{E}_2$ does. If $\mathcal{M}_{2.1}$ does not exist either, the optimal equilibrium is $\mathcal{E}_2$, only advisor 2 is hired, and Proposition 3 holds unchanged. If instead $\mathcal{M}_{2.1}$ exists, it is more informative than $\mathcal{E}_2$, so that $V_{M.2.1}(p, |b_1|) > V_2$, and the leader consults both advisors whenever $c < V_{M.2.1}(p, |b_1|) - V_2$. Because $\mu_{1,s_1}(s_1)$ and hence $V_{M.2.1}$ increase in $|b_1|$, there exists a threshold function $\eta_{M.1}(p, c)$, increasing in the consultation cost c , such that advisor 1 is consulted if and only if $|b_1| > \eta_{M.1}(p, c)$. Unexpectedly, consultation of advisor 1 requires that he is sufficiently biased.

We summarize our findings in the following result.

Proposition 5 *Suppose that $b_1 b_2 > 0$, $\eta_{M.1}(p, c) < |b_1| < \eta_{12.1}(p)$ and $\eta_{12.2}(p) < |b_2| \leq \eta_{M.2}(p, |b_1|)$, then the welfare optimal equilibrium is $\mathcal{M}_{2.1}$. Within the same-sign domain, no other mixed strategy equilibrium can ever be optimal when $|b_1| < \eta_{12.1}(p)$.*

For all p and all c , the results of Proposition 3 generalize. If (b_1, b_2) are such that both advisors are initially consulted and $|b_2|$ continuously increases without changing sign, then the leader first drops advisor 1 and retains 2, and eventually switches from 2 to 1.

5 Continuous Signals

The reversal in the optimal consultation choice that we identified in the binary setting also persists with continuous signals. In the class of games we study here, where the state and signals are continuous, equilibrium communication is partitional. Specifically, each advisor's signal space is divided into intervals, and each interval is associated with a distinct message to the leader (Crawford and Sobel, 1982). The advisor types whose signals

lie at the margins of adjacent intervals are indifferent between the corresponding messages.

Suppose that the leader considers consulting an additional expert. His presence gives her another source of information, as in the binary case. At the same time, any given consultant's report moves the leader's posterior belief about the state, and so her optimal decision, by less. For an expert whose signal lies close to the margin between any two messages, a deviation becomes less costly, making misreporting more attractive. To restore indifference at the margins, equilibrium communication requires a coarser partition of the signal space, resulting in less informative communication. In other words, we are facing the same trade-off as in the binary model.

Continuous signals strengthen the results from the binary model in one respect. In the latter, the leader turns to a single-advisor consultation whenever an informative two-advisor equilibrium ceases to exist. In contrast, under our continuous-signal specification, advisor 2 can be optimally consulted alone even though informative communication with two experts remains feasible. In such instances, adding advisor 1 worsens the quality of information provided by the better-informed expert by more than the informational gain from the additional consultant's messages.

The environment As before, the state is uniformly distributed on the unit interval, $x \sim U[0, 1]$. Advisor $i \in \{1, 2\}$ observes a noisy private signal $s_i = x + \varepsilon_i$, where the noise terms are mutually independent and jointly independent of x , with $\varepsilon_i \sim U[-\delta_i, \delta_i]$. Signals may therefore lie outside the unit interval, whereas the state remains in $[0, 1]$. We assume $0 < \delta_2 < \delta_1 < 1/2$, so that advisor 2 receives the more precise signal. As in the baseline model, $0 < b_1 < b_2$, so advisor 1 is more aligned with the leader. The leader and advisor i have quadratic-loss payoffs, $-(y - x)^2$ and $-(y - x - b_i)^2$, respectively.¹⁵

To keep the analysis tractable, we study equilibria where each advisor uses two intervals of his signal space, i.e., sends two informative messages in equilibrium. When only advisor i is consulted, he sends a low message m_L^i when $s_i < t_i$ and a high message m_H^i otherwise, where t_i denotes advisor i 's cutoff signal. We denote this equilibrium by E_i . When both advisors are consulted, each uses such a two-interval partition, and we denote the resulting equilibrium by E_{12} .

We call the interval $[t_i - \delta_i, t_i + \delta_i]$ advisor i 's *switching region*. Outside this region the state pins down his message, whereas inside it either message may arise depending on the

¹⁵The environment combines the uniform-quadratic communication model of Crawford and Sobel (1982) with imperfectly informed senders, as in Austen-Smith (1993); with perfect signals, the two-advisor game is related to multiple-expert models such as Krishna and Morgan (2001).




Figure 5: An illustration of two-interval communication by both advisors. Here, the switching regions are disjoint since $t_2 + \delta_2 < t_1 - \delta_1$.

realization of the noise. A switching region may extend beyond the state space, and the two switching regions may or may not overlap. Figure 5 illustrates one possible configuration.

Consider first the case in which only advisor i is consulted. Let $\theta_i = E[x|s_i = t_i]$ denote the posterior mean of the marginal signal, $\alpha(\theta_i) = \Pr(s_i < t_i)$ the probability of the low message, and $\gamma(\theta_i) = E[x|s_i \geq t_i] - E[x|s_i < t_i]$ the distance between the leader's induced actions. Define

$$\Psi_{\delta_i}(\theta) = 1/2 - \theta - [1 - 2\alpha(\theta)]\gamma(\theta)/2.$$

The marginal advisor must be indifferent between his low and high messages. The following result gives the corresponding existence condition.

Proposition 6 *Suppose that only advisor i is consulted. An informative two-interval equilibrium E_i exists if and only if $|b_i| < 1/4$, and is unique within this equilibrium class. Its threshold is the unique solution to $b_i = \Psi_{\delta_i}(\theta_i)$, where the function Ψ_{δ_i} is continuous and strictly decreasing from $1/4$ to $-1/4$. The leader's actions and value are*

$$y_L^i(b_i) = \frac{1}{2} - [1 - \alpha(\theta_i)]\gamma(\theta_i), \quad y_H^i(b_i) = \frac{1}{2} + \alpha(\theta_i)\gamma(\theta_i), \quad V_i = \alpha(\theta_i)[1 - \alpha(\theta_i)]\gamma(\theta_i)^2 - \frac{1}{12}.$$

Consider next the case in which both advisors are consulted. Given thresholds t_1 and t_2 , define

$$G_i(x) = \min\{1, \max\{0, (t_i + \delta_i - x)/(2\delta_i)\}\},$$

the probability that consultant i sends the low message conditional on state $x \in [0, 1]$, and write

$$\begin{aligned} A_i &= \int_0^1 G_i(x) dx, & B_i &= \int_0^1 x G_i(x) dx, \\ C &= \int_0^1 G_1(x) G_2(x) dx, & D &= \int_0^1 x G_1(x) G_2(x) dx. \end{aligned}$$

Proposition 7 *Suppose that both advisors are consulted and use thresholds t_1 and t_2 . Following each positive-probability message profile, the leader chooses*

$$y_{LL} = \frac{D}{C}, \quad y_{LH} = \frac{B_1 - D}{A_1 - C}, \quad y_{HL} = \frac{B_2 - D}{A_2 - C}, \quad y_{HH} = \frac{\frac{1}{2} - B_1 - B_2 + D}{1 - A_1 - A_2 + C}.$$

Her value is

$$V_{12} = Cy_{LL}^2 + (A_1 - C)y_{LH}^2 + (A_2 - C)y_{HL}^2 + (1 - A_1 - A_2 + C)y_{HH}^2 - \frac{1}{3},$$

with zero-probability terms omitted. The threshold pair (t_1, t_2) supports an equilibrium E_{12} if and only if there are sequentially rational beliefs and actions following off-path profiles such that, for each $i \in \{1, 2\}$, every signal below t_i weakly prefers L and every signal above t_i weakly prefers H . In particular, each type $s_i = t_i$ must be indifferent between his messages.

If the two switching regions overlap, all four message profiles occur with positive probability. If they are disjoint, one of the profiles in which the advisors send different messages occurs with probability zero. The leader's action following that profile must nevertheless be specified, since it may be reached after a unilateral deviation.

Further, marginal indifference is only a necessary condition for the existence of the two-advisor equilibrium. This is because the effect of advisor i 's report depends on the message sent by advisor j , and the probability of receiving each of j 's messages varies with i 's own signal. Marginal indifference can therefore hold even though some signals away from the threshold prefer the other message.¹⁶

We therefore impose the incentive constraints for every signal of both experts when identifying the two-advisor equilibria used below. Part C of the Supplementary Appendix gives these constraints explicitly.

The scope of the analysis Two-interval equilibria need not be the most informative equilibria of a continuous cheap-talk game. Since the leader's expected payoff rises with the informativeness of communication, we restrict attention to parameter values for which two-interval equilibria are the most informative equilibria available.

To characterize this region, let $\bar{b}_{\text{gen}}^{(N)}(\delta_i)$ denote the supremum bias at which advisor i , when consulted alone, can sustain an N -interval partition. Proposition 6 implies $\bar{b}_{\text{gen}}^{(2)}(\delta_i) = 1/4$.

We next characterize the corresponding boundary for three-interval communication.

¹⁶To see why, consider for example advisor 1 who compares $EU(H|s_1) - EU(L|s_1)$, averaging over what advisor 2 will report. As s_1 rises, two things happen. First, advisor 1's posterior about x rises, which pushes him toward H . Second, advisor 2 becomes more likely to report H . And when advisor 2 reports H , changing advisor 1's own report from L to H may move the leader's action much more, creating a larger risk of overshooting his preferred action.

Lemma 1 For each $\delta_i \in (0, 1/2)$,

$$\bar{b}_{\text{gen}}^{(3)}(\delta_i) = \begin{cases} \frac{3a_i^2 + \delta_i^2}{12a_i}, & \delta_i \leq \frac{3-\sqrt{3}}{4}, \quad a_i = \frac{2-\sqrt{1+\delta_i^2}}{3}, \\ \frac{a_i}{3}, & \delta_i > \frac{3-\sqrt{3}}{4}, \quad 8a_i^3 - 12\delta_i a_i + 3\delta_i = 0, \end{cases}$$

where in the second case a_i is the unique root in $(1/4, \delta_i)$.

Under the regularity conditions stated in Part C, the maximal sustainable bias decreases with the number of intervals. Therefore, $b_i > \bar{b}_{\text{gen}}^{(3)}(\delta_i)$ rules out all finer single-advisor partitions as well. Part C also searches directly for richer two-advisor partitions. For configurations with at most four intervals per advisor, our numerical analysis finds that consulting both advisors does not enable either advisor to communicate more finely than he could alone, except in near-degenerate cases.

Hence, we conduct the numerical comparison on the restricted parameter region

$$\mathcal{R} = \{(\delta_1, \delta_2, b_1, b_2) : 0 < \delta_2 < \delta_1 < \frac{1}{2}, 0 < b_1 < b_2 < \frac{1}{4}, b_i > \bar{b}_{\text{gen}}^{(3)}(\delta_i), i = 1, 2\}.$$

When comparing the informational value of the different equilibria, we abstract from consultation costs. This is conservative for the organizational result: a positive incremental cost of consulting a second advisor can only make single-advisor consultation more attractive.

Comparing consultation choices For the equilibria described above, we derive the posterior actions, equilibrium conditions and the leader's expected payoffs analytically, and then compare the leader's values V_{12} , V_2 , and V_1 numerically over $\mathcal{R}$. When several two-advisor equilibria exist at the same bias pair, we select the one giving the leader the highest value.

Figure 6 provides a simple illustration of the consultation reversal. Holding the two precisions and b_1 fixed, E_{12} gives the leader the highest value up to $b_2 = 0.1089$, E_2 on approximately $(0.1090, 0.1118)$, and E_1 thereafter.¹⁷

Figure 7 shows how this pattern varies with both biases. Fixing b_1 and increasing b_2 corresponds to moving upward in the figure. Close to the diagonal, where the two experts

¹⁷The change from consulting both experts to consulting only advisor 2 also illustrates why the incentive conditions must hold for every signal. At $b_2 = 0.10903$, the two marginal-indifference conditions continue to have a solution, but the resulting partition is not an equilibrium: some signals of advisor 1 away from his threshold prefer the opposite message. Thus joint communication can break down even though the usual indifference conditions at the two partition boundaries continue to hold.




Figure 6: The leader's values in E_1 , E_2 , and E_{12} as functions of b_2 , for $(\delta_1, \delta_2, b_1) = (0.20, 0.15, 0.10)$.




Figure 7: The equilibrium yielding the highest leader value as (b_1, b_2) vary over the restricted parameter region, for $(\delta_1, \delta_2) = (0.20, 0.15)$. Blue: E_{12} . Purple: E_2 , although an informative E_{12} exists. Crimson: E_2 , with no retained informative E_{12} . Gold: E_1 .

have similar biases, E_{12} often gives the leader the highest value. As advisor 2 becomes more biased, E_2 tends first to dominate, and eventually E_1 does.

The purple region highlights the additional reason for dismissing the more aligned advisor that arises with continuous signals. *The leader can obtain informative reports from both experts, yet she gains more information by relying on advisor 2 alone.* An informative equilibrium E_{12} exists here, but its value is lower than that of E_2 .

The reason is that adding advisor 1 changes what the better-informed expert is willing to communicate. Advisor 2 then uses a coarser partition of his signal space, so his report conveys less information. The additional consultant does contribute information of his own, but this gain is smaller than the information lost from advisor 2's report. Consulting a second expert therefore need not improve the leader's information. With a positive additional consultation cost, she strictly prefers to rely on advisor 2 alone.

In the crimson region, the reason for relying on a single expert is different: no informative two-advisor equilibrium in the retained class is available. The two regions thus distinguish *inability to obtain informative reports from both experts* from *a preference for relying on one even when both can communicate informatively*. In the purple region, the leader dismisses the more aligned consultant to preserve the quality of the better-informed expert's advice.

How the difference in expertise affects consultation The informational cost of dismissing advisor 1 depends on how much less precise his signal is than advisor 2's. Figure 8 compares the leader's preferred equilibrium for several pairs of noise parameters (δ_1, δ_2) . Each panel fixes these parameters and varies the two biases. A larger δ_1 means that the more aligned consultant has a noisier signal; a smaller δ_2 means that the better-informed expert has a more precise one.

This leads to two conclusions. First, joint communication is more valuable when the experts' precisions are similar. When advisor 2 has a large informational advantage, preserving the informativeness of his communication is especially valuable, making E_2 more attractive. As the precision gap narrows, the informational cost of excluding the more aligned consultant rises and E_{12} becomes more prevalent.

Second, the continuous model does not imply that the preferred equilibrium changes monotonically with b_2 for every precision pair. For some parameters, E_{12} can cease to give the highest value and later become preferred again as b_2 rises. The availability and value of joint communication need not vary monotonically once equilibrium thresholds can approach the boundaries of the signal space. Thus, for some fixed signal precisions and




Figure 8: The equilibrium yielding the highest leader value over the restricted parameter region for different precision pairs. Blue: E_{12} ; purple: E_2 , although an informative E_{12} exists; crimson: E_2 , with no retained informative E_{12} ; gold: E_1 . White denotes parameter values outside the restricted parameter region.

values of b_1 , increasing b_2 does not simply lead from consulting both experts to consulting only advisor 2 and then only advisor 1: the leader can return to consulting both along the way.

6 Uncertain Advisors Trade-Off

We generalize the model of Section 3 by assuming that advisor 2's bias is private information. Specifically, advisor 2 has bias equal to $b_1 > 0$ with probability $q \in (0, 1)$, and a higher bias $b_2 > b_1$ with probability $1 - q$. In words, advisor 2 is either equally aligned with or less aligned than advisor 1, and the leader does not know which. Private biases are mutually independent and independent of the state, all signals, and all other types. The

information structure satisfies the posterior-monotonicity condition of Appendix A: raising any advisor's signal from zero to one raises the leader's posterior mean at every realization of the other information available to her, including after the other advisors' signals have been garbled into equilibrium messages. The baseline signal structure satisfies the primitive likelihood-ratio condition given there. The remainder of the model is unchanged. For brevity, we focus on pure strategy equilibria.

Communication equilibria Two types of communication equilibria arise in this setting. First, both types of advisor 2 follow the same strategy, yielding equilibria $\mathcal{E}_{12}$, $\mathcal{E}_1$, $\mathcal{E}_2$, and $\mathcal{E}_0$. Second, the two types of advisor 2 behave differently, in which case (as shown in the appendix) two equilibria arise. In an equilibrium $\mathcal{E}_{12A}$ both advisor 1 and the aligned type of 2 are truthful, while the b_2 -type of advisor 2 reports the high message $\hat{m}_2 = 1$ whatever his signal; in the other equilibrium $\mathcal{E}_{2A}$ only the b_1 -type (i.e., the aligned type) of advisor 2 is truthful, the b_2 -type again always reports $\hat{m}_2 = 1$, and advisor 1 babbles.

The two equilibria exist under the following bias thresholds. $\mathcal{E}_{12A}$ exists if and only if $b_1 \leq \eta_{12A}(p, q) \equiv \min\{\eta_{12A.1}(p, q), \eta_{12A.2}(p, q)\}$ —the truth-telling bounds for advisor 1 and for advisor 2's b_1 -type—and $b_2 \geq \eta_{12A.2}(p, q)$. The last condition has no counterpart in Section 4: as the leader cannot tell the two types apart, her action responds to $\hat{m}_2$, so the b_2 -type must prefer the high message also when $s_2 = 0$, which he does exactly when his bias exceeds the bound up to which the b_1 -type reports $s_2 = 0$ truthfully. Likewise, $\mathcal{E}_{2A}$ exists if and only if $b_1 \leq \eta_{2A}(q) \leq b_2$. We note that $\eta_{12A.1}(p, q)$ increases in p and decreases in q , whereas $\eta_{12A.2}(p, q)$ decreases in p and increases in q : this mirrors the exact logic of the baseline model illustrated in Figure 2, and $\eta_{2A}(q)$ increases in q (verified in Part D of the Supplementary Appendix).

We now turn to the ranking of communication equilibria in terms of consultation value.

Lemma 2 *For all $(p, q) \in (1/2, 1) \times (0, 1)$, the communication equilibrium consultation value functions are ranked as follows: $V_{12} > \{V_2, V_{12A}\} > \{V_{2A}, V_1\}$. There exist functions $g_1, g_2 : q \mapsto p$ such that:*

(i) $V_2 > V_{12A}(p, q)$ if and only if $p < g_2(q)$;

(ii) $V_{2A}(q) > V_1(p)$ if and only if $p < g_1(q)$.

The function g_1 strictly increases in q with $g_1(0) = 1/2$ and $g_1(1) = 1$, while g_2 strictly decreases in q with $g_2(0) = 1$ and $g_2(1) = 1/2$.

Naturally, the fully revealing equilibrium $\mathcal{E}_{12}$ yields the highest consultation value for all (p, q) —this is the maximum amount of information the leader can receive. The equilibria




Figure 9: Value and existence region comparisons.

$\mathcal{E}_2$ and $\mathcal{E}_{12A}$ follow—recall that in $\mathcal{E}_2$ advisor 2 is truthful whereas in $\mathcal{E}_{12A}$ advisor 1 and the aligned type of 2 are truthful. Their relative ranking depends on (p, q) : if both parameters are low, $\mathcal{E}_2$ dominates $\mathcal{E}_{12A}$, since the signal of advisor 1 is not too informative and advisor 2 is unlikely to be aligned. Further, Lemma 2 shows that $\mathcal{E}_2$ and $\mathcal{E}_{12A}$ dominate both $\mathcal{E}_{2A}$ and $\mathcal{E}_1$, as the latter deliver less information in expectation to the leader.

Lemma 2 partitions the (p, q) space into four regions via the functions g_1 and g_2 (see Figure 9; ignore g_3 for now). Above g_1 , equilibrium $\mathcal{E}_1$ dominates $\mathcal{E}_{2A}$ because the precision p of advisor 1's signal s_1 is high, making his truthful message more valuable than that of advisor 2's aligned type. Right of g_2 , equilibrium $\mathcal{E}_{12A}$ dominates $\mathcal{E}_2$ because the probability q that advisor 2 is aligned is high. Since advisor 2 reports truthfully in both equilibria when aligned, receiving a truthful message in the unlikely event that he is misaligned is less valuable than truthful communication from advisor 1.

We now turn to the existence conditions for the equilibria $\mathcal{E}_{12A}$, $\mathcal{E}_1$, $\mathcal{E}_2$ and $\mathcal{E}_{2A}$. The bias thresholds satisfy $\eta_{12.1} < \eta_{12A.1} < \eta_1$, $\eta_{12A.2} < \eta_{12.2}$ and $\eta_{12A.2} < \eta_{2A} \leq \eta_2$, so that $\eta_{12A} < \{\eta_1, \eta_{2A}, \eta_{12.2}\} < \eta_2$; the comparison between η_{12A} and $\eta_{12.1}$ depends on the parameters: $\eta_{12A} = \eta_{12A.2}$ falls below $\eta_{12.1}$ when p exceeds $(\sqrt{7} - 1)/2 \approx 0.82$ and q lies below a bound increasing in p (0.23 at $p = 0.85$, 0.83 at $p = 0.95$), on 21 percent of the square. Hence some of the existence ranges are ordered in terms of inclusion. $\mathcal{E}_{12}$ requires a smaller value

of b_2 than $\mathcal{E}_2$, and, outside that set, a smaller value of b_1 than $\mathcal{E}_{12A}$. The equilibrium $\mathcal{E}_{12A}$ is supported by a smaller range of the bias b_1 than $\mathcal{E}_{2A}$ and $\mathcal{E}_1$, and its condition on b_2 fails only when $b_1 < b_2 < \eta_{12A.2}$.

The existence ranges of $\mathcal{E}_{2A}$ and $\mathcal{E}_1$, however, are not ordered in terms of inclusion. We now define a third threshold function g_3 , based on the equality $\eta_1(p) = \eta_{2A}(q)$, that identifies whether the existence condition of $\mathcal{E}_1$ is tighter or looser in terms of b_1 , compared to the existence condition for $\mathcal{E}_{2A}$.

Lemma 3 *There exists a strictly increasing function $g_3 : q \mapsto p$ such that $\eta_1(p) > \eta_{2A}(q)$ if and only if $p > g_3(q)$. Moreover, for low q , we have $g_1(q) < g_3(q) < g_2(q)$. As q increases, g_3 crosses g_2 and eventually coincides with g_1 at $q = 1$, with $g_3(1) = 1$.*

Consider Figure 9 again. Similar to g_1 , the function g_3 partitions the (p, q) space in two regions. Above g_3 , the existence conditions for equilibrium $\mathcal{E}_1$ are weaker than for $\mathcal{E}_{2A}$, and vice versa. The reason is connected to why $\mathcal{E}_1$ dominates $\mathcal{E}_{2A}$ above g_1 and vice versa. When the precision p of advisor 1's signal s_1 is higher, a lie by 1 moves the leader's decision more. By Proposition 1, this makes his truth-telling condition less demanding.

This result concludes the characterization of communication equilibria. We now determine the leader's consultation choice. For clarity, we first consider the case of small cost $c > 0$. We consider arbitrary consultation cost $c > 0$ later.

Advisor consultation choice with small cost Suppose that, fixing the other parameters in the game, one takes the consultation cost $c > 0$ to zero. In the limit, every advisor who is consulted communicates informatively for at least one type in the selected equilibrium; otherwise that advisor could be omitted at a strictly lower cost without changing the leader's information. Among the remaining equilibria, the leader selects the one with the highest ex-ante consultation value V_A , because $W_A \rightarrow V_A$ as $c \downarrow 0$. We call this the leader-preferred communication equilibrium. Lemmas 2 and 3 deliver its ranking. We first report the findings formally and then explain them with Figures 9 and 10.

Proposition 8 *The optimal communication equilibrium $\mathcal{E}^*$ in the game with uncertain trade-off is as follows:*

(i) *If $b_1 \leq \eta_{12.1}(p)$ and $b_2 \leq \eta_{12.2}(p)$, then $\mathcal{E}^* = \mathcal{E}_{12}$. Otherwise, if $b_1 > \eta_{12A}(p, q)$ and $b_2 \leq \eta_2$, then $\mathcal{E}^* = \mathcal{E}_2$, and if $b_1 \leq \eta_{12A}(p, q)$ and $b_2 > \eta_2$, then $\mathcal{E}^* = \mathcal{E}_{12A}$.*

(ii) *If $b_1 \leq \eta_{12A}(p, q)$ and $b_2 \leq \eta_2$, but either $b_1 > \eta_{12.1}(p)$ or $b_2 > \eta_{12.2}(p)$, then $\mathcal{E}^* = \mathcal{E}_{12A}$ if $p > g_2(q)$ and $b_2 \geq \eta_{12A.2}(p, q)$, and $\mathcal{E}^* = \mathcal{E}_2$ otherwise, the leader being indifferent between the two when $p = g_2(q)$ and $b_2 \geq \eta_{12A.2}(p, q)$.*

(iii) If $\eta_{12A}(p, q) < b_1 \leq \max\{\eta_{2A}(q), \eta_1(p)\}$ and $b_2 > \eta_2$, then $\mathcal{E}^* = \mathcal{E}_{2A}$ when either $b_1 \leq \min\{\eta_{2A}(q), \eta_1(p)\}$ and $p < g_1(q)$, or $\eta_1(p) < b_1 \leq \eta_{2A}(q)$; and $\mathcal{E}^* = \mathcal{E}_1$ when either $b_1 \leq \min\{\eta_{2A}(q), \eta_1(p)\}$ and $p > g_1(q)$, or $\eta_{2A}(q) < b_1 \leq \eta_1(p)$. If $b_1 > \max\{\eta_{2A}(q), \eta_1(p)\}$ and $b_2 > \eta_2$, no advisor is truthful and $\mathcal{E}^* = \mathcal{E}_0$.




Figure 10: Uncertain trade-off: leader-preferred equilibrium at $(p, q) = (3/4, 1/2)$. The thresholds are $\eta_{12.1} = 2/35$, $\eta_{12A} = 5952/79135$, $\eta_1 = 1/12$, $\eta_{12A.2} = 3509/33915$, $\eta_{2A} = 1/9$, $\eta_{12.2} = 11/70$, and $\eta_2 = 1/6$. At these parameters $V_2 > V_{12A}$ and $V_{2A} > V_1$.

Figure 10 illustrates the optimal communication equilibrium $\mathcal{E}^*$, and hence determines the leader's equilibrium consultation choice. In the south-west area marked $\mathcal{E}_{12}$, both advisors are consulted because neither is highly biased. In the north-west region marked $\mathcal{E}_{12A}$, the leader also consults both advisors; here, advisor 1 and the more aligned type of advisor 2 are truthful, while his less aligned type always reports the high message. In contrast, in the south-east region marked $\mathcal{E}_2$, where b_1 is sufficiently large, only advisor 2 is hired.

Consider next the “inverse L-shaped” region where $b_1 \leq \eta_{12A}(p, q)$ and $b_2 \leq \eta_2$, but either $b_1 > \eta_{12.1}(p)$ or $b_2 > \eta_{12.2}(p)$. Here, equilibrium $\mathcal{E}_2$ exists, and so does $\mathcal{E}_{12A}$ except in the small triangle where $\eta_{12.1}(p) < b_1 < b_2 < \eta_{12A.2}(p, q)$, marked $\mathcal{E}_2$ in the figure, where the less aligned type of advisor 2 would rather report his signal. Where both exist, the leader chooses between them based on whether advisor 1 is sufficiently informative or advisor 2 likely enough to be aligned with the leader, i.e., whether (p, q) lies above or below g_2 in

Figure 9. As a result, within this region of Figure 10, as q rises, the leader may shift from hiring only advisor 2 to consulting both advisors.

Finally, consider the north-east region, with $\eta_{12A}(p, q) < b_1 \leq \max\{\eta_{2A}(q), \eta_1(p)\}$ and $b_2 > \eta_2$. When $b_1 \leq \min\{\eta_{2A}(q), \eta_1(p)\}$, both $\mathcal{E}_1$ and $\mathcal{E}_{2A}$ exist, and the leader's choice depends on whether advisor 1 is more informative or advisor 2 more aligned—determined by the location of (p, q) relative to g_1 in Figure 9. When $\min\{\eta_{2A}(q), \eta_1(p)\} < b_1 \leq \max\{\eta_{2A}(q), \eta_1(p)\}$ only one of the equilibria $\mathcal{E}_{2A}$ and $\mathcal{E}_1$ exists, namely $\mathcal{E}_1$ if (p, q) is above g_3 , and $\mathcal{E}_{2A}$ otherwise. Therefore, if (p, q) is above g_1 and below g_3 , the leader must resort to hiring only advisor 2 because the equilibrium $\mathcal{E}_1$ does not exist.

We now turn to comparative statics results. As in the case of a certain trade-off (Proposition 3), increasing b_2 can lead to the dismissal of advisor 1 where $\mathcal{E}_2$ is preferred to $\mathcal{E}_{12A}$, or where $\mathcal{E}_{12A}$ is unavailable. This is not always the case: in parameter regions where $\mathcal{E}_{12A}$ remains available and leader-preferred, both advisors remain consulted as b_2 increases. We next highlight the additional comparative statics.

Proposition 9 *The main additional comparative statics results for the case of uncertain trade-off are as follows:*

(i) *Suppose that $b_1 \leq \eta_{12A}(p, q)$ and $b_2 > \eta_2$, so that the leader hires both advisor 1 and advisor 2, and the common aligned bias parameter b_1 increases, raising both advisor 1's bias and the bias of advisor 2's aligned type. If $p > g_1(q)$, then advisor 2 is dismissed when $b_1 > \eta_{12A}(p, q)$. If, additionally, $p < g_3(q)$ —which happens for intermediate p and low q —then 2 is hired back, and 1 discharged, when $\eta_1(p) < b_1 \leq \eta_{2A}(q)$.*

(ii) *Suppose that $\eta_{12.1}(p) < b_1 \leq \eta_{12A}(p, q)$, $\eta_{12.2}(p) < b_2 \leq \eta_2$, and $p < g_2(q)$, so that the leader only consults advisor 2. An increase in the probability q that advisor 2 is aligned leads to hiring advisor 1 as well, when $p > g_2(q)$ and $b_1 \leq \eta_{12A}(p, q)$ at the new value of q .*

Unlike in the baseline model where advisor 2's bias is always b_2 , an increase in the common aligned bias parameter b_1 may now lead to the dismissal of advisor 2, rather than just advisor 1. This comparative static simultaneously raises advisor 1's bias and the bias of advisor 2's aligned type. It occurs in the region above g_1 in Figure 9, when $b_2 > \eta_2$, as the increasing b_1 crosses $\eta_{12A}(p, q)$ (cf. Figure 10). In this scenario, because $b_2 > \eta_2$, the equilibrium $\mathcal{E}_2$ does not exist. When $b_1 < \eta_{12A}(p, q)$, the optimal equilibrium is $\mathcal{E}_{12A}$, so that both advisors are consulted; advisor 1 and advisor 2's aligned type report truthfully. However, as b_1 increases and crosses $\eta_{12A}(p, q)$, the equilibrium $\mathcal{E}_{12A}$ ceases to exist. Since $p > g_1(q)$, the optimal equilibrium becomes $\mathcal{E}_1$, and the leader dismisses advisor 2 to gather information only from advisor 1. As in the case of a certain trade-off, increasing one

advisor’s bias may lead to dismissing the other due to their strategic interaction. Here, however, this occurs not only when advisor 2’s bias increases, but also when the common aligned bias parameter is raised.

We also examine comparative statics as advisor 2 becomes more likely to be aligned with the leader, i.e., when q increases. Here, the leader may switch from consulting only advisor 2 to consulting *both* advisors. To see why, consider the “inverse-L” shaped region of Figure 10, where $b_1 \leq \eta_{12A}(p, q)$ and $b_2 \leq \eta_2$, but either $b_1 > \eta_{12.1}(p)$ or $b_2 > \eta_{12.2}(p)$. Suppose that (p, q) is located to the left of g_2 in Figure 9. As advisor 2 becomes more likely to be aligned, that is, as q increases, the function g_2 is crossed from the left and the equilibrium $\mathcal{E}_{12A}$ now dominates $\mathcal{E}_2$, leading the leader to hire advisor 1 as well, provided that $\mathcal{E}_{12A}$ still exists at the higher q (recall that $\eta_{12A.1}$ decreases in q). The leader is willing to forgo truthful advice from the b_2 -type of advisor 2 to gain the benefit of hiring advisor 1 and receiving his truthful report.

These findings naturally relate to our results in Proposition 3. There, we characterized when a newly-elected leader should primarily rely on pre-existing experienced experts, rather than bringing in her own less competent, politically aligned appointees. Here, we complement this result by considering the implications of uncovering information about the views of these appointees over time. The second part of Proposition 9 suggests that if the leader finds out that the views of experienced consultants are not as different from hers as expected, she can improve decision-making by consulting them alongside her own appointees, even if they are less experienced.

These results resonate with historical accounts of political appointments. It is not uncommon that leaders initially rely on potentially misaligned, yet experienced experts, and then add more aligned consultants over time. Relatedly, Spenskuch et al. (2023) document large partisan cycles in the stock of U.S. political appointees. The partisan composition of presidential appointees changes gradually after presidential transitions, which they suggest reflects delays in confirmation.

Experienced advisors are retained when they demonstrate stronger-than-expected alignment, and dismissed when alignment falls short of expectations. The second case is illustrated by the Nixon administration, which inherited a bureaucracy shaped by Johnson-era Democrats, viewed it as potentially obstructive, but initially retained experts for their competence. As Nixon stated after his inauguration, “We will do what we think will best serve the interest of effective government, and if the individual who has been frozen in can do the job, we are going to keep him” (Nixon, 1969). Over time, however, officials lacking

ideological alignment were increasingly sidelined.¹⁸ The result was a gradual realignment of the senior civil service, shifting toward greater support for Nixon’s agenda, without overt partisan purges.¹⁹

Advisor consultation choice with any cost In this section, we allow for any positive consultation cost $c > 0$ and focus on the most interesting case for comparative statics. Specifically, we assume that $p > g_2(q)$, $b_1 < \eta_{12A}(p, q)$, and $b_2 > \eta_{12.2}(p)$. These conditions imply that, without consultation costs, equilibrium $\mathcal{E}_{12A}$ exists and dominates $\mathcal{E}_2$, so both advisors are consulted.

