

Investment committee communication and institutional portfolio choice

Scott Condie*

This version: August 14, 2026

Abstract

I embed the [DeGroot \(1974\)](#) model of opinion formation into mean–variance portfolio selection to study how an investment committee’s internal communication shapes the portfolio it produces. Members hold noisy views about expected returns, revise them over a row-stochastic influence network, and a manager aggregates the post-deliberation views into the active tilt that drives the portfolio. This maps the committee’s communication style—its influence matrix, aggregation rule, and length of deliberation—in closed form into its performance, as measured by tracking error, information ratio, and alpha. Concentrating influence into a committee member transforms that member’s idiosyncratic estimation error into uncompensated active risk. Informationally efficient committees instead weight each member by the precision of her information. When influence is matched to skill, committee performance improves. Deliberation helps or hurts according to whether it moves effective influence toward skill. I quantify these effects in Monte Carlo simulations and state the resulting empirical predictions.

Keywords: committees, institutional investors, portfolio choice, governance, decision-making under uncertainty

JEL Classification: G11, G23, D71, D83, D85, G41

*Department of Economics, Brigham Young University, Provo, UT 84602. Email: scott_condie@byu.edu.

1 Introduction

Most institutional capital is managed by committees. Pension boards, endowment investment committees, mutual-fund management teams, and the oversight bodies of sovereign wealth funds decide collectively which assets to hold and in what proportion. The standard theory of portfolio choice, descending from [Markowitz \(1952\)](#), nonetheless treats the investor as a single decision-maker with a single set of beliefs. This paper takes up a question that sits in the gap between that theory and the institutional reality it is meant to describe: to what extent do the human interactions among the members of an investment committee—who listens to whom, whose views carry weight, and whether the group ultimately agrees—affect the performance of the portfolio the committee oversees? That the composition of investment teams matters is by now well documented—diverse teams, for example, outperform homogeneous ones ([Evans et al. 2025](#))—but the dynamics of who influences whom inside the room are largely invisible in the data. We observe the holdings a committee produces, not the deliberation that produced them, and so we have little sense of how much of measured performance is attributable to the committee’s internal structure rather than to the raw skill of its members.

That these dynamics should matter is something many will sense from experience. Consider two committees identical on paper—the same members, expertise, and access to information—but run in opposite ways. The first is led by a domineering chair who states his view first and treats dissent as an obstacle to be managed; the quieter members learn that their reservations change nothing and fall silent, so whatever private information they hold never enters the decision. The second is led by a chair who treats her own view as one input among several, drawing out reluctant members and protecting disagreement long enough for it to be weighed. Beginning from identical raw beliefs, the two committees arrive at different decisions—and, this paper argues, at different portfolios—purely because of how their members listen to one another. The claim is that committee communication is not a matter of collegiality but a measurable determinant of performance.

To study the question I embed a canonical model of group opinion formation, the [DeGroot \(1974\)](#) updating process, into the mean-variance portfolio selection framework of [Markowitz \(1952\)](#). Each member of the committee begins with a private, noisy belief about the vector of expected asset returns, expressed as a view relative to market-implied values. The members then deliberate. In each round, members replace their beliefs with a weighted average of their colleagues’ beliefs and their own. The extent to which each committee member is influenced by her colleagues on the committee is modeled by a row-stochastic influence matrix that encodes the committee’s communication structure—who listens to whom, and how much. After deliberation, an aggregation rule—the act of the portfolio manager or

chair who must turn the committee’s views into a single decision—combines the members’ post-discussion beliefs into one *effective view*. Added to the market-implied prior, that view feeds a mean-variance optimization, producing the institution’s portfolio. This portfolio can be broken down into the benchmark the institution would hold absent any view plus an active tilt driven by the committee’s view. I evaluate this portfolio using the metrics of professional asset management: the Sharpe ratio, tracking error, information ratio, and Jensen’s alpha. Because the influence matrix is what distinguishes one committee from another, I study a family of stylized archetypes—among them a committee dominated by a single voice, one split into rival factions, one in which a stubborn member is heeded by all, and one in which she is ignored—and trace each through the model both analytically and in Monte Carlo simulations.

The analysis yields four sets of results. The first concerns when a committee reaches a shared view.¹ Deliberation drives the committee to a single consensus if and only if its communication network contains one closed group toward which influence flows, directly or indirectly, from all members. Intuitively, this occurs when the committee holds “one conversation” in which everyone’s view is ultimately traceable to a common set of voices. What prevents consensus is not the presence of a stubborn member but *mutual deafness*: the committee fails to agree precisely when it fractures into subgroups that neither listen to nor are heard by one another. A single immovable member who listens to no one does not, by herself, block agreement. If her colleagues heed her, she draws the entire committee to her view. If they ignore her, however, she is isolated and the committee fragments. Consensus, in short, is a property of the structure of attention within the committee rather than of the stubbornness of individuals.

The second set of results concerns the portfolio manager who must aggregate the committee’s views. When the committee reaches consensus, all members hold the same belief, and any aggregation rule that respects this unanimity will use this belief in forming portfolios. In these settings the manager is a bystander, and the portfolio is determined by the deliberation of the committee. When the committee does not agree, members hold different beliefs and the aggregation rule must adjudicate among them, so the portfolio depends on the manager’s balancing of these differing beliefs. If the portfolio manager chooses a portfolio that weights committee members’ post-discussion beliefs equally, the portfolio represents an averaging of these beliefs. In other settings, the chair of the committee holds final say on committee

¹This first set of results restates, in the committee setting, the classical characterization of consensus in DeGroot-style repeated-averaging models. DeGroot (1974) gave a sufficient condition and Berger (1981) the necessary and sufficient one, while DeMarzo et al. (2003) and Golub and Jackson (2010) analyze the same dynamics in economic settings and Golub and Sadler (2016) survey them. This paper takes those consensus conditions and studies their consequences for portfolio choice and performance.

decisions, so the portfolio manager’s choice may correspond to the post-discussion opinions of the chair. The manager’s discretion, and with it the institution’s choice of governance rule, therefore matters most precisely when the committee is internally divided.

The third set of results concerns performance, and several of the findings may be surprising. Concentrating influence on a few members imports their idiosyncratic estimation error rather than averaging it away, so realized risk-adjusted performance falls as influence becomes more concentrated. At the far end of this continuum, the committee that adopts a single stubborn member’s unfiltered view performs worst of all. In this way, committee dynamics that overweight one individual’s perspective have the potential to subject the institution’s portfolio to a kind of “idiosyncratic risk”—not in the asset selection, but in the signals that are incorporated into the final portfolio selection. Deferring to a *biased* chair is worse still. This deference imports both the chair’s idiosyncratic noise and her systematic error, and because averaging cancels noise but not bias, the bias is irreducible. The only way to avoid this two-dimensional performance hit is to spread influence back onto the other committee members. These differences, moreover, are differences in risk-adjusted return rather than in alpha. When the market is efficient no committee earns alpha, because tilting away from a correct market buys active risk without return, so a mistaken view is paid for in a lower Sharpe ratio and the committee earns true alpha only to the extent its view is genuinely informed.

Concentrated committees are also systematically overconfident, perceiving information ratios several times larger than they realize, because one strong, unaveraged opinion appears attractive before the fact and underperforms after it. Notably, whether a committee reaches consensus is by itself a poor predictor of performance. Under equal-weight pooling the committees that fail to agree match the best-performing committee that does, because averaging re-mixes their fragmentary consensuses back into the wisdom of the crowd. That result does not hold, however, under a chair rule. In that case the same divided committees vary widely with the identity of the chair.

The model further speaks to the extent to which deliberation improves the decisions made by the committee. Under equal-weight pooling, additional rounds of discussion can *erode* performance by concentrating influence on dominant members, so that a committee which deliberates briefly outperforms the same committee that deliberates to consensus. Under a chair rule the opposite holds, because discussion improves the belief of a chair who begins knowing only her own view. The length of deliberation, like its structure, is not neutral for the portfolio it produces.

The fourth set of results concerns the *optimal* communication structure—how a committee *should* allocate influence. A committee is no better informed than its best-informed member

and no worse informed than its worst. Communication and aggregation only redistribute influence among the views already in the room and cannot manufacture insight no member brings to the committee’s deliberations. Good performance therefore demands two complementary things—*selecting* able members, which sets the ceiling on what the committee can achieve, and *attending* to them, which determines how close it comes to that ceiling. When members are equally skilled the best structure is the fully diffuse, egalitarian committee, which averages away the most idiosyncratic error, and the costs of concentration described above are simply the penalty for departing from it. More generally, when members differ in skill the performance-maximizing committee weights each member by the precision of her information. This decision rule coincides with beliefs that a Bayesian who observed all members’ signals would form. Concentration is then not costly in itself. What is costly is concentration *misaligned* with skill. Influence matched to a genuine expert can beat the equal-weight “wisdom of crowds,” whereas influence that accrues to whoever is most central or most vocal—the natural tendency, under naive updating, of the persuasion bias of [DeMarzo et al. \(2003\)](#)—falls short of it. The gap between the influence a committee actually pays and the influence its members’ skill warrants is the cost of deliberating naively, and it is also why deliberation helps only when it moves effective influence toward skill. In a calibrated example, precision weighting raises the realized information ratio by more than a fifth over equal weighting, concentrating on a lone expert outperforms equal weighting, and concentrating on a weak member decreases performance substantially.

This paper makes four methodological contributions. First, it derives a closed-form map from a committee’s primitives—its influence matrix $\mathbf{W}$, aggregation rule f , and number of deliberation rounds k —to the standard active-management metrics, expressing tracking error, information ratio, and Jensen’s alpha as explicit functions of the committee’s deliberation. Second, it shows that the variance of the committee’s accuracy is governed by a single scalar, the Herfindahl concentration of effective influence, so that concentrating influence raises active risk without raising expected return. Third, it characterizes the performance-maximizing committee using precision (inverse-variance) weighting and identifies the cost of naive deliberation with the gap between realized and precision-optimal influence. This shows the value of an extra round of discussion to be the movement discussion induces toward or away from skill. Fourth, it gives a spectral characterization of consensus and the limiting influence weights, and quantifies all of these effects in Monte Carlo simulations.

2 Related Literature and Hypotheses

This paper sits at the intersection of three literatures: social learning on networks, the theory of group and committee decision-making in finance, and mean-variance portfolio choice. Its contribution is to join them—to my knowledge for the first time—by treating the investment committee’s internal deliberation as a [DeGroot \(1974\)](#) updating process whose aggregate view drives an *active* portfolio that is then evaluated with standard active-management metrics.

The belief-updating block of the model comes from the opinion-dynamics process of [DeGroot \(1974\)](#), in which agents repeatedly replace their beliefs with a weighted average of their neighbors’. [Golub and Jackson \(2010\)](#) characterize when this process delivers the “wisdom of crowds”², and show that each agent’s weight in that consensus equals her eigenvector centrality in the communication network. These results also apply to the committees studied in this paper—they determine the extent to which each individual’s beliefs influence the aggregated belief μ^* and how the structure of the communication matrix $\mathbf{W}$ shapes it. Whereas this literature studies society-wide learning about a single state, I apply the DeGroot-style updating process to the repeated deliberation of a small professional committee choosing portfolio tilts. For this reason, the convergence results for a given committee size are more central than the results for large populations.