The relevant payoffs include the outside option $V_0 = -1/12$. Define the best available one-advisor value by

$$V_S = \begin{cases} V_2, & b_2 \leq \eta_2, \\ \max\{V_1(p), V_{2A}(q)\}, & b_2 > \eta_2, \end{cases}$$

where V_1 is selected above g_1 and V_{2A} below g_1 . Let $c_S = V_{12A} - V_S$ and $c_{0S} = V_S - V_0$. Part D proves that $0 < c_S < c_{0S}$ under the maintained assumptions.

Proposition 10 *Suppose $p > g_2(q)$, $b_1 < \eta_{12A}(p, q)$, and $b_2 > \eta_{12.2}(p)$. Then the leader consults both advisors for $0 \leq c < c_S$, only the advisor yielding V_S for $c_S < c < c_{0S}$, and neither for $c > c_{0S}$, with indifference between adjacent choices at the cutoffs.*

When c is small, the information gain justifies consulting both advisors even though advisor 2’s less aligned type babbles. At c_S , the leader begins consulting only the expert who provides her with the higher value. At c_{0S} , even that information gain no longer covers the consultation cost, and she stops consulting either advisor.

The consultation cost c captures the political capital required for advisor appointments. It rises when leaders are politically weakened (e.g., scandals, low approval) and falls when mandates are strong. While an immediate result given our previous analysis, Proposition 10 is useful to shed some light on how leaders navigate high-cost scenarios, showing they either consolidate around a highly competent expert or retreat to a compliant loyalist, depending on perceived alignment.

Reagan’s 1987 appointment of Howard Baker as Chief of Staff illustrates the first part of Proposition 10. Before the Iran-Contra scandal, Reagan balanced competent and more aligned advisors, including James Baker and Donald Regan (Baker and Glasser, 2020;

¹⁸This reflected Nixon’s “administrative presidency” strategy, which prioritized loyalty over expertise (Aberbach and Rockman, 1976; Lewis, 2008).

¹⁹Cole and Caputo (1979) find that career executives identifying as Independents became more likely to support Nixon’s domestic policies.

Weinraub, 1986).²⁰ After the scandal, Reagan dismissed the “loyalist” Regan and consolidated authority under Howard Baker, a bipartisan crisis manager (Gerstenzang, 1987). This demonstrates how rising political costs can lead a leader to rely solely on the most competent expert, even when he may be less politically aligned.

French President Chirac’s 2005 appointment of Dominique de Villepin illustrates the second part of Proposition 10. Before the failed EU referendum, Chirac had drawn on both competent figures like Nicolas Sarkozy and loyalists (Buchan, 2025). As the referendum raised political costs (Henley, 2005; de Boisgrollier, 2005), he chose between the experienced but independent Sarkozy and the more aligned but inexperienced Villepin, opting for the latter. This shows how high costs may push leaders toward reliable loyalists over more capable but less aligned alternatives.

7 Advisor Alignment vs. Diversity

This section considers the trade-off between advisors’ political alignment and diversity, modeled as access to information, or information sources, different from the leader’s. The leader has nothing to learn from an expert whose information is exactly the same as her own. We show how the previous results on the trade-off between alignment and competence carry over.

We modify the model of Section 3 as follows. First, we say that the leader holds a signal $s_0 \in \{0, 1\}$ such that $\Pr(s_0 = 1|x) = x$. Conditional on x , s_0 , s'_1 , and s_2 are mutually independent. Second, we stipulate that, with probability $\rho \in (0, 1)$ and unbeknownst to advisor 1, the signal s_1 observed by advisor 1 coincides with s_0 ; and otherwise $s_1 = s'_1$. Hence, while advisor 2’s signal is conditionally independent of the leader’s signal given the state, the signal s_1 of advisor 1 is correlated, and hence less informative to the leader. The remainder of the model of Section 3 is unmodified. Fix $\rho \in (0, 1)$ and the bias vector. We study the limit as $c \downarrow 0$; equivalently, the stated classification holds for every sufficiently small positive c at each fixed parameter vector. The statement is not uniform in ρ .

As in the analysis of the trade-off between alignment and competence, there are four equilibria in pure strategies: the equilibrium $\mathcal{E}_{12}$ where both advisors tell the truth, the two equilibria $\mathcal{E}_1$ and $\mathcal{E}_2$ where only one advisor (either 1 or 2) is truthful, and the babbling equilibrium.

Consider the equilibrium $\mathcal{E}_{12}$, in which both advisors are truthful and hence consulted. The thresholds $\eta_{12.1}(\rho)$ and $\eta_{12.2}(\rho)$ determine the maximum biases under which full revela-

²⁰James Baker had managed George H.W. Bush’s 1980 campaign, highlighting Reagan’s pragmatism.




Figure 11: Equilibrium bias thresholds and consultation values as functions of ρ .

tion is sustainable. Figure 11 depicts these thresholds, along with the leader's equilibrium payoff $V_{12}(\rho)$.

The threshold $\eta_{12.1}(\rho)$ lies below $\eta_{12.2}(\rho)$ for all values of ρ , and is decreasing in ρ . Conversely, $\eta_{12.2}(\rho)$ increases in ρ . That is, as signal s_1 becomes less correlated with the leader's private signal (i.e., ρ decreases), advisor 1's bias threshold $\eta_{12.1}$ becomes less stringent, while advisor 2's threshold $\eta_{12.2}$ becomes more demanding. The leader's equilibrium payoff $V_{12}(\rho)$ decreases in ρ and is convex, reflecting the diminishing value of signal s_1 as it becomes more redundant with the leader's own information.

In the equilibrium $\mathcal{E}_1$, where only advisor 1 is consulted, the threshold $\eta_1(\rho)$ decreases in ρ , reflecting that he must be more aligned with the leader when his signal is highly correlated with the leader's private signal. The threshold $\eta_1(\rho)$ always lies above $\eta_{12.1}(\rho)$, and crosses $\eta_{12.2}(\rho)$: truth-telling by advisor 1 requires a less strict alignment when he is the sole informant (and it may be less stringent than advisor 2's condition under full revelation).

The leader's payoff in this case, $V_1(\rho)$, also decreases in ρ and is convex. It is always strictly lower than $V_{12}(\rho)$ for all ρ – naturally, consulting both advisors dominates single-advisor consultation.

In the equilibria $\mathcal{E}_2$ and $\mathcal{E}_0$, in which only advisor 2 or no advisor is consulted, the leader holds two Bernoulli signals that are conditionally independent given the state, or one signal, respectively, so that, by the result of Galeotti et al. (2013) recalled in Section 4, advisor 2 alone is truthful if and only if $|b_2| \leq \eta_2 = 1/8$ (the threshold is lower than the $1/6$ of Section 4 because the leader now holds a signal of her own), and the payoff levels are the constants $V_2 = -1/24$ and $V_0 = -1/18$. The properties stated above and the proposition that follows are proved in Part E of the Supplementary Appendix.

The optimal equilibrium is characterized in the following proposition.

Proposition 11 *Fix $\rho \in (0, 1)$ and the bias vector, and take $c \downarrow 0$. If $|b_1| \leq \eta_{12.1}(\rho)$ and $|b_2| \leq \eta_{12.2}(\rho)$, then the leader optimally consults both advisors. If either $|b_1| > \eta_{12.1}(\rho)$ and $|b_2| \leq \eta_{12.2}(\rho)$, or $\eta_{12.2}(\rho) < |b_2| \leq \eta_2$, then she hires only advisor 2, the independently-informed advisor. If $|b_2| > \eta_2$ and $|b_1| \leq \eta_1(\rho)$, then the leader consults only advisor 1, the less biased advisor. If $|b_1| > \eta_1(\rho)$ and $|b_2| > \eta_2$, then the leader does not consult any advisor.*

Simultaneous reflection of signals and actions covers negative biases, which is why the classification depends on absolute biases.

We observe that this characterization of the optimal consultation choice is identical to Proposition 2 in the baseline model. The same logic applies despite the different information structure: whether an advisor’s information is less valuable due to low precision (p) or high redundancy (ρ), the leader follows the same threshold-based strategy. Therefore, qualitatively the same comparative statics hold as in the baseline model.

8 Conclusion

Political leaders face difficult trade-offs when selecting advisors, balancing political alignment, informational competence, and the diversity of perspectives offered by the experts they consult. We analyzed this problem in a model in which consulting one or two experts entails costs. One expert is more closely aligned with the leader’s preferences but provides less precise (or more correlated) information, and the other possesses higher-quality information but is less politically aligned.

Our analysis revealed that the leader’s optimal advisory choices can display subtle reversals. If both advisors are initially consulted, increasing ideological distance from the better-informed advisor can make it optimal to dismiss the more aligned one; if that distance grows too large, trust breaks down and the leader switches to the more loyal advisor. These reversals are not an artifact of binary signals. With continuous signals, within the retained numerical class an equilibrium in which only the better-informed advisor communicates may give the leader a higher gross payoff than a feasible equilibrium in which both communicate. The same decision rule can also be obtained after consulting both experts by having the other send an uninformative message. The consultation sets therefore tie at zero cost; our rule favoring fewer consultants selects the smaller set, and any positive cost makes

that choice strict. We related these patterns to historical episodes of advisor turnover, particularly during leadership transitions.

The experts in our model care only about influencing the leader’s decision. In practice, political consultants also care about their reputation for competence and integrity, which existing research has analyzed in settings with one expert and one decision maker. Extending reputational concerns to a leader choosing among multiple experts who differ in alignment and expertise is a promising direction: repeated interactions could enable the leader to evaluate past advice, encouraging even misaligned consultants to communicate truthfully. When multiple experts are consulted simultaneously, comparing their reports may also strengthen incentives for honest reporting.

Another dimension worth exploring is how a leader’s personal style shapes her choice of advisors and their influence on her decisions. Leaders more open to conflict and decentralized information flow often cultivate informal, task-oriented structures that encourage diverse views, whereas those who prioritize control and consensus tend to prefer hierarchical, loyalty-driven arrangements with tightly managed access and influence (Hermann and Preston, 1994). These styles affect not only who is consulted, but also how their input is processed and weighted. Distinct advisory roles emerge, and complementarities between “idea persons” and implementation-focused consultants become key to effective leadership.

References

- Joel D. Aberbach and Bert A. Rockman. Clashing beliefs within the executive branch: The Nixon administration bureaucracy. *American Political Science Review*, 70(2):456–468, 1976.
- Attila Ambrus and Shih En Lu. Almost fully revealing cheap talk with imperfectly informed senders. *Games and Economic Behavior*, 88:174–189, 2014. doi: 10.1016/j.geb.2014.09.001.
- David Austen-Smith. Interested experts and policy advice: Multiple referrals under open rule. *Games and Economic Behavior*, 5(1):3–43, 1993.
- Peter Baker and Susan Glasser. *The Man Who Ran Washington: The Life and Times of James A. Baker III*. Doubleday, 2020.
- Bill Barrow. Jimmy Carter: Family affair to the White House and beyond. AP News, February 28, 2023. URL <https://apnews.com/article/politics-united-states-government-libya-jimmy-carter-georgia-23768d6b322704b46535154a2f9fb519>.
- Marco Battaglini. Multiple referrals and multidimensional cheap talk. *Econometrica*, 70(4):1379–1401, 2002.

- Luca Bellodi, Massimo Morelli, Jörg L. Spenkuch, Edoardo Teso, Matia Vannoni, and Guo Xu. Personnel is policy: Delegation and political misalignment in the rulemaking process. Cambridge Working Papers in Economics 2622, Faculty of Economics, University of Cambridge, 2026. URL <https://www.econ.cam.ac.uk/publications/cwpe/2622>.
- Anthony Bertelli and Sven E. Feldmann. Strategic appointments. *Journal of Public Administration Research and Theory*, 17(1):19–38, 2007.
- David Buchan. Nicolas Sarkozy. Encyclopaedia Britannica, 2025. URL <https://www.britannica.com/biography/Nicolas-Sarkozy>.
- Peter Buisseret and Carlo Prato. Competing principals? Legislative representation in list proportional representation systems. *American Journal of Political Science*, 66(1):156–170, 2022.
- Yeon-Koo Che and Navin Kartik. Opinions as incentives. *Journal of Political Economy*, 117(5):815–860, 2009.
- David M. Cohen. Amateur government. *Journal of Public Administration Research and Theory*, 8(4):450–497, 1998.
- Richard L. Cole and David A. Caputo. Presidential control of the senior civil service: Assessing the strategies of the Nixon years. *American Political Science Review*, 73(2):399–413, 1979.
- Vincent P. Crawford and Joel Sobel. Strategic information transmission. *Econometrica*, 50(6):1431–1451, 1982.
- Nicolas de Boisgrollier. Dominique de Villepin: The new man on the right. Technical report, The Brookings Institution, 2005. URL <https://www.brookings.edu/wp-content/uploads/2016/06/boisgrollier20051103.pdf>.
- Anna Denisenko, Catherine Hafer, and Dimitri Landa. Competence and advice. *American Journal of Political Science*, 70(1):21–39, 2026.
- Torun Dewan and Francesco Squintani. Leadership with trustworthy associates. *American Political Science Review*, 112(4):844–859, 2018.
- Mathias Dewatripont and Jean Tirole. Advocates. *Journal of Political Economy*, 107(1):1–39, 1999.
- George C. Edwards, III. Why not the best? The loyalty–competence trade-off in presidential appointments. *The Brookings Review*, 19(2):12–16, 2001.
- Georgy Egorov and Konstantin Sonin. Dictators and their viziers: Endogenizing the loyalty–competence trade-off. *Journal of the European Economic Association*, 9(5):903–930, 2011.
- James Fallows. The passionless presidency: The trouble with Jimmy Carter’s administration. *The Atlantic*, 243(5), 1979.
- Manuel Foerster and Daniel Habermacher. Policy-advising competition and endogenous

- lobbies. *Journal of Public Economics*, 245:105354, 2025.
- Andrea Galeotti, Christian Ghiglino, and Francesco Squintani. Strategic information transmission networks. *Journal of Economic Theory*, 148(5):1751–1769, 2013.
- Robert M. Gates. Helping others defend themselves: The future of US security assistance. *Foreign Affairs*, 89(3):2–6, 2010.
- James Gerstenzang. Regan resigns, Reagan names Baker to top post: President hopes to regain momentum, end White House ‘feuding’. Los Angeles Times, February 28, 1987. URL <https://www.latimes.com/archives/la-xpm-1987-02-28-mn-6487-story.html>.
- Maria Goltsman and Gregory Pavlov. How to talk to multiple audiences. *Games and Economic Behavior*, 72(1):100–122, 2011. doi: 10.1016/j.geb.2010.08.007.
- Jeanne Hagenbach and Frédéric Koessler. Strategic communication networks. *The Review of Economic Studies*, 77(3):1072–1099, 2010.
- Jon Henley. Chirac opts for security with new PM. The Guardian, June 1, 2005. URL <https://www.theguardian.com/world/2005/jun/01/eu.france2>.
- Margaret G. Hermann and Thomas Preston. Presidents, advisers, and foreign policy: The effect of leadership style on executive arrangements. *Political Psychology*, 15(1):75–96, 1994.
- Patricia W. Ingraham, James R. Thompson, and Elliot F. Eisenberg. Political management strategies and political/career relationships: Where are we now in the federal government? *Public Administration Review*, 55(3):263–272, 1995.
- Burton Ira Kaufman and Scott Kaufman. *The Presidency of James Earl Carter, Jr.* University Press of Kansas, 2006.
- Vijay Krishna and John Morgan. A model of expertise. *Quarterly Journal of Economics*, 116(2):747–775, 2001.
- David E. Lewis. *The Politics of Presidential Appointments: Political Control and Bureaucratic Performance*. Princeton University Press, 2008.
- David E. Lewis. The contemporary presidency: The personnel process in the modern presidency. *Presidential Studies Quarterly*, 42(3):577–596, 2012.
- Ming Li and Kristóf Madarász. When mandatory disclosure hurts: Expert advice and conflicting interests. *Journal of Economic Theory*, 139(1):47–74, 2008.
- Dean E. Mann. The selection of federal political executives. *American Political Science Review*, 58(1):81–99, 1964.
- Terry M. Moe. The politicized presidency. In John E. Chubb and Paul E. Peterson, editors, *The New Direction in American Politics*, pages 235–271. Brookings Institution, Washington, DC, 1985.
- John Morgan and Phillip C. Stocken. An analysis of stock recommendations. *The RAND*

- Journal of Economics*, 34(1):183–203, 2003.
- Stephen Morris. Political correctness. *Journal of Political Economy*, 109(2):231–265, 2001.
- Jay Newton-Small. Ray LaHood: Secretary of Transportation. Time, December 19, 2008. URL https://content.time.com/time/specials/packages/article/0,28804,1863062_1863058_1868036,00.html.
- Richard Nixon. The President’s news conference, March 4, 1969. Available at The American Presidency Project, University of California, Santa Barbara, 1969. URL <https://www.presidency.ucsb.edu/documents/the-presidents-news-conference-153>.
- Marco Ottaviani and Peter Norman Sørensen. Information aggregation in debate: Who should speak first? *Journal of Public Economics*, 81(3):393–421, 2001.
- Marco Ottaviani and Peter Norman Sørensen. Reputational cheap talk. *The RAND Journal of Economics*, 37(1):155–175, 2006.
- Robert D. Putnam. The political attitudes of senior civil servants in Western Europe: A preliminary report. *British Journal of Political Science*, 3(3):257–290, 1973.
- Elizabeth N. Saunders. Leaders, advisers, and the political origins of elite support for war. *Journal of Conflict Resolution*, 62(10):2118–2149, 2018.
- Jörg L. Spenkuch, Edoardo Teso, and Guo Xu. Ideology and performance in public organizations. *Econometrica*, 91(4):1171–1203, 2023.
- Jeremi Suri. Liberal internationalism, law, and the first African American president. In Julian E. Zelizer, editor, *The Presidency of Barack Obama: A First Historical Assessment*, pages 195–212. Princeton University Press, 2018.
- Bernard Weinraub. How Donald Regan runs the White House. The New York Times Magazine, January 5, 1986. URL <https://www.nytimes.com/1986/01/05/magazine/how-donald-regan-runs-the-white-house.html>.

Appendix A: General Results

We prove a general result from which Proposition 1 and our subsequent communication equilibrium characterizations follow as corollaries. The framework covers the baseline model of Section 3, as well as the extensions considered in Sections 6 and 7.

Suppose there is a leader, player 0, and a set A of consulted advisors. Each player i 's payoff is $u_i(y, x) = -(y - x - B_i)^2$, with $B_0 = 0$, where the leader chooses $y \in \mathbb{R}$ and the state x is uniform on $[0, 1]$. Each player, including the leader, observes a binary signal $s_i \in \{0, 1\}$.²¹ The consultants simultaneously send binary messages, after which the leader chooses y on the basis of s_0 and the message profile.

We impose the following restrictions for the general result. Let $\mathbf{B}_A = (B_i)_{i \in A}$ denote the vector of advisors' private bias types. Its components are mutually independent, and $\mathbf{B}_A$ is independent of the state and all signals. For each advisor, the support of B_i is finite and lies entirely in $(0, \infty)$ or entirely in $(-\infty, 0)$: the magnitude may be private, but the direction is known. Every relevant signal profile has positive probability.

The information structure also satisfies posterior monotonicity. For every advisor i , every profile of possibly mixed communication strategies of the other advisors, and every on-path realization

$$z_i = (s_0, \hat{\mathbf{m}}_{A \setminus \{i\}})$$

for which both values of s_i have positive conditional probability,

$$\mathbb{E}[x | s_i = 1, z_i] > \mathbb{E}[x | s_i = 0, z_i]. \quad (2)$$

Thus raising an advisor's signal from zero to one raises the leader's posterior mean at every realization of the other information available to her. The requirement is stated after the coarsenings induced by the other advisors' messages: pointwise monotonicity conditional only on their complete raw signals need not hold after such pooling.

A primitive sufficient condition for (2) is that, conditional on (x, s_0) , the advisors' signals are independent and, for every i and s_0 ,

$$\frac{\Pr(s_i = 1 | x, s_0)}{\Pr(s_i = 0 | x, s_0)}$$

is strictly increasing in x . The other experts' messages change the prior to which this likelihood ratio is applied but leave the ratio itself unchanged. This sufficient condition holds in the baseline model, the private-bias extension, and the correlated-information specification of Section 7.

As we show below, up to relabeling messages, every informative report has the following form. Each type reports truthfully at the signal conforming to his bias. Let T_i be the set of types who also report truthfully with positive probability at the contrary signal: $s_i = 0$ for a rightward bias and $s_i = 1$ for a leftward bias. Types outside T_i pool on the message conforming to their bias. If a report is uninformative, we use the outcome-equivalent representation in which all types pool on that message. Write

$$\Delta_i(s_0, s_i, \hat{\mathbf{m}}_{-i}) = y(s_0, 1 - s_i, \hat{\mathbf{m}}_{-i}) - y(s_0, s_i, \hat{\mathbf{m}}_{-i})$$

²¹The leader's signal is uninformative in the baseline model and informative in Section 7.

for the change in the leader's decision caused by a false report. The communication equilibria (μ_A, y) , allowing for mixed strategies, and their values are characterized as follows.

Theorem 1 *Fix a non-empty consultation set A . Every communication equilibrium (μ_A, y) has the representation above. For each $i \in A$ and $b_i \in T_i$, the truth-telling condition at the signal contrary to his bias is*

$$|b_i| \leq \left| \frac{\sum_{z_i} \Delta_i(s_0, s_i, \hat{\mathbf{m}}_{-i})^2 \Pr(z_i | s_i)}{2 \sum_{z_i} \Delta_i(s_0, s_i, \hat{\mathbf{m}}_{-i}) \Pr(z_i | s_i)} \right|, \quad (3)$$

with equality if the type mixes. For types outside T_i , the inequality is reversed whenever the report is informative. The leader's decision rule is $y(s_0, \hat{\mathbf{m}}_A) = \mathbb{E}[x | s_0, \hat{\mathbf{m}}_A]$, and communication equilibria are ranked by her ex-ante payoff

$$U_0(\mu_A, y) = -\mathbb{E}[\text{Var}(x | s_0, \hat{\mathbf{m}}_A)]. \quad (4)$$

Proof. The proof is in Part A of the Supplementary Appendix. ■

Writing $Y = y(s_0, \hat{\mathbf{m}}_A)$, an expert's ex-ante payoff is

$$U_i(\mu_A, y) = U_0(\mu_A, y) - \mathbb{E}[B_i^2] + 2\mathbb{E}[B_i(Y - x)]. \quad (5)$$

For a fixed, commonly known bias, this reduces to $U_i = U_0 - b_i^2$, because $\mathbb{E}[Y - x] = 0$. With a private bias, the last term need not vanish, so the leader's ranking need not coincide with each expert's.

We continue with derivations that will be useful to calculate the consultation value expression (4), as well as the terms $\Delta_i(s_0, s_i, \hat{\mathbf{m}}_{A \setminus \{i\}})$ in condition (3). For any profile of information $(s_0, \hat{\mathbf{m}}_A)$, we let the n -th moment of x be:

$$P_n(s_0, \hat{\mathbf{m}}_A) = \int_0^1 x^n \Pr(s_0, \hat{\mathbf{m}}_A | x) dx.$$

Hence, we obtain:

$$\begin{aligned} E[x | s_0, \hat{\mathbf{m}}_A] &= \frac{P_1(s_0, \hat{\mathbf{m}}_A)}{P_0(s_0, \hat{\mathbf{m}}_A)}, \\ \Delta_i(s_0, s_i, \hat{\mathbf{m}}_{A \setminus \{i\}}) &= \frac{P_1(s_0, \hat{m}_i = 1 - s_i, \hat{\mathbf{m}}_{A \setminus \{i\}})}{P_0(s_0, \hat{m}_i = 1 - s_i, \hat{\mathbf{m}}_{A \setminus \{i\}})} - \frac{P_1(s_0, \hat{m}_i = s_i, \hat{\mathbf{m}}_{A \setminus \{i\}})}{P_0(s_0, \hat{m}_i = s_i, \hat{\mathbf{m}}_{A \setminus \{i\}})}. \end{aligned}$$

Likewise, expression (4) will be written as:

$$\begin{aligned} U_0(\mu_A, y) &= - \sum_{(s_0, \hat{\mathbf{m}}_A) \in \{0,1\}^{|A|+1}} \mathbb{E}_x \left[\left(\frac{P_1(s_0, \hat{\mathbf{m}}_A)}{P_0(s_0, \hat{\mathbf{m}}_A)} - x \right)^2 \middle| s_0, \hat{\mathbf{m}}_A \right] \Pr(s_0, \hat{\mathbf{m}}_A) \\ &= - \sum_{(s_0, \hat{\mathbf{m}}_A) \in \{0,1\}^{|A|+1}} v(s_0, \hat{\mathbf{m}}_A), \text{ s.t. } v(s_0, \hat{\mathbf{m}}_A) \equiv \left[P_2(s_0, \hat{\mathbf{m}}_A) - \frac{P_1(s_0, \hat{\mathbf{m}}_A)^2}{P_0(s_0, \hat{\mathbf{m}}_A)} \right]. \end{aligned}$$

We also use the equality of $\Pr(s_0, s_i, \hat{\mathbf{m}}_{A \setminus \{i\}})$ and $P_0(s_0, s_i, \hat{\mathbf{m}}_{A \setminus \{i\}})$.

We conclude the section by expanding the above terms $P_n(s_0, \hat{m}_i, \hat{\mathbf{m}}_{A \setminus \{i\}})$ for the models of Sections 3 and 6, where the set of advisors is $\{1, 2\}$ and the informativeness of signal s_1 is defined using the Bernoulli draw s'_1 . The corresponding derivations for Section 7 are in the Supplementary Appendix.

First, suppose that both advisors are consulted: $A = \{1, 2\}$, so that, omitting s_0 which is uninformative, P_n takes the form:

$$P_n(\hat{m}_1, \hat{m}_2) = \int_0^1 \sum_{s_1 \in \{0,1\}} \sum_{(s'_1, s_2) \in \{0,1\}^2} x^n \Pr(\hat{m}_1, \hat{m}_2 | s_1, s_2) \Pr(s_1 | s'_1) \Pr(s'_1, s_2 | x) dx.$$

Letting μ_{1,s_1} , μ_{b_1,s_2} and μ_{b_2,s_2} be mixed strategies of advisor 1 and of the two types of advisor 2, the first probability term simplifies:

$$\Pr(\hat{m}_1, \hat{m}_2 | s_1, s_2) = \mu_{1,s_1}(\hat{m}_1) [\mu_{b_1,s_2}(\hat{m}_2) q + \mu_{b_2,s_2}(\hat{m}_2) (1 - q)].$$

Turning to the second term, by construction $\Pr(s_1 | s'_1) = p$ if $s_1 = s'_1$ and $\Pr(s_1 | s'_1) = 1 - p$ if $s_1 \neq s'_1$.

For the final term, as s'_1 and s_2 are i.i.d. Bernoulli draws,

$$\Pr(s'_1, s_2 | x) = x^{s'_1 + s_2} (1 - x)^{2 - s'_1 - s_2}.$$

Hence, we obtain:

$$\begin{aligned} P_n(\hat{m}_1, \hat{m}_2) &= \sum_{(s_1, s_2) \in \{0,1\}^2} \mu_{1,s_1}(\hat{m}_1) [\mu_{b_1,s_2}(\hat{m}_2) q + \mu_{b_2,s_2}(\hat{m}_2) (1 - q)] \cdot \\ &\cdot \left[p \frac{(n + s_1 + s_2)!(2 - s_1 - s_2)!}{(n + 3)!} + (1 - p) \frac{(n + 1 - s_1 + s_2)!(1 + s_1 - s_2)!}{(n + 3)!} \right]. \end{aligned} \quad (6)$$

Proceeding in the same fashion, we see that in case only advisor 1 or 2 is consulted, $A = \{1\}$ or $A = \{2\}$:

$$P_n(\hat{m}_1) = \sum_{s_1 \in \{0,1\}} \mu_{1,s_1}(\hat{m}_1) \left[p \frac{(n + s_1)!(1 - s_1)!}{(n + 2)!} + (1 - p) \frac{(n + 1 - s_1)!s_1!}{(n + 2)!} \right], \quad (7)$$

$$P_n(\hat{m}_2) = \sum_{s_2 \in \{0,1\}} [\mu_{b_1,s_2}(\hat{m}_2) q + \mu_{b_2,s_2}(\hat{m}_2) (1 - q)] \frac{(n + s_2)!(1 - s_2)!}{(n + 2)!}. \quad (8)$$

For brevity, whenever considering an equilibrium where an advisor i plays the pure truth-telling strategy, $\hat{m}_i = s_i$ for $s_i = 0, 1$, we will interchange $\hat{m}_i$ with s_i in the expressions of P_n .

Appendix B: Precision vs. Alignment

Equilibrium $\mathcal{E}_{12}$: both advisors are truthful

Suppose first that $b_1 > 0$. Using condition (3) in Theorem 1, together with $\hat{m}_i = s_i$ for $s_i = 0, 1$ and $i = 1, 2$, the equilibrium condition for advisor 1 can be written here as follows:

$$b_1 \leq \eta_{12.1}(p) \equiv \frac{\sum_{s_2=0,1} \Delta_1(s_2)^2 P_0(s_1 = 0, s_2)}{2 \sum_{s_2=0,1} \Delta_1(s_2) P_0(s_1 = 0, s_2)}, \text{ with } \Delta_1(s_2) = \frac{P_1(s_1 = 1, s_2)}{P_0(s_1 = 1, s_2)} - \frac{P_1(s_1 = 0, s_2)}{P_0(s_1 = 0, s_2)}. \quad (9)$$

Using formula (6), we have

$$P_n(s_1, s_2) = p \frac{(s_1 + s_2 + n)! (2 - s_1 - s_2)!}{(n + 3)!} + (1 - p) \frac{(n + 1 - s_1 + s_2)! (1 + s_1 - s_2)!}{(n + 3)!}.$$

We can then immediately calculate $P_0(s_1, s_2)$ and $P_1(s_1, s_2)$, for $(s_1, s_2) \in \{0, 1\}^2$:

$$\begin{aligned} P_0(0, 0) &= \frac{1}{6}(p + 1) & P_1(0, 0) &= \frac{1}{12}, & P_0(0, 1) &= \frac{1}{6}(2 - p), & P_1(0, 1) &= \frac{1}{12}(3 - 2p), \\ P_0(1, 0) &= \frac{1}{6}(2 - p), & P_1(1, 0) &= \frac{1}{12}, & P_0(1, 1) &= \frac{1}{6}(p + 1), & P_1(1, 1) &= \frac{1}{12}(2p + 1). \end{aligned}$$

Plugging these expressions into $\Delta_1(s_2)$, for $s_2 \in \{0, 1\}$, and simplifying, we obtain:

$$\Delta_1(0) = \frac{1}{2} \frac{2p - 1}{(p + 1)(2 - p)} = \Delta_1(1) = \frac{1}{2} \frac{2p - 1}{(p + 1)(2 - p)} =: \Delta_1$$

Hence, expression (9) is simply:

$$\eta_{12.1}(p) = \frac{1}{2} \frac{(\Delta_1)^2}{\Delta_1} = \frac{1}{2} \Delta_1 = \frac{1}{4} \frac{2p - 1}{(p + 1)(2 - p)}.$$

We obtain that advisor 1's truth-telling condition when $b_1 > 0$ is $b_1 \leq \eta_{12.1}(p)$. Due to the model's ex-ante symmetry for signals $s_1 = 0$ and $s_1 = 1$, analogous algebra implies that, when $b_1 < 0$, his equilibrium condition is $-b_1 \leq \eta_{12.1}(p)$. Thus, his truth-telling condition for equilibrium $\mathcal{E}_{12}$ is $|b_1| \leq \eta_{12.1}(p)$.

Likewise, the truth-telling constraint of advisor 2 in equilibrium $\mathcal{E}_{12}$ is $|b_2| \leq \eta_{12.2}(p)$, with, again simplifying (1),

$$\eta_{12.2}(p) \equiv \frac{\sum_{s_1=0,1} \Delta_2(s_1)^2 P_0(s_1, s_2 = 0)}{2 \sum_{s_1=0,1} \Delta_2(s_1) P_0(s_1, s_2 = 0)}, \text{ with } \Delta_2(s_1) = \frac{P_1(s_1, s_2 = 1)}{P_0(s_1, s_2 = 1)} - \frac{P_1(s_1, s_2 = 0)}{P_0(s_1, s_2 = 0)}.$$

Plugging in $P_0(s_1, s_2)$ and $P_1(s_1, s_2)$ derived above into $\Delta_2(s_1)$, for $s_1 \in \{0, 1\}$, and simplifying,

$$\Delta_2(0) = \frac{1}{2} \frac{1 + 2p - 2p^2}{(p + 1)(2 - p)} = \Delta_2(1) = \frac{1}{2} \frac{1 + 2p - 2p^2}{(p + 1)(2 - p)} =: \Delta_2.$$

Hence, we obtain:

$$\eta_{12.2}(p) = \frac{1}{2} \Delta_2 = \frac{1}{4} \frac{1 + 2p - 2p^2}{(p + 1)(2 - p)}.$$

To calculate the equilibrium value $V_{12}(p)$, we simplify formula (4) and use the symmetry of the prior $\Pr(s_1, s_2)$, $\Pr(s_1, s_2) = \Pr(1 - s_1, 1 - s_2)$ for all $(s_1, s_2) \in \{0, 1\}^2$, to obtain:

$$V_{12}(p) = -2 \sum_{s_2=0,1} v(0, s_2), \text{ where } v(0, s_2) = P_2(0, s_2) - \frac{P_1(0, s_2)^2}{P_0(0, s_2)} \text{ for } s_2 = 0, 1. \quad (10)$$

Using the above expressions for $P_0(0, s_2)$ and $P_1(0, s_2)$ for $s_2 = 0, 1$, together with:

$$P_2(0, 0) = \frac{1}{60}(3 - p), \quad P_2(0, 1) = \frac{1}{20}(4 - 3p),$$

we obtain, after simplification:

$$v(0, 0) = \frac{1}{120} \frac{1 + 4p - 2p^2}{p + 1}, \quad v(0, 1) = \frac{1}{120} \frac{3 - 2p^2}{2 - p}.$$

Plugging these expressions into (10) and simplifying, we obtain:

$$V_{12}(p) = -\frac{1}{12} \frac{1 + 2p - 2p^2}{(p+1)(2-p)}.$$

Equilibrium $\mathcal{E}_1$: only advisor 1 is truthful

Condition (1) here is simply:

$$|b_1| \leq \frac{1}{2} \Delta_1 \equiv \eta_1(p), \text{ where } \Delta_1 = \frac{P_1(1)}{P_0(1)} - \frac{P_1(0)}{P_0(0)}.$$

Formula (7) yields here, for $\hat{m}_1 = s_1 \in \{0, 1\}$:

$$P_n(s_1) = p \frac{(s_1 + n)!(1 - s_1)!}{(2 + n)!} + (1 - p) \frac{(1 - s_1 + n)!s_1!}{(2 + n)!}.$$

Plugging $P_n(s_1)$ for $n = 0, 1$ and $s_1 = 0, 1$ into the expression of Δ_1 , we obtain:

$$\eta_1(p) = \frac{1}{2} \Delta_1 = \frac{1}{2} \left(\frac{\frac{1}{6}(p+1)}{\frac{1}{2}} - \frac{\frac{1}{6}(2-p)}{\frac{1}{2}} \right) = \frac{1}{6} (2p - 1).$$

Further, the equilibrium value is:

$$\begin{aligned} V_1(p) &= - \sum_{s_1=0,1} v(s_1), \text{ where } v(s_1) = P_2(s_1) - \frac{P_1(s_1)^2}{P_0(s_1)} \text{ for } s_1 = 0, 1 \\ &= -\frac{1}{36} (1 + 2p - 2p^2) - \frac{1}{36} (1 + 2p - 2p^2) = -\frac{1 + 2p - 2p^2}{18}. \end{aligned}$$

Simple algebra shows that $\eta_{12.1}(p) < \eta_{12.2}(p) < \eta_2$ and $\eta_{12.1}(p) < \eta_1(p) < \eta_2$, and that $V_0 < V_1(p) < V_2 < V_{12}(p)$ for all $p \in (1/2, 1)$.

Proofs of Propositions 2 and 3. First, we compare the value of equilibria $\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_{12}$, presuming they are available, with the babbling equilibrium $\mathcal{E}_0$, whose value is $V_0 = -1/12$. Simple algebra shows:

$$\begin{aligned} W_{12}(p) = -2c - \frac{1}{12} \frac{1 + 2p - 2p^2}{(p+1)(2-p)} \geq V_0 &\iff c \leq \tilde{c}_{12}(p) \equiv \frac{1}{24} \frac{1 - p + p^2}{(1+p)(2-p)}, \\ W_1(p) = -c - \frac{1}{18} (1 + 2p - 2p^2) \geq V_0 &\iff c \leq \bar{c}_1(p) \equiv \frac{1}{36} (2p - 1)^2, \\ W_2 = -\frac{1}{18} - c \geq V_0 &\iff c \leq \bar{c} \equiv \frac{1}{36}. \end{aligned}$$

Here $\tilde{c}_{12}(p)$ is the threshold of $\mathcal{E}_{12}$ against $\mathcal{E}_0$, as distinct from $\bar{c}_{12}(p)$, its threshold against $\mathcal{E}_2$. Since $\bar{c} > \max\{\tilde{c}_{12}(p), \bar{c}_1(p)\}$ for $p \in (1/2, 1)$, we conclude that for $c > \bar{c} = 1/36$, the leader hires no advisor as $V_0 > \max\{W_{12}(p), W_1(p), W_2\}$.

For the following, assume $c < 1/36$. In this range, we have $W_2 \geq \max\{W_1(p), V_0\}$: equilibrium $\mathcal{E}_2$ outperforms $\mathcal{E}_1$ and $\mathcal{E}_0$.

(i) Suppose that $|b_i| \leq \eta_{12,i}(p)$ for both $i = 1, 2$. Hence, equilibrium $\mathcal{E}_{12}$ exists, and hiring both advisors is optimal if

$$W_{12}(p) = -2c - \frac{1}{12} \frac{1+2p-2p^2}{(p+1)(2-p)} \geq W_2 = -\frac{1}{18} - c.$$

Therefore, for each $p \in (1/2, 1)$ there exists a cost threshold

$$\bar{c}_{12}(p) = \frac{1}{36} \frac{(2p-1)^2}{(p+1)(2-p)}$$

weakly increasing in p , such that for all $c \leq \bar{c}_{12}(p)$ it is optimal to consult two advisors; and for all $c \in (\bar{c}_{12}(p), 1/36]$ it is optimal to consult advisor 2. As $p \rightarrow 1$,

$$V_1(p) \rightarrow V_2, \quad W_{12}(p) - W_2 \rightarrow \frac{1}{72} - c.$$

Thus the leader is indifferent between the consultation choices in this limit only when $c = 1/72$.

(ii) Next, suppose that $|b_2| \leq \eta_2$ and, either $|b_1| > \eta_{12,1}(p)$, or $|b_2| > \eta_{12,2}(p)$, or both. Then, equilibrium $\mathcal{E}_{12}$ is not available, and $\mathcal{E}_2$ is the optimal equilibrium, since $|b_2| \leq \eta_2$.

(iii) Finally, suppose that $|b_2| > \eta_2$ and $|b_1| \leq \eta_1(p)$, so that neither equilibrium $\mathcal{E}_{12}$ nor $\mathcal{E}_2$ is available, while $\mathcal{E}_1$ exists. It is optimal to hire only advisor 1 and obtain equilibrium $\mathcal{E}_1$ if $c \leq \bar{c}_1(p)$; otherwise, it is optimal to consult no one and obtain $\mathcal{E}_0$.

Proposition 3 follows immediately from the above results, because $\eta_{12,2}(p) < \eta_2$ for all $p < 1$. Suppose that b_1, b_2, p, c are such that the leader hires both advisors to achieve equilibrium $\mathcal{E}_{12}$. Equilibrium $\mathcal{E}_{12}$ existence requires that $|b_1| \leq \eta_{12,1}(p)$ and $|b_2| \leq \eta_{12,2}(p)$ and welfare dominance that $c < \bar{c}_{12}$. Suppose that $|b_2|$ increases continuously. As $|b_2|$ crosses $\eta_{12,2}(p)$, equilibrium $\mathcal{E}_{12}$ ceases to exist and the optimal equilibrium is $\mathcal{E}_2$: the leader dismisses 1 and retains 2. As $|b_2|$ grows further, it crosses η_2 and then the optimal equilibrium is $\mathcal{E}_1$: the leader dismisses advisor 2 and rehires 1.

The mixed strategy equilibria—the equilibria $\mathcal{M}_i$ in which advisor i randomizes and j babbles, the equilibria $\mathcal{M}_{i,j}$ in which one advisor is truthful and the other mixes, and $\mathcal{M}_{12}$ in which both randomize—are derived in Part B of the Supplementary Appendix, which also contains the proofs of Propositions 4 and 5.

Appendix C: Advisors with Uncertain Trade-off

Equilibrium $\mathcal{E}_{12A}$ The equilibrium $\mathcal{E}_{12A}$ in which advisor 1 and type b_1 of advisor 2 are truthful, while type b_2 reports $\hat{m}_2 = 1$ for both signals, exists if and only if $|b_1| \leq \eta_{12A,1}(p, q)$ and $|b_1| \leq \eta_{12A,2}(p, q)$, i.e., $|b_1| \leq \eta_{12A}(p, q) \equiv \min\{\eta_{12A,1}(p, q), \eta_{12A,2}(p, q)\}$, and $b_2 \geq \eta_{12A,2}(p, q)$. The last condition is that type b_2 prefers $\hat{m}_2 = 1$ when $s_2 = 0$: it is condition (1) for $s_2 = 0$ with the inequality reversed, with the same right-hand side as for type b_1 , since both types face the same actions and the same conditional distribution of s_1 . Simplifying condition (1), we here obtain:

$$\eta_{12A,1}(p, q) = \frac{\sum_{\hat{m}_2=0,1} \Delta_1(\hat{m}_2)^2 P_0(0, \hat{m}_2)}{2 \sum_{\hat{m}_2=0,1} \Delta_1(\hat{m}_2) P_0(0, \hat{m}_2)} \text{ with } \Delta_1(\hat{m}_2) = \frac{P_1(1, \hat{m}_2)}{P_0(1, \hat{m}_2)} - \frac{P_1(0, \hat{m}_2)}{P_0(0, \hat{m}_2)}. \quad (11)$$

Using (6), together with the fact that, here, $\mu_{b_1, s_2}(s_2) = 1$ and $\mu_{b_2, s_2}(1) = 1$ for $s_2 = 0, 1$, we calculate $P_n(s_1, \hat{m}_2)$ for $n = 0, 1, 2$, and $(s_1, \hat{m}_2) \in \{0, 1\}^2$:

$$P_n(s_1, 0) = q \left(p \frac{(s_1 + n)!(2 - s_1)!}{(n + 3)!} + (1 - p) \frac{(n + 1 - s_1)!(1 + s_1)!}{(n + 3)!} \right) \quad (12)$$

$$\begin{aligned} P_n(s_1, 1) &= q \left(p \frac{(s_1 + 1 + n)!(1 - s_1)!}{(n + 3)!} + (1 - p) \frac{(2 + n - s_1)!s_1!}{(n + 3)!} \right) \\ &\quad + (1 - q) \left(p \frac{(s_1 + 1 + n)!(1 - s_1)!}{(n + 3)!} + (1 - p) \frac{(2 + n - s_1)!s_1!}{(n + 3)!} \right) \\ &\quad + (1 - q) \left(p \frac{(s_1 + n)!(2 - s_1)!}{(n + 3)!} + (1 - p) \frac{(1 + n - s_1)!(1 + s_1)!}{(n + 3)!} \right). \end{aligned} \quad (13)$$

Substituting the resulting moments into (11) gives, after simplification (Part D of the Supplementary Appendix):

$$\eta_{12A.1}(p, q) = \frac{1}{4} (2p - 1) \frac{\frac{q}{(p+1)(2-p)^2} + \frac{(6-6q+q^2)^2}{(3-2q+pq)^2(3-q-pq)}}{\frac{q}{2-p} + \frac{6-6q+q^2}{3-2q+pq}}.$$

In the same fashion, we obtain:

$$\eta_{12A.2}(p, q) = \frac{1}{4} (1 + 2p - 2p^2) \frac{\frac{1}{(p+1)(3-q-pq)^2} + \frac{1}{(2-p)(3-2q+pq)^2}}{\frac{1}{3-q-pq} + \frac{1}{3-2q+pq}}.$$

Simplifying expression (4), $V_{12A}(p, q) = -\sum v(s_1, \hat{m}_2)$, where the sum runs over $(s_1, \hat{m}_2) \in \{0, 1\}^2$ and $v(s_1, \hat{m}_2)$ equals $P_2(s_1, \hat{m}_2)$ minus $P_1(s_1, \hat{m}_2)^2/P_0(s_1, \hat{m}_2)$, so that, substituting the moments and simplifying (Part D of the Supplementary Appendix), we obtain

$$\begin{aligned} V_{12A} &= \frac{1}{120} q \frac{-4p + 2p^2 - 1}{p + 1} - \frac{1}{120} q \frac{2p^2 - 3}{p - 2} \\ &\quad - \frac{1}{120} \frac{-20p + 8q - 4pq^2 - 20p^2q + 2p^2q^2 + 24pq + 20p^2 - q^2 - 10}{q + pq - 3} \\ &\quad + \frac{1}{120} \frac{-20p + 12q + 16pq - 20p^2q + 2p^2q^2 + 20p^2 - 3q^2 - 10}{3 - 2q + pq} \end{aligned}$$

Equilibrium $\mathcal{E}_{2A}$ Only advisor 2 of type b_1 is truthful, while type b_2 reports $\hat{m}_2 = 1$ for both signals and advisor 1 babbles. Simplifying expression (1) with the moments $P_n(\hat{m}_2)$ of the leader's information gives the condition $|b_1| \leq \eta_{2A}(q) = 1/(6(2 - q))$ for type b_1 and, by the same reversal as above, $b_2 \geq \eta_{2A}(q)$ for type b_2 ; (4) gives $V_{2A}(q) = -(3 - 2q)/(18(2 - q))$. The steps are in Part D of the Supplementary Appendix.

Proof of Lemma 2. Three comparisons follow from the closed forms. First, $V_{12}(p) > V_2$ because $3(1 + 2p - 2p^2) < 2(p + 1)(2 - p)$ reduces to $(2p - 1)^2 > 0$. Second, $V_2 > V_{2A}(q)$ because $(3 - 2q)/(2 - q) > 1$ for $q < 1$. Third, $V_2 > V_1(p)$ because $1 + 2p - 2p^2 > 1$ for $p \in (1/2, 1)$. The remaining comparisons use the fact that the leader's payoff is minus the expected conditional

variance of the state, which cannot increase when her information is coarsened: if I' is a function of I and of an independent randomization, then $\mathbb{E}[\text{Var}(x|I')] \geq \mathbb{E}[\text{Var}(x|I)]$ by the law of total variance, strictly whenever $\mathbb{E}[x|I] \neq \mathbb{E}[x|I']$ with positive probability. In $\mathcal{E}_{12A}$ advisor 2's message equals s_2 with probability q and 1 otherwise, so $(s_1, \hat{m}_2)$ coarsens (s_1, s_2) and $V_{12A}(p, q) < V_{12}(p)$ for $q < 1$; in $\mathcal{E}_{2A}$ and $\mathcal{E}_1$ the leader observes $\hat{m}_2$ or s_1 alone, so $V_{12A}(p, q) > V_{2A}(q)$ for $p > 1/2$ and $V_{12A}(p, q) > V_1(p)$ for $q > 0$.

For g_1 , $V_{2A}(q) > V_1(p)$ if and only if $(3 - 2q)/(2 - q) < 1 + 2p - 2p^2$, that is, if and only if $2p(1 - p) > (1 - q)/(2 - q)$. As $2p(1 - p)$ is strictly decreasing in p on $(1/2, 1)$, this holds if and only if $p < g_1(q) \equiv \left(1 + \sqrt{q/(2 - q)}\right)/2$, a strictly increasing function with $g_1(0) = 1/2$ and $g_1(1) = 1$.

For g_2 , let $h(p, q) = V_{12A}(p, q) - V_2$. V_{12A} is continuous and strictly increasing in each argument, because lowering q replaces $\hat{m}_2$ by 1 and lowering p flips s_1 , each with an independent probability and so a coarsening in the sense above. For fixed $q \in (0, 1)$, $h \rightarrow V_{2A}(q) - V_2 < 0$ as $p \downarrow 1/2$, where s_1 becomes uninformative, while $h(1, q) > 0$ because $V_{12A}(1, 0) = V_1(1) = V_2$ and V_{12A} increases strictly in q ; hence $h(\cdot, q)$ has a unique root $g_2(q) \in (1/2, 1)$ and $V_2 > V_{12A}(p, q)$ if and only if $p < g_2(q)$. Monotonicity in q makes g_2 strictly decreasing, continuity of h makes it continuous, and $h(p, 0) = V_1(p) - V_2$, $h(p, 1) = V_{12}(p) - V_2$ give $g_2(0) = 1$ and $g_2(1) = 1/2$. ■

Proof of Lemma 3. $\eta_1(p) > \eta_{2A}(q)$ if and only if $2p - 1 > 1/(2 - q)$, that is, if and only if $p > g_3(q) \equiv (3 - q)/(2(2 - q))$, strictly increasing with $g_3(0) = 3/4$ and $g_3(1) = 1$. Since $g_3(q) - g_1(q) = \left[1/(2 - q) - \sqrt{q/(2 - q)}\right]/2 \geq 0$, because $q(2 - q) = 1 - (1 - q)^2 \leq 1$, with equality only at $q = 1$, $g_1 < g_3$ on $[0, 1)$ and $g_3(1) = g_1(1)$. As $g_3(0) < g_2(0) = 1$ and $g_3(1) > g_2(1) = 1/2$, while g_3 increases and g_2 decreases, both continuously, $g_3 - g_2$ changes sign exactly once. ■

The proof below uses the orderings of the bias thresholds stated in Section 6, verified in Part D.

Proof of Proposition 8. The optimal communication equilibrium is the one with the highest consultation value among those that exist, ranked by Lemma 2. If $b_1 \leq \eta_{12.1}(p)$ and $b_2 \leq \eta_{12.2}(p)$, then $\mathcal{E}_{12}$ exists and is top ranked. Otherwise, if $b_1 > \eta_{12A}(p, q)$ and $b_2 \leq \eta_2$, then $\mathcal{E}_2$ exists and $\mathcal{E}_{12A}$ does not, and if $b_1 \leq \eta_{12A}(p, q)$ and $b_2 > \eta_2$, then $\mathcal{E}_{12A}$ exists, as $b_2 > \eta_2 > \eta_{12A.2}(p, q)$, and $\mathcal{E}_2$ does not; in either case the existing one dominates $\mathcal{E}_1$ and $\mathcal{E}_{2A}$, which proves (i). If $b_1 \leq \eta_{12A}(p, q)$ and $b_2 \leq \eta_2$, then $\mathcal{E}_2$ exists, $\mathcal{E}_{12A}$ exists if and only if $b_2 \geq \eta_{12A.2}(p, q)$, both dominate $\mathcal{E}_1$ and $\mathcal{E}_{2A}$, and when both exist g_2 settles the choice between them, which proves (ii). If $b_1 > \eta_{12A}(p, q)$ and $b_2 > \eta_2$, only $\mathcal{E}_1$, which exists if and only if $b_1 \leq \eta_1(p)$, and $\mathcal{E}_{2A}$, which exists if and only if $b_1 \leq \eta_{2A}(q)$, as $b_2 > \eta_2 \geq \eta_{2A}(q)$, remain: g_1 settles the choice when both exist, the existing one is optimal when only one does, and no advisor is truthful when neither does, which proves (iii). ■

Part D of the Supplementary Appendix contains the derivations behind these closed forms, the threshold orderings and the proofs of Propositions 9 and 10; Part E, the results for Section 7.

Supplementary Appendix

Part A: Proof of Theorem 1

Proof. After every on-path history, the leader chooses y to maximize

$$-\mathbb{E}[(y - x)^2 | s_0, \hat{\mathbf{m}}_A].$$

Her unique best response is therefore

$$y(s_0, \hat{\mathbf{m}}_A) = \mathbb{E}[x | s_0, \hat{\mathbf{m}}_A].$$

Fix advisor i and write

$$\alpha_s = \Pr(\hat{m}_i = 1 | s_i = s), \quad s \in \{0, 1\},$$

where the probability integrates over i 's bias type and private randomization. Independence of B_i from the state, signals, and other types implies that α_s does not depend on $z_i = (s_0, \hat{\mathbf{m}}_{-i})$.

If $\alpha_1 = \alpha_0$, advisor i 's message contains no information about his signal. Because his bias is also independent of the state and the other players' information, his message contains no information about the state. The leader therefore ignores it on the equilibrium path. Replacing the possibly type-contingent strategies by pooling on the message conforming to the common direction of bias leaves all equilibrium actions and payoffs unchanged. This proves the stated outcome-equivalent canonical representation.

Suppose instead that $\alpha_1 \neq \alpha_0$. Relabel the messages so that $\alpha_1 > \alpha_0$. Let

$$\lambda_m(z_i) = \Pr(s_i = 1 | \hat{m}_i = m, z_i).$$

Bayes' rule gives, with the usual limiting interpretation when one of the probabilities is zero,

$$\frac{\frac{\lambda_1(z_i)}{1-\lambda_1(z_i)}}{\frac{\lambda_0(z_i)}{1-\lambda_0(z_i)}} = \frac{\alpha_1(1-\alpha_0)}{\alpha_0(1-\alpha_1)} > 1.$$

Hence $\lambda_1(z_i) > \lambda_0(z_i)$.

Set

$$\theta_s(z_i) = \mathbb{E}[x | s_i = s, z_i].$$

Conditional on (s_i, z_i) , advisor i 's message depends only on his bias type and private randomization. Independence of B_i therefore implies

$$\mathbb{E}[x | s_i = s, \hat{m}_i = m, z_i] = \theta_s(z_i).$$

It follows that

$$y(s_0, \hat{m}_i = m, \hat{\mathbf{m}}_{-i}) = \lambda_m(z_i)\theta_1(z_i) + (1 - \lambda_m(z_i))\theta_0(z_i).$$

Combining $\lambda_1(z_i) > \lambda_0(z_i)$ with (2) gives

$$\theta_0(z_i) \leq y(s_0, 0, \hat{\mathbf{m}}_{-i}) < y(s_0, 1, \hat{\mathbf{m}}_{-i}) \leq \theta_1(z_i).$$

Consider a type $b > 0$ who observes $s_i = 1$. Conditional on z_i , his expected loss from inducing action y is

$$\text{Var}(x|s_i = 1, z_i) + (y - \theta_1(z_i) - b)^2.$$

Both available actions are weakly below $\theta_1(z_i)$, and the action following message one is larger. Message one is therefore his strict best reply. Symmetrically, if $b < 0$ and $s_i = 0$, message zero is the strict best reply given the equilibrium decision rule and beliefs.

Let s_i^a denote the signal contrary to the bias. The truthful message at this signal is sent by no type who observes the opposite signal. Independence of the bias type from the state and signals therefore gives

$$y(s_0, \hat{m}_i = s_i^a, \hat{\mathbf{m}}_{-i}) = \mathbb{E}[x|s_i = s_i^a, z_i].$$

Let $d_i(z_i) > 0$ denote the absolute change in the leader's action when the advisor sends the conforming message instead. This change has the same direction as b . Conditional on z_i , the difference between the expert's expected payoff from the truthful message and from the conforming message is

$$d_i(z_i)^2 - 2|b|d_i(z_i).$$

Averaging over z_i conditional on $s_i = s_i^a$, truth-telling is optimal if and only if

$$\sum_{z_i} d_i(z_i)^2 \Pr(z_i|s_i = s_i^a) - 2|b| \sum_{z_i} d_i(z_i) \Pr(z_i|s_i = s_i^a) \geq 0.$$

Since $d_i(z_i) = |\Delta_i(s_0, s_i^a, \hat{\mathbf{m}}_{-i})|$ and Δ_i has the sign of b at every z_i , this is equivalent to (3). The conforming report is optimal if and only if the inequality is reversed, and the type can mix only when equality holds. The distribution of z_i conditional on s_i does not depend on B_i , so the same cutoff applies to every type of advisor i . This proves the strategy characterization.

Conversely, for an informative advisor, the preceding comparisons show that a canonical strategy profile satisfying the cutoff conditions is a best response for every type at both signals. Together with the leader's posterior-mean decision rule, these conditions are sufficient for equilibrium. If every type of an advisor pools, off-path beliefs can assign the unused message the same posterior as the used message at the same z_i , so deviating does not change the leader's action.

Finally, let $I = (s_0, \hat{\mathbf{m}}_A)$ and $Y = \mathbb{E}[x|I]$. Then

$$U_0 = -\mathbb{E}[(Y - x)^2] = -\mathbb{E}[\text{Var}(x|I)].$$

Expanding advisor i 's loss gives

$$U_i = -\mathbb{E}[(Y - x - B_i)^2] = U_0 - \mathbb{E}[B_i^2] + 2\mathbb{E}[B_i(Y - x)].$$

The last term need not vanish because the privately observed bias affects communication. Hence the fixed-bias Pareto-ranking argument does not extend, and equilibria must instead be ranked by the leader's ex-ante payoff. ■

Part B: Mixed Strategy Equilibria

Throughout Part B we maintain $b_1 b_2 > 0$. By simultaneous reflection it is sufficient to derive the results for $0 < b_1 < b_2$; the all-negative case $b_2 < b_1 < 0$ follows by the same reflection.

Equilibria $\mathcal{M}_i$: advisor i randomizes and j babbles

Suppose that $b_i > 0$ and write μ_i for the probability $\mu_{i,0}(0)$ that i reports truthfully when $s_i = 0$; the case $b_i < 0$ follows from the ex-ante symmetry of the signals, with μ_i reading as the probability $\mu_{i,1}(1)$ that i is truthful when $s_i = 1$. Because j babbles, the leader conditions on $\hat{m}_i$ alone, and expressions (7) and (8) yield

$$\begin{aligned} i = 1: & P_0 = \left(\frac{\mu_1}{2}, \frac{2-\mu_1}{2} \right), P_1 = \left(\frac{(2-p)\mu_1}{6}, \frac{3-(2-p)\mu_1}{6} \right), P_2 = \left(\frac{(3-2p)\mu_1}{12}, \frac{4-(3-2p)\mu_1}{12} \right), \\ i = 2: & P_0 = \left(\frac{\mu_2}{2}, \frac{2-\mu_2}{2} \right), P_1 = \left(\frac{\mu_2}{6}, \frac{3-\mu_2}{6} \right), P_2 = \left(\frac{\mu_2}{12}, \frac{4-\mu_2}{12} \right), \end{aligned}$$

where each pair reports the value at $\hat{m}_i = 0$ and at $\hat{m}_i = 1$. In both cases the difference between the posterior means is:

$$\Delta_i = \frac{P_1(1)}{P_0(1)} - \frac{P_1(0)}{P_0(0)} = \frac{2\eta_i}{2 - \mu_i}, \text{ so that, by (1), } |b_i| = \frac{1}{2}\Delta_i = \frac{\eta_i}{2 - \mu_i},$$

where $\eta_1 = (2p - 1)/6$ and $\eta_2 = 1/6$ are the existence thresholds of $\mathcal{E}_1$ and $\mathcal{E}_2$. As the right-hand side increases in μ_i and equals η_i at $\mu_i = 1$, equilibrium $\mathcal{M}_i$ exists only for $|b_i| < \eta_i$, i.e., only if $\mathcal{E}_i$ also exists.

Turning to the value of equilibrium, expression (4) gives $V_{\mathcal{M}_i} = -\sum_{\hat{m}_i=0,1} v(\hat{m}_i)$ with $v(\hat{m}_i) = P_2(\hat{m}_i) - P_1(\hat{m}_i)^2 / P_0(\hat{m}_i)$, and substituting the expressions above yields, again in both cases,

$$V_{\mathcal{M}_i} = \frac{\mu_i V_i - \frac{1-\mu_i}{6}}{2 - \mu_i}, \text{ so that } V_{\mathcal{M}_i} - V_i = -\frac{(1 - \mu_i)(1 + 12V_i)}{6(2 - \mu_i)} < 0$$

for every $\mu_i \in (0, 1)$, where $V_1 = -[1 + 2p(1 - p)]/18$ and $V_2 = -1/18$. The value interpolates between the two pure strategy values, equaling V_i at $\mu_i = 1$ and the babbling value $-1/12$ at $\mu_i = 0$, and the sign of the difference is immediate because $1 + 12V_i > 0$ is precisely the statement that $\mathcal{E}_i$ is more informative than babbling; explicitly, $1 + 12V_1 = (2p - 1)^2/3$ and $1 + 12V_2 = 1/3$. We conclude that equilibrium $\mathcal{M}_i$ is dominated by $\mathcal{E}_i$: its existence requires $\mathcal{E}_i$'s existence, and it yields a lower value.

Equilibrium $\mathcal{M}_{2,1}$: advisor 1 randomizes and advisor 2 tells the truth (Proof of Proposition 4)

Suppose that $b_1 > 0$. The case where $b_1 < 0$ is handled as above. Denote by μ_1 the probability $\mu_{1,0}(0)$ that 1 tells the truth when $s_1 = 0$. Because $\mu_{1,1}(1) = 1$, and $\hat{m}_2 = s_2$ for $s_2 = 0, 1$,

$$\frac{P_1(\hat{m}_1 = 0, \hat{m}_2)}{P_0(\hat{m}_1 = 0, \hat{m}_2)} = \frac{P_1(s_1 = 0, s_2)}{P_0(s_1 = 0, s_2)}. \quad (14)$$

Hence, the calculation of

$$\Delta_1(s_2) = \frac{P_1(\hat{m}_1 = 1, s_2)}{P_0(\hat{m}_1 = 1, s_2)} - \frac{P_1(\hat{m}_1 = 0, s_2)}{P_0(\hat{m}_1 = 0, s_2)}, \text{ for } s_2 = 0, 1$$

requires only using expression (6) to determine:

$$P_n(\hat{m}_1 = 1, s_2) = P_n(s_1 = 1, s_2) + (1 - \mu_1)P_n(s_1 = 0, s_2), \text{ for } n = 0, 1, 2. \quad (15)$$

Plugging in the earlier calculated $P_n(s_1, s_2)$ with $s_1 = 0, 1$ in the above expression and simplifying, we obtain:

$$\begin{aligned} P_0(\hat{m}_1 = 1, 0) &= -\frac{1}{6}(\mu_1 + p\mu_1 - 3), & P_1(\hat{m}_1 = 1, 0) &= -\frac{1}{12}(\mu_1 - 2), \\ P_0(\hat{m}_1 = 1, 1) &= \frac{1}{6}(-2\mu_1 + p\mu_1 + 3), & P_1(\hat{m}_1 = 1, 1) &= \frac{1}{12}(-3\mu_1 + 2p\mu_1 + 4). \end{aligned}$$

From these expressions, we obtain $\Delta_1(s_2)$ for $s_2 = 0, 1$. Plugging them in expression (1) and simplifying, we obtain that advisor 1's equilibrium condition is:

$$b_1 = \eta_{2.1.1}(p, \mu_1) \equiv \frac{1}{4}(2p - 1) \frac{\frac{1}{(2-p)(3-2\mu_1+p\mu_1)^2} + \frac{1}{(p+1)(3-\mu_1-p\mu_1)^2}}{\frac{1}{3-2\mu_1+p\mu_1} + \frac{1}{3-\mu_1-p\mu_1}}.$$

It can be checked that $\eta_{2.1.1}(p, \mu_1)$ increases in μ_1 , that $\eta_{2.1.1}(p, \mu_1) = \eta_{12.1}(p)$ for $\mu_1 = 1$ and that $\eta_{2.1.1}(p, \mu_1) = \eta_{12.1}(p)/2$ for $\mu_1 = 0$. The implication is that, for any $b_1 \notin [\eta_{12.1}(p)/2, \eta_{12.1}(p)]$, the equilibrium $\mathcal{M}_{2.1}$ cannot exist.

Turning to advisor 2, again with $b_2 > 0$, note that (14), implied by $\mu_{1,1}(1) = 1$, yields

$$\Delta_2(\hat{m}_1 = 0) = \Delta_2(s_1 = 0),$$

and the latter has already been calculated for equilibrium $\mathcal{E}_{12}$. Using $P_n(\hat{m}_1 = 1, s_2)$ for $s_2 = 0, 1$ and $n = 0, 1$ calculated above, we obtain $\Delta_2(\hat{m}_1 = 1)$. Plugging $\Delta_2(\hat{m}_1)$, $\hat{m}_1 = 0, 1$, into expression (1) and simplifying, we obtain advisor 2's equilibrium condition:

$$b_2 \leq \eta_{2.1.2}(p, \mu_1) \equiv \frac{1}{4} \frac{\mu_1 \frac{(1+2p-2p^2)^2}{(p+1)(2-p)^2} + \frac{(6-6\mu_1-2p^2\mu_1^2+\mu_1^2+2p\mu_1^2)^2}{(3-2\mu_1+p\mu_1)^2(3-\mu_1-p\mu_1)}}{\mu_1 \frac{1+2p-2p^2}{2-p} + \frac{6-6\mu_1-2p^2\mu_1^2+\mu_1^2+2p\mu_1^2}{3-2\mu_1+p\mu_1}}.$$

It can be checked that $\eta_{2.1.2}(p, \mu_1)$ strictly decreases in μ_1 , $\eta_{2.1.2}(p, \mu_1) = \eta_{12.2}(p)$ for $\mu_1 = 1$ and $\eta_{2.1.2}(p, \mu_1) = \eta_2$ for $\mu_1 = 0$.

We introduce the function $\eta_{M.2} : (p, b_1) \mapsto b_2$ according to the rule that, for every $p \in (1/2, 1)$, $b_1 \in (\eta_{12.1}(p)/2, \eta_{12.1}(p))$ and $\mu_1 \in (0, 1)$, it is the case that $\eta_{M.2}(p, b_1) = \eta_{2.1.2}(p, \mu_1)$ if and only if $b_1 = \eta_{2.1.1}(p, \mu_1)$. We conclude that the existence region of equilibrium $\mathcal{M}_{2.1}$ is $\eta_{12.1}(p)/2 < b_1 < \eta_{12.1}(p)$ and $b_2 \leq \eta_{M.2}(p, b_1)$. The function $\eta_{M.2}$ strictly decreases in b_1 , because $\eta_{2.1.2}$ is strictly decreasing in μ_1 , whereas $\eta_{2.1.1}$ is strictly increasing in μ_1 . For all p , it is further the case that $\eta_{M.2}(p, b_1) = \eta_2$ for $b_1 = \eta_{12.1}(p)/2$ and $\eta_{M.2}(p, b_1) = \eta_{12.2}(p)$ for $b_1 = \eta_{12.1}(p)$.

This concludes the proof of Proposition 4.

To continue, expression (4) gives $V_{M.2.1}(p, b_1) = -\sum v(\hat{m}_1, s_2)$, the sum running over $(\hat{m}_1, s_2) \in \{0, 1\}^2$ and v being as above with the pair $(\hat{m}_1, s_2)$ in place of $\hat{m}_i$. Again, expression (15) gives $P_2(\hat{m}_1 = 1, s_2) = P_2(s_1 = 1, s_2) + (1 - \mu_1)P_2(s_1 = 0, s_2)$ for $s_2 = 0, 1$, and we calculated $P_2(s_1, s_2)$ for equilibrium $\mathcal{E}_{12}$. We obtain:

$$P_2(\hat{m}_1 = 1, 0) = \frac{1}{60}(-3\mu_1 + p\mu_1 + 5), \quad P_2(\hat{m}_1 = 1, 1) = \frac{1}{20}(-4\mu_1 + 3p\mu_1 + 5),$$

where recall that μ_1 is such that $b_1 = \eta_{2.1.1}(p, \mu_1)$, and note that $\eta_{2.1.1}(p, \mu_1)$ depends on p but not on b_2 .