Why model the deliberations of sophisticated investment professionals as naive DeGroot averaging rather than fully Bayesian inference? The justification is [DeMarzo et al. \(2003\)](#), who micro-found exactly this updating rule. Agents who repeatedly average their neighbors’ opinions but fail to undo the redundancy of information that reaches them through many paths—“persuasion bias”—update precisely as in [DeGroot \(1974\)](#), double-counting the views of well-connected members; [DeMarzo et al. \(2003\)](#) show that the resulting long-run influence is governed by a member’s network position rather than by the quality of her information, and that the distortion is hard to detect from the inside. The DeGroot process is therefore not an assumption of stupidity but a tractable model of a familiar committee failure: the views of central, senior, or simply more voluble members are heard repeatedly and credited each time, so they come to carry weight out of proportion to their accuracy. This is the friction the present paper traces into portfolio performance. It also makes the fully Bayesian committee—which weights each member by the precision of her signal, irrespective of her standing in the room—the natural normative benchmark against which deliberative distortions are measured. [Proposition 2](#) below shows that this Bayesian benchmark is exactly the committee that DeGroot deliberation reaches only when influence happens to track skill, so that the cost of

²Specifically, when the societal belief vector converges in probability to the true state as the number of agents grows large.

naive updating is the wedge between the two.

In addition to the network literature in [Golub and Jackson \(2010\)](#), a growing literature embeds network-based belief formation into asset markets. [Pedersen \(2022\)](#) derives closed-form prices, portfolios, and beliefs in an economy populated by naive (DeGroot) learners alongside rational traders and “fanatics,” generating influencers, bubbles, momentum, and reversal. That paper studies a single asset and looks at price dynamics, whereas I study a multi-asset setting and look at the portfolio implications of the belief dynamics of committee members.

[Hatcher and Hellmann \(2025\)](#) let DeGroot-style influence weights depend on past trading performance and show how network centrality and diameter shape price dynamics and the limiting distribution of belief types in a market with a single risky asset. The current paper differs in that I seek to understand the committee dynamics within a single institution, modeling its *internal* belief aggregation and tracing the effect on the institution’s portfolio and relative performance, holding prices fixed.

The economic motivation for the communication matrix $\mathbf{W}$ and the aggregator f comes from the theory of collective choice under heterogeneous beliefs. [Garlappi et al. \(2017\)](#) study dynamic corporate investment decided by a group of agents with heterogeneous beliefs and fixed influence weights, showing that the aggregation rule itself can generate dynamic inconsistency and underinvestment—direct evidence that *how* beliefs are combined, not merely their average, matters for outcomes. On the empirical side, [Evans et al. \(2025\)](#) find that diverse investment teams outperform homogeneous ones through better decision making and monitoring, with the benefit eroded by polarization. These findings motivate treating committee structure as a first-order determinant of portfolio choice and connect to the broader social-finance agenda surveyed by [Kuchler and Stroebe \(2021\)](#).

The closest antecedent to this paper’s framework in corporate finance is [Garlappi et al. \(2017\)](#). They study a group of managers with heterogeneous, ambiguous beliefs who choose a sequence of real investments, and show that aggregating those beliefs through a Pareto-optimal group rule makes the firm behave as a single agent with time-varying, pessimistically tilted beliefs—generating dynamic inconsistency and underinvestment. This paper shares their central premise—that the rule for combining heterogeneous beliefs is itself an economic primitive with first-order consequences—but differs in three respects that make the two analyses complementary rather than overlapping. First, where their aggregator is a static, ambiguity-averse Pareto weighting, this paper studies the outcome of an explicit *deliberative* process. Beliefs are not pooled once but revised through repeated communication on a network, so the effective weights $\mathbf{w}_k = (\mathbf{W}')^k \mathbf{a}$ evolve with the length of discussion, and—as the simulations show—the value of deliberation can change sign with the governance

rule. Second, where their object is the *level* of investment and its dynamic consistency, this paper looks at the *risk profile* of a portfolio—tracking error, information ratio, and the diversification of estimation error—so the cost of committee dynamics is not underinvestment but uncompensated active risk imported by concentrated influence. Third, in their model disagreement under ambiguity resolves into conservatism, whereas in the unbiased case studied here disagreement is, under the right aggregator, *harmless*: averaging re-mixes a divided committee’s fragmentary consensuses back into the wisdom of the crowd (Section 4.5). The two papers thus identify distinct costs of belief aggregation.

The portfolio allocation mechanism used in this paper is the standard mean-variance framework of Markowitz (1952), with the benchmark-relative structure of Black and Litterman (1992): the institution holds the benchmark plus an active tilt, and the committee’s deliberation determines that tilt by fixing the active view $\boldsymbol{\mu}^*$ behind it. This active view is an endogenous output of the committee’s deliberation, so the communication structure $\mathbf{W}$ and portfolio manager’s decision f map directly into the performance metrics of the institution’s portfolio: the Sharpe ratio, tracking error, and information ratio.

3 Theory

The investment committee is composed of I members, each of whom begins the investment process with beliefs about the expected returns of the N assets in the market. These beliefs are expressed as a vector $\boldsymbol{\mu}_i^{(0)} \in \mathbb{R}^N$ for each member $i = 1, \dots, I$, normalized as a view relative to market-implied beliefs about excess returns. That is, the sum of the entries of $\boldsymbol{\mu}_i^{(0)}$ is zero, and the entries can be positive (overweight) or negative (underweight) relative to the market portfolio. The committee’s beliefs are stacked as the rows of the $I \times N$ matrix $\mathbf{M}^{(0)}$, so that the i -th row of $\mathbf{M}^{(0)}$ is member i ’s belief $\boldsymbol{\mu}_i^{(0)'}.$

The deliberations of the committee in deciding on their optimal portfolio are modeled in the style of DeGroot (1974). Each committee member’s inclination to be influenced by the members in the committee is represented by an I -dimensional vector $\mathbf{w}_i$. This vector, the entries of which sum to 1, represents the extent to which investor i ’s beliefs are influenced by the beliefs of each other member of the committee. A committee member who is supremely confident in their own beliefs about the market places all weight on their own entry in the vector $\mathbf{w}_i$ and zero weight on other committee members. A committee member who perceives that each committee member brings useful information to the decision making process will put positive weight on each member’s view. In each round of discussion, member i updates her belief by taking a weighted average of her own belief and those of her peers, with weights defined in $\mathbf{w}_i$. The weights that each member places on each other’s beliefs ($\mathbf{w}_i$ for each i)

are captured by a communication matrix $\mathbf{W}$, which is row-stochastic and where $\mathbf{W}_{ij}$ captures the weight that member i places on member j 's beliefs. The matrix $\mathbf{W}$ encodes the aggregate structure of influence within the committee.

In a one-shot DeGroot updating process, committee members update their beliefs by taking a weighted average of their own beliefs and those of their peers. Formally, the updated beliefs after one round of communication are given by:

$$\mathbf{M}^{(1)} = \mathbf{W}\mathbf{M}^{(0)}. \quad (1)$$

This process can be iterated, leading to beliefs after k rounds of discussion given by:

$$\mathbf{M}^{(k)} = \mathbf{W}^k \mathbf{M}^{(0)}. \quad (2)$$

Following the committee discussion, a portfolio manager aggregates the final beliefs of the committee members into a single *effective view*, $\boldsymbol{\mu}^*$ —the committee's effective beliefs, the object from which the portfolio is built. The term *consensus* is reserved throughout for the distinct event that deliberation itself carries all members to a common belief (Section 4.5); the effective view exists whether or not the committee agrees, and when consensus obtains the two coincide. In this model, the portfolio manager is represented by an aggregator function $f : \mathbb{R}^{I \times N} \rightarrow \mathbb{R}^N$, such that:

$$\boldsymbol{\mu}^* = f(\mathbf{M}^{(k)}). \quad (3)$$

The portfolio manager's method of aggregating committee beliefs, represented by the function f , can reflect different decision rules, such as taking the average belief across members, selecting the belief of a particular member (e.g., the chair), or applying a more complex weighting scheme based on expertise or seniority. Because each member's belief is a zero-sum view relative to market-implied values, the aggregated vector $\boldsymbol{\mu}^*$ is itself zero-sum: it is the committee's effective *view*, the active deviation from market-implied expected returns that the institution's deliberation produces.

The institution holds the benchmark portfolio $\boldsymbol{\theta}_b$ —the capitalization-weighted market—as its strategic core, and the committee's deliberation determines an *active tilt* away from this benchmark. Throughout, a subscript b denotes a quantity associated with the benchmark. To make this precise, treat the benchmark as the portfolio the institution would hold absent any view. By reverse optimization, $\boldsymbol{\theta}_b$ is mean-variance optimal at risk aversion γ for the equilibrium (market-implied) expected excess returns

$$\boldsymbol{\mu}_b = \gamma \boldsymbol{\Sigma} \boldsymbol{\theta}_b, \quad (4)$$

the standard reverse-optimized prior of [Black and Litterman \(1992\)](#). The committee's effective view $\boldsymbol{\mu}^*$ is added to this prior, so the institution's effective belief about *total* expected excess returns is $\boldsymbol{\mu}_b + \boldsymbol{\mu}^*$. The investor solves the mean-variance problem

$$\begin{aligned}\boldsymbol{\theta}_p &= \arg \max_{\boldsymbol{\theta}} \mathbb{E}[R_p] - \frac{\gamma}{2} \text{Var}[R_p] \\ &= \arg \max_{\boldsymbol{\theta}} \boldsymbol{\theta}'(\boldsymbol{\mu}_b + \boldsymbol{\mu}^*) - \frac{\gamma}{2} \boldsymbol{\theta}'\boldsymbol{\Sigma}\boldsymbol{\theta},\end{aligned}\tag{5}$$

where γ is the risk aversion parameter. The solution decomposes into the benchmark plus an active tilt driven entirely by the committee's view:

$$\boldsymbol{\theta}_p = \frac{1}{\gamma} \boldsymbol{\Sigma}^{-1}(\boldsymbol{\mu}_b + \boldsymbol{\mu}^*) = \boldsymbol{\theta}_b + \underbrace{\frac{1}{\gamma} \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}^*}_{\text{active tilt } \boldsymbol{\theta}_a}.\tag{6}$$

When the committee holds no view ($\boldsymbol{\mu}^* = \mathbf{0}$) the institution holds exactly the benchmark, since the deliberation enters only through the tilt $\boldsymbol{\theta}_a = \frac{1}{\gamma} \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}^*$. This framework makes it possible to analyze how the structure of communication within the committee (captured by $\mathbf{W}$) and the choice of aggregator function f influence the institution's portfolio $\boldsymbol{\theta}_p$ and its performance relative to the benchmark.

The committee dynamics shape the institution's portfolio and therefore its performance. To make this precise, the standard performance metrics can be written as functions of the committee's belief but *evaluated against the true distribution of returns*, which need not coincide with that belief. Let $\bar{\boldsymbol{\mu}} \equiv \boldsymbol{\mu}_b + \bar{\boldsymbol{\mu}}^*$ denote the true total expected excess returns, where $\bar{\boldsymbol{\mu}}^*$ is the true optimal view relative to the market-implied prior. Crucially, I do not assume that the committee is correct: the portfolio is built on the committee's belief $\boldsymbol{\mu}_p \equiv \boldsymbol{\mu}_b + \boldsymbol{\mu}^*$, through $\boldsymbol{\theta}_p = \frac{1}{\gamma} \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}_p$, but its performance is measured under the true returns $\bar{\boldsymbol{\mu}}$. The two coincide only when the committee happens to be right ($\boldsymbol{\mu}^* = \bar{\boldsymbol{\mu}}^*$). Under rational expectations the market itself is right and $\bar{\boldsymbol{\mu}}^* = \mathbf{0}$, an important benchmark case I return to but do not impose *ex ante*. The risk premium of the portfolio is its true expected excess return,

$$\text{Risk Premium} = \boldsymbol{\theta}_p' \bar{\boldsymbol{\mu}} = \frac{1}{\gamma} \boldsymbol{\mu}_p' \boldsymbol{\Sigma}^{-1} \bar{\boldsymbol{\mu}}.\tag{7}$$

The portfolio's variance is given by

$$\text{Variance} = \boldsymbol{\theta}_p' \boldsymbol{\Sigma} \boldsymbol{\theta}_p = \frac{1}{\gamma^2} \boldsymbol{\mu}_p' \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}_p.\tag{8}$$

The Sharpe ratio of the portfolio is

$$\text{Sharpe Ratio} = \frac{\text{Risk Premium}}{\sqrt{\text{Variance}}} = \frac{\boldsymbol{\mu}'_p \boldsymbol{\Sigma}^{-1} \bar{\boldsymbol{\mu}}}{\sqrt{\boldsymbol{\mu}'_p \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}_p}}. \quad (9)$$

Because the institution holds the benchmark plus the active tilt, these three metrics describe the *total* portfolio and are dominated by the benchmark. The metrics that isolate the committee's contribution are those measured relative to the benchmark. The tracking error is the risk of the active tilt $\boldsymbol{\theta}_a = \boldsymbol{\theta}_p - \boldsymbol{\theta}_b$,

$$\text{Tracking Error} = \sqrt{(\boldsymbol{\theta}_p - \boldsymbol{\theta}_b)' \boldsymbol{\Sigma} (\boldsymbol{\theta}_p - \boldsymbol{\theta}_b)} = \sqrt{\boldsymbol{\theta}'_a \boldsymbol{\Sigma} \boldsymbol{\theta}_a} = \frac{1}{\gamma} \sqrt{\boldsymbol{\mu}^{*'} \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}^*}, \quad (10)$$

which depends only on the committee's view $\boldsymbol{\mu}^*$ and not on the benchmark directly. The information ratio scales the portfolio's active return by this tracking error,

$$\text{Information Ratio} = \frac{\boldsymbol{\theta}'_p \bar{\boldsymbol{\mu}} - \boldsymbol{\theta}'_b \bar{\boldsymbol{\mu}}}{\sqrt{(\boldsymbol{\theta}_p - \boldsymbol{\theta}_b)' \boldsymbol{\Sigma} (\boldsymbol{\theta}_p - \boldsymbol{\theta}_b)}} = \frac{\boldsymbol{\theta}'_a \bar{\boldsymbol{\mu}}}{\text{Tracking Error}}. \quad (11)$$

Finally, Jensen's alpha measures the portfolio's expected excess return net of its systematic exposure to the benchmark. The portfolio's beta with respect to the benchmark is

$$\beta = \frac{\boldsymbol{\theta}'_p \boldsymbol{\Sigma} \boldsymbol{\theta}_b}{\boldsymbol{\theta}'_b \boldsymbol{\Sigma} \boldsymbol{\theta}_b} = 1 + \frac{\boldsymbol{\mu}^{*'} \boldsymbol{\theta}_b}{\gamma \sigma_b^2}, \quad \sigma_b^2 \equiv \boldsymbol{\theta}'_b \boldsymbol{\Sigma} \boldsymbol{\theta}_b. \quad (12)$$

Because the institution starts with the benchmark, beta is centered around one but departs from it to the extent that the active tilt is aligned with the benchmark. When the committee's zero-sum view and the benchmark weights are orthogonal, the overlap $\boldsymbol{\mu}^{*'} \boldsymbol{\theta}_b$ vanishes and the second term in this expression is zero.

Jensen's alpha is the difference between the portfolio's risk premium and the beta-scaled risk premium of the benchmark,

$$\alpha = \boldsymbol{\theta}'_p \bar{\boldsymbol{\mu}} - \beta \boldsymbol{\theta}'_b \bar{\boldsymbol{\mu}}. \quad (13)$$

Since $\beta \approx 1$, alpha is essentially the active risk premium $\boldsymbol{\theta}'_a \bar{\boldsymbol{\mu}}$ that the tilt earns *under the true returns*, beyond benchmark exposure—the value the committee's view actually adds. Alpha is zero when the committee holds no view ($\boldsymbol{\mu}^* = \mathbf{0}$) and, more pointedly, whenever the committee's view is uninformative about the truth. Under rational expectations ($\bar{\boldsymbol{\mu}}^* = \mathbf{0}$) the true alpha is exactly zero for *any* view the committee might hold, because tilting away from a correct market adds risk without adding return. This is a property of the beta adjustment, not a claim that the tilt is harmless. Because the benchmark is mean-variance efficient under

$\bar{\mu} = \mu_b$, the security market line holds exactly, so every portfolio’s expected return equals its beta times the benchmark premium and its alpha is zero. The spurious tilt nonetheless adds tracking error that earns nothing, and so lowers the portfolio’s *Sharpe ratio* strictly below the efficient benchmark’s. Under rational expectations the cost of a mistaken view is borne in risk-adjusted return, not in alpha.

3.1 Performance as a function of committee primitives

The metrics above depend on the committee only through the effective view μ^* —which the fixed prior μ_b shifts to the total belief $\mu_p = \mu_b + \mu^*$ —together with the market primitives μ_b , Σ , and the true view $\bar{\mu}^*$. Composing the belief-updating process $M^{(k)} = W^k M^{(0)}$ with the aggregator $\mu^* = f(M^{(k)})$ expresses this view—and hence each performance metric—directly in terms of the committee’s primitives: the initial beliefs $M^{(0)}$, the influence matrix W , the number of discussion rounds k , and the aggregation rule f . Write the effective view after k rounds as

$$\mu_k^* = f(W^k M^{(0)}), \quad (14)$$

where $f : \mathbb{R}^{I \times N} \rightarrow \mathbb{R}^N$ is left as a general, possibly nonlinear, map.³ The corresponding portfolio is the benchmark plus the active tilt,

$$\theta_{p,k} = \theta_b + \frac{1}{\gamma} \Sigma^{-1} \mu_k^* = \theta_b + \frac{1}{\gamma} \Sigma^{-1} f(W^k M^{(0)}). \quad (15)$$

These metrics are clearest in the geometry that the precision matrix Σ^{-1} places on beliefs. For any two belief vectors $\mathbf{x}$ and $\mathbf{y}$, write

$$\langle \mathbf{x}, \mathbf{y} \rangle \equiv \mathbf{x}' \Sigma^{-1} \mathbf{y}, \quad \|\mathbf{x}\| \equiv \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle}, \quad (16)$$

the inner product and norm in which mean–variance distances are measured. Consider three important vectors written in this notation: the benchmark’s market-implied belief $\mu_b = \gamma \Sigma \theta_b$, so that $\theta_b' \mathbf{x} = \frac{1}{\gamma} \langle \mu_b, \mathbf{x} \rangle$; the committee’s view μ_k^* ; and the true view $\bar{\mu}^*$. The portfolio is built from the committee’s total belief $\mu_{p,k} = \mu_b + \mu_k^*$ and judged against the truth $\bar{\mu} = \mu_b + \bar{\mu}^*$,

³Linear pooling rules such as equal weighting or chair-take-all are special cases but are not assumed for now.

and each performance metric is a simple combination of these vectors:

$$\begin{aligned}
\text{Risk Premium}(k) &= \frac{1}{\gamma} \langle \boldsymbol{\mu}_{p,k}, \bar{\boldsymbol{\mu}} \rangle, & \text{Tracking Error}(k) &= \frac{1}{\gamma} \|\boldsymbol{\mu}_k^*\|, \\
\text{Variance}(k) &= \frac{1}{\gamma^2} \|\boldsymbol{\mu}_{p,k}\|^2, & \text{Information Ratio}(k) &= \frac{\langle \boldsymbol{\mu}_k^*, \bar{\boldsymbol{\mu}} \rangle}{\|\boldsymbol{\mu}_k^*\|}, \\
\text{Sharpe Ratio}(k) &= \frac{\langle \boldsymbol{\mu}_{p,k}, \bar{\boldsymbol{\mu}} \rangle}{\|\boldsymbol{\mu}_{p,k}\|}.
\end{aligned} \tag{17}$$

The total-portfolio metrics in the left column—risk premium, variance, and the Sharpe ratio—depend on the committee’s full belief $\boldsymbol{\mu}_{p,k}$ and are dominated by the benchmark component $\boldsymbol{\mu}_b$. The benchmark-relative metrics in the right column strip the benchmark away and isolate the committee: tracking error is simply the committee’s view measured in this norm, and the information ratio is the alignment of that view with the truth.

What the committee contributes is therefore controlled by two attributes of its view $\boldsymbol{\mu}_k^*$: its *magnitude* $\|\boldsymbol{\mu}_k^*\|$, the conviction with which the committee bets and hence the active risk it takes, and its *direction relative to the truth* $\bar{\boldsymbol{\mu}}^*$. Both are summarized by the gap between what the committee believes and what is true,

$$\boldsymbol{\delta}_k \equiv \boldsymbol{\mu}_k^* - \bar{\boldsymbol{\mu}}^*, \tag{18}$$

the committee’s *estimation error*. Its squared norm decomposes as

$$\|\boldsymbol{\delta}_k\|^2 = \underbrace{\|\boldsymbol{\mu}_k^*\|^2}_{\text{conviction}} - 2 \underbrace{\langle \boldsymbol{\mu}_k^*, \bar{\boldsymbol{\mu}}^* \rangle}_{\text{accuracy}} + \|\bar{\boldsymbol{\mu}}^*\|^2, \tag{19}$$

which ties the committee’s conviction to its accuracy. For a fixed truth $\bar{\boldsymbol{\mu}}^*$, the committee adds value only by shrinking the gap $\boldsymbol{\delta}_k$, and a large but mistaken view is worse than none at all. This gap between the effective view and the true view is the object the rest of the paper traces back to the committee’s communication structure.

It is convenient to collect the scalar functionals of the effective view $\boldsymbol{\mu}_k^*$ through which these metrics act—the building blocks that resolve the gap $\boldsymbol{\delta}_k$ into conviction, benchmark overlap, and accuracy. Three depend on the committee’s view,

$$q(k) = \boldsymbol{\mu}_k^{*\prime} \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}_k^*, \quad r(k) = \frac{1}{\gamma} \boldsymbol{\mu}_b' \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}_k^*, \quad s(k) = \boldsymbol{\mu}_k^{*\prime} \boldsymbol{\Sigma}^{-1} \bar{\boldsymbol{\mu}}^*, \tag{20}$$

and two are fixed primitives of the market,

$$\sigma_b^2 = \frac{1}{\gamma^2} \boldsymbol{\mu}_b' \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}_b, \quad \rho = \frac{1}{\gamma} \boldsymbol{\mu}_b' \boldsymbol{\Sigma}^{-1} \bar{\boldsymbol{\mu}}^*. \tag{21}$$

Here $q(k)$ is the squared active belief, which determines the portfolio's active risk. The expression

$$\begin{aligned} r(k) &= \frac{1}{\gamma} \boldsymbol{\mu}_b' \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}_k^* \\ &= \frac{1}{\gamma} \cos(\boldsymbol{\mu}_b, \boldsymbol{\mu}_k^*) \|\boldsymbol{\mu}_b\| \|\boldsymbol{\mu}_k^*\| \end{aligned} \quad (22)$$

gives the degree of collinearity of the committee's view with the benchmark. For given magnitudes of the benchmark and the committee's view, $r(k)$ is larger when the two are more aligned. This alignment then determines the portfolio's beta.