This allows us to calculate $v(\hat{m}_1, s_2)$ for $(\hat{m}_1, s_2) \in \{0, 1\}^2$ and hence the value of equilibrium $\mathcal{M}_{2.1}$ as a function of μ_1 :

$$\tilde{V}_{M.2.1}(p, \mu_1) = -\frac{1}{24}\mu_1 \frac{1+2p-2p^2}{(p+1)(2-p)} - \frac{1}{24}(2-\mu_1) \frac{6-6\mu_1-2p^2\mu_1^2+2p\mu_1^2+\mu_1^2}{(3-\mu_1-p\mu_1)(3-2\mu_1+p\mu_1)}.$$

It can be checked that $\tilde{V}_{M.2.1}$ increases in μ_1 . As $\eta_{2.1.1}$ increases in μ_1 , we conclude that $V_{M.2.1}(p, b_1)$, the value of equilibrium $\mathcal{M}_{2.1}$, increases in b_1 . As μ_1 is pinned down by p and b_1 through the condition $b_1 = \eta_{2.1.1}(p, \mu_1)$, whereas the condition $b_2 \leq \eta_{2.1.2}(p, \mu_1)$ is a weak inequality, the value $V_{M.2.1}(p, b_1)$ is independent of b_2 generically, i.e., when $b_2 \leq \eta_{2.1.2}(p, \mu_1)$ holds as a strict inequality.

Equilibrium $\mathcal{M}_{1.2}$: advisor 2 randomizes and advisor 1 tells the truth.

The derivation retraces the steps taken for $\mathcal{M}_{2.1}$ with the roles of the two advisors interchanged; because the model is not symmetric across advisors, the resulting expressions differ, and we record them. Suppose that $b_2 > 0$ and let $\mu_2 = \mu_{2,0}(0)$ be the probability that 2 tells the truth when $s_2 = 0$. Since $\mu_{2,1}(1) = 1$, the analogue of (14) holds at $\hat{m}_2 = 0$, so that $\Delta_2(s_1)$ requires only $P_n(s_1, \hat{m}_2 = 1) = P_n(s_1, s_2 = 1) + (1 - \mu_2)P_n(s_1, s_2 = 0)$ from (6), which yields:

$$\begin{aligned} P_1(0, \hat{m}_2 = 1) &= \frac{1}{12}(3-2p) + (1-\mu_2)\frac{1}{12}, \quad P_1(1, \hat{m}_2 = 1) = \frac{1}{12}(2p+1) + (1-\mu_2)\frac{1}{12}, \\ P_0(0, \hat{m}_2 = 1) &= \frac{1}{6}(2-p) + (1-\mu_2)\frac{1}{6}(p+1), \quad P_0(1, \hat{m}_2 = 1) = \frac{1}{6}(p+1) + (1-\mu_2)\frac{1}{6}(2-p). \end{aligned}$$

Substituting the resulting $\Delta_2(s_1)$ in (1), advisor 2's equilibrium condition is:

$$b_2 = \eta_{1.2.2}(p, \mu_2) \equiv \frac{1}{4}(1+2p-2p^2) \frac{\frac{1}{(p+1)(3-\mu_2-p\mu_2)^2} + \frac{1}{(2-p)(3-2\mu_2+p\mu_2)^2}}{\frac{1}{3-\mu_2-p\mu_2} + \frac{1}{3-2\mu_2+p\mu_2}},$$

which increases in μ_2 and equals $\eta_{12.2}(p)/2$ at $\mu_2 = 0$ and $\eta_{12.2}(p)$ at $\mu_2 = 1$, so that existence requires $\eta_{12.2}(p)/2 < b_2 < \eta_{12.2}(p)$. For advisor 1, $\mu_{2,1}(1) = 1$ gives $\Delta_1(\hat{m}_2 = 0) = \Delta_1(s_2 = 0)$, calculated for $\mathcal{E}_{12}$, while $\Delta_1(\hat{m}_2 = 1)$ follows from the $P_n(s_1, \hat{m}_2 = 1)$ above; his condition is:

$$b_1 \leq \eta_{1.2.1}(p, \mu_2) = \frac{1}{4}(2p-1) \frac{\frac{\mu_2}{(p+1)(p-2)^2} + \frac{(6-6\mu_2+\mu_2^2)^2}{(3-2\mu_2+p\mu_2)^2(3-\mu_2-p\mu_2)}}{\frac{\mu_2}{2-p} + \frac{6-6\mu_2+\mu_2^2}{3-2\mu_2+p\mu_2}}.$$

Here $\eta_{1.2.1}(p, 1) = \eta_{12.1}(p)$, $\eta_{1.2.1}(p, 0) = \eta_1(p)$ and $\eta_{1.2.1}$ is decreasing in μ_2 . Define $\eta_{1.M} : (p, b_2) \mapsto b_1$ —not to be confused with $\eta_{M.1}$ of Section 4, which bounds b_1 as a function of the consultation cost—by $\eta_{1.M}(p, b_2) = \eta_{1.2.1}(p, \mu_2)$ whenever $b_2 = \eta_{1.2.2}(p, \mu_2)$. Equilibrium $\mathcal{M}_{1.2}$ exists for $\eta_{12.2}(p)/2 < |b_2| < \eta_{12.2}(p)$ and $|b_1| \leq \eta_{1.M}(p, |b_2|)$. Taking the negative of the sum of v over $(s_1, \hat{m}_2) \in \{0, 1\}^2$ in (4), we obtain its value:

$$\tilde{V}_{M.1.2}(p, \mu_2) = -\frac{1}{24}\mu_2 \frac{1+2p-2p^2}{(p+1)(2-p)} - \frac{1}{24}(2-\mu_2) \frac{(6-6\mu_2+\mu_2^2)(1+2p-2p^2)}{(3-\mu_2-p\mu_2)(3-2\mu_2+p\mu_2)}$$

which increases in μ_2 , as does $\eta_{1.2.2}$; hence $V_{M.1.2}$ increases in b_2 and is constant in b_1 whenever $b_1 \leq \eta_{1.2.1}(p, \mu_2)$ holds strictly.

Equilibrium $\mathcal{M}_{12}$: both advisors randomize.

Suppose again that $b_1 > 0$ and $b_2 > 0$. The case in which both are negative follows by simultaneous reflection. Use again notation $\mu_{1,0}(0) = \mu_1$ and $\mu_{2,0}(0) = \mu_2$, and note that $\mu_{1,1}(1) = 1$, $\mu_{2,1}(1) = 1$. Using notation $\hat{P}_n(\hat{m}_1, \hat{m}_2) = P_n(\hat{m}_1, \hat{m}_2)$ to distinguish from $P_n(s_1, s_2)$, for $(s_1, s_2, \hat{m}_1, \hat{m}_2) \in \{0, 1\}^4$, expression (6) yields:

$$\begin{aligned}\hat{P}_n(0, 0) &= \mu_1 \mu_2 P_n(0, 0), \quad \hat{P}_n(1, 0) = \mu_2 P_n(1, 0) + (1 - \mu_1) \mu_2 P_n(0, 0), \\ \hat{P}_n(0, 1) &= \mu_1 P_n(0, 1) + \mu_1 (1 - \mu_2) P_n(0, 0), \\ \hat{P}_n(1, 1) &= P_n(1, 1) + (1 - \mu_1) P_n(0, 1) + (1 - \mu_2) P_n(1, 0) + (1 - \mu_1)(1 - \mu_2) P_n(0, 0).\end{aligned}$$

Substituting in the expressions of $P_n(s_1, s_2)$ for $(s_1, s_2) \in \{0, 1\}^2$ calculated for equilibrium $\mathcal{E}_{12}$, we obtain, after simplification:

$$\begin{aligned}\hat{P}_1(0, 0) &= \mu_1 \mu_2 \frac{1}{12}, \quad \hat{P}_1(1, 0) = \mu_2 \frac{1}{12} + (1 - \mu_1) \mu_2 \frac{1}{12}, \\ \hat{P}_1(0, 1) &= \mu_1 \frac{1}{12} (3 - 2p) + \mu_1 (1 - \mu_2) \frac{1}{12}, \\ \hat{P}_1(1, 1) &= \frac{1}{12} (2p + 1) + (1 - \mu_1) \frac{1}{12} (3 - 2p) + (1 - \mu_2) \frac{1}{12} + (1 - \mu_1)(1 - \mu_2) \frac{1}{12}.\end{aligned}$$

These expressions give $\hat{\Delta}_1(\hat{m}_2)$ for $\hat{m}_2 = 0, 1$ by the usual formula, with $\hat{P}$ in place of P . Plugging them, together with

$$\begin{aligned}P_0(s_1 = 0, \hat{m}_2 = 0) &= \mu_2 P_0(0, 0), \quad P_0(s_1 = 0, \hat{m}_2 = 1) = P_0(0, 1) + (1 - \mu_2) P_0(0, 0), \\ P_0(s_1 = 0, \hat{m}_2 = 0) &= \mu_2 \frac{1}{6} (p + 1), \quad P_0(s_1 = 0, \hat{m}_2 = 1) = \frac{1}{6} (2 - p) + (1 - \mu_2) \frac{1}{6} (p + 1),\end{aligned}$$

in condition (1), and substituting, we obtain the equilibrium condition for advisor 1:

$$b_1 = \bar{\eta}_{12.1}(p, \mu_1, \mu_2) \equiv \frac{1}{4} (2p - 1) \frac{\frac{\mu_2}{(p+1)(3-\mu_1-p\mu_1)^2} + \frac{(-6\mu_2+\mu_2^2+6)^2}{(3-\mu_2-p\mu_2)(6-3\mu_1-3\mu_2+\mu_1\mu_2+p\mu_1\mu_2)^2}}{\frac{\mu_2}{3-\mu_1-p\mu_1} + \frac{(-6\mu_2+\mu_2^2+6)}{6-3\mu_1-3\mu_2+\mu_1\mu_2+p\mu_1\mu_2}}.$$

We note that $\partial \bar{\eta}_{12.1} / \partial \mu_1 > 0$ and $\partial \bar{\eta}_{12.1} / \partial \mu_2 < 0$, while $\bar{\eta}_{12.1}(p, 1, 1) = \eta_{12.1}(p)$, $\bar{\eta}_{12.1}(p, 0, 1) = \eta_{12.1}(p)/2$, $\bar{\eta}_{12.1}(p, 1, 0) = \eta_1(p)$ and $\bar{\eta}_{12.1}(p, 0, 0) = \eta_1(p)/2$. In other terms, the existence interval $B_{12.1}(\mu_2) \equiv \{b_1 : b_1 = \bar{\eta}_{12.1}(p, \mu_1, \mu_2) \text{ for some } \mu_1 \in (0, 1)\}$ in b_1 of $\mathcal{M}_{12}$ goes from $(\eta_{12.1}(p)/2, \eta_{12.1}(p))$ to $(\eta_1(p)/2, \eta_1(p))$ as μ_2 decreases from 1 to 0. Both the lower end and the upper end of $B_{12.1}(\mu_2)$ increase. Since we maintain $b_1 < \eta_{12.1}(p)$, the relevant interval is $B_{12.1}(\mu_2) \cap (0, \eta_{12.1}(p))$, and this intersection shrinks by inclusion as μ_2 decreases from 1 to 0.

Turning to $\bar{\eta}_{12.2}(\mu_1, \mu_2)$, proceeding as above, we obtain:

$$b_2 = \bar{\eta}_{12.2}(p, \mu_1, \mu_2) = \frac{1}{4} \frac{\frac{\mu_1(1+2p-2p^2)^2}{(p+1)(\mu_2+p\mu_2-3)^2} + \frac{(6\mu_1+2p^2\mu_1^2-\mu_1^2-2p\mu_1^2-6)^2}{(3-\mu_1-p\mu_1)(6-3\mu_1-3\mu_2+\mu_1\mu_2+p\mu_1\mu_2)^2}}{\frac{\mu_1(1+2p-2p^2)}{3-\mu_2-p\mu_2} + \frac{6-6\mu_1-2p^2\mu_1^2+\mu_1^2+2p\mu_1^2}{6-3\mu_1-3\mu_2+\mu_1\mu_2+p\mu_1\mu_2}}.$$

We note that $\partial \bar{\eta}_{12.2} / \partial \mu_2 > 0$ and $\partial \bar{\eta}_{12.2} / \partial \mu_1 < 0$, while $\bar{\eta}_{12.2}(p, 1, 1) = \eta_{12.2}(p)$, $\bar{\eta}_{12.2}(p, 1, 0) = \eta_{12.2}(p)/2$, $\bar{\eta}_{12.2}(p, 0, 1) = \eta_2$ and $\bar{\eta}_{12.2}(p, 0, 0) = \eta_2/2 = 1/12$.

The comparison must hold b_1 , rather than μ_1 , fixed. Write $F_i(\mu_1, \mu_2) = \bar{\eta}_{12.i}(p, \mu_1, \mu_2)$ for fixed p . Along the indifference locus $b_1 = F_1(\mu_1, \mu_2)$, the implicit-function theorem gives

$$\frac{d\mu_1}{d\mu_2} = -\frac{\frac{\partial F_1}{\partial \mu_2}}{\frac{\partial F_1}{\partial \mu_1}} > 0.$$

Along the same locus,

$$\frac{db_2}{d\mu_2} = \frac{J}{\frac{\partial F_1}{\partial \mu_1}} > 0, \quad J = \frac{\partial F_2}{\partial \mu_2} \frac{\partial F_1}{\partial \mu_1} - \frac{\partial F_2}{\partial \mu_1} \frac{\partial F_1}{\partial \mu_2}.$$

Direct differentiation of the displayed closed forms establishes $J > 0$ on $p \in (1/2, 1)$ and $(\mu_1, \mu_2) \in (0, 1)^2$. One algebraic sign certificate sets $r = 2p - 1$: after clearing the positive denominator, the numerator of J has 960 positive and 64 zero Bernstein coefficients on the unit cube $(r, \mu_1, \mu_2) \in [0, 1]^3$.

Starting from any interior $\mathcal{M}_{12}$ with $b_1 < \eta_{12.1}(p)$, continue this fixed- b_1 locus until $\mu_2 = 1$, and denote the terminal mixing probability by μ_1^* . Then $\mu_1^* > \mu_1$, the endpoint is the equilibrium $\mathcal{M}_{2.1}$, and the original b_2 lies strictly below its advisor 2 incentive bound. Thus, on this strict domain, every parameter vector supporting $\mathcal{M}_{12}$ also supports $\mathcal{M}_{2.1}$, and the former existence region is a proper subset of the latter. At $b_1 = \eta_{12.1}(p)$, the continuation reaches $\mu_1^* = 1$ and ends at the pure equilibrium $\mathcal{E}_{12}$.

We are left to check that, for any (b_1, b_2) such that $\mathcal{M}_{12}$ and $\mathcal{M}_{2.1}$ coexist, the former is welfare dominated by the latter: $W_{M.12}(p, b_1, b_2) = V_{M.12}(p, b_1, b_2) - 2c < W_{M.2.1}(p, b_1) = V_{M.2.1}(p, b_1) - 2c$.

Proceeding as above, we calculate the value of equilibrium $\mathcal{M}_{12}$ as a function of μ_1 and μ_2 and obtain:

$$\begin{aligned} \tilde{V}_{M.12}(p, \mu_1, \mu_2) = & -\frac{1}{120}\mu_1\mu_2\frac{1+4p-2p^2}{p+1} \\ & -\frac{1}{120}\mu_2\frac{10-2p^2\mu_1^2+4p\mu_1^2-4p\mu_1+\mu_1^2-8\mu_1}{3-\mu_1-p\mu_1} \\ & -\frac{1}{120}\mu_1\frac{\left(\begin{array}{c} 10-2p^2\mu_2^2+20p^2\mu_2-20p^2 \\ +4p\mu_2^2-24p\mu_2+20p+\mu_2^2-8\mu_2 \end{array}\right)}{3-\mu_2-p\mu_2} \\ & -\frac{1}{120}\frac{\left(\begin{array}{c} 60-2p^2\mu_1^2\mu_2^2+20p^2\mu_1^2\mu_2-20p^2\mu_1^2 \\ +4p\mu_1^2\mu_2^2-24p\mu_1^2\mu_2+20p\mu_1^2-4p\mu_1\mu_2^2 \\ +8p\mu_1\mu_2+\mu_1^2\mu_2^2-8\mu_1^2\mu_2+10\mu_1^2 \\ -8\mu_1\mu_2^2+56\mu_1\mu_2-60\mu_1+10\mu_2^2-60\mu_2 \end{array}\right)}{6-3\mu_1-3\mu_2+\mu_1\mu_2+p\mu_1\mu_2}, \end{aligned}$$

For fixed μ_1 , this expression is strictly increasing in μ_2 and satisfies $\tilde{V}_{M.12}(p, \mu_1, 1) = \tilde{V}_{M.2.1}(p, \mu_1)$. At the fixed- b_1 endpoint above, $\mu_1^* > \mu_1$, and increasing either truth-telling probability is a strict Blackwell improvement. Therefore

$$\tilde{V}_{M.12}(p, \mu_1, \mu_2) < \tilde{V}_{M.2.1}(p, \mu_1) < \tilde{V}_{M.2.1}(p, \mu_1^*).$$

The last term is the value of $\mathcal{M}_{2.1}$ at the same p and b_1 as the original $\mathcal{M}_{12}$. Hence, whenever the two equilibria coexist at a parameter vector, $\mathcal{M}_{12}$ has lower value for the leader and lower welfare. This comparison adjusts both mixing probabilities as required by advisor 1's indifference condition.

Proof of Proposition 5 We earlier concluded that the only mixed strategy equilibria that are not dominated by another equilibrium are $\mathcal{M}_{2.1}$ and $\mathcal{M}_{1.2}$. Consider the former. If it coexists with $\mathcal{E}_{12}$, it is dominated by $\mathcal{E}_{12}$; both configurations nevertheless consult both advisors. Further, we know that $V_{M.2.1}(p, b_1) > V_2$ for all b_1, b_2 and p such that equilibrium $\mathcal{M}_{2.1}$ exists, and that $V_{M.2.1}(p, b_1)$ increases in b_1 . Hence, for every s_1 precision p and consultation cost c , there exists a threshold $\eta_{M.1}(p, c)$ such that $W_{M.2.1}(p, b_1) = V_{M.2.1}(p, b_1) - 2c > (<)W_2 = V_2 - c$ if $|b_1| > (<)\eta_{M.1}(p, c)$. Provided that equilibrium $\mathcal{M}_{2.1}$ exists, the leader hires both advisors 1 and 2 for $|b_1| > \eta_{M.1}(p, c)$ and only advisor 2 if $|b_1| < \eta_{M.1}(p, c)$. The threshold $\eta_{M.1}$ increases in c and eventually crosses $\eta_{12.1}(p)$ so that $W_{M.2.1}(p, b_1) < W_2$ for all b_1, b_2 and p such that equilibrium $\mathcal{M}_{2.1}$ exists. Beyond that high-cost crossing, the leader never hires advisor 1 to achieve equilibrium $\mathcal{M}_{2.1}$. Together with the existence conditions $\eta_{12.1}(p)/2 < |b_1| < \eta_{12.1}(p)$ and $|b_2| \leq \eta_{M.2}(p, b_1)$, this concludes the optimality characterization of equilibrium $\mathcal{M}_{2.1}$. As for $\mathcal{M}_{1.2}$, when $|b_1| < \eta_{12.1}(p)$ it exists only where $\mathcal{E}_{12}$ does, and its value $\tilde{V}_{M.1.2}(p, \mu_2)$ increases in μ_2 up to $V_{12}(p)$ at $\mu_2 = 1$, so that it is dominated by $\mathcal{E}_{12}$.

For the second result, suppose that b_1, b_2, p, c are such that the leader hires both advisors to achieve equilibrium $\mathcal{E}_{12}$, i.e., $|b_1| \leq \eta_{12.1}(p)$, $|b_2| \leq \eta_{12.2}(p)$ and $c < \bar{c}_{12}(p)$. Say that $|b_2|$ increases continuously. If it is the case that $|b_1| < \eta_{M.1}(p, c)$, then as $|b_2|$ crosses $\eta_{12.2}(p)$, the optimal equilibrium is $\mathcal{E}_2$ as $\mathcal{M}_{2.1}$ either is not available, or it is welfare dominated by $\mathcal{E}_2$. In either case, the leader dismisses 1 and retains 2. If instead it is the case that $|b_1| \geq \eta_{M.1}(p, c)$, then as $|b_2|$ crosses $\eta_{12.2}(p)$ the leader keeps consulting both 1 and 2. However, as $|b_2|$ increases it will eventually cross $\eta_{M.2}(p, b_1)$ which is smaller than η_2 . When it does, again, the leader dismisses 1 and retains 2. Regardless of the value of b_1 , as b_2 grows further, it will cross η_2 and then the leader will dismiss advisor 2 and rehire 1.

Part C: Continuous Signals

This part contains the derivations and proofs for the continuous-signal analysis in Section 5. We first characterize two-message communication with one and two advisors, then establish the global incentive-compatibility conditions, and finally present the finer-partition calculations and numerical analysis used to define the comparison region $\mathcal{R}$.

Throughout, $x \sim U[0, 1]$, $s_i = x + \varepsilon_i$, and the noise terms are mutually independent and jointly independent of x , with $\varepsilon_i \sim U[-\delta_i, \delta_i]$. For a signal s with noise parameter δ , Bayes' rule gives

$$x|s \sim U[\max\{0, s - \delta\}, \min\{1, s + \delta\}]$$

and posterior mean

$$\theta_\delta(s) = \frac{\max\{0, s - \delta\} + \min\{1, s + \delta\}}{2}.$$

The posterior mean is nondecreasing in s and equals s whenever $\delta < s < 1 - \delta$. The truncation at the boundaries of the state space is incorporated directly in the message probabilities and posterior moments below.

Single-advisor communication. Let advisor i send m_L^i when $s_i < t_i$ and m_H^i otherwise, with $t_i \in (-\delta_i, 1 + \delta_i)$. Parametrize the threshold by the posterior mean of the marginal signal,

$$\theta = \theta_{\delta_i}(t_i) \in (0, 1), \quad \nu = 1 - \theta.$$

Let $\alpha(\theta)$ denote the probability of the low message and let $\gamma(\theta) = y_H^i - y_L^i$ denote the distance between the leader's two induced actions. Direct integration gives

$$\alpha(\theta) = \begin{cases} \frac{\theta^2}{\delta_i}, & \theta \leq \delta_i, \\ \theta, & \delta_i \leq \theta \leq 1 - \delta_i, \\ 1 - \frac{\nu^2}{\delta_i}, & \theta \geq 1 - \delta_i, \end{cases} \quad \gamma(\theta) = \begin{cases} \frac{\delta_i(3 - 4\theta)}{6(\delta_i - \theta^2)}, & \theta \leq \delta_i, \\ \frac{3\theta\nu - \delta_i^2}{6\theta\nu}, & \delta_i \leq \theta \leq 1 - \delta_i, \\ \frac{\delta_i(3 - 4\nu)}{6(\delta_i - \nu^2)}, & \theta \geq 1 - \delta_i. \end{cases} \quad (16)$$

Because $\mathbb{E}[y] = \mathbb{E}[x] = 1/2$, the leader's actions and value are

$$y_L^i = \frac{1}{2} - (1 - \alpha)\gamma, \quad y_H^i = \frac{1}{2} + \alpha\gamma, \quad V_i = \alpha(1 - \alpha)\gamma^2 - \frac{1}{12}. \quad (17)$$

The marginal type is indifferent if and only if $b_i = (y_L^i + y_H^i)/2 - \theta$, or, equivalently, $(1 - 2\alpha)\gamma = 1 - 2\theta - 2b_i$. Define

$$\Psi_{\delta_i}(\theta) = \frac{1}{2} - \theta - \frac{1 - 2\alpha(\theta)}{2}\gamma(\theta). \quad (18)$$

Proof of Proposition 6. Write

$$G_i(x) = \Pr(s_i < t_i | x) = \min \left\{ 1, \max \left\{ 0, \frac{t_i + \delta_i - x}{2\delta_i} \right\} \right\}.$$

Integrating $G_i(x)$ and $xG_i(x)$ over $[0, 1]$ for the three possible positions of the switching region yields (16); (17) then follows from $\mathbb{E}[y] = 1/2$. Substitution of the marginal-indifference condition gives $b_i = \Psi_{\delta_i}(\theta)$. By the argument in the proof of Proposition 7 below, marginal indifference is also sufficient when only one advisor is consulted.

It remains to establish the properties of Ψ_{δ_i} . For $\delta_i < \theta < 1 - \delta_i$, we have $\alpha = \theta$ and $\gamma = 1/2 - \delta_i^2/(6\theta\nu)$. Hence

$$\Psi_{\delta_i}(\theta) = \frac{1}{4} - \frac{\theta}{2} + \frac{(1 - 2\theta)\delta_i^2}{12\theta\nu},$$

with derivative

$$\Psi'_{\delta_i}(\theta) = -\frac{1}{2} - \frac{\delta_i^2(\theta^2 + \nu^2)}{12(\theta\nu)^2} < -\frac{1}{2}.$$

For $0 < \theta < \delta_i$, we have $\alpha = \theta^2/\delta_i$, which gives

$$\Psi_{\delta_i}(\theta) = \frac{1}{2} - \theta - \frac{(\delta_i - 2\theta^2)(3 - 4\theta)}{12(\delta_i - \theta^2)}$$

and

Profile	Probability	$\mathbb{E}[x\mathbf{1}\{\text{profile}\}]$	Posterior action
LL	C	D	$\frac{D}{C}$
LH	$A_1 - C$	$B_1 - D$	$\frac{B_1 - D}{A_1 - C}$
HL	$A_2 - C$	$B_2 - D$	$\frac{B_2 - D}{A_2 - C}$
HH	$1 - A_1 - A_2 + C$	$\frac{1}{2} - B_1 - B_2 + D$	$\frac{\frac{1}{2} - B_1 - B_2 + D}{1 - A_1 - A_2 + C}$

Table 1: Message probabilities and posterior actions with two advisors. An action is defined by Bayes' rule only when the corresponding profile has positive probability.

$$\Psi'_{\delta_i}(\theta) = -\frac{4\delta_i^2 - 3\delta_i\theta - 2\delta_i\theta^2 + 2\theta^4}{6(\delta_i - \theta^2)^2}.$$

For $0 < \theta < \delta_i < 1/2$, the numerator is greater than $(\delta_i - \theta^2)^2/2$, so $\Psi'_{\delta_i}(\theta) \leq -1/12$. The corresponding result for $1 - \delta_i < \theta < 1$ follows by symmetry. The formulas agree at $\theta = \delta_i$ and $\theta = 1 - \delta_i$, so Ψ_{δ_i} is continuous and strictly decreasing on $(0, 1)$. Moreover, $\Psi_{\delta_i}(0^+) = 1/4$, $\Psi_{\delta_i}(1^-) = -1/4$. It therefore maps $(0, 1)$ one-to-one onto $(-1/4, 1/4)$. Hence an informative two-interval equilibrium exists if and only if $|b_i| < 1/4$, and, when it exists, it is unique within the class of informative two-interval equilibria.

■

Two-advisor communication. Let both advisors use two-interval partitions with unrestricted thresholds t_1, t_2 , and define

$$G_i(x) = \Pr(s_i < t_i | x), \quad A_i = \int_0^1 G_i(x) dx, \quad B_i = \int_0^1 x G_i(x) dx,$$

together with

$$C = \int_0^1 G_1(x) G_2(x) dx, \quad D = \int_0^1 x G_1(x) G_2(x) dx.$$

Conditional independence of the two noise terms given x gives the message-profile probabilities, first moments, and posterior actions in Table 1.

The leader's value is therefore

$$V_{12} = -\frac{1}{3} + \sum_{j,k \in \{L,H\}} \Pr(jk) y_{jk}^2 = \text{Var}(y) - \frac{1}{12}. \quad (19)$$

Proof of Proposition 7. The probability and first-moment columns of Table 1 follow immediately from conditional independence and subtraction. For example,

$$\Pr(LL) = \int_0^1 G_1(x)G_2(x) dx = C, \quad \mathbb{E}[x\mathbf{1}\{LL\}] = D,$$

so $y_{LL} = D/C$ whenever $C > 0$; the other three profiles follow analogously. Equation (19) uses $\mathbb{E}[y] = \mathbb{E}[x] = 1/2$.

Because $C \leq \min\{A_1, A_2\}$, one mixed profile has probability zero exactly when the two switching regions have disjoint interiors. The action following such a profile is off path and is therefore not pinned down by on-path Bayes' rule. Indifference of consultant i 's marginal type supplies one equation in (t_1, t_2) for given (b_1, b_2) .

Global incentive compatibility. The marginal conditions identify candidate threshold pairs, but with two advisors they need not guarantee equilibrium. The following expression makes the additional incentive constraints explicit and is also useful for the finer-partition calculations below.

Suppose advisor 1 uses N intervals and advisor 2 uses M . Let

$$\pi_k(s) = \Pr(m_2^k | s_1 = s), \quad \xi_k(s) = \mathbb{E}[x\mathbf{1}\{m_2^k\} | s_1 = s],$$

and let $y_{j,k}$ be the leader's action after advisor 1's j -th and advisor 2's k -th messages. Up to a term independent of advisor 1's message, his expected payoff from message j equals

$$\sum_k [2y_{j,k}(\xi_k(s) + b_1\pi_k(s)) - y_{j,k}^2\pi_k(s)].$$

Hence the gain from sending m_1^{j+1} instead of m_1^j is

$$\sum_k (y_{j+1,k} - y_{j,k}) [2\xi_k(s) + (2b_1 - y_{j,k} - y_{j+1,k})\pi_k(s)]. \quad (20)$$

The corresponding expression for advisor 2 is symmetric. Evaluating (20) at a partition threshold and setting it equal to zero gives the marginal-indifference condition at that threshold.

Given the thresholds, Bayes' rule determines the leader's actions after all positive-probability profiles. For a zero-probability profile, sequential rationality requires the specified action to be a posterior mean in $[0, 1]$. The thresholds and these actions form an equilibrium if and only if no signal of either advisor gains by switching message. This is exactly the global condition stated in the proposition.

When only advisor i is consulted, (20) reduces to

$$(y_H^i - y_L^i) [2\theta_{\delta_i}(s) + 2b_i - y_L^i - y_H^i].$$

Since $y_H^i - y_L^i > 0$ and $\theta_{\delta_i}(s)$ is nondecreasing in s , the gain is nondecreasing in the signal. It vanishes at t_i by marginal indifference, and is therefore nonpositive below the threshold and nonnegative above it. Thus marginal indifference is sufficient with one advisor.

With two advisors, both the conditional target and the probabilities $\pi_k(s)$ in (20) vary with the expert's signal, so marginal indifference alone does not establish the required incentive inequalities. ■

Finer partitions with one advisor. An N -interval equilibrium of advisor i consulted alone is described by thresholds

$$t^1 < \dots < t^{N-1}$$

in $(-\delta_i, 1 + \delta_i)$, with every interval having positive probability and every threshold type indifferent between the adjacent actions. Since the posterior mean is monotone, these marginal conditions imply global incentive compatibility. Let $\bar{b}_{\text{gen}}^{(N)}(\delta_i)$ be the supremum of $b_i > 0$ over such equilibria. Proposition 6 gives $\bar{b}_{\text{gen}}^{(2)} = 1/4$ for every δ_i .

Let $(\theta(b), y_L^i(b), y_H^i(b))$ denote the unique two-interval equilibrium at bias $b \in (0, 1/4)$.

Proof of Lemma 1.

Step 1: the boundary equation.

Define $F(b) = y_L^i(b)/2 - b$, $b \in (0, 1/4)$.

We first show that F has a unique zero.

Along the two-interval equilibrium, $b = \Psi_{\delta_i}(\theta)$. Since $b > 0$ and Ψ_{δ_i} is strictly decreasing, $\theta < 1 - \delta_i$. Hence only the ranges $0 < \theta \leq \delta_i$ and $\delta_i \leq \theta < 1 - \delta_i$ are relevant. In these two ranges,

$$y_L^i = \begin{cases} \frac{2\theta}{3}, & 0 < \theta \leq \delta_i, \\ \frac{\theta}{2} + \frac{\delta_i^2}{6\theta}, & \delta_i \leq \theta < 1 - \delta_i. \end{cases}$$

For $\theta \neq \delta_i$ in these ranges, $dy_L^i/d\theta \geq 1/3$.

From the proof of Proposition 6, $\Psi'_{\delta_i}(\theta) \leq -1/12$ at these points. It follows that

$$\frac{d}{d\theta} \left[\frac{y_L^i}{2} - \Psi_{\delta_i}(\theta) \right] > 0.$$

The expression in brackets is continuous at $\theta = \delta_i$, so it is strictly increasing throughout these ranges. Since Ψ_{δ_i} is strictly decreasing, θ falls as b rises. Hence $F(b)$ is strictly decreasing in b .

As $b \downarrow 0$, $\theta \rightarrow 1/2$ and $y_L^i \rightarrow 1/4 + \delta_i^2/3$, whereas as $b \uparrow 1/4$, $\theta \downarrow 0$ and $y_L^i \downarrow 0$. Thus

$$F(0^+) = \frac{1}{8} + \frac{\delta_i^2}{6} > 0, \quad F\left(\frac{1}{4}\right) = -\frac{1}{4}.$$

Hence the boundary equation $F(b) = 0$ has a unique solution, at which $b = y_L^i/2$. Combining this condition with marginal indifference, $b_i = (y_L^i + y_H^i)/2 - \theta$, gives $\theta = y_H^i/2$.

Step 2: the vanishing-interval limit.

Consider a sequence of three-interval equilibria with thresholds $t_k^1 < t_k^2$ in which the probability of the lowest interval converges to zero. Then $t_k^1 \rightarrow -\delta_i$.

Consequently, both the leader's action following the lowest message and the posterior mean of the marginal type at t_k^1 converge to zero.

Let y_k^1, y_k^2, y_k^3 denote the three induced actions. Indifference at the first threshold requires $\theta_{\delta_i}(t_k^1) + b_k = (y_k^1 + y_k^2)/2$. Since $\theta_{\delta_i}(t_k^1) \rightarrow 0$ and $y_k^1 \rightarrow 0$, $b_k - y_k^2/2 \rightarrow 0$. As the lowest interval disappears, the remaining two intervals converge to a two-interval equilibrium. By Proposition 6, this equilibrium is uniquely determined within the two-interval class by the limiting bias b , and $y_k^2 \rightarrow y_L^i(b)$. Therefore $b = y_L^i(b)/2$, so $F(b) = 0$. By Step 1, the limiting bias must be $\bar{b}_{\text{gen}}^{(3)}(\delta_i)$.

Step 3: closed form.

Let a_i denote the posterior mean θ at $b = \bar{b}_{\text{gen}}^{(3)}(\delta_i)$. From Step 1, $a_i = y_H^i/2$. Suppose first that $a_i \geq \delta_i$. In this range,

$$y_H^i = \frac{1 + a_i}{2} - \frac{\delta_i^2}{6(1 - a_i)}.$$

The condition $a_i = y_H^i/2$ becomes $9a_i^2 - 12a_i + 3 - \delta_i^2 = 0$, whose relevant root is $a_i = (2 - \sqrt{1 + \delta_i^2})/3$. The condition $a_i \geq \delta_i$ is equivalent to $\delta_i \leq (3 - \sqrt{3})/4$.