The term

$$\begin{aligned} s(k) &= \boldsymbol{\mu}_k^* \boldsymbol{\Sigma}^{-1} \bar{\boldsymbol{\mu}}^* = \langle \boldsymbol{\mu}_k^*, \bar{\boldsymbol{\mu}}^* \rangle \\ &= \cos(\boldsymbol{\mu}_k^*, \bar{\boldsymbol{\mu}}^*) \|\boldsymbol{\mu}_k^*\| \|\bar{\boldsymbol{\mu}}^*\| \end{aligned} \quad (23)$$

is the alignment of the committee's view with the *true* view $\bar{\boldsymbol{\mu}}^*$ —the accuracy term of the error decomposition (19), and the channel through which deliberation adds value. Holding the committee's conviction $\|\boldsymbol{\mu}_k^*\|$ fixed, $s(k)$ is largest—and the gap $\boldsymbol{\delta}_k$ smallest—when the committee's view points in the direction of the truth.

The constants σ_b^2 and ρ describe the benchmark and the truth and do not depend on the committee, while q , r , and s inherit their dependence on the primitives $(\boldsymbol{M}^{(0)}, \boldsymbol{W}, k, f)$ entirely through $\boldsymbol{\mu}_k^*$. Two further combinations summarize the portfolio: the squared total belief $Q(k) = \gamma^2 \sigma_b^2 + 2\gamma r(k) + q(k) = \boldsymbol{\mu}_{p,k}' \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}_{p,k}$, which governs its risk, and the cross term $\bar{Q}(k) = \gamma^2 \sigma_b^2 + \gamma \rho + \gamma r(k) + s(k) = \boldsymbol{\mu}_{p,k}' \boldsymbol{\Sigma}^{-1} \bar{\boldsymbol{\mu}}$ between the committee's total belief and the true returns, which governs its true expected return. Substituting these building blocks, each

metric of (17) becomes an explicit scalar function of the committee’s deliberation:

$$\text{Risk Premium}(k) = \frac{\overline{Q}(k)}{\gamma}, \quad (24)$$

$$\text{Variance}(k) = \frac{Q(k)}{\gamma^2}, \quad (25)$$

$$\text{Sharpe Ratio}(k) = \frac{\overline{Q}(k)}{\sqrt{Q(k)}}, \quad (26)$$

$$\text{Tracking Error}(k) = \frac{\sqrt{q(k)}}{\gamma}, \quad (27)$$

$$\text{Information Ratio}(k) = \frac{\gamma r(k) + s(k)}{\sqrt{q(k)}}, \quad (28)$$

$$\beta(k) = 1 + \frac{r(k)}{\gamma \sigma_b^2}, \quad (29)$$

$$\alpha(k) = \frac{1}{\gamma} \left(s(k) - \frac{r(k) \rho}{\sigma_b^2} \right). \quad (30)$$

This decomposition separates the committee’s *conviction*, the magnitude $\|\boldsymbol{\mu}_k^*\|^2 = q(k)$, from its *accuracy* $s(k)$ —the two terms of the gap $\boldsymbol{\delta}_k$ in (19). The portfolio’s risk—variance and tracking error—depends on the squared active belief $q(k)$ alone, so a committee that takes a large active position bears large active risk whether or not it is right. What the committee *earns* for that risk is governed instead by $s(k)$, the alignment of its view with the truth—equivalently, by how small it makes the gap $\boldsymbol{\delta}_k$ to the true view. This is clearest in Jensen’s alpha, $\alpha(k) = \frac{1}{\gamma} \left(s(k) - r(k) \rho / \sigma_b^2 \right)$, which is benchmark-neutral and depends on the committee’s view only through its accuracy $s(k)$ and the benchmark overlaps $r(k)$ and ρ : under rational expectations ($\bar{\boldsymbol{\mu}}^* = \mathbf{0}$, so $s(k) = \rho = 0$) it is exactly zero for every committee, however confident, because with a correct market active deviation buys risk but not return. The information ratio, $(\gamma r(k) + s(k)) / \sqrt{q(k)}$, adds to the accuracy term $s(k)$ a benchmark-overlap term $\gamma r(k)$ —the active return the tilt inherits from its correlation with the benchmark—and so need not vanish under rational expectations even though its skill component does. The total metrics—risk premium, variance, and Sharpe ratio—remain dominated by the fixed benchmark term $\gamma^2 \sigma_b^2$ and leave the committee only a secondary role. Communication enters each metric only through the single argument $\mathbf{W}^k \mathbf{M}^{(0)}$ passed to the aggregator f : the number of rounds k reshapes performance solely by transforming the matrix of initial beliefs into $\mathbf{W}^k \mathbf{M}^{(0)}$ before f maps it to an effective view. Consider two limiting cases. With no deliberation ($k = 0$), $\boldsymbol{\mu}_0^* = f(\mathbf{M}^{(0)})$ aggregates members’ priors directly.

As deliberation continues ($k \rightarrow \infty$), and provided the communication network is strongly connected and aperiodic, the iterated beliefs $\mathbf{W}^k \mathbf{M}^{(0)}$ converge. Then, if f is continuous, $\boldsymbol{\mu}_k^*$ converges to a limiting consensus belief in which each member's influence is governed by their eigenvector centrality in $\mathbf{W}$ (Golub and Jackson 2010), and the performance metrics converge to the corresponding limits.

3.2 Accuracy and the concentration of influence

The decomposition above isolates accuracy $s(k)$ as the channel through which deliberation adds value, but leaves it as an implicit function of the committee's primitives through the aggregator f . When f is a linear pooling rule—as are both aggregators used in the simulations, equal-weight pooling and chair-take-all—accuracy can be written explicitly in terms of $\mathbf{W}$, and its dependence on the committee's communication structure becomes transparent. Throughout this subsection let f be linear pooling with a fixed weight vector $\mathbf{a} \in \mathbb{R}^I$, $\mathbf{a} \geq \mathbf{0}$, $\mathbf{a}'\mathbf{1} = 1$, so that $f(\mathbf{M}) = \mathbf{M}'\mathbf{a}$; equal-weight pooling is $\mathbf{a} = \frac{1}{I}\mathbf{1}$ and a chair who decides is $\mathbf{a} = \mathbf{e}_c$, the c -th standard basis vector.

Proposition 1 (Accuracy as influence-weighted individual accuracy). *Under linear pooling with weights $\mathbf{a}$, define the effective influence vector after k rounds and the vector of individual accuracies by*

$$\mathbf{w}_k \equiv (\mathbf{W}')^k \mathbf{a}, \quad \mathbf{c} \equiv \mathbf{M}^{(0)} \boldsymbol{\Sigma}^{-1} \bar{\boldsymbol{\mu}}^*, \quad i.e. \quad c_i = \boldsymbol{\mu}_i^{(0)'} \boldsymbol{\Sigma}^{-1} \bar{\boldsymbol{\mu}}^*. \quad (31)$$

Then $\mathbf{w}_k$ is a probability vector and the committee's accuracy is the $\mathbf{w}_k$ -weighted average of its members' individual accuracies,

$$s(k) = \mathbf{a}' \mathbf{W}^k \mathbf{c} = \mathbf{w}_k' \mathbf{c} = \sum_{i=1}^I w_{k,i} c_i. \quad (32)$$

Consequently accuracy is bounded by the most and least accurate members,

$$\min_i c_i \leq s(k) \leq \max_i c_i, \quad (33)$$

and, when $\mathbf{W}$ is strongly connected and aperiodic, $\mathbf{w}_k \rightarrow \boldsymbol{\pi}$ and $s(k) \rightarrow \boldsymbol{\pi}' \mathbf{c}$ as $k \rightarrow \infty$, independently of the aggregator $\mathbf{a}$.

Proof. With $f(\mathbf{M}) = \mathbf{M}'\mathbf{a}$ the effective view after k rounds is $\boldsymbol{\mu}_k^* = (\mathbf{W}^k \mathbf{M}^{(0)})' \mathbf{a} = \mathbf{M}^{(0)'} (\mathbf{W}^k)' \mathbf{a} = \mathbf{M}^{(0)'} \mathbf{w}_k$. Substituting into the definition $s(k) = \boldsymbol{\mu}_k^{*'} \boldsymbol{\Sigma}^{-1} \bar{\boldsymbol{\mu}}^*$ in (23) gives $s(k) = \mathbf{w}_k' \mathbf{M}^{(0)} \boldsymbol{\Sigma}^{-1} \bar{\boldsymbol{\mu}}^* = \mathbf{w}_k' \mathbf{c}$, and $\mathbf{w}_k' \mathbf{c} = \mathbf{a}' \mathbf{W}^k \mathbf{c}$ by definition of $\mathbf{w}_k$. The vector $\mathbf{w}_k$ is

nonnegative because $\mathbf{W}^k \geq \mathbf{0}$ and $\mathbf{a} \geq \mathbf{0}$, and it sums to one because $\mathbf{1}'\mathbf{w}_k = \mathbf{1}'(\mathbf{W}')^k\mathbf{a} = (\mathbf{W}^k\mathbf{1})'\mathbf{a} = \mathbf{1}'\mathbf{a} = 1$, using $\mathbf{W}^k\mathbf{1} = \mathbf{1}$ (row-stochasticity). Equation (32) thus expresses $s(k)$ as a convex combination of the c_i , which yields (33). Finally, when $\mathbf{W}$ is strongly connected and aperiodic, $\mathbf{W}^k \rightarrow \mathbf{1}\pi'$, so $\mathbf{w}_k = (\mathbf{W}^k)'\mathbf{a} \rightarrow \pi\mathbf{1}'\mathbf{a} = \pi$ and $s(k) \rightarrow \pi'\mathbf{c}$, which does not involve $\mathbf{a}$. $\square$

The content of Proposition 1 is that the communication structure enters accuracy only by reweighting members: $\mathbf{W}$ redistributes influence among the c_i but cannot push accuracy beyond the best individual signal or below the worst. The effective influence vector $\mathbf{w}_k = (\mathbf{W}')^k\mathbf{a}$ is the aggregator's weights propagated backward through k rounds of communication; it interpolates between the raw aggregator $\mathbf{a}$ at $k = 0$ and the eigenvector-centrality weights π as $k \rightarrow \infty$, recovering the limiting consensus of Section 3.3 as the special case in which deliberation has run to completion.

This places a sharp limit on what committee dynamics can accomplish. Because $s(k)$ is a convex combination of the members' individual accuracies $c_i = \langle \mu_i^{(0)}, \bar{\mu}^* \rangle$, it can never exceed the accuracy of the committee's most accurate member nor fall below that of its least accurate: in this setup a committee is no better informed than its best-informed member and no worse informed than its worst-informed. Communication and aggregation only move the committee *within* this envelope—they reweight the views already in the room—and cannot manufacture insight that no member brought. In the extreme, a committee that places all effective influence on its best-informed member reaches the ceiling $\max_i c_i$, while one captured by its worst-informed member sinks to the floor $\min_i c_i$. The influence structure $\mathbf{W}$ and the aggregator f thus determine *where within the envelope* the committee lands, but the envelope itself is fixed by who sits on the committee.

Two distinct requirements for a good investment committee can be seen here. The first is *selection*: the ceiling $\max_i c_i$ is set by the ablest member's signal, so recruiting genuinely skilled members—those whose views align with the truth—raises the best the committee can possibly do, and no amount of deliberative process substitutes for it. The second is *attention*: a committee of able members still underperforms if its effective influence $\mathbf{w}_k$ rests on its weaker voices, so the structure must place weight on the members worth listening to. Choosing well-informed people and actually heeding them are separate things and both of them are important for committee performance. What the committee *can* do that no single member can is average away idiosyncratic error, lowering the *variance* of its accuracy rather than raising its ceiling—the subject of the results that follow.