Using $\bar{b}_{\text{gen}}^{(3)} = y_L^i/2$ then gives

$$\bar{b}_{\text{gen}}^{(3)}(\delta_i) = \frac{3a_i^2 + \delta_i^2}{12a_i}. \quad (21)$$

For $\delta_i \geq (3 - \sqrt{3})/4$, we instead have $a_i \leq \delta_i$. In this range, the condition $a_i = y_H^i/2$ reduces to

$$8a_i^3 - 12\delta_i a_i + 3\delta_i = 0. \quad (22)$$

This equation has a unique relevant root $a_i \in (1/4, \delta_i]$, and since $y_L^i = 2a_i/3$,

$$\bar{b}_{\text{gen}}^{(3)}(\delta_i) = \frac{a_i}{3}. \quad (23)$$

Step 4: comparative statics.

For $\delta_i \leq (3 - \sqrt{3})/4$,

let $\tau = \sqrt{1 + \delta_i^2}$, so that $a_i = (2 - \tau)/3$. Then

$$\bar{b}_{\text{gen}}^{(3)}(\delta_i) - \frac{1}{12} = \frac{(\tau - 1)(4\tau + 1)}{36a_i} > 0,$$

and $\bar{b}_{\text{gen}}^{(3)}(\delta_i)$ is strictly increasing in δ_i . For $\delta_i \geq (3 - \sqrt{3})/4$, we have $\bar{b}_{\text{gen}}^{(3)}(\delta_i) = a_i/3 > 1/12$. Implicit differentiation of (22) gives

$$\frac{da_i}{d\delta_i} = \frac{3(4a_i - 1)}{24a_i^2 - 12\delta_i} < 0,$$

so $\bar{b}_{\text{gen}}^{(3)}(\delta_i)$ is strictly decreasing in δ_i .

The two expressions coincide at $\delta_i = (3 - \sqrt{3})/4$, where $\bar{b}_{\text{gen}}^{(3)}(\delta_i) = (3 - \sqrt{3})/12$. This is therefore the maximum of $\bar{b}_{\text{gen}}^{(3)}(\delta_i)$.

■

Lemma 4 *Under these conditions, for every $N \geq 3$, $\bar{b}_{\text{gen}}^{(N+1)} < \bar{b}_{\text{gen}}^{(N)}$, and the boundary $\bar{b}_{\text{gen}}^{(N+1)}$ satisfies $\bar{b}_{\text{gen}}^{(N+1)} = y_1^{(N)}(\bar{b}_{\text{gen}}^{(N+1)})/2$. Moreover, $\bar{b}_{\text{gen}}^{(3)} < \bar{b}_{\text{gen}}^{(2)} = 1/4$.*

Proof.

Consider a sequence of $(N + 1)$ -interval equilibria approaching $\bar{b}_{\text{gen}}^{(N+1)}$ in which the lowest interval vanishes. By assumption, the remaining N intervals converge to an N -interval equilibrium at the same limiting bias. Hence, $\bar{b}_{\text{gen}}^{(N+1)} \leq \bar{b}_{\text{gen}}^{(N)}$.

Suppose, to the contrary, that $\bar{b}_{\text{gen}}^{(N+1)} = \bar{b}_{\text{gen}}^{(N)}$. As the lowest interval of the $(N + 1)$ -interval equilibrium vanishes, its action and the posterior mean of its marginal type both converge to zero. The corresponding marginal-indifference condition therefore implies

$$\bar{b}_{\text{gen}}^{(N+1)} = \frac{1}{2} \lim_{b \uparrow \bar{b}_{\text{gen}}^{(N)}} y_1^{(N)}(b).$$

	$M = 2$	$M = 3$	$M = 4$
$N = 3$	1.002162	0.781261	0.426761
$N = 4$	1.002603	0.862351	0.716093

Table 2: Largest ratio $b_1/\bar{b}_{\text{gen}}^{(N)}(\delta_1)$ among globally incentive-compatible (N, M) equilibria with $0 < b_1 < b_2 < 1/4$, maximized over $(\delta_1, \delta_2) \in \{(0.05, 0.01), (0.10, 0.05), (0.15, 0.05), (0.20, 0.10), (0.25, 0.005), (0.30, 0.10), (0.35, 0.20), (0.40, 0.05), (0.45, 0.30), (0.49, 0.25)\}$. A ratio below one means that advisor 1 can sustain N intervals when consulted alone at the same bias.

By assumption, the limit on the right-hand side is zero. This would imply $\bar{b}_{\text{gen}}^{(N+1)} = 0$, a contradiction. Therefore, $\bar{b}_{\text{gen}}^{(N+1)} < \bar{b}_{\text{gen}}^{(N)}$.

Since $\bar{b}_{\text{gen}}^{(N+1)}$ lies strictly below $\bar{b}_{\text{gen}}^{(N)}$, the limiting N -interval equilibrium is nondegenerate. The same marginal-indifference condition then gives

$$\bar{b}_{\text{gen}}^{(N+1)} = \frac{y_1^{(N)}(\bar{b}_{\text{gen}}^{(N+1)})}{2}.$$

For $N = 2$,

$$\bar{b}_{\text{gen}}^{(3)}(\delta_i) = \frac{y_L^i(\bar{b}_{\text{gen}}^{(3)}(\delta_i))}{2} < \frac{1}{4},$$

because $y_L^i < 1/2$. Hence $\bar{b}_{\text{gen}}^{(3)}(\delta_i) < \bar{b}_{\text{gen}}^{(2)}(\delta_i)$. ■

Finer partitions with two advisors. We next examine numerically whether consulting a second advisor can support a finer partition than an advisor could sustain alone. If an (N, M) two-advisor equilibrium exists, the relevant comparison for advisor 1 is whether the same bias satisfies

$$b_1 < \bar{b}_{\text{gen}}^{(N)}(\delta_1).$$

For $N = 2$ this holds immediately because $\bar{b}_{\text{gen}}^{(2)} = 1/4$ and the analysis restricts attention to $b_1 < 1/4$. For $N \geq 3$ we verify the comparison numerically; the corresponding exercise for advisor 2 is symmetric.

For a given (N, M) , there are $(N - 1) + (M - 1)$ thresholds and $(N - 2) + (M - 2)$ consistency conditions. The two lowest thresholds therefore parametrize a two-dimensional equilibrium family, while the biases are recovered from the marginal-indifference conditions. We grid the two lowest thresholds, refine the grid near vanishing-interval boundaries, solve the remaining consistency conditions to 10^{-12} , retain solutions with $0 < b_1 < b_2 < 1/4$, and impose the global incentive constraints (20) for every signal of both advisors.

Table 2 reports the largest ratio $b_1/\bar{b}_{\text{gen}}^{(N)}(\delta_1)$ among globally incentive-compatible solutions for $N \in \{3, 4\}$ and $M \in \{2, 3, 4\}$, maximized over ten precision pairs spanning $0 < \delta_2 < \delta_1 < 1/2$.

The two ratios in Table 2 that slightly exceed one occur only at near-degenerate boundary configurations: in every such case, one message has probability between 1.3×10^{-12} and 3.1×10^{-12} . Table 3 summarizes how the maximal ratios change when these boundary configurations are progressively excluded. The advisor 2 calculations repeat the same exercise after interchanging the advisors' roles.

Minimum message probability	$\max b_1/\bar{b}_{\text{gen}}^{(N)}(\delta_1)$	$\max b_2/\bar{b}_{\text{gen}}^{(M)}(\delta_2)$
No floor	1.002603	1.0075
10^{-6}	0.998744	1.0022
One percent	0.773428	0.757

Table 3: Sensitivity of the finer-partition comparison to a minimum message-probability requirement. The maxima are taken over the same precision grid and partition sizes described in the text.

At the one-percent screen used for the comparisons in Section 5, both maximal ratios are well below one. Thus, over the numerically examined configurations, an advisor who could not sustain more than two intervals alone does not acquire a nondegenerate finer partition merely because the other advisor is also consulted. This is the numerical basis for restricting the main comparison to two-interval equilibria on $\mathcal{R}$.

Numerical implementation. Three features of the computation are worth recording.

1. *Equilibrium parametrization.* With two advisors and two intervals each, there are two thresholds and no additional consistency conditions. Hence

$$(t_1, t_2) \mapsto (b_1, b_2, V_{12})$$

is a map. Sweeping the threshold plane rather than solving separately at each bias pair finds equilibria with thresholds anywhere in $(-\delta_i, 1 + \delta_i)$, including boundary configurations that a solver initialized near the interior can miss. The threshold grid is refined geometrically down to 10^{-9} from either end. When several equilibria correspond to the same bias pair, we retain the one giving the leader the highest value.

2. *Posterior moments.* Only C and D depend jointly on the geometry of the two switching regions. The integrands $G_1 G_2$ and $x G_1 G_2$ are piecewise polynomials of degree at most two and three, respectively. After splitting the integrals at the sorted switching-region endpoints, three-point Gaussian quadrature is exact.

3. *Incentive compatibility.* A candidate satisfying the marginal conditions can fail global incentive compatibility on a narrow set of signals. The unnormalized gain in (20) is piecewise cubic in the advisor’s signal. We therefore split the signal space at all support, threshold, clipping, and induced switching-region knots and evaluate each polynomial piece at its endpoints and real stationary points. For general (N, M) partitions, the prescribed message is compared with every alternative message in the same way. When a message profile is off path, the leader’s action following it is searched over a grid in $[0, 1]$.

In the reported main-text comparisons we require every message to occur with probability at least one percent, thereby excluding near-degenerate partitions. Table 3 reports the corresponding finer-partition calculation under smaller probability floors as well.

Finally, the implementation was cross-checked against an independent code path written for general (N, M) partitions. Across all 397 configurations tested, message probabilities agree to 3×10^{-15} and implied biases to 7×10^{-10} . The scripts are reproduced in Part F.

Part D: Derivations for Section 6

Equilibrium $\mathcal{E}_{12A}$. We derive the closed forms of $\eta_{12A.1}$, $\eta_{12A.2}$ and V_{12A} stated in Appendix C. Using expressions (12) and (13), we obtain, after simplification:

$$\begin{aligned} P_0(0,0) &= \frac{q}{6} + \frac{pq}{6}, & P_1(0,0) &= \frac{q}{12}, \\ P_0(1,0) &= \frac{q}{3} - \frac{pq}{6}, & P_1(1,0) &= \frac{q}{12}, \\ P_0(0,1) &= -\frac{q}{6} - \frac{pq}{6} + \frac{1}{2}, & P_1(0,1) &= -\frac{p}{6} - \frac{q}{12} + \frac{1}{3}, \\ P_0(1,1) &= \frac{1}{2} - \frac{q}{3} + \frac{pq}{6}, & P_1(1,1) &= \frac{1}{6} + \frac{p}{6} - \frac{q}{12}. \end{aligned}$$

Plugging these expressions into the formula for

$$\Delta_1(\hat{m}_2) = \frac{P_1(1, \hat{m}_2)}{P_0(1, \hat{m}_2)} - \frac{P_1(0, \hat{m}_2)}{P_0(0, \hat{m}_2)}, \text{ for } \hat{m}_2 = 0, 1,$$

and simplifying, we get:

$$\begin{aligned} \Delta_1(0) &= \frac{\frac{1}{12}q}{\frac{1}{3}q - \frac{1}{6}pq} - \frac{\frac{1}{12}q}{\frac{1}{6}q + \frac{1}{6}pq} = -\frac{1}{2} \frac{2p-1}{(p+1)(p-2)} \\ \Delta_1(1) &= \frac{\frac{1}{6} + \frac{1}{6}p - \frac{1}{12}q}{\frac{1}{2} - \frac{1}{3}q + \frac{1}{6}pq} - \frac{-\frac{1}{6}p - \frac{1}{12}q + \frac{1}{3}}{-\frac{1}{6}q - \frac{1}{6}pq + \frac{1}{2}} \\ &= -\frac{1}{2} \frac{(2p-1)(6-6q+q^2)}{(3-2q+pq)(q+pq-3)}. \end{aligned}$$

The threshold $\eta_{12A.1}(p, q)$ is obtained by plugging $\Delta_1(\hat{m}_2)$ and $P_0(0, \hat{m}_2)$, $\hat{m}_2 \in \{0, 1\}$, in the formula:

$$\begin{aligned} \eta_{12A.1} &= \frac{1}{2} \frac{\Delta_1(0)^2 P_0(0,0) + \Delta_1(1)^2 P_0(0,1)}{\Delta_1(0) P_0(0,0) + \Delta_1(1) P_0(0,1)} \\ &= \frac{1}{2} \frac{\left(-\frac{1}{2} \frac{2p-1}{(p+1)(p-2)}\right)^2 \left(\frac{1}{6}q + \frac{1}{6}pq\right) + \left(-\frac{1}{2} \frac{(2p-1)(6-6q+q^2)}{(3-2q+pq)(q+pq-3)}\right)^2 \left(-\frac{1}{6}q - \frac{1}{6}pq + \frac{1}{2}\right)}{\left(-\frac{1}{2} \frac{2p-1}{(p+1)(p-2)}\right) \left(\frac{1}{6}q + \frac{1}{6}pq\right) + \left(-\frac{1}{2} \frac{(2p-1)(6-6q+q^2)}{(3-2q+pq)(q+pq-3)}\right) \left(-\frac{1}{6}q - \frac{1}{6}pq + \frac{1}{2}\right)}. \end{aligned}$$

Similarly, $\eta_{12A.2}(p, q)$ is obtained with the formula:

$$\eta_{12A.2} = \frac{\sum_{s_1=0,1} \Delta_2(s_1)^2 \Pr(s_1, s_2=0)}{2 \sum_{s_1 \in \{0,1\}} \Delta_2(s_1) \Pr(s_1, s_2=0)}, \text{ where } \Delta_2(s_1) = \frac{P_1(s_1, 1)}{P_0(s_1, 1)} - \frac{P_1(s_1, 0)}{P_0(s_1, 0)} \text{ for } s_1 = 0, 1,$$

Hence, using again the above expressions $P_0(s_1, \hat{m}_2)$, $s_1 = 0, 1$, $\hat{m}_2 = 0, 1$,

$$\begin{aligned} \Delta_2(0) &= \frac{-\frac{1}{6}p - \frac{1}{12}q + \frac{1}{3}}{-\frac{1}{6}q - \frac{1}{6}pq + \frac{1}{2}} - \frac{\frac{1}{12}q}{\frac{1}{6}q + \frac{1}{6}pq} = \frac{1}{2} \frac{-2p + 2p^2 - 1}{(p+1)(q+pq-3)} \\ \Delta_2(1) &= \frac{\frac{1}{6} + \frac{1}{6}p - \frac{1}{12}q}{\frac{1}{2} - \frac{1}{3}q + \frac{1}{6}pq} - \frac{\frac{1}{12}q}{\frac{1}{3}q - \frac{1}{6}pq} = \frac{1}{2} \frac{-2p + 2p^2 - 1}{(p-2)(3-2q+pq)} \end{aligned}$$

Further, we obtain:

$$\Pr(0,0) = \frac{1}{6}(p+1), \quad \Pr(1,0) = \frac{1}{6}(2-p),$$

and hence:

$$\eta_{12A.2} = \frac{1}{2} \frac{\left(\frac{1}{2} \frac{-2p+2p^2-1}{(p+1)(q+pq-3)}\right)^2 \frac{1}{6}(p+1) + \left(\frac{1}{2} \frac{-2p+2p^2-1}{(p-2)(3-2q+pq)}\right)^2 \frac{1}{6}(2-p)}{\left(\frac{1}{2} \frac{-2p+2p^2-1}{(p+1)(q+pq-3)}\right) \frac{1}{6}(p+1) + \left(\frac{1}{2} \frac{-2p+2p^2-1}{(p-2)(3-2q+pq)}\right) \frac{1}{6}(2-p)},$$

In terms of the consultation value, for each $(s_1, \hat{m}_2) \in \{0,1\}^2$, let

$$v(s_1, \hat{m}_2) = P_2(s_1, \hat{m}_2) - \frac{P_1(s_1, \hat{m}_2)^2}{P_0(s_1, \hat{m}_2)},$$

$$V_{12A} = - \sum_{(s_1, \hat{m}_2) \in \{0,1\}^2} v(s_1, \hat{m}_2).$$

This yields, using also

$$\begin{aligned} P_2(0,0) &= \frac{1}{20}q - \frac{1}{60}pq, \quad P_2(1,0) = \frac{1}{30}q + \frac{1}{60}pq, \quad P_2(0,1) = -\frac{1}{6}p - \frac{1}{20}q + \frac{1}{60}pq + \frac{1}{4}, \\ P_2(1,1) &= \frac{1}{12} + \frac{1}{6}p - \frac{1}{30}q - \frac{1}{60}pq, \end{aligned}$$

the following expressions:

$$\begin{aligned} v(0,0) &= \frac{q}{20} - \frac{pq}{60} - \frac{\left(\frac{q}{12}\right)^2}{\frac{q}{6} + \frac{pq}{6}} = -\frac{q}{120} \frac{-4p+2p^2-1}{p+1}, \\ v(1,0) &= \frac{q}{30} + \frac{pq}{60} - \frac{\left(\frac{q}{12}\right)^2}{\frac{q}{3} - \frac{pq}{6}} = \frac{q}{120} \frac{2p^2-3}{p-2}, \\ v(0,1) &= -\frac{1}{6}p - \frac{1}{20}q + \frac{1}{60}pq + \frac{1}{4} - \frac{\left(-\frac{1}{6}p - \frac{1}{12}q + \frac{1}{3}\right)^2}{-\frac{1}{6}q - \frac{1}{6}pq + \frac{1}{2}} \\ &= \frac{1}{120} \frac{-20p+8q-4pq^2-20p^2q+2p^2q^2+24pq+20p^2-q^2-10}{q+pq-3} \\ v(1,1) &= \frac{1}{12} + \frac{1}{6}p - \frac{1}{30}q - \frac{1}{60}pq - \frac{\left(\frac{1}{6} + \frac{1}{6}p - \frac{1}{12}q\right)^2}{\frac{1}{2} - \frac{1}{3}q + \frac{1}{6}pq} \\ &= -\frac{1}{120} \frac{-20p+12q+16pq-20p^2q+2p^2q^2+20p^2-3q^2-10}{3-2q+pq} \end{aligned}$$

Taking the negative of the sum of the four terms $v(s_1, \hat{m}_2)$ gives the expression of $V_{12A}(p, q)$ reported in Appendix C.

Equilibrium $\mathcal{E}_{2A}$. Only advisor 2 of type b_1 is truthful. Simplifying expression (1), the equilibrium condition is:

$$|b_1| \leq \eta_{2A}(q) = \frac{1}{2}\Delta_{2A}, \quad \text{where } \Delta_{2A} = \frac{P_1(1)}{P_0(1)} - \frac{P_1(0)}{P_0(0)}.$$

Hence using $\mu_{b_1, s_2}(s_2) = 1$ and $\mu_{b_2, s_2}(1) = 1$ for $s_2 = 0, 1$, and equation (8), we obtain:

$$P_n(0) = q \frac{n!1!}{(2+n)!} \text{ and } P_n(1) = q \frac{(1+n)!0!}{(2+n)!} + (1-q) \frac{n!1!}{(n+2)!} + (1-q) \frac{(n+1)!0!}{(n+2)!}.$$

Plugging $P_n(\hat{m}_2)$ for $n = 0, 1$ and $\hat{m}_2 = 0, 1$ into Δ_{2A} , we obtain, after simplification, $\eta_{2A}(q) = 1/(6(2 - q))$. Likewise, simplifying expression (4), and plugging in $P_n(\hat{m}_2)$ for $n = 0, 1, 2$ and $\hat{m}_2 = 0, 1$, we obtain:

$$\begin{aligned} V_{2A}(q) &= - \sum_{\hat{m}_2=0,1} v(\hat{m}_2), \text{ where } v(\hat{m}_2) = P_2(\hat{m}_2) - \frac{P_1(\hat{m}_2)^2}{P_0(\hat{m}_2)} \text{ for } \hat{m}_2 = 0, 1, \\ &= -\frac{1}{18} \frac{3-2q}{2-q}. \end{aligned}$$

Threshold orderings. From Appendix B, $\eta_{12.1}(p) < \eta_{12.2}(p)$ because $2p - 1 < 1 + 2p - 2p^2$; $\eta_{12.2}(p) < \eta_2$ because $3(1 + 2p - 2p^2) < 2(p + 1)(2 - p)$; $\eta_{12.1}(p) < \eta_1(p)$ because $3 < 2(p + 1)(2 - p)$ for $p \in (1/2, 1)$; and $\eta_1(p) < \eta_2$. Further, $\eta_{2A}(q) = 1/(6(2 - q)) \leq \eta_2$ with equality only at $q = 1$.

By (11), $\eta_{12A.1}$ is $1/2$ of the average of $\Delta_1(0)$ and $\Delta_1(1)$ with the positive weights $\Delta_1(\hat{m}_2)P_0(0, \hat{m}_2)$, where $\Delta_1(0) = 2\eta_{12.1}(p)$, because $\hat{m}_2 = 0$ reveals $s_2 = 0$, and $\Delta_1(1) = (2p - 1)f(q)/2$ with $f(q) = (6 - 6q + q^2)/((3 - 2q + pq)(3 - q - pq))$. Since $2(3 - 2q + pq)(3 - q - pq) - 3(6 - 6q + q^2) = q^2(1 + 2p - 2p^2)$ and $(6 - 6q + q^2)(p + 1)(2 - p) - (3 - 2q + pq)(3 - q - pq) = 3(1 - q)(1 + 2p - 2p^2)$ are nonnegative, $2\eta_{12.1}(p) \leq \Delta_1(1) \leq 2\eta_1(p)$, strictly for $q \in (0, 1)$, and the weighted average lies strictly between $2\eta_{12.1}(p)$ and $2\eta_1(p)$: $\eta_{12.1}(p) < \eta_{12A.1}(p, q) < \eta_1(p)$ for $q \in (0, 1)$, with equalities at $q = 1$ and $q = 0$ respectively.

Consider $\eta_{12A.2}$. It equals $1/2$ of the average of $\Delta_2(0) = (1 + 2p - 2p^2)/(2(p + 1)(3 - q - pq))$ and $\Delta_2(1) = (1 + 2p - 2p^2)/(2(2 - p)(3 - 2q + pq))$ with the positive weights $\Delta_2(s_1)P_0(s_1, s_2 = 0)$. Both lie below $2\eta_{12.2}(p) = (1 + 2p - 2p^2)/(2(p + 1)(2 - p))$, because $3 - q - pq \geq 2 - p$ and $3 - 2q + pq \geq 1 + p$, strictly for $q < 1$; hence $\eta_{12A.2}(p, q) < \eta_{12.2}(p)$ for $q < 1$, with equality at $q = 1$. Finally, direct computation gives

$$\eta_{2A}(q) - \eta_{12A.2}(p, q) = \frac{(2p - 1)^2 [9 - 6q - q^2 + p(1 - p)q(6 - 5q)]}{12(2 - p)(2 - q)(p + 1)(3 - 2q + pq)(3 - q - pq)},$$

whose bracket is at least 2 on $q \in [0, 1]$, so $\eta_{12A.2}(p, q) < \eta_{2A}(q)$ for $p > 1/2$. In particular $\eta_{12A} < \{\eta_1, \eta_{2A}, \eta_{12.2}\} < \eta_2$ on $(1/2, 1) \times (0, 1)$, whereas η_{12A} may fall below $\eta_{12.1}$: $\eta_{12A.2}(p, 0) = \eta_{12.2}(p)/2 < \eta_{12.1}(p)$ whenever $3 - 2p - 2p^2 < 0$, that is, for $p > (\sqrt{7} - 1)/2 \approx 0.82$, and by continuity $\eta_{12A} < \eta_{12.1}$ for such p and small q .

Monotonicity of the thresholds. Section 6 states that $\eta_{12A.1}$ increases in p and decreases in q , that $\eta_{12A.2}$ decreases in p and increases in q , and that η_{2A} increases in q . The last is immediate from $\eta_{2A}(q) = 1/(6(2 - q))$. For the other four, each partial derivative of the closed forms of $\eta_{12A.1}(p, q)$ and $\eta_{12A.2}(p, q)$ is a ratio whose denominator is positive on $(1/2, 1) \times (0, 1)$, being a product of squares and of the factors $2 - p$ and $p + 1$, and whose numerator is a polynomial in (p, q) . The substitutions $p = (1 + t)/(2 + t)$ and $q = s/(1 + s)$ map $t, s > 0$ onto the domain, and after them, once the positive factors $(2 + t)^a(1 + s)^b$ that clear the denominators are multiplied in, each numerator becomes a polynomial in (t, s) all of whose coefficients have the same sign: positive for $\partial\eta_{12A.1}/\partial p$ and $\partial\eta_{12A.2}/\partial q$, negative for $\partial\eta_{12A.1}/\partial q$ and $\partial\eta_{12A.2}/\partial p$ (the four polynomials have 63, 35, 35 and 30 terms; the signs were verified by symbolic computation). The four monotonicity properties follow.

Proof of Proposition 9. For the first result, Proposition 8(i) gives $\mathcal{E}_{12A}$ initially and (iii) applies once $b_1 > \eta_{12A}(p, q)$: just above η_{12A} both $\mathcal{E}_1$ and $\mathcal{E}_{2A}$ exist, and $p > g_1(q)$ selects $\mathcal{E}_1$,

so advisor 2 is dismissed; if moreover $p < g_3(q)$, then $\eta_1(p) < \eta_{2A}(q)$ by Lemma 3, and on $\eta_1(p) < b_1 \leq \eta_{2A}(q)$ only $\mathcal{E}_{2A}$ exists, so advisor 2 is hired back and advisor 1 discharged. For the second result, Proposition 8(ii) gives $\mathcal{E}_2$ at the initial q , where $p < g_2(q)$, and $\mathcal{E}_{12A}$ at any higher q at which $p > g_2(q)$ and $\mathcal{E}_{12A}$ still exists. ■

Proof of Proposition 10. Under $p > g_2(q)$, $b_1 < \eta_{12A}(p, q)$ and $b_2 > \eta_{12.2}(p) > \eta_{12A.2}(p, q)$, the equilibria $\mathcal{E}_{12A}$ and $\mathcal{E}_1$ exist, $\mathcal{E}_2$ exists if and only if $b_2 \leq \eta_2$, $\mathcal{E}_{2A}$ exists whenever $b_2 \geq \eta_{2A}(q)$, in particular when $b_2 > \eta_2$, and $\mathcal{E}_{12}$ does not, because $b_2 > \eta_{12.2}(p)$. Lemma 2 identifies V_S as the best available one-advisor value and gives $V_{12A} > V_S > V_0$, hence $c_S, c_{0S} > 0$.

It remains to prove that consulting one advisor is optimal over a nonempty range of costs. If $V_S = V_2$, then $V_{12A} < V_{12} \leq -1/24$, while $V_2 = -1/18$ and $V_0 = -1/12$. Therefore

$$V_{12A} + V_0 < -\frac{1}{8} < -\frac{1}{9} = 2V_2,$$

which is equivalent to $c_S < c_{0S}$. If $b_2 > \eta_2$, then $2V_S \geq V_1 + V_{2A}$. Put $r = 2p - 1 \in (0, 1)$. Direct substitution of the closed forms gives

$$V_1 + V_{2A} - V_{12A} - V_0 = \frac{qr^2 S(r, q)}{576(2-p)(p+1)(2-q)(pq-2q+3)(3-q-pq)},$$

where

$$S(r, q) = 180 - 162q + 27q^2 + (-36 + 24q - 2q^2)r^2 + (2q - q^2)r^4.$$

The polynomial S decreases in $q \in [0, 1]$, and $S(r, 1) = 45 - 14r^2 + r^4 \geq 32$. Thus the displayed difference is positive, so $V_{12A} + V_0 < 2V_S$, again equivalent to $c_S < c_{0S}$.

The three relevant net payoffs are $V_{12A} - 2c$, $V_S - c$, and V_0 . Their slopes and the strict ordering $c_S < c_{0S}$ give exactly the two changes in the optimal consultation choice stated in the proposition, with indifference at the cutoffs. ■

Part E: Alignment vs. Diversity

We begin by expanding P_n for the model of Section 7. First suppose that $A = \{1, 2\}$, so that we can write:

$$P_n(s_0, \hat{m}_1, \hat{m}_2) = \int_0^1 \sum_{s_1 \in \{0,1\}} \sum_{(s'_1, s_2) \in \{0,1\}^2} x^n \Pr(\hat{m}_1, \hat{m}_2 | s_1, s_2) \Pr(s_1 | s_0, s'_1) \Pr(s_0, s'_1, s_2 | x) dx.$$

Here, $\Pr(s_1 | s_0, s'_1) = 1$ if $s_1 = s'_1 = s_0$, $\Pr(s_1 | s_0, s'_1) = 1 - \rho$ if $s_1 = s'_1 \neq s_0$ and $\Pr(s_1 | s_0, s'_1) = \rho$ if $s_1 = s_0 \neq s'_1$. Further, $\Pr(s_0, s'_1, s_2 | x) = x^{s_0+s'_1+s_2}(1-x)^{3-s_0-s'_1-s_2}$. Hence, we get:

$$\begin{aligned} P_n(s_0, \hat{m}_1, \hat{m}_2) &= \sum_{(s_1, s_2) \in \{0,1\}^2} \mu_{1,s_1}(\hat{m}_1) \mu_{2,s_2}(\hat{m}_2) \\ &\times \left[\rho \mathbf{1}\{s_1 = s_0\} \frac{(n+s_0+s_2)!(2-s_0-s_2)!}{(n+3)!} \right. \\ &\quad \left. + (1-\rho) \frac{(n+s_0+s_1+s_2)!(3-s_0-s_1-s_2)!}{(n+4)!} \right]. \end{aligned} \quad (24)$$

In the same fashion, when $A = \{1\}$ and $A = \{2\}$, we obtain:

$$P_n(s_0, \hat{m}_1) = \sum_{s_1 \in \{0,1\}} \mu_{1,s_1}(\hat{m}_1) \left[\rho \mathbf{1}_{\{s_1 = s_0\}} \frac{(n+s_0)!(1-s_0)!}{(n+2)!} + (1-\rho) \frac{(n+s_0+s_1)!(2-s_0-s_1)!}{(n+3)!} \right],$$

$$P_n(s_0, \hat{m}_2) = \sum_{s_2 \in \{0,1\}} \mu_{2,s_2}(\hat{m}_2) \frac{(n+s_0+s_2)!(2-s_0-s_2)!}{(n+3)!}.$$

The full information equilibrium with correlated signal s_1 In order to find $\eta_{12.1}(\rho)$, $\eta_{12.2}(\rho)$ and $V_{12}(\rho)$, we begin by calculating $P_n(s_0, s_1, s_2)$ for $n = 0, 1, 2$, $s_0 = 0, 1$, $s_1 = 0, 1$ and $s_2 = 0, 1$. Using (24), we obtain:

$$\begin{aligned} P_0(0, 0, 0) &= \rho \frac{0!2!}{3!} + (1-\rho) \frac{0!3!}{4!} = \frac{\rho+3}{12}, & P_1(0, 0, 0) &= \rho \frac{1!2!}{4!} + (1-\rho) \frac{1!3!}{5!} = \frac{2\rho+3}{60}, \\ P_2(0, 0, 0) &= \rho \frac{2!2!}{5!} + (1-\rho) \frac{2!3!}{6!} = \frac{\rho+1}{60}, & P_0(0, 0, 1) &= \rho \frac{1!1!}{3!} + (1-\rho) \frac{1!2!}{4!} = \frac{\rho+1}{12}, \\ P_1(0, 0, 1) &= \rho \frac{2!1!}{4!} + (1-\rho) \frac{2!2!}{5!} = \frac{3\rho+2}{60}, & P_2(0, 0, 1) &= \rho \frac{3!1!}{5!} + (1-\rho) \frac{3!2!}{6!} = \frac{2\rho+1}{60}, \\ P_0(0, 1, 0) &= (1-\rho) \frac{1!2!}{4!} = \frac{1-\rho}{12}, & P_1(0, 1, 0) &= (1-\rho) \frac{2!2!}{5!} = \frac{1-\rho}{30}, \\ P_2(0, 1, 0) &= (1-\rho) \frac{3!2!}{6!} = \frac{1-\rho}{60}, & P_0(0, 1, 1) &= (1-\rho) \frac{2!1!}{4!} = \frac{1-\rho}{12}, \\ P_1(0, 1, 1) &= (1-\rho) \frac{3!1!}{5!} = \frac{1-\rho}{20}, & P_2(0, 1, 1) &= (1-\rho) \frac{4!1!}{6!} = \frac{1-\rho}{30}. \end{aligned}$$

We also obtain

$$\begin{aligned} P_0(1, 1, 1) &= \rho \frac{2!0!}{3!} + (1-\rho) \frac{3!0!}{4!} = \frac{\rho+3}{12}, & P_1(1, 1, 1) &= \rho \frac{3!0!}{4!} + (1-\rho) \frac{4!0!}{5!} = \frac{\rho+4}{20}, \\ P_0(1, 1, 0) &= \rho \frac{1!1!}{3!} + (1-\rho) \frac{2!1!}{4!} = \frac{1+\rho}{12}, & P_1(1, 1, 0) &= \rho \frac{2!1!}{4!} + (1-\rho) \frac{3!1!}{5!} = \frac{3+2\rho}{60}, \\ P_0(1, 0, 1) &= (1-\rho) \frac{2!1!}{4!} = \frac{1-\rho}{12}, & P_1(1, 0, 1) &= (1-\rho) \frac{3!1!}{5!} = \frac{1-\rho}{20}, \\ P_0(1, 0, 0) &= (1-\rho) \frac{1!2!}{4!} = \frac{1-\rho}{12}, & P_1(1, 0, 0) &= (1-\rho) \frac{2!2!}{5!} = \frac{1-\rho}{30}. \end{aligned}$$