The next result shows that when members hold unbiased signals of the true view, the communication structure has *no* effect on expected accuracy and acts purely on its variance, through a single scalar: the concentration of effective influence.

Corollary 1 (Concentration determines the risk of accuracy). *Suppose members' initial beliefs are unbiased signals of the true view, $\boldsymbol{\mu}_i^{(0)} = \bar{\boldsymbol{\mu}}^* + \boldsymbol{\varepsilon}_i$, with $\boldsymbol{\varepsilon}_i$ mean-zero and independent across members, and let $\eta_i \equiv \boldsymbol{\varepsilon}_i' \boldsymbol{\Sigma}^{-1} \bar{\boldsymbol{\mu}}^*$ have common variance $\text{Var}(\eta_i) = \tau^2$. Write $\kappa \equiv \bar{\boldsymbol{\mu}}^{*'} \boldsymbol{\Sigma}^{-1} \bar{\boldsymbol{\mu}}^*$ and let $H(\mathbf{w}_k) \equiv \sum_i w_{k,i}^2$ denote the Herfindahl concentration of effective influence, with $N_{\text{eff}}(\mathbf{w}_k) \equiv 1/H(\mathbf{w}_k)$ the effective number of members. Then for every $\mathbf{W}$, every k , and every linear aggregator $\mathbf{a}$,*

$$\mathbb{E}[s(k)] = \kappa, \quad \text{Var}[s(k)] = \tau^2 H(\mathbf{w}_k) = \frac{\tau^2}{N_{\text{eff}}(\mathbf{w}_k)}. \quad (34)$$

Expected accuracy is therefore invariant to the committee's communication structure, while the variance of accuracy is increasing in the concentration of effective influence: it is minimized, at τ^2/I , by equal effective influence $\mathbf{w}_k = \frac{1}{I} \mathbf{1}$ and maximized, at τ^2 , when all effective influence rests on one member.

Proof. Unbiasedness gives $c_i = (\bar{\boldsymbol{\mu}}^* + \boldsymbol{\varepsilon}_i)' \boldsymbol{\Sigma}^{-1} \bar{\boldsymbol{\mu}}^* = \kappa + \eta_i$, so by (32) and $\mathbf{w}_k' \mathbf{1} = 1$,

$$s(k) = \mathbf{w}_k' \mathbf{c} = \kappa (\mathbf{w}_k' \mathbf{1}) + \sum_i w_{k,i} \eta_i = \kappa + \sum_i w_{k,i} \eta_i.$$

Taking expectations and using $\mathbb{E}[\eta_i] = 0$ gives $\mathbb{E}[s(k)] = \kappa$. Independence of the η_i gives $\text{Var}[s(k)] = \sum_i w_{k,i}^2 \text{Var}(\eta_i) = \tau^2 \sum_i w_{k,i}^2 = \tau^2 H(\mathbf{w}_k)$. The Herfindahl index is Schur-convex on the simplex, hence minimized at the uniform vector, where $H = 1/I$, and maximized at a vertex, where $H = 1$. $\square$

The same scalar controls how far the committee's view sits from the truth, tying [Corollary 1](#) to the gap $\boldsymbol{\delta}_k$ that [Section 3.1](#) places at the center of the analysis.

Corollary 2 (Concentration increases the expected gap to the truth). *Under the assumptions of [Corollary 1](#), with members' signal errors identically distributed across i , each committee's view is unbiased, $\mathbb{E}[\boldsymbol{\mu}_k^*] = \bar{\boldsymbol{\mu}}^*$, so the estimation error $\boldsymbol{\delta}_k = \boldsymbol{\mu}_k^* - \bar{\boldsymbol{\mu}}^*$ of (18) has mean zero, and its expected squared size in the precision norm is*

$$\mathbb{E} \|\boldsymbol{\delta}_k\|^2 = \sigma_\varepsilon^2 H(\mathbf{w}_k) = \frac{\sigma_\varepsilon^2}{N_{\text{eff}}(\mathbf{w}_k)}, \quad \sigma_\varepsilon^2 \equiv \mathbb{E}[\boldsymbol{\varepsilon}_i' \boldsymbol{\Sigma}^{-1} \boldsymbol{\varepsilon}_i], \quad (35)$$

the per-member signal error scaled by the same concentration $H(\mathbf{w}_k)$ that governs $\text{Var}[s(k)]$ in (34).

Proof. Unbiased signals and $\mathbf{w}_k' \mathbf{1} = 1$ give $\boldsymbol{\mu}_k^* = \mathbf{M}^{(0)'} \mathbf{w}_k = \sum_i w_{k,i} (\bar{\boldsymbol{\mu}}^* + \boldsymbol{\varepsilon}_i) = \bar{\boldsymbol{\mu}}^* + \sum_i w_{k,i} \boldsymbol{\varepsilon}_i$,

so $\mathbb{E}[\boldsymbol{\mu}_k^*] = \bar{\boldsymbol{\mu}}^*$ and $\boldsymbol{\delta}_k = \sum_i w_{k,i} \boldsymbol{\epsilon}_i$. Hence, using independence across members,

$$\mathbb{E}\|\boldsymbol{\delta}_k\|^2 = \sum_{i,j} w_{k,i} w_{k,j} \mathbb{E}[\boldsymbol{\epsilon}_i' \boldsymbol{\Sigma}^{-1} \boldsymbol{\epsilon}_j] = \sum_i w_{k,i}^2 \mathbb{E}[\boldsymbol{\epsilon}_i' \boldsymbol{\Sigma}^{-1} \boldsymbol{\epsilon}_i] = \sigma_\epsilon^2 H(\mathbf{w}_k). \quad \square$$

Through the decomposition (19), concentration thus acts on two faces of the same per-member noise: it inflates the committee’s expected gap $\mathbb{E}\|\boldsymbol{\delta}_k\|^2$ to the truth and, equivalently, the variance of its realized accuracy $s(k)$, both by the factor $H(\mathbf{w}_k) = 1/N_{\text{eff}}(\mathbf{w}_k)$. A diffuse committee that pools all I signals shrinks both by a factor of I —the wisdom of crowds—whereas a committee that relies on one voice inherits the full per-member error σ_ϵ^2 .

Corollary 1 is the analytical core of the performance results that follow. The Schur-convexity of $H(\mathbf{w}_k)$ implies that any change in $\mathbf{W}$ that concentrates effective influence—a transfer of weight toward an already-influential member—weakly raises the variance of accuracy while leaving its mean unchanged. Concentration thus imports a dominant member’s idiosyncratic estimation error as risk that earns nothing on average, exactly the mechanism behind the higher tracking error and lower information ratios of the concentrated committees in Section 4.2. The reciprocal $N_{\text{eff}}(\mathbf{w}_k) = 1/H(\mathbf{w}_k)$ is the “effective number of members whose signals enter $\boldsymbol{\mu}^*$ ” that governs the dispersion of realized performance: a diffuse committee pools many signals and concentrates the distribution of accuracy near κ , whereas a committee that rides one voice has $N_{\text{eff}} \approx 1$ and the full per-member variance τ^2 .

The result also pins down the effect of *deliberation*. Since $\mathbf{w}_k = (\mathbf{W}')^k \mathbf{a}$ drifts from the aggregator $\mathbf{a}$ toward eigenvector centrality $\boldsymbol{\pi}$, additional rounds raise or lower the variance of accuracy to the extent that $\mathbf{w}_k$ moves away from or toward the uniform vector. Under equal-weight pooling $\mathbf{a} = \frac{1}{I} \mathbf{1}$ sits at the minimum-concentration point, so deliberation can only (weakly) increase $H(\mathbf{w}_k)$ and erode performance—and leaves it exactly flat when $\boldsymbol{\pi} = \frac{1}{I} \mathbf{1}$. Under a chair rule $\mathbf{a} = \mathbf{e}_c$ sits at the maximum, so deliberation lowers $H(\mathbf{w}_k)$ as the chair’s effective weight diffuses toward her colleagues, improving accuracy. This sign reversal in the value of deliberation, documented in the simulations of Section 4.2, is thus a direct consequence of where each aggregator starts on the concentration scale.

The efficient committee. Corollaries 1 and 2 hold the members symmetric: all signals carry the same precision, so equal influence is best and any concentration decreases portfolio performance. When members differ in skill this is no longer so. Let member i ’s signal have precision-norm error variance $\sigma_i^2 \equiv \mathbb{E}[\boldsymbol{\epsilon}_i' \boldsymbol{\Sigma}^{-1} \boldsymbol{\epsilon}_i]$ and precision $p_i \equiv 1/\sigma_i^2$, with the $\boldsymbol{\epsilon}_i$ independent across members but no longer identically distributed.

Proposition 2 (The efficient committee). *Among all committees whose effective influence*

$\mathbf{w}_k$ lies on the simplex, the expected squared estimation error

$$\mathbb{E}\|\boldsymbol{\delta}_k\|^2 = \sum_{i=1}^I w_{k,i}^2 \sigma_i^2 \quad (36)$$

is minimized by precision (inverse-variance) weighting,

$$w_i^* = \frac{p_i}{\sum_j p_j}, \quad (37)$$

at which the committee's precision is the sum of its members' and its error the reciprocal of that sum,

$$\left(\mathbb{E}\|\boldsymbol{\delta}\|_{\min}^2\right)^{-1} = \sum_{i=1}^I p_i. \quad (38)$$

For any influence vector $\mathbf{w}_k$, the excess error above this floor is a precision-weighted measure of misallocated attention,

$$\mathbb{E}\|\boldsymbol{\delta}_k\|^2 = \frac{1}{\sum_j p_j} + \sum_{i=1}^I \sigma_i^2 (w_{k,i} - w_i^*)^2. \quad (39)$$

Because all unbiased committees earn the same expected accuracy $\mathbb{E}[s(k)] = \kappa$ and bear active risk $\mathbb{E}[q(k)] = \|\bar{\boldsymbol{\mu}}^*\|^2 + \mathbb{E}\|\boldsymbol{\delta}_k\|^2$, the precision-weighted committee attains the common expected active return at the lowest expected tracking error: it is the minimum-active-risk unbiased committee.

Proof. Independence and unbiasedness give $\boldsymbol{\delta}_k = \sum_i w_{k,i} \boldsymbol{\varepsilon}_i$, so as in [Corollary 2](#) $\mathbb{E}\|\boldsymbol{\delta}_k\|^2 = \sum_{i,j} w_{k,i} w_{k,j} \mathbb{E}[\boldsymbol{\varepsilon}_i' \boldsymbol{\Sigma}^{-1} \boldsymbol{\varepsilon}_j] = \sum_i w_{k,i}^2 \sigma_i^2$. By the Cauchy–Schwarz inequality,

$$1 = \left(\sum_i w_{k,i}\right)^2 = \left(\sum_i (w_{k,i} \sigma_i) \frac{1}{\sigma_i}\right)^2 \leq \left(\sum_i w_{k,i}^2 \sigma_i^2\right) \left(\sum_i \frac{1}{\sigma_i^2}\right),$$

so $\sum_i w_{k,i}^2 \sigma_i^2 \geq 1/\sum_i p_i$, with equality iff $w_{k,i} \sigma_i \propto 1/\sigma_i$, i.e. $w_{k,i} \propto p_i$; this gives (37)–(38), and the weights are nonnegative and hence attainable. For (39), expand $\sum_i \sigma_i^2 (w_{k,i} - w_i^*)^2 = \sum_i \sigma_i^2 w_{k,i}^2 - 2 \sum_i \sigma_i^2 w_i^* w_{k,i} + \sum_i \sigma_i^2 (w_i^*)^2$ and use $\sigma_i^2 w_i^* = 1/\sum_j p_j$, so the middle term equals $-2/\sum_j p_j$ and the last $1/\sum_j p_j$, leaving $\sum_i \sigma_i^2 w_{k,i}^2 - 1/\sum_j p_j = \mathbb{E}\|\boldsymbol{\delta}_k\|^2 - 1/\sum_j p_j$. The final claim uses $\mathbb{E}[q(k)] = \|\bar{\boldsymbol{\mu}}^*\|^2 + \mathbb{E}\|\boldsymbol{\delta}_k\|^2$, immediate from $\boldsymbol{\mu}_k^* = \bar{\boldsymbol{\mu}}^* + \boldsymbol{\delta}_k$ and $\mathbb{E}[\boldsymbol{\delta}_k] = \mathbf{0}$. $\square$

Three implications follow. First, the result gives a situation where there is a benefit to concentration. In this case, what is costly is not concentration as such but concentration *misaligned* with skill. The penalty borne by any committee is exactly the precision-weighted distance (39) between the attention it pays and the attention members deserve, so a committee

that concentrates on its most precise member is efficient, not wasteful, while one that spreads influence equally over members of unequal skill is itself misallocated. The blanket superiority of equal weighting in [Corollary 1](#) is the special case $\sigma_i^2 \equiv \sigma^2$, where skill is uniform and $w_i^* = 1/I$; the diffuse nonpolar committee is the efficient committee only when members are interchangeable.

Second, the result formalizes the *selection-versus-attention* distinction of [Proposition 1](#). Selection sets the precisions p_i and, through the additivity (38), the ceiling on what any committee can achieve; attention sets $\mathbf{w}_k$ and, through (39), how close the committee comes to that ceiling. A committee of able members (large $\sum_i p_i$) still squanders its talent if its influence rests on the wrong voices, and a perfectly aligned committee of weak members cannot manufacture precision it does not have. The two levers are complementary, and either without the other reduces the committee’s effectiveness.

Third, precision weighting (37) is exactly the aggregate a Bayesian would form from the I independent signals, given a diffuse prior and signal-error covariances proportional to Σ . The efficient committee is thus the *Bayesian* committee. Naive DeGroot deliberation reaches it only when effective influence coincides with skill, $\mathbf{w}_k = \mathbf{w}^*$ —in the limit, only when eigenvector centrality $\boldsymbol{\pi}$ equals the precision shares. In this case, the value of an extra round of discussion can be either positive or negative depending on the movement it induces in effective influence—toward skill or away from it—which is the deeper reason the deliberation effect reverses between pooling and a chair rule in [Section 4.2](#).

The efficient committee is, moreover, *reachable* by either instrument the institution controls. Because effective influence is $\mathbf{w}_k = (\mathbf{W}')^k \mathbf{a}$, a committee can implement precision weighting through its governance rule if a chair or portfolio manager sets the aggregator to $\mathbf{a} = \mathbf{w}^*$. It can also reach this through its communication structure. Any influence matrix $\mathbf{W}$ whose eigenvector centrality is $\boldsymbol{\pi} = \mathbf{w}^*$ attains it in the limit, and then for *every* unanimity-respecting aggregator, since deliberation has already done the weighting. What neither lever can overcome is poor selection: by the additivity (38) the error floor $1/\sum_i p_i$ falls only by bringing more precision into the room, never by rearranging the influence over a fixed set of members.

The results so far assume every member’s signal is unbiased, so the committee errs only through noise that averaging can cancel. A committee that defers to a dominant chair, however, inherits whatever *systematic* error she carries. Suppose the chair, member c , holds a biased view

$$\boldsymbol{\mu}_c^{(0)} = \bar{\boldsymbol{\mu}}^* + \mathbf{b} + \boldsymbol{\varepsilon}_c, \quad (40)$$

where $\mathbf{b}$ is a fixed bias—the difference between the chair’s systematic view and the truth. The remaining members are unbiased, $\boldsymbol{\mu}_i^{(0)} = \bar{\boldsymbol{\mu}}^* + \boldsymbol{\varepsilon}_i$, and all $\boldsymbol{\varepsilon}_i$ are mean-zero, independent, and

identically distributed with $\mathbb{E}\|\boldsymbol{\varepsilon}_i\|^2 = \sigma_\varepsilon^2$. Let $\omega \equiv w_{k,c}$ denote the chair's effective influence, which is large in exactly the two cases of interest: when the committee is heavily swayed by her, so her eigenvector centrality π_c is high and $\omega \rightarrow \pi_c$ as deliberation lengthens, or when the aggregator defers to her, as under chair-take-all ($\mathbf{a} = \mathbf{e}_c$), where $\omega = 1$ at $k = 0$.

Proposition 3 (A biased chair). *Under (40), the effective view inherits a fraction ω of the chair's bias, $\mathbb{E}[\boldsymbol{\mu}_k^*] = \bar{\boldsymbol{\mu}}^* + \omega \mathbf{b}$, and the committee's expected performance is*

$$\mathbb{E}\|\boldsymbol{\delta}_k\|^2 = \underbrace{\omega^2 \|\mathbf{b}\|^2}_{\text{bias}} + \underbrace{\sigma_\varepsilon^2 H(\mathbf{w}_k)}_{\text{variance}}, \quad (41)$$

$$\mathbb{E}[s(k)] = \kappa + \omega \langle \mathbf{b}, \bar{\boldsymbol{\mu}}^* \rangle, \quad (42)$$

$$\mathbb{E}[q(k)] = \|\bar{\boldsymbol{\mu}}^* + \omega \mathbf{b}\|^2 + \sigma_\varepsilon^2 H(\mathbf{w}_k), \quad (43)$$

$$\mathbb{E}[\alpha(k)] = \frac{1}{\gamma} \left(\|\bar{\boldsymbol{\mu}}_\perp^*\|^2 + \omega \langle \mathbf{b}, \bar{\boldsymbol{\mu}}_\perp^* \rangle \right), \quad (44)$$

where $\bar{\boldsymbol{\mu}}_\perp^* \equiv \bar{\boldsymbol{\mu}}^* - (\langle \bar{\boldsymbol{\mu}}^*, \boldsymbol{\mu}_b \rangle / \|\boldsymbol{\mu}_b\|^2) \boldsymbol{\mu}_b$ is the component of the true view orthogonal to the benchmark in the precision metric. Setting $\mathbf{b} = \mathbf{0}$ recovers [Corollaries 1 and 2](#).

Proof. Writing $\boldsymbol{\zeta} \equiv \sum_i w_{k,i} \boldsymbol{\varepsilon}_i$, the effective view is $\boldsymbol{\mu}_k^* = \sum_i w_{k,i} \boldsymbol{\mu}_i^{(0)} = \bar{\boldsymbol{\mu}}^* + \omega \mathbf{b} + \boldsymbol{\zeta}$, since the bias enters only through the chair's term, of weight $\omega = w_{k,c}$, and $\mathbf{w}'_k \mathbf{1} = 1$. Hence $\mathbb{E}[\boldsymbol{\mu}_k^*] = \bar{\boldsymbol{\mu}}^* + \omega \mathbf{b}$ and $\boldsymbol{\delta}_k = \omega \mathbf{b} + \boldsymbol{\zeta}$, with $\mathbb{E}[\boldsymbol{\zeta}] = \mathbf{0}$ and $\mathbb{E}\|\boldsymbol{\zeta}\|^2 = \sum_i w_{k,i}^2 \sigma_\varepsilon^2 = \sigma_\varepsilon^2 H(\mathbf{w}_k)$ by independence. Expanding $\|\boldsymbol{\delta}_k\|^2$, $s(k) = \langle \boldsymbol{\mu}_k^*, \bar{\boldsymbol{\mu}}^* \rangle$, and $q(k) = \|\boldsymbol{\mu}_k^*\|^2$ and taking expectations—the terms linear in $\boldsymbol{\zeta}$ vanishing—gives (41)–(43). For (44), substituting the precision-form building blocks of (20) into $\alpha(k) = \frac{1}{\gamma} (s(k) - r(k) \rho / \sigma_b^2)$ gives $\alpha(k) = \frac{1}{\gamma} \langle \boldsymbol{\mu}_k^*, \bar{\boldsymbol{\mu}}_\perp^* \rangle$; taking expectations and using $\langle \bar{\boldsymbol{\mu}}^*, \bar{\boldsymbol{\mu}}_\perp^* \rangle = \|\bar{\boldsymbol{\mu}}_\perp^*\|^2$ yields the result. $\square$

The two benchmark-relative metrics follow explicitly. They are nonlinear in $\boldsymbol{\mu}_k^*$, so the tracking error is reported through its expected square and the information ratio through its value at the representative view $\mathbb{E}[\boldsymbol{\mu}_k^*] = \bar{\boldsymbol{\mu}}^* + \omega \mathbf{b}$.

Corollary 3 (Tracking error and information ratio under a biased chair). *Let $\bar{\boldsymbol{\mu}} = \boldsymbol{\mu}_b + \bar{\boldsymbol{\mu}}^*$ be the true total belief. The expected squared tracking error is*

$$\mathbb{E}[TE(k)^2] = \frac{\mathbb{E}[q(k)]}{\gamma^2} = \frac{1}{\gamma^2} \left(\|\bar{\boldsymbol{\mu}}^* + \omega \mathbf{b}\|^2 + \sigma_\varepsilon^2 H(\mathbf{w}_k) \right), \quad (45)$$

so the bias raises squared active risk by $2\omega \langle \bar{\boldsymbol{\mu}}^*, \mathbf{b} \rangle + \omega^2 \|\mathbf{b}\|^2$ above the noise floor $\sigma_\varepsilon^2 H(\mathbf{w}_k) / \gamma^2$. The information ratio has the compact form

$$IR(k) = \frac{\langle \boldsymbol{\mu}_k^*, \bar{\boldsymbol{\mu}} \rangle}{\|\boldsymbol{\mu}_k^*\|} = \|\bar{\boldsymbol{\mu}}\| \cos(\boldsymbol{\mu}_k^*, \bar{\boldsymbol{\mu}}), \quad (46)$$

the magnitude of the true total belief times the cosine of the angle between the committee's view and the truth; conviction $\|\boldsymbol{\mu}_k^*\|$ cancels, so only the direction of $\boldsymbol{\mu}_k^*$ matters. At the representative view—exact as $\omega \rightarrow 1$, when the dominant chair's view swamps the residual noise—it is

$$\widetilde{IR}(k) = \frac{\langle \bar{\boldsymbol{\mu}}^* + \omega \mathbf{b}, \bar{\boldsymbol{\mu}} \rangle}{\|\bar{\boldsymbol{\mu}}^* + \omega \mathbf{b}\|} = \frac{\kappa + \gamma\rho + \omega \langle \mathbf{b}, \bar{\boldsymbol{\mu}} \rangle}{\sqrt{\kappa + 2\omega \langle \bar{\boldsymbol{\mu}}^*, \mathbf{b} \rangle + \omega^2 \|\mathbf{b}\|^2}}. \quad (47)$$