Using these expressions, we calculate $\Delta_1(s_0, s_2)$ for $s_0 = 0, 1$ and $s_2 = 0, 1$:

$$\begin{aligned} \Delta_1(0, 0) &= \frac{P_1(0, 1, 0)}{P_0(0, 1, 0)} - \frac{P_1(0, 0, 0)}{P_0(0, 0, 0)} = \frac{\frac{1-\rho}{30}}{\frac{1-\rho}{12}} - \frac{\frac{2\rho+3}{60}}{\frac{\rho+3}{12}} = \frac{3}{5(3+\rho)}, \\ \Delta_1(0, 1) &= \frac{P_1(0, 1, 1)}{P_0(0, 1, 1)} - \frac{P_1(0, 0, 1)}{P_0(0, 0, 1)} = \frac{\frac{1-\rho}{20}}{\frac{1-\rho}{12}} - \frac{\frac{3\rho+2}{60}}{\frac{\rho+1}{12}} = \frac{1}{5(\rho+1)}, \\ \Delta_1(1, 0) &= \frac{P_1(1, 1, 0)}{P_0(1, 1, 0)} - \frac{P_1(1, 0, 0)}{P_0(1, 0, 0)} = \frac{\frac{3+2\rho}{60}}{\frac{1+\rho}{12}} - \frac{\frac{1-\rho}{30}}{\frac{1-\rho}{12}} = \frac{1}{5(1+\rho)}, \\ \Delta_1(1, 1) &= \frac{P_1(1, 1, 1)}{P_0(1, 1, 1)} - \frac{P_1(1, 0, 1)}{P_0(1, 0, 1)} = \frac{\frac{\rho+4}{20}}{\frac{\rho+3}{12}} - \frac{\frac{1-\rho}{20}}{\frac{1-\rho}{12}} = \frac{3}{5(\rho+3)}. \end{aligned}$$

Using condition (3) in Theorem 1, together with $\hat{m}_i = s_i$ for $s_i = 0, 1$ and $i = 1, 2$, the equilibrium condition for advisor 1 can be written here as follows:

$$|b_1| \leq \eta_{12.1}(\rho) \equiv \frac{\sum_{(s_0, s_2) \in \{0,1\}^2} \Delta_1(s_0, s_2)^2 P_0(s_0, s_1 = 0, s_2)}{2 \sum_{(s_0, s_2) \in \{0,1\}^2} \Delta_1(s_0, s_2) P_0(s_0, s_1 = 0, s_2)}.$$

Plugging $\Delta_1(s_0, s_2)$ and $P_0(s_0, 0, s_2)$ for $s_0 = 0, 1$ and $s_2 = 0, 1$, into the above expression, we obtain:

$$\begin{aligned} \eta_{12.1}(\rho) &= \frac{1 \left(\frac{3}{5(3+\rho)}\right)^2 \frac{\rho+3}{12} + \left(\frac{1}{5(\rho+1)}\right)^2 \frac{\rho+1}{12} + \left(\frac{1}{5(1+\rho)}\right)^2 \frac{1-\rho}{12} + \left(\frac{3}{5(3+\rho)}\right)^2 \frac{1-\rho}{12}}{2 \left(\frac{3}{5(3+\rho)}\right) \frac{\rho+3}{12} + \left(\frac{1}{5(\rho+1)}\right) \frac{\rho+1}{12} + \left(\frac{1}{5(1+\rho)}\right) \frac{1-\rho}{12} + \left(\frac{3}{5(3+\rho)}\right) \frac{1-\rho}{12}} \\ &= \frac{\frac{1}{2} \frac{3}{5(\rho+3)^2} + \frac{1}{30(\rho+1)^2}}{\frac{1}{2} \frac{1}{(\rho+3)} + \frac{1}{6(\rho+1)}}. \end{aligned}$$

Turning to the equilibrium condition for advisor 2, we calculate $\Delta_2(s_0, s_1)$, for $s_0 = 0, 1$, $s_1 = 0, 1$ and obtain:

$$\begin{aligned} \Delta_2(0, 0) &= \frac{P_1(0, 0, 1)}{P_0(0, 0, 1)} - \frac{P_1(0, 0, 0)}{P_0(0, 0, 0)} = \frac{\frac{3\rho+2}{60}}{\frac{\rho+1}{12}} - \frac{\frac{2\rho+3}{60}}{\frac{\rho+3}{12}} = \frac{1}{5} \frac{6\rho + \rho^2 + 3}{(\rho+3)(\rho+1)}, \\ \Delta_2(0, 1) &= \frac{P_1(0, 1, 1)}{P_0(0, 1, 1)} - \frac{P_1(0, 1, 0)}{P_0(0, 1, 0)} = \frac{\frac{1-\rho}{20}}{\frac{1-\rho}{12}} - \frac{\frac{1-\rho}{30}}{\frac{1-\rho}{12}} = \frac{1}{5}, \\ \Delta_2(1, 0) &= \frac{P_1(1, 0, 1)}{P_0(1, 0, 1)} - \frac{P_1(1, 0, 0)}{P_0(1, 0, 0)} = \frac{\frac{1-\rho}{20}}{\frac{1-\rho}{12}} - \frac{\frac{1-\rho}{30}}{\frac{1-\rho}{12}} = \frac{1}{5}, \\ \Delta_2(1, 1) &= \frac{P_1(1, 1, 1)}{P_0(1, 1, 1)} - \frac{P_1(1, 1, 0)}{P_0(1, 1, 0)} = \frac{\frac{\rho+4}{20}}{\frac{\rho+3}{12}} - \frac{\frac{3+2\rho}{60}}{\frac{1+\rho}{12}} = \frac{1}{5} \frac{6\rho + \rho^2 + 3}{(\rho+3)(\rho+1)}. \end{aligned}$$

By plugging these expressions and $P_0(s_0, s_1, 0)$ for $s_0 = 0, 1$ and $s_1 = 0, 1$ in the formula:

$$|b_2| \leq \eta_{12.2}(\rho) \equiv \frac{\sum_{(s_0, s_1) \in \{0,1\}^2} \Delta_2(s_0, s_1)^2 P_0(s_0, s_1, s_2 = 0)}{2 \sum_{(s_0, s_1) \in \{0,1\}^2} \Delta_2(s_0, s_1) P_0(s_0, s_1, s_2 = 0)},$$

we obtain:

$$\begin{aligned} \eta_{12.2}(\rho) &= \frac{1 \left(\frac{1}{5} \frac{6\rho+\rho^2+3}{(\rho+3)(\rho+1)}\right)^2 \frac{\rho+3}{12} + \left(\frac{1}{5}\right)^2 \frac{1-\rho}{12} + \left(\frac{1}{5}\right)^2 \frac{1-\rho}{12} + \left(\frac{1}{5} \frac{6\rho+\rho^2+3}{(\rho+3)(\rho+1)}\right)^2 \frac{1+\rho}{12}}{2 \left(\frac{1}{5} \frac{6\rho+\rho^2+3}{(\rho+3)(\rho+1)}\right) \frac{\rho+3}{12} + \frac{1}{5} \frac{1-\rho}{12} + \frac{1}{5} \frac{1-\rho}{12} + \left(\frac{1}{5} \frac{6\rho+\rho^2+3}{(\rho+3)(\rho+1)}\right) \frac{1+\rho}{12}} \\ &= \frac{1}{10} \frac{(\rho+2) \frac{(6\rho+\rho^2+3)^2}{(\rho+3)^2(\rho+1)^2} + (1-\rho)}{(\rho+2) \frac{6\rho+\rho^2+3}{(\rho+3)(\rho+1)} + (1-\rho)}. \end{aligned}$$

Turning to calculating the value $V_{12}(\rho)$, formula (4) together with ex-ante signal symmetry gives us: For every $(s_1, s_2) \in \{0, 1\}^2$, let

$$\begin{aligned} v(0, s_1, s_2) &= P_2(0, s_1, s_2) - \frac{P_1(0, s_1, s_2)^2}{P_0(0, s_1, s_2)}, \\ V_{12}(\rho) &= -2 \sum_{(s_1, s_2) \in \{0,1\}^2} v(0, s_1, s_2). \end{aligned}$$

We thus calculate:

$$\begin{aligned}
v(0,0,0) &= \frac{\rho+1}{60} - \frac{\left(\frac{2\rho+3}{60}\right)^2}{\frac{\rho+3}{12}} = \frac{1}{300} \frac{8\rho+\rho^2+6}{\rho+3}, \quad v(0,1,0) = \frac{1-\rho}{60} - \frac{\left(\frac{1-\rho}{30}\right)^2}{\frac{1-\rho}{12}} = \frac{1-\rho}{300}, \\
v(0,0,1) &= \frac{2\rho+1}{60} - \frac{\left(\frac{3\rho+2}{60}\right)^2}{\frac{\rho+1}{12}} = \frac{1}{300} \frac{3\rho+\rho^2+1}{\rho+1}, \quad v(0,1,1) = \frac{1-\rho}{30} - \frac{\left(\frac{1-\rho}{20}\right)^2}{\frac{1-\rho}{12}} = \frac{1-\rho}{300}.
\end{aligned}$$

We then obtain:

$$\begin{aligned}
V_{12}(\rho) &= -2 \left(\frac{1-\rho}{300} + \frac{1-\rho}{300} + \frac{1}{300} \frac{3\rho+\rho^2+1}{\rho+1} + \frac{1}{300} \frac{8\rho+\rho^2+6}{\rho+3} \right) \\
&= -\frac{1}{150} \frac{26\rho+9\rho^2+15}{(\rho+3)(\rho+1)}.
\end{aligned}$$

The equilibrium with only advisor 1 truthful Let us turn to the equilibrium in which only advisor 1 is truthful. Proceeding as before, we use (24) to calculate $P_1(s_0, s_1)$ and $P_0(s_0, s_1)$ for $s_0 = 0, 1$, and $s_1 = 0, 1$,

$$\begin{aligned}
P_0(0,0) &= \rho \frac{0!1!}{2!} + (1-\rho) \frac{0!2!}{3!} = \frac{2+\rho}{6}, \quad P_1(0,0) = \rho \frac{1!1!}{3!} + (1-\rho) \frac{1!2!}{4!} = \frac{1+\rho}{12}, \\
P_2(0,0) &= \rho \frac{2!1!}{4!} + (1-\rho) \frac{2!2!}{5!} = \frac{2+3\rho}{60}, \quad P_0(0,1) = P_0(1,0) = (1-\rho) \frac{1!1!}{3!} = \frac{1-\rho}{6}, \\
P_1(0,1) &= P_1(1,0) = (1-\rho) \frac{2!1!}{4!} = \frac{1-\rho}{12}, \quad P_2(0,1) = (1-\rho) \frac{3!1!}{5!} = \frac{1-\rho}{20}, \\
P_0(1,1) &= \rho \frac{1!0!}{2!} + (1-\rho) \frac{2!0!}{3!} = \frac{2+\rho}{6}, \quad P_1(1,1) = \rho \frac{2!0!}{3!} + (1-\rho) \frac{3!0!}{4!} = \frac{3+\rho}{12}.
\end{aligned}$$

Plugging these expressions in the formulas for $\Delta_1(s_0)$, for $s_0 = 0, 1$, we obtain:

$$\begin{aligned}
\Delta_1(0) &= \frac{P_1(0,1)}{P_0(0,1)} - \frac{P_1(0,0)}{P_0(0,0)} = \frac{\frac{1-\rho}{12}}{\frac{1-\rho}{6}} - \frac{\frac{1+\rho}{12}}{\frac{2+\rho}{6}} = \frac{1}{2(2+\rho)}, \\
\Delta_1(1) &= \frac{P_1(1,1)}{P_0(1,1)} - \frac{P_1(1,0)}{P_0(1,0)} = \frac{\frac{3+\rho}{12}}{\frac{2+\rho}{6}} - \frac{\frac{1-\rho}{12}}{\frac{1-\rho}{6}} = \frac{1}{2(2+\rho)}.
\end{aligned}$$

Combining these expressions, we obtain:

$$\begin{aligned}
\eta_1(\rho) &= \frac{\frac{1}{2} \Delta_1(0)^2 P_0(0,0) + \Delta_1(1)^2 P_0(1,0)}{\frac{1}{2} \Delta_1(0) P_0(0,0) + \Delta_1(1) P_0(1,0)} = \frac{1}{2} \frac{\left(\frac{1}{2(2+\rho)}\right)^2 \frac{2+\rho}{6} + \left(\frac{1}{2(2+\rho)}\right)^2 \frac{1-\rho}{6}}{\frac{1}{2(2+\rho)} \frac{2+\rho}{6} + \frac{1}{2(2+\rho)} \frac{1-\rho}{6}} \\
&= \frac{1}{4(2+\rho)}.
\end{aligned}$$

Turning to calculating the value of advice, we use formula (4) and ex-ante symmetry again:

$$V_1(\rho) = -2 \sum_{s_1=0,1} v(0, s_1), \quad \text{where } v(0, s_1) = P_2(0, s_1) - \frac{(P_1(0, s_1))^2}{P_0(0, s_1)}.$$

After calculating

$$v(0,0) = \frac{2+3\rho}{60} - \frac{\left(\frac{1+\rho}{12}\right)^2}{\frac{2+\rho}{6}} = \frac{1}{120} \frac{3+6\rho+\rho^2}{2+\rho}, \quad v(0,1) = \frac{1-\rho}{20} - \frac{\left(\frac{1-\rho}{12}\right)^2}{\frac{1-\rho}{6}} = \frac{1-\rho}{120},$$

we obtain:

$$V_1(\rho) = -2 \left(\frac{1}{120} \frac{3+6\rho+\rho^2}{2+\rho} + \frac{1-\rho}{120} \right) = -\frac{1}{12} \frac{1+\rho}{2+\rho}.$$

Properties of the thresholds and values, and proof of Proposition 11 Simplifying the expressions above, $\eta_{12.1}(\rho) = (27 + 42\rho + 19\rho^2)/(10(1+\rho)(3+\rho)(9+7\rho))$ and $\eta_{12.2}(\rho) = (27 + 96\rho + 118\rho^2 + 52\rho^3 + 7\rho^4)/(10(1+\rho)(3+\rho)(9+16\rho+5\rho^2))$; at $\rho = 0$ both equal $1/10$, the threshold with three independent signals. Their difference is $\eta_{12.2}(\rho) - \eta_{12.1}(\rho) = \rho(81 + 225\rho + 185\rho^2 + 49\rho^3)/(10(1+\rho)(9+7\rho)(9+16\rho+5\rho^2))$, positive for $\rho > 0$. The derivatives are $\eta'_{12.1}(\rho) = -(405 + 972\rho + 1038\rho^2 + 588\rho^3 + 133\rho^4)/(10(1+\rho)^2(3+\rho)^2(9+7\rho)^2) < 0$ and $\eta'_{12.2}(\rho) = 2(81 + 405\rho + 690\rho^2 + 510\rho^3 + 163\rho^4 + 13\rho^5 - 2\rho^6)/(5(1+\rho)^2(3+\rho)^2(9+16\rho+5\rho^2)^2) > 0$ on $[0, 1]$, because $2\rho^6 < 81$. For the value, $V'_{12}(\rho) = -(9 + 12\rho + 5\rho^2)/(75(1+\rho)^2(3+\rho)^2) < 0$ and $V''_{12}(\rho) = 2(18 + 27\rho + 18\rho^2 + 5\rho^3)/(75(1+\rho)^3(3+\rho)^3) > 0$, so V_{12} decreases and is convex.

For advisor 1 alone, $\eta'_1(\rho) = -1/(4(2+\rho)^2) < 0$, $V'_1(\rho) = -1/(12(2+\rho)^2) < 0$ and $V''_1(\rho) = 1/(6(2+\rho)^3) > 0$. Further, $\eta_1(\rho) - \eta_{12.1}(\rho) = (27 + 63\rho + 25\rho^2 - 3\rho^3)/(20(1+\rho)(2+\rho)(3+\rho)(9+7\rho)) > 0$ on $[0, 1]$, while $\eta_1(\rho) - \eta_{12.2}(\rho) = (27 - 18\rho - 224\rho^2 - 264\rho^3 - 107\rho^4 - 14\rho^5)/(20(1+\rho)(2+\rho)(3+\rho)(9+16\rho+5\rho^2))$ has a numerator that is strictly decreasing in ρ , positive at $\rho = 0$ and negative at $\rho = 1$: η_1 crosses $\eta_{12.2}$ exactly once, from above, at $\rho \approx 0.270$. Finally, $V_{12}(\rho) - V_1(\rho) = (15 + 41\rho + 37\rho^2 + 7\rho^3)/(300(1+\rho)(2+\rho)(3+\rho)) > 0$, and the constants $V_2 = -1/24$ and $V_0 = -1/18$ satisfy $V_{12}(\rho) - V_2 = (1-\rho)(15+11\rho)/(600(1+\rho)(3+\rho)) > 0$, $V_2 - V_1(\rho) = \rho/(24(2+\rho)) > 0$ and $V_1(\rho) - V_0 = (1-\rho)/(36(2+\rho)) > 0$ for $\rho \in (0, 1)$, so that $V_0 < V_1(\rho) < V_2 < V_{12}(\rho)$. As for the threshold $\eta_2 = 1/8$ of advisor 2 alone, $\eta_2 - \eta_{12.2}(\rho) = (1-\rho)(27+63\rho+31\rho^2+3\rho^3)/(40(1+\rho)(3+\rho)(9+16\rho+5\rho^2)) > 0$ and $\eta_2 - \eta_1(\rho) = \rho/(8(2+\rho)) > 0$ for $\rho \in (0, 1)$.

The proof of Proposition 11 is that of Proposition 2 with these orderings. Fix ρ and the bias vector. Since the displayed value differences are strict whenever the relevant alternatives exist, the same ranking holds for every sufficiently small positive c ; this is the asserted $c \downarrow 0$ limit. If $|b_1| \leq \eta_{12.1}(\rho)$ and $|b_2| \leq \eta_{12.2}(\rho)$, then $\mathcal{E}_{12}$ exists and has the highest value. Otherwise $\mathcal{E}_{12}$ does not exist, and $\mathcal{E}_2$, which exists if and only if $|b_2| \leq \eta_2$, dominates $\mathcal{E}_1$ and $\mathcal{E}_0$; both the case $|b_1| > \eta_{12.1}(\rho)$ with $|b_2| \leq \eta_{12.2}(\rho) < \eta_2$ and the case $\eta_{12.2}(\rho) < |b_2| \leq \eta_2$ fall here. If $|b_2| > \eta_2$, neither $\mathcal{E}_{12}$ nor $\mathcal{E}_2$ exists, and $\mathcal{E}_1$, which exists if and only if $|b_1| \leq \eta_1(\rho)$, dominates $\mathcal{E}_0$; when also $|b_1| > \eta_1(\rho)$, only $\mathcal{E}_0$ remains.

Part F: Numerical Code

This part reproduces the scripts behind the figures and the numerical statements of Section 5 and Part C. They are plain Python and depend only on `numpy`, `scipy` and `matplotlib`. Long lines have been broken at bracket boundaries so that they fit the page; the code is otherwise exactly as run.

cs_core.py — closed forms for the single-advisor and two-advisor two-interval equilibria, the incentive check, and the solver at a given bias pair.

```

"""
Closed forms for the continuous-signal section: general two-interval equilibria,
one advisor and two advisors.

Model
-----
x ~ U[0,1]; s_i = x + eps_i, where eps_1 and eps_2 are mutually independent
and jointly independent of x, and eps_i ~ U[-delta_i, delta_i].
Leader payoff -(y-x)^2, advisor i payoff -(y-x-b_i)^2.

Everything here is "general": thresholds may lie anywhere in (-delta, 1+delta),
so a switching region may protrude outside the state space, and with two advisors
the two switching regions may be disjoint in either order or overlap.

Contents
-----
bgen3(delta)          general three-interval single-advisor bound
bbar_int(delta)       largest bias with an INTERIOR two-interval threshold
single_general(b, delta) the unique two-interval equilibrium at bias b
two_by_two(t1, t2, d1, d2) profile probabilities and leader actions
implied_biases(...)   the biases that make each threshold type indifferent
    global_ic_22(...) incentive compatibility at every signal (certified
                        over the polynomial pieces)
value_22(P, y)        the leader's value

The two-advisor objects were checked against run_NM_16.py on 400 random
configurations: max |dP| = 3e-15, max |d(implied bias)| = 7e-10, and the
incentive verdict agreed at all 397 admissible points (about 75x faster).
"""
import numpy as np
from scipy.optimize import brentq

TOL = 1e-13

# -----
# one advisor
# -----
def bgen3(d):
    """Largest bias at which one advisor alone sustains three intervals,
    with no restriction on the position of his thresholds.

    Closed form: for d <= (3-sqrt(3))/4 the binding configuration is on the
    middle branch, bgen3 = (3a^2 + d^2)/12a with a = (2-sqrt(1+d^2))/3;
    beyond that it is on the lower outer branch, bgen3 = mu/3 with mu the
    root in (1/4, d] of 8 mu^3 - 12 d mu + 3 d = 0.
    """
    dstar = (3 - np.sqrt(3)) / 4
    if d <= dstar:
        a = (2 - np.sqrt(1 + d * d)) / 3
        return (3 * a * a + d * d) / (12 * a)
    mu = brentq(lambda m: 8 * m ** 3 - 12 * d * m + 3 * d, 0.25 + 1e-13, d, xtol=1e-15)
    return mu / 3

def bbar_int(d):
    """Largest bias with an interior two-interval threshold: 1/4 - 5d/12 - d^2/12(1-d)."""
    return 0.25 - 5 * d / 12 - d * d / (12 * (1 - d))

def alpha_Delta(mu, d):
    """Low-message probability alpha and action spread Delta as functions of the
    posterior mean mu at the threshold signal (the three branches)."""
    nu = 1 - mu

```

```

    if mu <= d:
        return mu * mu / d, d * (3 - 4 * mu) / (6 * (d - mu * mu))
    if mu >= 1 - d:
        return 1 - nu * nu / d, d * (3 - 4 * nu) / (6 * (d - nu * nu))
    return mu, (3 * mu * (1 - mu) - d * d) / (6 * mu * (1 - mu))

def Psi(mu, d):
    """Implied-bias map: mu supports a two-interval equilibrium at bias b iff b = Psi(mu).
    Strictly decreasing, onto (-1/4, 1/4)."""
    a, D = alpha_Delta(mu, d)
    return 0.5 - mu - 0.5 * (1 - 2 * a) * D

def single_general(b, d):
    """The unique two-interval equilibrium of a single advisor at bias b.

    Returns a dict with the posterior mean mu at the threshold, the low-message
    probability alpha, the spread Delta, the two actions, the leader's value
    V = alpha(1-alpha)Delta^2 - 1/12, the threshold signal t, and which branch
    it is on ('interior' iff |b| < bbar_int(d)). None if |b| >= 1/4.
    """
    if abs(b) >= 0.25:
        return None
    mu = brentq(lambda m: Psi(m, d) - b, 1e-14, 1 - 1e-14, xtol=1e-15)
    a, D = alpha_Delta(mu, d)
    yL, yH = 0.5 - (1 - a) * D, 0.5 + a * D
    V = a * (1 - a) * D * D - 1 / 12
    if mu <= d:
        t, branch = 2 * mu - d, 'low'
    elif mu >= 1 - d:
        t, branch = 2 * mu - 1 + d, 'high'
    else:
        t, branch = mu, 'interior'
    return dict(mu=mu, alpha=a, Delta=D, yL=yL, yH=yH, V=V, t=t, branch=branch)

def V_single(bs, d):
    """Vectorised leader value of consulting one advisor with noise d at biases bs."""
    bs = np.atleast_1d(np.asarray(bs, dtype=float))
    out = np.full(bs.shape, np.nan)
    flat = out.ravel()
    for k, b in enumerate(bs.ravel()):
        if 0 < b < 0.25:
            flat[k] = single_general(b, d)['V']
    return out

# -----
# two advisors: exact integrals of the clipped-linear message likelihoods
# -----
def _int_G(lo, hi, t, d):
    """int_lo^hi G dx and int_lo^hi x G dx, with G(x) = clip((t+d-x)/2d, 0, 1).
    Vectorised in lo, hi."""
    lo = np.asarray(lo, dtype=float)
    hi = np.asarray(hi, dtype=float)
    hi1 = np.minimum(hi, t - d) # flat part where G = 1
    e1 = np.clip(hi1 - lo, 0.0, None)
    m1 = np.where(e1 > 0, (hi1 ** 2 - lo ** 2) / 2, 0.0)
    A = np.maximum(lo, t - d) # linear part
    B = np.minimum(hi, t + d)
    ok = B > A
    Ac, Bc = np.where(ok, A, 0.0), np.where(ok, B, 0.0)
    p2 = np.where(ok, ((t + d) * (Bc - Ac) - (Bc ** 2 - Ac ** 2) / 2) / (2 * d), 0.0)
    m2 = np.where(ok, ((t + d) * (Bc ** 2 - Ac ** 2) / 2 - (Bc ** 3 - Ac ** 3) / 3) / (2 * d), 0.0)
    return e1 + p2, m1 + m2

```

```

def _int_GG(t1, d1, t2, d2):
    """The joint terms  $C = \int_0^1 G_1 G_2 dx$  and  $D = \int_0^1 x G_1 G_2 dx$ .
    Their integrands have degree  $\leq 2$  and  $\leq 3$  in  $x$ , respectively, so
    3-point Gauss on each piece is exact."""
    pts = np.array(sorted({0.0, 1.0} | {p for p in (t1 - d1, t1 + d1, t2 - d2, t2 + d2) if 0 < p < 1}))
    a, b = pts[-1], pts[1:]
    xs = np.array([-np.sqrt(3 / 5), 0.0, np.sqrt(3 / 5)])
    ws = np.array([5 / 9, 8 / 9, 5 / 9])
    mid, half = (a + b) / 2, (b - a) / 2
    X = mid[:, None] + half[:, None] * xs[None, :]
    W = half[:, None] * ws[None, :]
    G1 = np.clip((t1 + d1 - X) / (2 * d1), 0, 1)
    G2 = np.clip((t2 + d2 - X) / (2 * d2), 0, 1)
    return float(np.sum(W * G1 * G2)), float(np.sum(W * X * G1 * G2))

def two_by_two(t1, t2, d1, d2):
    """Profile probabilities  $P$  and leader actions  $y$  for thresholds  $(t1, t2)$ .

    Index  $[j, k]$ :  $j = 0/1$  is advisor 1's low/high message,  $k$  likewise for advisor 2.
     $y$  is nan at profiles of zero probability (off path). The switching regions
    have disjoint interiors iff
     $P[0,0] = \min(P[0,0] + P[0,1], P[0,0] + P[1,0])$ .
    """
    A1, B1 = _int_G(0.0, 1.0, t1, d1)
    A2, B2 = _int_G(0.0, 1.0, t2, d2)
    C, D = _int_GG(t1, d1, t2, d2)
    P = np.clip(np.array([[C, A1 - C], [A2 - C, 1 - A1 - A2 + C]]), 0.0, None)
    M = np.array([[D, B1 - D], [B2 - D, 0.5 - B1 - B2 + D]])
    y = np.where(P > TOL, M / np.maximum(P, 1e-300), np.nan)
    return P, y

def _cond_other(s, d_own, t_o, d_o):
    """Given own signal  $s$  (vector): the probability of the other advisor's low
    message conditional on  $s$ , the corresponding partial mean of  $x$ , and the same
    two objects for his high message."""
    lo = np.maximum(0.0, s - d_own)
    hi = np.minimum(1.0, s + d_own)
    L = hi - lo
    pL, mL = _int_G(lo, hi, t_o, d_o)
    pL, mL = pL / L, mL / L
    return pL, mL, 1 - pL, (lo + hi) / 2 - mL

def implied_biases(t1, t2, d1, d2, P=None, y=None):
    """The biases  $(b1, b2)$  at which each threshold type is exactly indifferent.

    This inverts the marginal conditions: rather than solving for thresholds at
    a given bias, it reads the bias off a given threshold pair, which turns the
    equilibrium set into a graph over the threshold plane.
    """
    if P is None:
        P, y = two_by_two(t1, t2, d1, d2)
    out = []
    for adv in (1, 2):
        if adv == 1:
            tau, d_own, t_o, d_o = t1, d1, t2, d2
            yA, yB, PA, PB = y[0, :], y[1, :], P[0, :], P[1, :]
        else:
            tau, d_own, t_o, d_o = t2, d2, t1, d1
            yA, yB, PA, PB = y[:, 0], y[:, 1], P[:, 0], P[:, 1]
        pL, mL, pH, mH = _cond_other(np.array([tau]), d_own, t_o, d_o)
        pk, mk = np.array([pL[0], pH[0]]), np.array([mL[0], mH[0]])
        num = den = 0.0
        for k in range(2):
            if pk[k] < 1e-12 or PA[k] < TOL or PB[k] < TOL:
                continue

```

```

        m = mk[k] / pk[k]
        Dk = yB[k] - yA[k]
        num += pk[k] * Dk * (0.5 * (yA[k] + yB[k]) - m)
        den += pk[k] * Dk
        out.append(num / den if abs(den) > 1e-14 else np.nan)
    return out[0], out[1]

def _signal_knots(t_own, d_own, t_other, d_other):
    """Breakpoints of the signal-wise deviation gain.

    Between consecutive points the gain, multiplied by the positive length of
    the own-signal conditional state support, is a polynomial of degree at most
    three. The points include changes in the prescribed message, clipping of
    that support at 0 or 1, and entry or exit of the other advisor's switching
    region.
    """
    lo, hi = -d_own, 1 + d_own
    candidates = [lo, hi, d_own, 1 - d_own, t_own]
    for x_break in (t_other - d_other, t_other + d_other):
        candidates.extend((x_break - d_own, x_break + d_own))
    return np.array(sorted({float(z) for z in candidates if lo <= z <= hi}))

def _cubic_maximum(fun, a, b):
    """Maximum of a function known to be cubic on [a,b].

    Four evaluations recover the cubic on a normalized interval. The returned
    value is evaluated from 'fun' itself at the endpoints and stationary
    points, rather than from the interpolant.
    """
    if b - a <= 1e-14:
        return -np.inf
    mid, half = 0.5 * (a + b), 0.5 * (b - a)
    z_fit = np.array([-0.8, -0.25, 0.3, 0.85])
    vals = np.array([fun(mid + half * z) for z in z_fit])
    poly = np.polynomial.Polynomial.fit(z_fit, vals, 3).convert()
    candidates = [-1.0, 1.0]
    candidates.extend(r.real for r in poly.deriv().roots()
                      if abs(r.imag) < 1e-8 and -1 < r.real < 1)
    return max(float(fun(mid + half * z)) for z in candidates)

def _report_gain_numerator(s, d_own, t_other, d_other, actions, bias,
                          prescribed, deviation):
    """Unnormalized payoff gain from 'deviation' at own signal 's'."""
    lo, hi = max(0.0, s - d_own), min(1.0, s + d_own)
    if hi <= lo:
        return 0.0
    p_low, m_low = _int_G(lo, hi, t_other, d_other)
    length = hi - lo
    p = np.array([float(p_low), length - float(p_low)])
    m = np.array([float(m_low), 0.5 * (hi * hi - lo * lo) - float(m_low)])

    def relevant_utility(report):
        action = actions[report, :]
        return float(np.sum(2 * action * (m + bias * p) - action * action * p))

    return relevant_utility(deviation) - relevant_utility(prescribed)

def _advisor_ic_22(t_own, d_own, t_other, d_other, actions, bias, tol):
    """Certified signal-wise check for one advisor in a two-by-two profile."""
    knots = _signal_knots(t_own, d_own, t_other, d_other)
    for a, b in zip(knots[:-1], knots[1:]):
        if b - a <= 1e-14:
            continue
        prescribed = int(0.5 * (a + b) >= t_own)

```

```

    deviation = 1 - prescribed

    def excess(s):
        length = min(1.0, s + d_down) - max(0.0, s - d_down)
        gain = _report_gain_numerator(s, d_down, t_other, d_other,
                                      actions, bias, prescribed, deviation)
        return gain - tol * max(length, 0.0)

    # The small slack absorbs roundoff in recovering a cubic from four
    # floating-point evaluations; it is far below the economic tolerance.
    if _cubic_maximum(excess, a, b) > 1e-13:
        return False
    return True

def global_ic_22(t1, t2, d1, d2, b1, b2, P, y, nS=None, nfill=11, tol=1e-11):
    """Check incentive compatibility at every signal, for both advisors.

    The check evaluates the endpoints and stationary points of every polynomial
    piece of every message-deviation gain. It therefore does not rely on a
    uniform signal grid. 'nS' is retained only for compatibility with older
    calls and is ignored. Off-path profiles (probability at most 'TOL') get a
    common fill action searched over 'nfill' values in [0,1].
    """
    miss = P <= TOL
    fills = np.array([0.5]) if not miss.any() else np.linspace(0.0, 1.0, nfill)
    for fv in fills:
        yy = np.where(miss, fv, y)
        if not _advisor_ic_22(t1, d1, t2, d2, yy, b1, tol):
            continue
        if _advisor_ic_22(t2, d2, t1, d1, yy.T, b2, tol):
            return True, float(fv)
    return False, None

def value_22(P, y):
    """Leader's gross value  $-1/3 + \sum P y^2$  over on-path profiles."""
    return -1 / 3 + float(np.nansum(P * np.where(np.isnan(y), 0.0, y) ** 2))

def in_region(b1, b2, d1, d2):
    """Restricted region:  $b1 < b2$  and  $b_i > \text{bgen3}(\text{delta}_i)$ ,  $i=1,2$ ."""
    return (b1 > bgen3(d1)) & (b2 > bgen3(d2)) & (b2 > b1)

def solve_22(d1, d2, b1, b2, nstart=6, tol=1e-11, starts=None, pmin=1e-9):
    """Threshold pairs supporting a two-advisor two-interval equilibrium at (b1, b2).

    Solves implied_biases(t1, t2) = (b1, b2) from several starts and returns the
    incentive-compatible solution with the highest leader value, or None.
    Cheaper than sweeping the whole threshold plane when the biases are given.
    'starts' prepends threshold pairs to try first, which is how a neighbouring
    solution is used as a continuation seed; 'pmin' discards solutions in which
    some message is sent with probability below it.
    """
    from scipy.optimize import fsolve

    def F(z):
        t1, t2 = z
        P, y = two_by_two(t1, t2, d1, d2)
        if min(P.sum(axis=1).min(), P.sum(axis=0).min()) < 1e-12:
            return [1e3, 1e3]
        a, b = implied_biases(t1, t2, d1, d2, P, y)
        if not (np.isfinite(a) and np.isfinite(b)):
            return [1e3, 1e3]
        return [a - b1, b - b2]

    best = None

```

```

grid = np.linspace(0.12, 0.88, nstart)
starts = [s for s in (starts or []) if s is not None] + [(u, v) for u in grid for v in grid]
seen = []
for z0 in starts:
    try:
        z, info, ier, _ = fsolve(F, z0, full_output=True, xtol=1e-13)
    except Exception:
        continue
    if ier != 1 or max(abs(np.asarray(F(z)))) > tol:
        continue
    if any(abs(z[0] - s[0]) < 1e-7 and abs(z[1] - s[1]) < 1e-7 for s in seen):
        continue
    seen.append(tuple(z))
    t1, t2 = z
    P, y = two_by_two(t1, t2, d1, d2)
    pm = min(P.sum(axis=1).min(), P.sum(axis=0).min())
    if pm < pmin:
        continue
    ok, fill = global_ic_22(t1, t2, d1, d2, b1, b2, P, y)
    if not ok:
        continue
    V = value_22(P, y)
    if best is None or V > best['V']:
        best = dict(t1=t1, t2=t2, V=V, pmin=pm, P=P, y=y, fill=fill)
return best

```

`cs_manifold.py` — sweeps the two-advisor threshold plane and rasterizes the leader-best value onto the bias plane.

```

"""
The general two-advisor two-interval equilibrium manifold at fixed precisions, and its
image in the bias plane.