Proof. Equation (45) is $\mathbb{E}[\text{TE}(k)^2] = \mathbb{E}[q(k)]/\gamma^2$ with $\mathbb{E}[q(k)]$ from (43). For (46), the numerator of $IR(k) = (\gamma r(k) + s(k))/\sqrt{q(k)}$ is $\langle \boldsymbol{\mu}_k^*, \boldsymbol{\mu}_b \rangle + \langle \boldsymbol{\mu}_k^*, \bar{\boldsymbol{\mu}}^* \rangle = \langle \boldsymbol{\mu}_k^*, \bar{\boldsymbol{\mu}} \rangle$, since $\gamma r(k) = \langle \boldsymbol{\mu}_k^*, \boldsymbol{\mu}_b \rangle$, and the denominator is $\|\boldsymbol{\mu}_k^*\|$, giving $IR(k) = \langle \boldsymbol{\mu}_k^*, \bar{\boldsymbol{\mu}} \rangle / \|\boldsymbol{\mu}_k^*\| = \|\bar{\boldsymbol{\mu}}\| \cos(\boldsymbol{\mu}_k^*, \bar{\boldsymbol{\mu}})$. Substituting $\boldsymbol{\mu}_k^* = \bar{\boldsymbol{\mu}}^* + \omega \mathbf{b}$ and using $\langle \bar{\boldsymbol{\mu}}^*, \bar{\boldsymbol{\mu}} \rangle = \kappa + \gamma\rho$ (because $\langle \bar{\boldsymbol{\mu}}^*, \boldsymbol{\mu}_b \rangle = \gamma\rho$) gives (47). $\square$

As the chair's influence ω grows, $\boldsymbol{\mu}_k^*$ rotates from the unbiased view $\bar{\boldsymbol{\mu}}^*$ toward $\bar{\boldsymbol{\mu}}^* + \mathbf{b}$ and generically away from the direction $\bar{\boldsymbol{\mu}}$ that maximizes the information ratio, so the cosine in (46), and with it $\widetilde{IR}(k)$, falls even as (45) has the committee taking ever-larger active positions. A bias that is large but uninformative— $\|\mathbf{b}\|$ large but $\langle \mathbf{b}, \bar{\boldsymbol{\mu}} \rangle$ small—drives the information ratio toward zero: a big, confident, and wrong-headed bet.

Proposition 3 sharpens the earlier results in one decisive way. For unbiased committees, concentration left expected accuracy untouched and acted only on variance (Corollaries 1 and 2); a biased chair makes concentration act on the *mean*. The expected gap (41) now splits into a bias term $\omega^2 \|\mathbf{b}\|^2$ and the familiar variance term $\sigma_\varepsilon^2 H(\mathbf{w}_k)$ —a bias–variance decomposition of the committee's error. Both worsen as the chair's sway grows: the variance term rises with the concentration $H(\mathbf{w}_k)$, and the bias term rises with the square of her effective influence ω . Concentrating on a biased chair is therefore doubly costly, importing both her idiosyncratic noise and her systematic error. The bias, moreover, is *irreducible*: averaging cancels noise but not bias, and it is the chair's own dominance that injects it, so deliberation that drives $\omega \rightarrow 1$ —the holdout, or chair-take-all—leaves the committee carrying her entire bias $\mathbf{b}$. Only spreading influence back onto the unbiased members, lowering ω , can dilute it.

The *direction* of the bias determines how it is paid for. Resolve $\mathbf{b}$ into its components along and orthogonal to the benchmark $\boldsymbol{\mu}_b$. The component along $\boldsymbol{\mu}_b$ moves the portfolio's beta but is invisible to alpha; the component along the true-alpha direction $\bar{\boldsymbol{\mu}}_\perp^*$ shifts expected alpha by $\frac{\omega}{\gamma} \langle \mathbf{b}, \bar{\boldsymbol{\mu}}_\perp^* \rangle$ in (44)—a chair whose bias happens to point toward the truth earns more, one whose bias points away earns less—while the leading term $\frac{1}{\gamma} \|\bar{\boldsymbol{\mu}}_\perp^*\|^2$ is just the alpha available to any unbiased committee. Whatever its direction, the bias adds active risk through the $\omega^2 \|\mathbf{b}\|^2$ in $\mathbb{E}[q(k)]$, so that even a bias orthogonal to the truth raises tracking error while contributing nothing to return, lowering the information and Sharpe ratios. Under

rational expectations ($\bar{\boldsymbol{\mu}}^* = \mathbf{0}$, so $\bar{\boldsymbol{\mu}}_{\perp}^* = \mathbf{0}$) expected alpha stays exactly zero even for a biased chair. The bias still inflates tracking error to $\frac{1}{\gamma}\sqrt{\omega^2\|\mathbf{b}\|^2 + \sigma_\varepsilon^2 H(\mathbf{w}_k)}$: the committee takes ever-larger uncompensated positions and, mistaking its biased view for the truth, grows more overconfident the more it defers to the chair.

3.3 The limiting consensus

The belief the committee approaches as deliberation lengthens has a clean characterization in terms of the spectrum of $\mathbf{W}$. Because $\mathbf{W}$ is row-stochastic, the all-ones vector $\mathbf{1} \in \mathbb{R}^I$ is a right eigenvector with eigenvalue one, $\mathbf{W}\mathbf{1} = \mathbf{1}$. When the communication network is strongly connected and aperiodic, this eigenvalue is simple and strictly dominant, $1 = \lambda_1 > |\lambda_2| \geq \dots \geq |\lambda_I|$, and $\mathbf{W}$ has a unique left eigenvector $\boldsymbol{\pi}$ satisfying

$$\boldsymbol{\pi}'\mathbf{W} = \boldsymbol{\pi}', \quad \boldsymbol{\pi}'\mathbf{1} = 1, \quad \boldsymbol{\pi} \geq \mathbf{0}, \quad (48)$$

the stationary distribution of the chain. The Perron–Frobenius theorem then gives $\lim_{k \rightarrow \infty} \mathbf{W}^k = \mathbf{1}\boldsymbol{\pi}'$, so that $\lim_{k \rightarrow \infty} \mathbf{M}^{(k)} = \mathbf{1}\boldsymbol{\pi}'\mathbf{M}^{(0)}$. The rows of this limit are identical: all members converge to the common consensus belief

$$\boldsymbol{\mu}_\infty^* = \mathbf{M}^{(0)'}\boldsymbol{\pi} = \sum_{i=1}^I \pi_i \boldsymbol{\mu}_i^{(0)}, \quad (49)$$

the influence-weighted average of the members' initial views. The weight π_i is member i 's *eigenvector centrality*: her sway over the committee's final view is set by her global position in the network $\mathbf{W}$ —whom she is heard by, and whom they in turn are heard by. Because all members hold the same belief in the limit, any unanimity-respecting aggregator uses this as the effective belief, so the limiting effective view $\boldsymbol{\mu}_\infty^*$ is the consensus belief (49), independently of the rule f .

More generally, the entire trajectory is governed by the spectrum of $\mathbf{W}$. If $\mathbf{W}$ is diagonalizable, with eigenvalues $\lambda_1 = 1, \lambda_2, \dots, \lambda_I$ and right and left eigenvectors $\mathbf{v}_j$ and $\mathbf{u}_j$ normalized so that $\mathbf{u}_i'\mathbf{v}_j = \delta_{ij}$ (with $\mathbf{v}_1 = \mathbf{1}$ and $\mathbf{u}_1 = \boldsymbol{\pi}$), the k -round beliefs admit the spectral expansion

$$\mathbf{M}^{(k)} = \mathbf{W}^k \mathbf{M}^{(0)} = \underbrace{\mathbf{1}\boldsymbol{\pi}'\mathbf{M}^{(0)}}_{\text{consensus}} + \sum_{j=2}^I \lambda_j^k \mathbf{v}_j \mathbf{u}_j' \mathbf{M}^{(0)}. \quad (50)$$

The leading term is the consensus of (49); the remaining terms are transient deviations that decay geometrically, the slowest at the rate set by the second-largest eigenvalue modulus $|\lambda_2|$.

The spectral gap $1 - |\lambda_2|$ thus measures how quickly the committee converges: a cohesive committee with a large gap reaches consensus in a few rounds, whereas one that is nearly reducible—two subgroups joined by only thin communication—has $|\lambda_2|$ close to one and converges slowly. This is the precise sense in which the number of discussion rounds k matters: each round contracts the transient by a factor of about $|\lambda_2|$, and the committee’s effective view equals the consensus belief only once enough rounds have passed for these terms to die away. When the network is instead *reducible*, the eigenvalue one is no longer simple—its multiplicity equals the number of closed communicating classes—and $\lim_k \mathbf{W}^k$ is block-structured rather than rank-one. Each closed class then converges to its own internal influence-weighted average, and no single committee-wide belief emerges: this is the failure of consensus analyzed in [Section 4.5](#).

4 Simulations

4.1 Stylized influence structures

Before turning to simulated portfolios, it is useful to build intuition for the influence matrix $\mathbf{W}$. Recall that row i records the weights member i places on each colleague’s beliefs, and on her own, when updating. Because $\mathbf{W}$ is row-stochastic, each row sums to one. For what follows, fix a committee of $I = 5$ members and consider seven stylized communication structures, ranging from a single dominant voice through fully diffuse influence and on to committees that splinter or are captured by an immovable member.

Unipolar. A single member dominates the conversation. Member 1 listens almost entirely to herself and barely to anyone else, while the other members place the bulk of their weight on member 1 and little on themselves or on each other:

$$\mathbf{W}_{\text{uni}} = \begin{bmatrix} 0.90 & 0.025 & 0.025 & 0.025 & 0.025 \\ 0.80 & 0.05 & 0.05 & 0.05 & 0.05 \\ 0.80 & 0.05 & 0.05 & 0.05 & 0.05 \\ 0.80 & 0.05 & 0.05 & 0.05 & 0.05 \\ 0.80 & 0.05 & 0.05 & 0.05 & 0.05 \end{bmatrix}. \quad (51)$$

Because member 1 scarcely revises her view while everyone else is pulled toward it, repeated communication drives the whole committee toward member 1’s initial belief. Iterating the

influence matrix makes this precise: its steady state is

$$\mathbf{W}_{\text{uni}}^{\infty} \equiv \lim_{k \rightarrow \infty} \mathbf{W}_{\text{uni}}^k = \begin{bmatrix} 0.8889 & 0.0278 & 0.0278 & 0.0278 & 0.0278 \\ 0.8889 & 0.0278 & 0.0278 & 0.0278 & 0.0278 \\ 0.8889 & 0.0278 & 0.0278 & 0.0278 & 0.0278 \\ 0.8889 & 0.0278 & 0.0278 & 0.0278 & 0.0278 \\ 0.8889 & 0.0278 & 0.0278 & 0.0278 & 0.0278 \end{bmatrix}, \quad (52)$$

whose identical rows are the long-run influence (left-eigenvector centrality) weights $\boldsymbol{\pi}$. Member 1 commands roughly 89% of the committee's influence (8/9, against 1/36 for each of the others), so the consensus belief $\sum_{i=1}^I \pi_i \boldsymbol{\mu}_i^{(0)}$ —and hence the portfolio—essentially reflects one person's signal. (Entries are rounded to four decimal places.)