```

Why a sweep of thresholds rather than a solve at each bias: with two advisors and two intervals each there are two thresholds and no consistency equations, so the equilibrium set is a two-dimensional manifold and $(t1, t2) \rightarrow (b1, b2, V12)$ is a map, not a fixed point. Sweeping the thresholds and reading the biases off the marginal conditions therefore finds every equilibrium of this class, including the ones with a threshold outside $[\delta, 1-\delta]$ that a solver seeded near the interior would miss. Folds of the map (several equilibria at one bias pair) are resolved by keeping the leader-best.

```

sweep()      evaluates the map on a threshold grid and checks incentive
              compatibility inside a bias window
rasterise()  pushes the compatible nodes onto a regular grid of the bias plane,
              interpolating over the triangles of the threshold grid and keeping
              the largest value at each pixel

```

```

"""

```

```

import time
import numpy as np

```

```

from cs_core import two_by_two, implied_biases, global_ic_22, value_22, TOL

```

```

def threshold_grid(d, n_uniform=1200, layer=1e-3, n_layer=30, umin=1e-9):
    """Uniform over the signal support  $(-d, 1+d)$ , refined geometrically down to
    umin at both ends, where one message becomes rare and the equilibria that
    approach the existence frontier live."""
    lo = -d + np.geomspace(umin, layer, n_layer)
    hi = 1 + d - np.geomspace(umin, layer, n_layer)[::-1]
    mid = np.linspace(-d + layer, 1 + d - layer, n_uniform)
    return np.unique(np.concatenate([lo, mid, hi]))

```

```

def sweep(d1, d2, window, n1=1200, n2=1000, pmin=1e-9, nS=2001, nfill=11, verbose=True):
    """Evaluate the two-advisor map on the threshold grid.

```

```

window = (b1lo, b1hi, b2lo, b2hi): incentive compatibility is checked only at
nodes whose implied biases fall in (a slightly enlarged) window, since the
check is the expensive step.
pmin: nodes at which some message has probability below pmin are discarded,
the marginal conditions being numerically unstable there.
"""
T1, T2 = threshold_grid(d1, n1), threshold_grid(d2, n2)
m, n = len(T1), len(T2)
B1 = np.full((m, n), np.nan)
B2 = np.full((m, n), np.nan)
V = np.full((m, n), np.nan)
IC = np.zeros((m, n), bool)
PMIN = np.full((m, n), np.nan)
OVL = np.zeros((m, n), np.int8)      # 0 overlap, 1 disjoint 2-below-1, 2 reversed
b1lo, b1hi, b2lo, b2hi = window
t0 = time.time()
for i, t1 in enumerate(T1):
    for j, t2 in enumerate(T2):
        P, y = two_by_two(t1, t2, d1, d2)
        m1, m2 = P.sum(axis=1).min(), P.sum(axis=0).min()
        if min(m1, m2) < pmin:
            continue
        b1, b2 = implied_biases(t1, t2, d1, d2, P, y)
        if not (np.isfinite(b1) and np.isfinite(b2)):
            continue
        B1[i, j], B2[i, j], V[i, j], PMIN[i, j] = b1, b2, value_22(P, y), min(m1, m2)
        OVL[i, j] = 1 if (t2 + d2 <= t1 - d1) else (2 if (t1 + d1 <= t2 - d2) else 0)
        if (b1lo - 0.005 <= b1 <= b1hi + 0.005 and b2lo - 0.005 <= b2 <= b2hi + 0.005
            and b2 > b1 - 0.005):
            IC[i, j], _ = global_ic_22(t1, t2, d1, d2, b1, b2, P, y, nS=nS, nfill=nfill)
        if verbose and i % 100 == 0:
            print(' row %d/%d %.0fs' % (i, m, time.time() - t0), flush=True)
return dict(T1=T1, T2=T2, B1=B1, B2=B2, V=V, IC=IC, PMIN=PMIN, OVL=OVL,
            d1=d1, d2=d2, window=window)

def drop_jumps(sw, tol=0.004):
    """Mark nodes adjacent to a discontinuity of the map, where linear
    interpolation over a grid triangle would be meaningless."""
    B1, B2 = sw['B1'], sw['B2']
    jump = np.zeros_like(B1, bool)
    for ax in (0, 1):
        step = np.abs(np.diff(B1, axis=ax)) + np.abs(np.diff(B2, axis=ax))
        big = step > tol
        if ax == 0:
            jump[:-1, :] |= big
            jump[1:, :] |= big
        else:
            jump[:, :-1] |= big
            jump[:, 1:] |= big
    out = dict(sw)
    out['IC'] = sw['IC'] & ~jump
    return out

def rasterise(sw, npix=600, require_ic=True):
    """Max-accumulate the leader value over the triangles of the threshold grid
    onto an npix x npix grid of the bias window.

    Returns (Vmax, PMbest, b1grid, b2grid); Vmax is nan where no admissible
    two-advisor equilibrium maps into the pixel, and PMbest carries the smallest
    message probability of the equilibrium achieving Vmax, which is what tells a
    substantive equilibrium from a degenerate one.
    """
    b1lo, b1hi, b2lo, b2hi = sw['window']
    B1, B2, V, IC, PM = sw['B1'], sw['B2'], sw['V'], sw['IC'], sw['PMIN']
    valid = np.isfinite(B1) & np.isfinite(B2) & np.isfinite(V)

```

```

if require_ic:
    valid &= IC
Vmax = np.full((npix, npix), -np.inf)
PMbest = np.full((npix, npix), np.nan)
sx = (npix - 1) / (b1hi - b1lo)
sy = (npix - 1) / (b2hi - b2lo)
m, n = B1.shape

def splat(x, y, v, p):
    px, py = int(round(x)), int(round(y))
    if 0 <= px < npix and 0 <= py < npix and v > Vmax[py, px]:
        Vmax[py, px] = v
        PMbest[py, px] = p

for tri in range(2):
    if tri == 0:
        va = np.s_[0:m - 1, 0:n - 1]
        vb = np.s_[1:m, 0:n - 1]
        vc = np.s_[0:m - 1, 1:n]
    else:
        va = np.s_[1:m, 1:n]
        vb = np.s_[1:m, 0:n - 1]
        vc = np.s_[0:m - 1, 1:n]
    ok = valid[va] & valid[vb] & valid[vc]
    xa, ya, Va, Pa = (B1[va] - b1lo) * sx, (B2[va] - b2lo) * sy, V[va], PM[va]
    xb, yb, Vb, Pb = (B1[vb] - b1lo) * sx, (B2[vb] - b2lo) * sy, V[vb], PM[vb]
    xc, yc, Vc, Pc = (B1[vc] - b1lo) * sx, (B2[vc] - b2lo) * sy, V[vc], PM[vc]
    xmin = np.minimum(np.minimum(xa, xb), xc); xmax = np.maximum(np.maximum(xa, xb), xc)
    ymin = np.minimum(np.minimum(ya, yb), yc); ymax = np.maximum(np.maximum(ya, yb), yc)
    for (i, j) in np.argwhere(ok & (xmax >= 0) & (xmin <= npix - 1) & (ymax >= 0) & (ymin <= npix - 1)):
        verts = ((xa[i, j], ya[i, j], Va[i, j], Pa[i, j]),
                  (xb[i, j], yb[i, j], Vb[i, j], Pb[i, j]),
                  (xc[i, j], yc[i, j], Vc[i, j], Pc[i, j]))
        x0, x1 = int(np.floor(max(xmin[i, j], 0))), int(np.ceil(min(xmax[i, j], npix - 1)))
        y0, y1 = int(np.floor(max(ymin[i, j], 0))), int(np.ceil(min(ymax[i, j], npix - 1)))
        if x1 < x0 or y1 < y0:
            continue
        det = ((xb[i, j] - xa[i, j]) * (yc[i, j] - ya[i, j])
               - (xc[i, j] - xa[i, j]) * (yb[i, j] - ya[i, j]))
        if abs(det) < 1e-14:
            for v in verts:
                splat(*v)
            continue
        X, Y = np.meshgrid(np.arange(x0, x1 + 1), np.arange(y0, y1 + 1))
        l2 = ((X - xa[i, j]) * (yc[i, j] - ya[i, j]) - (xc[i, j] - xa[i, j]) * (Y - ya[i, j])) / det
        l3 = ((xb[i, j] - xa[i, j]) * (Y - ya[i, j]) - (X - xa[i, j]) * (yb[i, j] - ya[i, j])) / det
        l1 = 1 - l2 - l3
        inside = (l1 >= -1e-9) & (l2 >= -1e-9) & (l3 >= -1e-9)
        if not inside.any():
            # triangle smaller than a pixel
            for v in verts:
                splat(*v)
            continue
        Vint = l1 * Va[i, j] + l2 * Vb[i, j] + l3 * Vc[i, j]
        Pint = l1 * Pa[i, j] + l2 * Pb[i, j] + l3 * Pc[i, j]
        Ys, Xs = Y[inside], X[inside]
        better = Vint[inside] > Vmax[Ys, Xs]
        Vmax[Ys[better], Xs[better]] = Vint[inside][better]
        PMbest[Ys[better], Xs[better]] = Pint[inside][better]
Vmax[~np.isfinite(Vmax)] = np.nan
return Vmax, PMbest, np.linspace(b1lo, b1hi, npix), np.linspace(b2lo, b2hi, npix)

def report(sw):
    """What the incentive check does to the candidates in the window: how many
    pass it, how degenerate they are, and which advisor breaks the rest."""
    B1, B2, IC, PM = sw['B1'], sw['B2'], sw['IC'], sw['PMIN']
    b1lo, b1hi, b2lo, b2hi = sw['window']
    win = np.isfinite(B1) & (B1 > b1lo) & (B1 < b1hi) & (B2 > B1) & (B2 < b2hi)

```

```

print('nodes in the window: %d; incentive compatible: %d' % (win.sum(), (win & IC).sum()))
if (win & IC).any():
    print('    compatible: b1 from %.4f, smallest message probability at most %.4f'
          % (B1[win & IC].min(), PM[win & IC].max()))
print('    nodes with every message above one percent: %d, of which compatible: %d'
      % ((win & (PM >= 1e-2)).sum(), (win & IC & (PM >= 1e-2)).sum()))

def run(d1=0.25, d2=0.005, window=(0.09689, 0.25, 0.09689, 0.27), npix=600,
        n1=1200, n2=1000, pmin=1e-9, out='manifold_%s_%s.npz', save_sweep=False):
    """Sweep, filter, rasterise and save. About 8 minutes at the default sizes.

    save_sweep=True also stores the raw threshold-plane arrays (about 50 MB),
    which is what you want to look at individual equilibria rather than the map.
    """
    sw = sweep(d1, d2, window, n1=n1, n2=n2, pmin=pmin)
    print('nodes with finite biases: %d; incentive compatible in window: %d'
          % (np.isfinite(sw['B1']).sum(), sw['IC'].sum()), flush=True)
    report(sw)
    sw = drop_jumps(sw)
    Vmax, PMbest, b1g, b2g = rasterise(sw, npix=npix)
    path = out % (d1, d2)
    fields = dict(Vmax=Vmax, PMbest=PMbest, b1=b1g, b2=b2g)
    if save_sweep:
        fields.update(T1=sw['T1'], T2=sw['T2'], B1=sw['B1'], B2=sw['B2'],
                      V=sw['V'], IC=sw['IC'], PMIN=sw['PMIN'], OVL=sw['OVL'])
    np.savez_compressed(path, **fields)
    print('pixels with an admissible two-advisor equilibrium: %d / %d -> %s'
          % (np.isfinite(Vmax).sum(), npix * npix, path))
    return path

if __name__ == '__main__':
    run()

```

`cs.fixed_bias.py` — traces the two-advisor branch at a fixed bias of advisor 1.

```

"""Comparison at a fixed first-expert bias: fix (delta1, delta2, b1) and follow the
leader's three values as b2 rises.

```

```

V1 is a number. V2 comes from the closed form. V12 is the leader-best
incentive-compatible two-advisor equilibrium at each b2, obtained by tracing the
level set {(t1, t2) : b1(t1, t2) = b1} in the threshold plane: for each t2 the
implied b1 is a function of t1, its roots are found by bisection on the sign
changes of a fine t1 grid, and each root gives one (b2, V12) pair.
"""

```

```

import numpy as np
from scipy.optimize import brentq

```

```

from cs_core import (two_by_two, implied_biases, global_ic_22, value_22,
                     single_general, bgen3)

```

```

def _b1_of(t1, t2, d1, d2):
    P, y = two_by_two(t1, t2, d1, d2)
    if min(P.sum(axis=1).min(), P.sum(axis=0).min()) < 1e-10:
        return np.nan
    b1, _ = implied_biases(t1, t2, d1, d2, P, y)
    return b1

```

```

def values_at_b1(d1, d2, b1, n_t2=900, n_t1=900, pmin=1e-6, nS=2001, nfill=11, check_ic=True):
    """Returns an array with columns (t1, t2, b2, V12, min message probability, IC)."""
    T2 = np.linspace(-d2 + 1e-9, 1 + d2 - 1e-9, n_t2)
    T1 = np.linspace(-d1 + 1e-9, 1 + d1 - 1e-9, n_t1)
    rows = []

```

```

for t2 in T2:
    vals = np.array([_b1_of(t, t2, d1, d2) for t in T1]) - b1
    ok = np.isfinite(vals)
    for i in range(len(T1) - 1):
        if not (ok[i] and ok[i + 1]) or vals[i] * vals[i + 1] > 0:
            continue
        try:
            t1 = brentq(lambda t: _b1_of(t, t2, d1, d2) - b1, T1[i], T1[i + 1], xtol=1e-13)
        except (ValueError, RuntimeError):
            continue
        P, y = two_by_two(t1, t2, d1, d2)
        pm = min(P.sum(axis=1).min(), P.sum(axis=0).min())
        if pm < pmin:
            continue
        bb1, b2 = implied_biases(t1, t2, d1, d2, P, y)
        if not np.isfinite(b2) or b2 <= b1 or b2 >= 0.25:
            continue
        ic = True
        if check_ic:
            ic, _ = global_ic_22(t1, t2, d1, d2, bb1, b2, P, y, nS=nS, nfill=nfill)
        rows.append((t1, t2, b2, value_22(P, y), pm, float(ic)))
    return np.array(rows) if rows else np.zeros((0, 6))

def upper_envelope(rows, b2grid, gap=None, tgap=8.0):
    """Leader-best V12 at each b2 on the grid.

    The traced points are a scatter along one or more branches, so binning them
    into pixels would leave holes; instead the points are sorted in b2 and
    interpolated between consecutive points that belong to the same branch, so
    that a genuine gap in the branch stays a gap.

    Adjacency is judged in the threshold plane, where the trace is a curve: a run
    breaks where the step in (t1, t2) exceeds 'tgap' times its median. Judging it
    by the gap in b2 instead would tie the answer to the plotting grid, and a fine
    grid would then cut a continuous branch into pieces; 'gap', an absolute bound
    on the b2 step, is only a backstop.
    """
    out = np.full(len(b2grid), np.nan)
    pm = np.full(len(b2grid), np.nan)
    if not len(rows):
        return out, pm
    gap = 0.02 if gap is None else gap
    o = np.argsort(rows[:, 2])
    b, v, p = rows[o, 2], rows[o, 3], rows[o, 4]
    t1, t2 = rows[o, 0], rows[o, 1]
    # keep the leader-best value at (numerically) coincident b2
    keep = np.ones(len(b), bool)
    for i in range(1, len(b)):
        if b[i] - b[i - 1] < 1e-12:
            j = i - 1 if v[i - 1] >= v[i] else i
            keep[j if j == i else i - 1] = False
    b, v, p, t1, t2 = b[keep], v[keep], p[keep], t1[keep], t2[keep]
    d = np.hypot(np.diff(t1), np.diff(t2)) if len(b) > 1 else np.zeros(0)
    dmax = tgap * np.median(d) if len(d) and np.median(d) > 0 else np.inf
    start = 0
    for i in range(1, len(b) + 1):
        if i == len(b) or d[i - 1] > dmax or b[i] - b[i - 1] > gap:
            if i - start >= 2:
                m = (b2grid >= b[start]) & (b2grid <= b[i - 1])
                cand = np.interp(b2grid[m], b[start:i], v[start:i])
                pcand = np.interp(b2grid[m], b[start:i], p[start:i])
                cur = out[m]
                better = ~np.isfinite(cur) | (cand > cur)
                idx = np.where(m)[0][better]
                out[idx] = cand[better]
                pm[idx] = pcand[better]
            start = i

```

```

return out, pm

def pattern(d1, d2, b1, n_t2=600, n_t1=600, nb=400):
    """Classify the leader's choice as the second expert's bias varies and report the sequence."""
    rows = values_at_b1(d1, d2, b1, n_t2=n_t2, n_t1=n_t1)
    rows = rows[rows[:, 5] > 0]
    b2g = np.linspace(b1, 0.25 - 1e-6, nb)
    V12, PM = upper_envelope(rows, b2g)
    V1 = single_general(b1, d1)['V']
    V2 = np.array([single_general(b, d2)['V'] for b in b2g])
    best = np.where(np.nan_to_num(V12, nan=-np.inf) > np.maximum(V1, V2), 0,
        np.where(V2 >= V1, 1, 2)) # 0 = E12, 1 = E2, 2 = E1
    seq, prev = [], None
    for k, c in enumerate(best):
        if c != prev:
            seq.append((c, b2g[k]))
            prev = c
    return dict(rows=rows, b2=b2g, V12=V12, PM=PM, V1=V1, V2=V2, best=best, seq=seq)

if __name__ == '__main__':
    import sys
    d1, d2, b1 = float(sys.argv[1]), float(sys.argv[2]), float(sys.argv[3])
    r = pattern(d1, d2, b1)
    names = {0: 'E12', 1: 'E2', 2: 'E1'}
    print('(d1,d2,b1)=(%.3f,%.3f,%.4f) bgen3(d1)=%.4f' % (d1, d2, b1, bgen3(d1)))
    print('sequence: ' + ' -> '.join('%s from b2=%.4f' % (names[c], b) for c, b in r['seq']))
    m = np.isfinite(r['V12'])
    if m.any():
        print('E12 traced on b2 in [%.4f, %.4f]; smallest message probability there up to %.3f'
            % (r['b2'][m].min(), r['b2'][m].max(), np.nanmax(r['PM'])))

```

make_figs_R.py — draws Figures 6 and 7.

```

"""
Figures 6 and 7 of the paper, on the restricted bias region, with the general
two-interval equilibria: at most two intervals per advisor, thresholds
unrestricted.

R = { (b_1, b_2) : bgen3(delta_i) < b_i for i=1,2,
      and b_1 < b_2 < 1/4 }

where bgen3(delta) is the largest bias at which one advisor alone can sustain a
three-interval partition, thresholds unrestricted; above it no advisor can
sustain more than two intervals, so the equilibria compared here are the
finest retained class under the numerical message-probability cutoff.

The leader's three options, each at its best equilibrium:
V1*, V2* the general two-interval equilibrium of one advisor, whose threshold
may lie anywhere in (-delta_i, 1+delta_i), against babbling;
V12* the leader-best two-advisor equilibrium, switching regions disjoint
in either order or overlapping, incentive compatible at every signal
and with every message sent with probability at least PMIN.

Both printed figures are at (delta_1, delta_2) = (0.20, 0.15); Figure 6 fixes
b_1 = 0.10. The bias-plane sweep behind Figure 7 is manifoldZ_R_0.2_0.15.npz,
produced by cs_manifold.run(d1=0.20, d2=0.15, window=(bgen3(0.20), 0.15,
bgen3(0.20), 0.15), npix=600, n1=2600, n2=2600, pmin=0.01); the comparison behind
Figure 6 is computed by cs_fixed_bias.values_at_b1 on first use and cached.

python3 make_figs_R.py

```

```

Figure 6 fig_values the three values against b_2 on (b_1, 0.12]
Figure 7 fig_regimes the choice over the bias plane
"""

```

```

import numpy as np
import matplotlib
matplotlib.use('Agg')
import matplotlib.pyplot as plt
from matplotlib.colors import ListedColormap
from matplotlib.patches import Patch

from cs_core import V_single, bgen3, bbar_int

COL = dict(E12='#2b5f8e', E2='#9b1b30', E2x='#7b4ea3', E1='#e0a020', out='#e2e2e2')
PMIN = 0.01
D1, D2 = 0.20, 0.15 # precisions of Figures 6 and 7
B1_FIXED = 0.10 # advisor 1's bias in Figure 6

def Vbest_single(bs, d):
    return np.maximum(V_single(bs, d), -1 / 12)

# -----
def fig_regimes(npz='manifoldZ_R_0.2_0.15.npz', d1=D1, d2=D2, path='continuous_regimes.pdf',
               b1_mark=B1_FIXED):
    z = np.load(npz)
    b1g, b2g, V12, PM = z['b1'], z['b2'], z['Vmax'], z['PMbest']
    n = len(b1g)
    B1, B2 = np.meshgrid(b1g, b2g)
    V1 = Vbest_single(b1g, d1)[None, :] * np.ones((n, 1))
    V2 = Vbest_single(b2g, d2)[:, None] * np.ones((1, n))
    found = np.isfinite(V12) & (PM >= PMIN)
    v12 = np.where(found, V12, -np.inf)
    inR = ((B1 > bgen3(d1)) & (B2 > bgen3(d2))) & (B2 > B1)
    cls = np.zeros((n, n), int)
    cls[inR & (V2 >= V1)] = 1 # advisor 2 alone, no two-advisor equilibrium
    cls[inR & (V1 > V2)] = 2 # advisor 1 alone
    cls[inR & found & (V2 > v12) & (V2 >= V1)] = 4
    cls[inR & found & (v12 > np.maximum(V1, V2))] = 3
    cmap = ListedColormap(['white', COL['E2'], COL['E1'], COL['E2x']])
    fig, ax = plt.subplots(figsize=(6.2, 6.2))
    ax.imshow(cls, origin='lower', extent=[b1g[0], b1g[-1], b2g[0], b2g[-1]], cmap=cmap,
              vmin=-0.5, vmax=4.5, aspect='auto', interpolation='nearest')
    ax.contour(B1, B2, np.where(inR, V1 - V2, np.nan), levels=[0.0], colors='k', linewidths=0.8)
    ax.plot([b1g[0], b1g[-1]], [b1g[0], b1g[-1]], 'k--', lw=0.8)
    ax.axhline(bbar_int(d2), color='k', ls=':', lw=0.7)
    if b1_mark is not None:
        # the b_1 of the companion value comparison
        ax.axvline(b1_mark, color='k', ls='-', lw=0.6, alpha=0.5)
    ax.set_xlabel(r'$b_1$'); ax.set_ylabel(r'$b_2$')
    ax.set_xlim(b1g[0], b1g[-1]); ax.set_ylim(b2g[0], b2g[-1])
    ax.legend(handles=[Patch(color=COL['E12'], label='both advisors'),
                      Patch(color=COL['E2x'], label=r'advisor 2 alone; retained $E_{12}$ found'),
                      Patch(color=COL['E2'], label=r'advisor 2 alone; none found at cutoff'),
                      Patch(color=COL['E1'], label='advisor 1 alone')],
              loc='upper left', fontsize=8, frameon=True, facecolor='white',
              framealpha=0.92, edgecolor='none')
    plt.tight_layout(); plt.savefig(path); plt.close(fig)
    tot = inR.sum()
    print('Figure 7 at (delta_1, delta_2) = (%.3f, %.3f), region starts at b_1 = %.4f:'
          % (d1, d2, bgen3(d1)))
    print(' both advisors %.3f; advisor 2 alone %.3f, of which %.3f where both could be consulted; '
          'advisor 1 alone %.3f'
          % ((cls == 3).sum() / tot, ((cls == 1) | (cls == 4)).sum() / tot, (cls == 4).sum() / tot,
            (cls == 2).sum() / tot))
    print(' a retained informative two-advisor equilibrium was found on %.3f of the region'
          % ((inR & found).sum() / tot))
    if (cls == 3).any():
        print(' where both are consulted the smallest message probability reaches %.3f'
              % np.nanmax(PM[cls == 3]))
    return b1g, b2g, cls

```

```

# -----
def fig_values(d1=D1, d2=D2, b1=B1_FIXED, step=0.0001, path='continuous_values.pdf',
              cache='valuesR_%s_%s_%s.npz', b2max=0.12):
    from cs_fixed_bias import values_at_b1, upper_envelope
    key = cache % (d1, d2, b1)
    try:
        rows = np.load(key)['rows']
    except IOError:
        rows = values_at_b1(d1, d2, b1, n_t2=1300, n_t1=1300, pmin=PMIN)
        np.savez(key, rows=rows)
    rows = rows[(rows[:, 5] > 0) & (rows[:, 4] >= PMIN)]
    b2 = np.arange(b1 + step, b2max, step)
    V12, PM = upper_envelope(rows, b2)
    # The branch traced on the threshold grid can miss the two-advisor equilibrium
    # where b_2 is within a few 10^-4 of b_1; solve directly at those biases.
    from cs_core import solve_22
    for k in np.where(~np.isfinite(V12))[0]:
        r = solve_22(d1, d2, b1, float(b2[k]), nstart=8)
        if r is not None and r['pmin'] >= PMIN:
            V12[k], PM[k] = r['V'], r['pmin']
    V1 = float(Vbest_single(np.array([b1]), d1)[0])
    V2 = Vbest_single(b2, d2)
    v12 = np.nan_to_num(V12, nan=-np.inf)
    best = np.where(v12 > np.maximum(V1, V2), 0, np.where(V2 >= V1, 1, 2))
    edges = [0] + [k for k in range(1, len(b2)) if best[k] != best[k - 1]] + [len(b2)]
    names = {0: r'$E_{12}$', 1: r'$E_2$', 2: r'$E_1$'}
    keys = {0: 'E12', 1: 'E2', 2: 'E1'}
    fig, ax = plt.subplots(figsize=(7.4, 4.6))
    ax.set_xlim(b1, b2max)
    for k in range(len(edges) - 1):
        a, b = edges[k], edges[k + 1] - 1
        ax.axvspan(b2[a], b2[b], color=COL[keys[best[a]]], alpha=0.07, lw=0)
        if b2[b] - b2[a] > 0.04 * (b2max - b1): # label a band only if it is wide enough to hold one
            ax.text(0.5 * (b2[a] + b2[b]), 0.965, names[best[a]], transform=ax.get_xaxis_transform(),
                    ha='center', va='top', fontsize=9, color=COL[keys[best[a]]])
    ax.plot(b2, np.full_like(b2, V1), color=COL['E1'], lw=1.5, label=r'$V_1$')
    ax.plot(b2, V2, color=COL['E2'], lw=1.5, label=r'$V_2$')
    m = np.isfinite(V12)
    ax.plot(b2[m], V12[m], color=COL['E12'], lw=1.7, label=r'$V_{12}$')
    for k in range(1, len(edges) - 1):
        ax.axvline(b2[edges[k]], color='k', lw=0.5, ls='--')
    ax.set_xlabel(r'$b_2$')
    ax.legend(loc='lower left', frameon=False, fontsize=9)
    plt.tight_layout(); plt.savefig(path); plt.close(fig)
    seq = [(keys[best[edges[k]]], b2[edges[k]], b2[edges[k + 1] - 1]) for k in range(len(edges) - 1)]
    print('Figure 6 at (delta_1, delta_2, b_1) = (%.3f, %.3f, %.4f): ' % (d1, d2, b1)
          + ' | '.join('%s on [%.4f, %.4f]' % s for s in seq))
    if m.any():
        print(' a retained informative two-advisor equilibrium was found for b_2 in '
              ' [%.4f, %.4f]; smallest message probability up to %.3f'
              % (b2[m].min(), b2[m].max(), np.nanmax(PM)))
    return b2, V1, V2, V12, best

if __name__ == '__main__':
    fig_regimes()
    fig_values()

```

make_filmstrip3.py — draws Figure 8.

```

"""
A filmstrip of bias planes: one panel per precision pair, each panel the leader's
consultation choice over the restricted bias region itself,

```

```

    bgen3(delta_i) < b_i for i=1,2, and b_1 < b_2 < 1/4,

```

so the axes of a panel begin at its own `bgen3(delta_1)` and end where an advisor alone stops being able to sustain two intervals.

Columns are `delta_1` in (0.10, 0.20, 0.32, 0.45); rows are the precision ratio `delta_2 / delta_1` in (0.25, 0.50, 0.75), so every panel satisfies $0 < \delta_2 < \delta_1 < 1/2$.

Within a panel:

```

white      b_2 <= b_1, outside the maintained ordering
blue       both advisors
violet     advisor 2 alone; retained informative E_12 found
crimson    advisor 2 alone; none found at the numerical cutoff
gold       advisor 1 alone

python3 make_filmstrip3.py      sweep (cached) and draw
python3 make_filmstrip3.py draw draw from the cached sweeps
"""
import os
import sys
import time

import numpy as np

import cs_manifold as M
from cs_core import bgen3, single_general

BBABBLE = 0.25          # no two-interval equilibrium at or above this bias
BMAX = BBABBLE          # the upper end of every panel's axes
BABBLE = -1.0 / 12.0
PMIN = 0.01
D1S = (0.10, 0.20, 0.32, 0.45)
RATIOS = (0.25, 0.50, 0.75)
NPIX = 500
NGRID = 1200
COL = dict(E12='#2b5f8e', E2='#9b1b30', E2x='#7b4ea3', E1='#e0a020',
           out='#e2e2e2', none='#8c8c8c')

def panels():
    return [(d1, round(r * d1, 6)) for r in RATIOS for d1 in D1S]

def sweep_one(args):
    d1, d2 = args
    path = 'fs3_%s_%s.npz' % (d1, d2)
    if os.path.exists(path):
        return path
    t0 = time.time()
    lo = bgen3(d1)
    sw = M.sweep(d1, d2, (lo - 0.002, BMAX + 0.002, lo - 0.002, BMAX + 0.002),
                  n1=NGRID, n2=NGRID, pmin=PMIN, verbose=False)
    sw = M.drop_jumps(sw)
    sw['window'] = (lo, BMAX, lo, BMAX)
    Vmax, PMbest, b1g, b2g = M.rasterise(sw, npix=NPIX)
    np.savez_compressed(path, Vmax=Vmax, PMbest=PMbest, b1=b1g, b2=b2g)
    print(' (%.3f, %.3f): %d of %d pixels carry an equilibrium  [%.1f min]'
          % (d1, d2, np.isfinite(Vmax).sum(), NPIX * NPIX, (time.time() - t0) / 60), flush=True)
    return path

def build(workers=2):
    import multiprocessing as mp
    todo = [p for p in panels() if not os.path.exists('fs3_%s_%s.npz' % p)]
    print('panels to sweep: %d of %d' % (len(todo), len(panels()))), flush=True
    if todo:
        with mp.Pool(workers) as pool:
            pool.map(sweep_one, todo, chunksize=1)

```

```

def classify(d1, d2):
    z = np.load('fs3_%s_%s.npz' % (d1, d2))
    b1g, b2g, V12, PM = z['b1'], z['b2'], z['Vmax'], z['PMbest']
    n = len(b1g)
    B1, B2 = np.meshgrid(b1g, b2g)
    V1 = np.array([max(single_general(b, d1)['V'], BABBLE) if b < BBABBLE and single_general(b, d1)
                    else BABBLE for b in b1g])[None, :] * np.ones((n, 1))
    V2 = np.array([max(single_general(b, d2)['V'], BABBLE) if b < BBABBLE and single_general(b, d2)
                    else BABBLE for b in b2g])[None, :] * np.ones((1, n))
    exists = np.isfinite(V12) & (PM >= PMIN)
    v12 = np.where(exists, V12, -np.inf)
    inR = ((B1 > bgen3(d1)) & (B2 > bgen3(d2)) & (B2 > B1))
    cls = np.zeros((n, n), int)
    cls[inR & (V2 >= V1)] = 4
    cls[inR & (V1 > V2)] = 5
    cls[inR & exists & (V2 > v12) & (V2 >= V1)] = 3
    cls[inR & exists & (v12 > np.maximum(V1, V2))] = 2
    return b1g, b2g, cls

def draw(path='continuous_filmstrip.pdf'):
    import matplotlib
    matplotlib.use('Agg')
    import matplotlib.pyplot as plt
    from matplotlib.colors import ListedColormap
    from matplotlib.patches import Patch

    cmap = ListedColormap(['white', COL['out'], COL['E12'], COL['E2x'], COL['E2'],
                           COL['E1'], COL['none']])
    nrow, ncol = len(RATIOS), len(D1S)
    fig, axes = plt.subplots(nrow, ncol, figsize=(3.0 * ncol, 3.1 * nrow))
    for k, (d1, d2) in enumerate(panels()):
        ax = axes.ravel()[k]
        b1g, b2g, cls = classify(d1, d2)
        lo = b1g[0]
        ax.imshow(cls, origin='lower', extent=[lo, BMAX, lo, BMAX], cmap=cmap,
                  vmin=-0.5, vmax=6.5, aspect='auto', interpolation='nearest')
        ax.plot([lo, BMAX], [lo, BMAX], 'k--', lw=0.6)
        if bgen3(d2) > lo:
            ax.axhline(bgen3(d2), color='k', ls=':', lw=0.5)
        tot = max((cls > 1).sum(), 1)
        sh = {c: (cls == c).sum() / tot for c in (2, 3, 4, 5)}
        ax.set_title(r'$\delta_1$=%.2f, \delta_2=%.3f'
                    % (d1, d2, lo, bgen3(d2)), fontsize=8.5)
        for i, (c, lab, col) in enumerate([(2, r'$E_{12}$', 'E12'),
                                           (3, r'$E_{2x}$; retained $E_{12}$ found', 'E2x'),
                                           (4, r'$E_{2x}$; none found at cutoff', 'E2'),
                                           (5, r'$E_{1x}$', 'E1')]):
            ax.text(0.97, 0.36 - 0.075 * i, '%s %.1f%' % (lab, 100 * sh[c]),
                    transform=ax.transAxes, ha='right', va='top', fontsize=7, color=COL[col])
        ax.set_xlim(lo, BMAX); ax.set_ylim(lo, BMAX)
        ax.set_xticks([0.10, 0.15, 0.20, 0.25]); ax.set_yticks([0.10, 0.15, 0.20, 0.25])
        ax.tick_params(labelsize=7)
        if k % ncol == 0:
            ax.set_ylabel(r'$b_2$', fontsize=8)
        if k >= len(panels()) - ncol:
            ax.set_xlabel(r'$b_1$', fontsize=8)
        print('%.2f, %.3f): both %.3f, advisor 2 %.3f (of which %.3f with a retained '
              'informative E12 found), advisor 1 %.3f'
              % (d1, d2, sh[2], sh[3] + sh[4], sh[3], sh[5]))
    fig.legend(handles=[Patch(color=COL['E12'], label='both advisors'),
                       Patch(color=COL['E2x'], label=r'advisor 2 alone; retained $E_{12}$ found'),
                       Patch(color=COL['E2'], label=r'advisor 2 alone; none found at cutoff'),
                       Patch(color=COL['E1'], label='advisor 1 alone')],
              loc='lower center', ncol=4, fontsize=8, frameon=False, bbox_to_anchor=(0.5, 0.0))

```

```

plt.tight_layout(rect=[0, 0.055, 1, 1.0])
plt.savefig(path); plt.close(fig)
print('wrote', path)

if __name__ == '__main__':
    if any(arg in ('-h', '--help') for arg in sys.argv[1:]):
        print('usage: python3 make_filmstrip3.py [draw]')
    elif len(sys.argv) > 1 and sys.argv[1] == 'draw':
        draw()
    elif len(sys.argv) == 1:
        build()
        draw()
    else:
        raise SystemExit('unknown argument; use --help for usage')

```

nm.py — general partitions: solves an (N, M) configuration, computes $\bar{b}_{\text{gen}}^{(N)}$, and verifies global incentive compatibility.

```

#!/usr/bin/env python3
"""
Statement (1) probe for general two-advisor equilibria, N1, N2 in {3, 4}.

Tests:  does any (N1, N2) equilibrium have  $b_1 > \bar{b}^{\text{gen}}(N1, \text{delta1})$ ?
        -- unrestricted, and restricted to the ordered region  $b_2 \geq b_1$ .

Design notes
-----
* b1 and b2 are NOT grid coordinates. Fixing (delta1, delta2), an (N1, N2)
  configuration has (N1-1)+(N2-1) thresholds and (N1-2)+(N2-2) consistency
  equations, so the equilibrium set is a 2-dimensional manifold for EVERY
  (N1, N2). The grid is over (delta1, delta2) and the two lowest thresholds;
  the remaining thresholds are solved to 1e-12 and the biases come out at
  solver tolerance.
* Thresholds are gridded logarithmically in the boundary-layer coordinate
   $u_i = a_i + \text{delta}_i$ . Known violations of the analogous (3,2) statement sit
  at  $u \sim 5e-6$  with lowest-interval probability  $\sim 1e-11$ ; a uniform grid misses them.
* Three stages: (A) coarse sweep, (B) adaptive refinement of the maximum for
  each parameter configuration, (C) global-IC verification of near-violations.
  Stage C is the item omitted from the earlier five-minute run.

Usage
-----
python3 nm.py --hours 24 --workers 8
python3 nm.py --hours 24 --resume          # continue a run
python3 nm.py --smoke --out NM_smoke      # short self-test
"""

import argparse
import os
import pickle
import time

import numpy as np
from scipy.optimize import brentq, fsolve

# -----
# quadrature
# -----
GLx, GLw = np.polynomial.legendre.leggauss(10)
TOL = 1e-13
STAGE_A_SHARE = 0.35 # each configuration's remaining budget goes to adaptive zoom
NRANGE = (2, 3, 4, 5) # N_i values swept

def G(x, t, d):

```

```

    """P(signal < t | state x) for signal = x + U[-d, d]."""
    return np.clip((t + d - x) / (2 * d), 0.0, 1.0)

def nodes(lo, hi, brk):
    """Gauss-Legendre nodes on [lo, hi], subdivided at the kink points brk."""
    pts = sorted({lo, hi} | {p for p in brk if lo < p < hi})
    a = np.array(pts[:-1])
    b = np.array(pts[1:])
    mid, half = (a + b) / 2, (b - a) / 2
    return ((mid[:, None] + half[:, None] * GLx[None, :]).ravel(),
            (half[:, None] * GLw[None, :]).ravel())

def WN(x, ts, d):
    """Interval-membership weights for thresholds ts: shape (len(ts)+1, len(x))."""
    gs = [G(x, t, d) for t in ts]
    rows = [gs[0]]
    for k in range(1, len(gs)):
        rows.append(gs[k] - gs[k - 1])
    rows.append(1 - gs[-1])
    return np.stack(rows)

def brks(ts, d):
    return [t + s * d for t in ts for s in (-1, 1)]

# -----
# two-advisor system
# -----
def sysNM(ts1, ts2, d1, d2):
    """Joint profile probabilities P[j,k] and leader actions y[j,k]."""
    x, w = nodes(0.0, 1.0, brks(ts1, d1) + brks(ts2, d2))
    A, B = WN(x, ts1, d1), WN(x, ts2, d2)
    P = np.einsum('i,ji,ki->jk', w, A, B)
    M = np.einsum('i,i,ji,ki->jk', w, x, A, B)
    return P, np.where(P > TOL, M / np.maximum(P, 1e-300), np.nan)

def impb(tau, d_own, brk_o, fWo, yA, yB, PA, PB):
    """Own bias implied by indifference at threshold tau between two messages."""
    lo, hi = max(0.0, tau - d_own), min(1.0, tau + d_own)
    if hi <= lo:
        return np.nan
    x, w = nodes(lo, hi, brk_o)
    w = w / (hi - lo)
    Wo = fWo(x)
    num = den = 0.0
    for k in range(Wo.shape[0]):
        pk = np.sum(w * Wo[k])
        if pk < 1e-12 or PA[k] < TOL or PB[k] < TOL:
            continue
        mk = np.sum(w * x * Wo[k]) / pk
        D = yB[k] - yA[k]
        num += pk * D * (0.5 * (yA[k] + yB[k]) - mk)
        den += pk * D
    return num / den if abs(den) > 1e-14 else np.nan

def biasesNM(ts1, ts2, d1, d2):
    P, y = sysNM(ts1, ts2, d1, d2)
    b2b, f2 = brks(ts2, d2), (lambda x: WN(x, ts2, d2))
    b1b, f1 = brks(ts1, d1), (lambda x: WN(x, ts1, d1))
    B1 = [impb(ts1[r], d1, b2b, f2, y[r, :], y[r + 1, :], P[r, :], P[r + 1, :])
           for r in range(len(ts1))]
    B2 = [impb(ts2[r], d2, b1b, f1, y[:, r], y[:, r + 1], P[:, r], P[:, r + 1])
           for r in range(len(ts2))]

```

```

return np.array(B1), np.array(B2), P, y

def solve_point(a1, a2, d1, d2, N1, N2, guess):
    """Given the two lowest thresholds, solve for the rest. Returns None on failure.

    For (N1, N2) = (2, 2) there are no consistency equations: the biases are an
    explicit function of the two thresholds, so no solve is needed (40x cheaper).
    """
    n1, n2 = N1 - 2, N2 - 2
    if n1 + n2 == 0:
        ts1, ts2 = np.array([a1]), np.array([a2])
        B1, B2, P, y = biasesNM(ts1, ts2, d1, d2)
        if not (np.isfinite(B1[0]) and np.isfinite(B2[0])):
            return None
        if P.sum(axis=1).min() < 1e-14 or P.sum(axis=0).min() < 1e-14:
            return None
        return ts1, ts2, float(B1[0]), float(B2[0]), P, y

    def F(z):
        ts1 = np.concatenate([[a1], z[:n1]])
        ts2 = np.concatenate([[a2], z[n1:]])
        if np.any(np.diff(ts1) <= 1e-9) or np.any(np.diff(ts2) <= 1e-9):
            return np.full(n1 + n2, 1e2)
        B1, B2, P, _ = biasesNM(ts1, ts2, d1, d2)
        if np.any(np.isnan(B1)) or np.any(np.isnan(B2)):
            return np.full(n1 + n2, 1e2)
        return np.concatenate([B1[1:] - B1[0], B2[1:] - B2[0]])

    try:
        z, _, ier, _ = fsolve(F, guess, full_output=True, xtol=1e-12)
    except Exception:
        return None
    if ier != 1:
        return None
    r = F(z)
    if not np.all(np.isfinite(r)) or np.max(np.abs(r)) > 1e-10:
        return None
    ts1 = np.concatenate([[a1], z[:n1]])
    ts2 = np.concatenate([[a2], z[n1:]])
    if np.any(np.diff(ts1) <= 1e-9) or np.any(np.diff(ts2) <= 1e-9):
        return None
    B1, B2, P, y = biasesNM(ts1, ts2, d1, d2)
    if P.sum(axis=1).min() < 1e-14 or P.sum(axis=0).min() < 1e-14:
        return None
    if not (np.isfinite(B1[0]) and np.isfinite(B2[0])):
        return None
    return ts1, ts2, float(B1[0]), float(B2[0]), P, y

# -----
# single-advisor bounds  bbar^(N)_gen(delta)
# -----
def single_expert_solve(b, d, N, guess):
    """General N-interval single-advisor equilibrium at bias b."""
    if N == 1:
        return np.array([]), np.array([0.5])

    def F(ts):
        ts = np.asarray(ts)
        if np.any(np.diff(ts) <= 1e-10):
            return np.full(N - 1, 1e2)
        x, w = nodes(0.0, 1.0, brks(ts, d))
        W = WN(x, ts, d)
        P = np.einsum('i,ji->j', w, W)
        M = np.einsum('i,i,ji->j', w, x, W)
        if P.min() < TOL:
            return np.full(N - 1, 1e2)

```

```

        y = M / P
        mu = np.array([0.5 * (max(0.0, t - d) + min(1.0, t + d)) for t in ts])
        return 0.5 * (y[:-1] + y[1:]) - mu - b

    z, _, ier, _ = fsolve(F, guess, full_output=True, xtol=1e-13)
    if ier != 1 or np.max(np.abs(F(z))) > 1e-11:
        return None
    ts = np.sort(z)
    x, w = nodes(0.0, 1.0, brks(ts, d))
    W = WN(x, ts, d)
    P = np.einsum('i,ji->j', w, W)
    M = np.einsum('i,i,ji->j', w, x, W)
    if P.min() < 1e-14:
        return None
    return ts, M / P

_BBAR_CACHE = {}

def bbar_gen(d, N, nb=400):
    """Existence bound for general N-interval single-advisor equilibria.

    N=2: 1/4 exactly. N=3: closed form. N>=4: recursion with a vanishing interval
    bbar^*(N) = root of b = y_1^(N-1)(b)/2, marched with continuation in b
    (a plain grid scan fails from N=5 on, where the inner solve is delicate).

    Cached: the uncached version is called once per (combo, delta) configuration,
    which at N=5 costs about 1.1 h of the run on its own.
    """
    key = (round(float(d), 12), N)
    if key in _BBAR_CACHE:
        return _BBAR_CACHE[key]
    if N <= 1:
        v = 0.0
    elif N == 2:
        v = 0.25
    elif N == 3:
        a = (2 - np.sqrt(1 + d * d)) / 3
        v = ((3 * a * a + d * d) / (12 * a) if a >= d else
            brentq(lambda m: 8 * m ** 3 - 12 * d * m + 3 * d,
                0.2500000001, d, xtol=1e-15) / 3)
    else:
        hi = bbar_gen(d, N - 1)
        M = N - 1
        prev = None
        xs, ys = [], []
        for b in np.linspace(1e-7, hi * (1 - 1e-9), nb):
            r = single_expert_solve(b, d, M, np.linspace(1.0 / M, 1 - 1.0 / M, M - 1))
            if prev is None else prev)
            if r is None:
                r = single_expert_solve(b, d, M, np.linspace(1.0 / M, 1 - 1.0 / M, M - 1))
            if r is None:
                continue
            prev = r[0]
            xs.append(b)
            ys.append(0.5 * r[1][0] - b)
        xs, ys = np.array(xs), np.array(ys)
        v = np.nan
        for i in range(len(xs) - 1):
            if ys[i] * ys[i + 1] < 0:
                lo, hi2 = xs[i], xs[i + 1]
                for _ in range(60):
                    mid = 0.5 * (lo + hi2)
                    r = single_expert_solve(mid, d, M, prev)
                    if r is None:
                        break
                if (0.5 * r[1][0] - mid) * ys[i] > 0:

```

```

        lo = mid
    else:
        hi2 = mid
        v = 0.5 * (lo + hi2)
        break
    _BBAR_CACHE[key] = v
    return v

def prime_bbar_cache(dg, nrange):
    """Precompute every bound once; workers inherit the cache by fork."""
    for d in dg:
        for N in nrange:
            bbar_gen(float(d), N)
    return dict(_BBAR_CACHE)

# -----
# stage C: global incentive compatibility
# -----
def _signal_knots(ts_own, d_own, ts_other, d_other):
    """Breakpoints of all signal-wise deviation-gain polynomials."""
    lo, hi = -d_own, 1 + d_own
    candidates = [lo, hi, d_own, 1 - d_own]
    candidates.extend(np.asarray(ts_own, dtype=float))
    for x_break in brks(ts_other, d_other):
        candidates.extend((x_break - d_own, x_break + d_own))
    return np.array(sorted({float(z) for z in candidates if lo <= z <= hi}))

def _cubic_maximum(fun, a, b):
    """Maximum on [a,b] of a function known to be cubic there."""
    if b - a <= 1e-14:
        return -np.inf
    mid, half = 0.5 * (a + b), 0.5 * (b - a)
    z_fit = np.array([-0.8, -0.25, 0.3, 0.85])
    vals = np.array([fun(mid + half * z) for z in z_fit])
    poly = np.polynomial.Polynomial.fit(z_fit, vals, 3).convert()
    candidates = [-1.0, 1.0]
    candidates.extend(r.real for r in poly.deriv().roots()
                      if abs(r.imag) < 1e-8 and -1 < r.real < 1)
    return max(float(fun(mid + half * z)) for z in candidates)

def _report_utility_numerator(s, d_own, other_breaks, other_weights,
                             actions, bias, report):
    """Expected report utility times the conditional-support length.

    The report-independent term is retained here; it cancels in differences.
    """
    lo, hi = max(0.0, s - d_own), min(1.0, s + d_own)
    if hi <= lo:
        return 0.0
    x, w = nodes(lo, hi, other_breaks)
    weights = other_weights(x)
    return -sum(np.sum(w * weights[k] * (actions[report, k] - x - bias) ** 2)
               for k in range(weights.shape[0]))

def _advisor_ic(ts_own, d_own, ts_other, d_other, actions, bias, tol=1e-11):
    """Check every deviation at endpoints and stationary points of each piece."""
    other_breaks = brks(ts_other, d_other)
    other_weights = lambda x: WN(x, ts_other, d_other)
    knots = _signal_knots(ts_own, d_own, ts_other, d_other)
    for a, b in zip(knots[:-1], knots[1:]):
        if b - a <= 1e-14:
            continue
        prescribed = int(np.searchsorted(ts_own, 0.5 * (a + b)))

```

```

    for deviation in range(actions.shape[0]):
        if deviation == prescribed:
            continue

        def excess(s):
            length = min(1.0, s + d_own) - max(0.0, s - d_own)
            gain = (_report_utility_numerator(s, d_own, other_breaks,
                                                other_weights, actions, bias, deviation)
                    - _report_utility_numerator(s, d_own, other_breaks,
                                                other_weights, actions, bias, prescribed))
            return gain - tol * max(length, 0.0)

        if _cubic_maximum(excess, a, b) > 1e-13:
            return False
    return True

def global_ic(ts1, ts2, d1, d2, b1, b2, P, y, nS=None, nfill=25):
    """Check that each advisor's message choice is optimal at every own signal.

    For each possible message deviation, the payoff gain multiplied by the
    positive conditional-support length is piecewise cubic. The check evaluates
    every endpoint and stationary point rather than sampling a uniform grid.
    'nS' is retained only for compatibility with older calls and is ignored.
    Off-path actions are searched over the same common fill grid as before.
    """
    miss = P <= TOL
    fills = [0.5] if not miss.any() else np.linspace(0.0, 1.0, nfill)
    for fv in fills:
        yy = np.where(miss, fv, y)
        if not _advisor_ic(ts1, d1, ts2, d2, yy, b1):
            continue
        if _advisor_ic(ts2, d2, ts1, d1, yy.T, b2):
            return True, float(fv)
    return False, None

# -----
# grids
# -----
def delta_grid(n):
    return np.geomspace(0.004, 0.48, n)

def u_grid(n, umin=1e-12, umax=0.60):
    return np.geomspace(umin, umax, n)

def default_guess(a1, a2, N1, N2):
    g1 = np.linspace(a1 + (1 - a1) / (N1 - 1), 1 - (1 - a1) / (N1 - 1), N1 - 2) \
        if N1 > 2 else np.array([])
    g2 = np.linspace(a2 + (1 - a2) / (N2 - 1), 1 - (1 - a2) / (N2 - 1), N2 - 2) \
        if N2 > 2 else np.array([])
    return np.concatenate([g1, g2])

# -----
# worker: one (N1, N2, delta1, delta2) configuration
# -----
def run_cell(task):
    """Stage A sweep + stage B refinement for one parameter configuration."""
    N1, N2, d1, d2, nA, nB, budget = task
    t0 = time.time()
    B1bar = bbar_gen(d1, N1)
    if not np.isfinite(B1bar):
        return dict(key=(N1, N2, d1, d2), error='bbar failed')

    rows = []

```

```

seed = None
U1, U2 = u_grid(nA), u_grid(nA)

# ---- stage A: coarse sweep -----
for u1 in U1:
    if time.time() - t0 > STAGE_A_SHARE * budget:
        break
    a1 = -d1 + u1
    for u2 in U2:
        if time.time() - t0 > STAGE_A_SHARE * budget:
            break
        a2 = -d2 + u2
        g = seed if seed is not None else default_guess(a1, a2, N1, N2)
        r = solve_point(a1, a2, d1, d2, N1, N2, g)
        if r is None and seed is not None:
            r = solve_point(a1, a2, d1, d2, N1, N2,
                           default_guess(a1, a2, N1, N2))
        if r is None:
            continue
        ts1, ts2, b1, b2, P, _ = r
        seed = np.concatenate([ts1[1:], ts2[1:]])
        rows.append((u1, u2, b1, b2, b1 / B1bar,
                     P.sum(axis=1).min(), P.sum(axis=0).min()))

if not rows:
    return dict(key=(N1, N2, d1, d2), bbar=B1bar, rows=[], best=None,
               best_ord=None, npts=0)

R = np.array(rows)

# ---- stage B: adaptive zoom on the ordered-region maximum -----
def zoom(mask, label):
    sub = R[mask]
    if not len(sub):
        return None
    i = int(np.argmax(sub[:, 4]))
    u1c, u2c = sub[i, 0], sub[i, 1]
    best = (sub[i, 4], sub[i])
    span = 8.0
    while time.time() - t0 < budget and span > 1.0000001:
        got = False
        for f1 in np.geomspace(1 / span, span, nB):
            if time.time() - t0 > budget:
                break
            for f2 in np.geomspace(1 / span, span, nB):
                a1, a2 = -d1 + u1c * f1, -d2 + u2c * f2
                if a1 <= -d1 or a2 <= -d2:
                    continue
                r = solve_point(a1, a2, d1, d2, N1, N2,
                               default_guess(a1, a2, N1, N2))
                if r is None:
                    continue
                ts1, ts2, b1, b2, P, _ = r
                if label == 'ord' and b2 < b1:
                    continue
                ratio = b1 / B1bar
                if ratio > best[0]:
                    best = (ratio, np.array([u1c * f1, u2c * f2, b1, b2, ratio,
                                             P.sum(axis=1).min(),
                                             P.sum(axis=0).min()]))
            got = True
        if got:
            u1c, u2c = best[1][0], best[1][1]
            span = np.sqrt(span)
    return best

best_any = zoom(np.ones(len(R), bool), 'any')
best_ord = zoom(R[:, 3] >= R[:, 2], 'ord')

```

```

        return dict(key=(N1, N2, d1, d2), bbar=B1bar, npts=len(R),
                    best=best_any, best_ord=best_ord,
                    near=R[R[:, 4] > 0.98], elapsed=time.time() - t0)

# -----
# driver
# -----
def main():
    ap = argparse.ArgumentParser()
    ap.add_argument('--hours', type=float, default=2.0)
    ap.add_argument('--workers', type=int, default=1)
    ap.add_argument('--ndelta', type=int, default=16)
    ap.add_argument('--nA', type=int, default=20, help='stage-A grid per axis')
    ap.add_argument('--nB', type=int, default=9, help='stage-B zoom grid per axis')
    ap.add_argument('--out', default=None)
    ap.add_argument('--resume', action='store_true')
    ap.add_argument('--smoke', action='store_true')
    args = ap.parse_args()

    if args.out is None:
        args.out = 'NM_smoke' if args.smoke else 'NM_day'
    if args.smoke:
        args.hours, args.ndelta, args.nA, args.nB = 120 / 3600, 2, 5, 3

    combos = [(N1, N2) for N1 in NRANGE for N2 in NRANGE]
    dg = delta_grid(args.ndelta)
    print('priming single-advisor bounds ...', flush=True)
    prime_bbar_cache(dg, NRANGE)
    tasks = [(N1, N2, float(d1), float(d2))
              for (N1, N2) in combos for d1 in dg for d2 in dg]

    total = args.hours * 3600
    reserve = 0.10 * total # stage C + reporting
    per = (total - reserve) / len(tasks)

    ckpt = args.out + '.pkl'
    done = {}
    if args.resume and os.path.exists(ckpt):
        with open(ckpt, 'rb') as f:
            done = pickle.load(f)
        print('resumed with %d configurations complete' % len(done))

    todo = [t for t in tasks if t not in done]
    payload = [(N1, N2, d1, d2, args.nA, args.nB, per) for (N1, N2, d1, d2) in todo]

    t0 = time.time()
    if args.workers > 1:
        import multiprocessing as mp
        with mp.Pool(args.workers) as pool:
            for k, res in enumerate(pool.imap_unordered(run_cell, payload)):
                done[res['key']] = res
                if k % 20 == 0:
                    with open(ckpt, 'wb') as f:
                        pickle.dump(done, f)
                    print(' %d/%d configurations, %.1f min elapsed'
                          % (len(done), len(tasks), (time.time() - t0) / 60), flush=True)
    else:
        for k, p in enumerate(payload):
            done[tuple(p[:4])] = run_cell(p)
            if k % 20 == 0:
                with open(ckpt, 'wb') as f:
                    pickle.dump(done, f)
                print(' %d/%d configurations, %.1f min elapsed'
                      % (len(done), len(tasks), (time.time() - t0) / 60), flush=True)
    with open(ckpt, 'wb') as f:
        pickle.dump(done, f)

```

```

# ---- report -----
print('\n=== STATEMENT (1): b1 <= bbar^(N1)_gen(delta1)? ===')
print('configurations: %d solving time: %.2f h'
      % (len(done), (time.time() - t0) / 3600))
for (N1, N2) in combos:
    sel = [v for k, v in done.items() if k[0] == N1 and k[1] == N2
           and v.get('best_ord')]
    if not sel:
        print(' (%d,%d): no ordered-region points' % (N1, N2))
        continue
    ro = max(v['best_ord'][0] for v in sel)
    ra = max(v['best'][0] for v in sel if v.get('best'))
    nv = sum(1 for v in sel if v['best_ord'][0] > 1.0)
    npts = sum(v.get('npts', 0) for v in sel)
    print(' (%d,%d): %7d pts | max ratio ordered %.6f | unrestricted %.6f '
          '| precision pairs violating (ordered) %d'
          % (N1, N2, npts, ro, ra, nv))
print('\nNA ratio strictly above 1.000000 in the ordered region would refute '
      'statement (1); stage C then re-checks those points under global IC.')

if __name__ == '__main__':
    main()

```

`nm_check.py` — the check behind Table 2.

```

#!/usr/bin/env python3
"""Does existence of a two-advisor N x M equilibrium imply that advisor 1 alone
can sustain N intervals?

For N in {3,4} and M in {2,3,4}, sweep the two lowest thresholds at a grid of
precision pairs, keep the solutions with  $0 < b_1 < b_2 < 1/4$  and every message
of positive probability, and report the largest ratio  $b_1 / b_{gen}^N(\delta_1)$  among
them -- first over all solutions of the marginal conditions, then over those that
are also globally incentive compatible. A ratio below one at every pair is the
numerical content of the claim. N = 2 is trivial, since  $b_{gen}^2 = 1/4$  and
 $b_1 < 1/4$  is maintained throughout.
"""
import sys
import time
import pickle

import numpy as np

from nm import solve_point, bbar_gen, u_grid, default_guess, global_ic

PAIRS = [(0.05, 0.01), (0.10, 0.05), (0.15, 0.05), (0.20, 0.10), (0.25, 0.005),
          (0.30, 0.10), (0.35, 0.20), (0.40, 0.05), (0.45, 0.30), (0.49, 0.25)]
TYPES = [(N, M) for N in (3, 4) for M in (2, 3, 4)]
NA = 20
NTOP = 40

def one(d1, d2, N, M):
    t0 = time.time()
    B1bar = bbar_gen(d1, N)
    rows, seed = [], None
    for u1 in u_grid(NA):
        for u2 in u_grid(NA):
            a1, a2 = -d1 + u1, -d2 + u2
            g = seed if seed is not None else default_guess(a1, a2, N, M)
            r = solve_point(a1, a2, d1, d2, N, M, g)
            if r is None and seed is not None:
                r = solve_point(a1, a2, d1, d2, N, M, default_guess(a1, a2, N, M))
            if r is None:

```



```

"""
import numpy as np
from nm import solve_point, bbar_gen, u_grid, default_guess, global_ic

PAIRS = [(0.05, 0.01), (0.10, 0.05), (0.15, 0.05), (0.20, 0.10), (0.25, 0.005),
         (0.30, 0.10), (0.35, 0.20), (0.40, 0.05), (0.45, 0.30), (0.49, 0.25)]
TYPES = [(3, 2), (4, 2)]
FLOORS = (1e-6, 1e-2)
NA = 20

def sweep(d1, d2, N, M):
    rows, seed = [], None
    for u1 in u_grid(NA):
        for u2 in u_grid(NA):
            a1, a2 = -d1 + u1, -d2 + u2
            g = seed if seed is not None else default_guess(a1, a2, N, M)
            r = solve_point(a1, a2, d1, d2, N, M, g)
            if r is None and seed is not None:
                r = solve_point(a1, a2, d1, d2, N, M, default_guess(a1, a2, N, M))
            if r is None:
                continue
            ts1, ts2, b1, b2, P, y = r
            seed = np.concatenate([ts1[1:], ts2[1:]])
            if not (0 < b1 < b2 < 0.25):
                continue
            rows.append((b1, min(P.sum(axis=1).min(), P.sum(axis=0).min()), a1, a2))
    return np.array(rows) if rows else np.zeros((0, 4))

if __name__ == '__main__':
    print('%-14s %-6s %12s %12s %12s'
          % ('(d1,d2)', '(N,M)', 'no floor', 'probability>=1e-6', 'probability>=1e-2'),
          flush=True)
    worst = {f: 0.0 for f in (0.0,) + FLOORS}
    for (d1, d2) in PAIRS:
        for (N, M) in TYPES:
            B = bbar_gen(d1, N)
            R = sweep(d1, d2, N, M)
            out = []
            for f in (0.0,) + FLOORS:
                S = R[R[:, 1] >= f]
                best = None
                for i in np.argsort(-S[:, 0])[:40]:
                    r = solve_point(S[i, 2], S[i, 3], d1, d2, N, M,
                                    default_guess(S[i, 2], S[i, 3], N, M))
                    if r is None:
                        continue
                    ts1, ts2, b1, b2, P, y = r
                    if not (0 < b1 < b2 < 0.25):
                        continue
                    if min(P.sum(axis=1).min(), P.sum(axis=0).min()) < f:
                        continue
                    if global_ic(ts1, ts2, d1, d2, b1, b2, P, y, nS=401, nfill=25)[0]:
                        best = b1 / B
                        break
                out.append(best)
                if best is not None:
                    worst[f] = max(worst[f], best)
            print('%(%.2f,%.3f) (%d,%d) %12s %12s %12s'
                  % (d1, d2, N, M,
                     *['---' if v is None else '%.6f' % v for v in out]), flush=True)
    print('\nlargest ratio anywhere: no floor %.6f; probability >= 1e-6 %.6f; '
          'probability >= 1e-2 %.6f'
          % (worst[0.0], worst[1e-6], worst[1e-2]))

```

validate.py — the cross-check of cs_core.py against the general-partition code and the internal

consistency checks reported in Part C.

```

"""
Check cs_core against nm.py, and the closed forms against each other.

Put nm.py in this directory and run:
    python3 validate.py
"""
import time
import numpy as np

from cs_core import (two_by_two, implied_biases, global_ic_22, single_general,
                     bgen3, bbar_int, Psi)

try:
    import nm
except ImportError:
    nm = None

def against_script(ntrial=400, seed=1):
    if nm is None:
        print('nm.py not found; skipping the cross-check')
        return
    rng = np.random.default_rng(seed)
    dP = dy = db = 0.0
    tested = agree = ic_true = 0
    t_new = t_old = 0.0
    for _ in range(ntrial):
        d1 = rng.uniform(0.01, 0.45)
        d2 = rng.uniform(0.002, min(d1, 0.49 - d1))
        # half the draws in the boundary layer, where the interesting equilibria are
        t1 = ((-d1 + 10 ** rng.uniform(-6, np.log10(2 * d1))) if rng.uniform() < 0.5
              else rng.uniform(-d1, 1 + d1))
        t2 = ((-d2 + 10 ** rng.uniform(-6, np.log10(2 * d2))) if rng.uniform() < 0.5
              else rng.uniform(-d2, 1 + d2))
        P, y = two_by_two(t1, t2, d1, d2)
        b1, b2 = implied_biases(t1, t2, d1, d2, P, y)
        B1, B2, Ps, ys = nm.biasesNM(np.array([t1]), np.array([t2]), d1, d2)
        dP = max(dP, np.max(np.abs(P - Ps)))
        m = (P > 1e-9) & (Ps > 1e-9)
        if m.any():
            dy = max(dy, np.max(np.abs(y[m] - ys[m])))
        for a, b in ((b1, B1[0]), (b2, B2[0])):
            if np.isfinite(a) and np.isfinite(b):
                db = max(db, abs(a - b))
        if not (np.isfinite(b1) and np.isfinite(b2)):
            continue
        if min(P.sum(axis=1).min(), P.sum(axis=0).min()) < 1e-12:
            continue
        tested += 1
        t = time.time()
        ok_new, _ = global_ic_22(t1, t2, d1, d2, b1, b2, P, y, nS=201, nfill=11)
        t_new += time.time() - t
        t = time.time()
        ok_old, _ = nm.global_ic(np.array([t1]), np.array([t2]), d1, d2, B1[0], B2[0], Ps, ys,
                                nS=201, nfill=11)
        t_old += time.time() - t
        ic_true += ok_new
        agree += (ok_new == ok_old)
    print('vs nm.py: max |dP| = %.1e, max |dy| = %.1e, max |d bias| = %.1e' % (dP, dy, db))
    print('incentive check: %d configurations, %d compatible, agreement %d/%d, %.2fs vs %.2fs'
          % (tested, ic_true, agree, tested, t_new, t_old))

def internal():
    # the general two-interval equilibrium coincides with the interior formulas
    worst = 0.0

```

```

for d in (0.005, 0.05, 0.1, 0.25, 0.4):
    for b in np.linspace(0.001, bbar_int(d) - 1e-6, 40):
        g = single_general(b, d)
        s = g['t']
        yL = (d * d + 3 * s * s) / (6 * s)
        yH = (3 - d * d - 3 * s * s) / (6 * (1 - s))
        V = -1 / 3 + s * yL ** 2 + (1 - s) * yH ** 2
        worst = max(worst, abs(V - g['V']), abs(yL - g['yL']), abs(yH - g['yH']))
print('interior branch vs Proposition 6 formulas: max discrepancy %.1e' % worst)

# Psi strictly decreasing onto (-1/4, 1/4)
bad = 0
for d in (0.005, 0.1, 0.3, 0.49):
    mus = np.linspace(1e-6, 1 - 1e-6, 20000)
    vals = np.array([Psi(m, d) for m in mus])
    bad += int((np.diff(vals) >= 0).sum())
print('Psi monotone on a 20000-point grid at four precisions: violations %d' % bad)

# V strictly decreasing in b, and bgen3 larger for the noisier advisor
bad = 0
for d in (0.005, 0.1, 0.25, 0.4):
    V = np.array([single_general(b, d)['V'] for b in np.linspace(0.001, 0.249, 2000)])
    bad += int((np.diff(V) >= 0).sum())
print('V decreasing in b: violations %d' % bad)
worst = -1
for d1 in np.linspace(0.002, 0.498, 200):
    for d2 in np.linspace(0.001, 0.498, 200):
        if d2 < d1 and d1 + d2 < 0.5:
            worst = max(worst, bgen3(d2) - bgen3(d1))
print('max over Delta of bgen3(delta2) - bgen3(delta1): %.1e (negative means the '
      'noisier advisor has the larger bound)' % worst)

if __name__ == '__main__':
    against_script()
    internal()

```

The check for advisor 2 reported in Part C uses `nm_check.py` and `nm_floor.py` with the roles of the two advisors interchanged: TYPES set to (2,3), (2,4), (3,3) and (3,4), the solutions ordered by `b2`, and the ratio computed as `b2 / bbar_gen(d2, M)`.