Bipolar. Two members dominate, each anchoring one of two poles. Members 1 and 2 each listen mostly to themselves and little to anyone else, while the three remaining members weight the two poles about equally and place only a small weight on themselves:

$$\mathbf{W}_{\text{bi}} = \begin{bmatrix} 0.90 & 0.025 & 0.025 & 0.025 & 0.025 \\ 0.025 & 0.90 & 0.025 & 0.025 & 0.025 \\ 0.45 & 0.45 & 0.10 & 0 & 0 \\ 0.45 & 0.45 & 0 & 0.10 & 0 \\ 0.45 & 0.45 & 0 & 0 & 0.10 \end{bmatrix}. \quad (53)$$

The long-run consensus is a blend of the two poles' initial beliefs, with the three followers' own priors washing out. Its steady state is

$$\mathbf{W}_{\text{bi}}^{\infty} \equiv \lim_{k \rightarrow \infty} \mathbf{W}_{\text{bi}}^k = \begin{bmatrix} 0.4615 & 0.4615 & 0.0256 & 0.0256 & 0.0256 \\ 0.4615 & 0.4615 & 0.0256 & 0.0256 & 0.0256 \\ 0.4615 & 0.4615 & 0.0256 & 0.0256 & 0.0256 \\ 0.4615 & 0.4615 & 0.0256 & 0.0256 & 0.0256 \\ 0.4615 & 0.4615 & 0.0256 & 0.0256 & 0.0256 \end{bmatrix}, \quad (54)$$

so members 1 and 2 jointly command about 92% of the long-run influence (6/13 each), while each of the three followers retains only about 2.6% (1/39). The consensus is essentially an equal blend of the two poles' initial beliefs, and committee performance then hinges on how far apart those two views are and on whether they happen to be correct.

Nonpolar. Influence is fully diffuse: every member places equal weight on everyone, including herself,

$$\mathbf{W}_{\text{non}} = \begin{bmatrix} 0.20 & 0.20 & 0.20 & 0.20 & 0.20 \\ 0.20 & 0.20 & 0.20 & 0.20 & 0.20 \\ 0.20 & 0.20 & 0.20 & 0.20 & 0.20 \\ 0.20 & 0.20 & 0.20 & 0.20 & 0.20 \\ 0.20 & 0.20 & 0.20 & 0.20 & 0.20 \end{bmatrix}. \quad (55)$$

This matrix is doubly stochastic and symmetric, and it is idempotent ($\mathbf{W}_{\text{non}}^k = \mathbf{W}_{\text{non}}$ for every $k \geq 1$), so it equals its own steady state,

$$\mathbf{W}_{\text{non}}^\infty \equiv \lim_{k \rightarrow \infty} \mathbf{W}_{\text{non}}^k = \mathbf{W}_{\text{non}}, \quad (56)$$

in which every member carries identical influence $\pi_i = 1/5$. The committee’s consensus is therefore simply the equal-weighted average of the five members’ initial beliefs—the “wisdom of crowds” aggregator. When members’ signals are noisy but unbiased, this structure cancels the most idiosyncratic error and yields the most diversified effective view of any committee that reaches a single consensus.

Faction. The committee splits into two self-contained factions that never listen across the divide. Members 1–3 form one faction and members 4–5 the other; within each faction members weight one another and themselves equally, but place no weight on the opposing faction:

$$\mathbf{W}_{\text{fac}} = \begin{bmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & 0 & 0 \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & 0 & 0 \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & 0 & 0 \\ 0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} \\ 0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} \end{bmatrix}. \quad (57)$$

Unlike the previous committees, this influence matrix is *reducible*: it contains two closed communicating classes, so deliberation never produces a single committee-wide consensus. Because each diagonal block is itself a uniform averaging matrix, $\mathbf{W}_{\text{fac}}$ is idempotent and equals its own steady state,

$$\mathbf{W}_{\text{fac}}^\infty \equiv \lim_{k \rightarrow \infty} \mathbf{W}_{\text{fac}}^k = \mathbf{W}_{\text{fac}}. \quad (58)$$

Communication drives the first faction to its internal average belief $\frac{1}{3}(\boldsymbol{\mu}_1^{(0)} + \boldsymbol{\mu}_2^{(0)} + \boldsymbol{\mu}_3^{(0)})$ and the second to its own average $\frac{1}{2}(\boldsymbol{\mu}_4^{(0)} + \boldsymbol{\mu}_5^{(0)})$, but the two camps stay permanently apart—there is no single influence vector, only one stationary distribution per faction. Any

committee-wide effective view $\boldsymbol{\mu}^*$ must therefore be supplied by the aggregator f rather than by deliberation. Under equal-weight pooling, for instance, $\boldsymbol{\mu}^*$ places weight $3/5$ on the larger faction’s consensus and $2/5$ on the smaller’s.

The dictator. One member listens only to herself, and the rest of the committee places no weight on her. Member 1’s row loads entirely on her own belief, while the other four ignore her and listen equally to one another:

$$\mathbf{W}_{\text{dic}} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0.25 & 0.25 & 0.25 & 0.25 \\ 0 & 0.25 & 0.25 & 0.25 & 0.25 \\ 0 & 0.25 & 0.25 & 0.25 & 0.25 \\ 0 & 0.25 & 0.25 & 0.25 & 0.25 \end{bmatrix}. \quad (59)$$

Like the faction committee, this matrix is reducible—it has two closed communicating classes, the dictator $\{1\}$ on her own and the deliberating bloc $\{2, 3, 4, 5\}$ —so there is no committee-wide consensus. Because each diagonal block is a uniform averaging matrix, $\mathbf{W}_{\text{dic}}$ is idempotent and equals its own steady state,

$$\mathbf{W}_{\text{dic}}^\infty \equiv \lim_{k \rightarrow \infty} \mathbf{W}_{\text{dic}}^k = \mathbf{W}_{\text{dic}}. \quad (60)$$

The dictator remains permanently at her own prior $\boldsymbol{\mu}_1^{(0)}$, while the other four converge to their internal average $\frac{1}{4}(\boldsymbol{\mu}_2^{(0)} + \boldsymbol{\mu}_3^{(0)} + \boldsymbol{\mu}_4^{(0)} + \boldsymbol{\mu}_5^{(0)})$; the two groups never reconcile, and there is no single influence vector, only one stationary distribution per class. Although she takes no input from her colleagues, the dictator also exerts no influence on them: she is isolated from the group, neither revising her own belief nor altering theirs. Any committee-wide effective view $\boldsymbol{\mu}^*$ is therefore supplied by the aggregator f rather than by deliberation; under equal-weight pooling, $\boldsymbol{\mu}^*$ places weight $1/5$ on the dictator’s view and $4/5$ on the majority bloc’s consensus.

The holdout. One member listens only to herself, while the rest of the committee remains fully open, including to her. Member 5’s row loads entirely on her own belief, whereas

members 1 through 4 weight all five members equally, exactly as in the nonpolar committee:

$$\mathbf{W}_{\text{hold}} = \begin{bmatrix} 0.20 & 0.20 & 0.20 & 0.20 & 0.20 \\ 0.20 & 0.20 & 0.20 & 0.20 & 0.20 \\ 0.20 & 0.20 & 0.20 & 0.20 & 0.20 \\ 0.20 & 0.20 & 0.20 & 0.20 & 0.20 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}. \quad (61)$$

Member 5 is an absorbing state, and because the other members place positive weight on her she is reachable from the entire committee. Repeated communication therefore drives the entire committee to the holdout's belief, so the steady state has identical rows,

$$\mathbf{W}_{\text{hold}}^{\infty} \equiv \lim_{k \rightarrow \infty} \mathbf{W}_{\text{hold}}^k = \begin{bmatrix} 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \quad (62)$$

with long-run influence vector $\boldsymbol{\pi} = (0, 0, 0, 0, 1)$ and consensus belief equal to the holdout's own prior, $\boldsymbol{\mu}^* = \boldsymbol{\mu}_5^{(0)}$. The holdout stands in sharp contrast to the dictator: both place all weight on their own belief, but the dictator is ignored and isolated, whereas the holdout is heeded and therefore determines the committee's consensus—a single member, attended to by all, fixes the committee's effective view. Because each colleague places only one-fifth of her weight on the holdout, the committee converges to her belief gradually, over many rounds of discussion, rather than at once.

The idealist. Like the holdout, the idealist occupies the bottom row and listens only to herself; unlike the holdout, the rest of the committee places no weight on her. Member 5's row loads entirely on her own belief, while members 1 through 4 place no weight on her and listen equally to one another:

$$\mathbf{W}_{\text{ide}} = \begin{bmatrix} 0.25 & 0.25 & 0.25 & 0.25 & 0 \\ 0.25 & 0.25 & 0.25 & 0.25 & 0 \\ 0.25 & 0.25 & 0.25 & 0.25 & 0 \\ 0.25 & 0.25 & 0.25 & 0.25 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}. \quad (63)$$

This is the dictator structure with the isolated member moved to the bottom row: it is reducible, with two closed classes—the idealist $\{5\}$ alone and the bloc $\{1, 2, 3, 4\}$ —so deliberation produces no committee-wide consensus. Each diagonal block is a uniform averaging matrix, so $\mathbf{W}_{\text{ide}}$ is idempotent and equals its own steady state,

$$\mathbf{W}_{\text{ide}}^\infty \equiv \lim_{k \rightarrow \infty} \mathbf{W}_{\text{ide}}^k = \mathbf{W}_{\text{ide}}. \quad (64)$$

The idealist holds permanently to her prior $\boldsymbol{\mu}_5^{(0)}$, while the four others converge to their internal average $\frac{1}{4}(\boldsymbol{\mu}_1^{(0)} + \boldsymbol{\mu}_2^{(0)} + \boldsymbol{\mu}_3^{(0)} + \boldsymbol{\mu}_4^{(0)})$. The idealist and the holdout differ only in whether the other members attend to the stubborn member: the holdout, heeded by all, draws the committee to her belief, whereas the idealist, heeded by none, has no effect on the others, who converge to their own internal average. As in the dictator committee, any committee-wide belief is supplied by the aggregator rather than by deliberation: under equal-weight pooling $\boldsymbol{\mu}^*$ is again the grand mean, while a chair drawn from the bloc adopts the four-member average and excludes the idealist entirely.

Together these cases span a spectrum from maximal concentration of influence (unipolar) through partial concentration (bipolar) to perfect diffusion (nonpolar); the faction, dictator, and idealist committees illustrate a qualitatively different possibility in which communication produces no single consensus at all, while the holdout shows the opposite extreme, in which one immovable member is heeded by everyone and captures the entire committee. In the simulations that follow, each structure is fed through the belief-updating and portfolio-choice mechanism of [Section 3](#) and the resulting performance metrics are compared.

4.2 Performance of the committee styles

This section quantifies how each communication structure translates into portfolio performance. A committee of $I = 5$ members oversees an active portfolio over $N = 10$ assets with covariance matrix $\boldsymbol{\Sigma}$, generated from a three-factor model with idiosyncratic noise. The benchmark $\boldsymbol{\theta}_b$ is the capitalization-weighted market, and $\boldsymbol{\Sigma}$ is scaled so that the benchmark’s Sharpe ratio under the equilibrium prior $\boldsymbol{\mu}_b = \gamma \boldsymbol{\Sigma} \boldsymbol{\theta}_b$ is 0.50; at risk aversion $\gamma = 3$ the institution holds the benchmark plus the active tilt, $\boldsymbol{\theta}_p = \boldsymbol{\theta}_b + \frac{1}{\gamma} \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}^*$. Each member observes an independent noisy signal of the true view, $\boldsymbol{\mu}_i^{(0)} = \bar{\boldsymbol{\mu}}^* + \boldsymbol{\varepsilon}_i$ (renormalized to a zero-sum view), with signal-noise standard deviation 0.04, so the members’ equal-weighted average is an unbiased estimate of the truth. I aggregate by equal-weight *pooling* and by a *chair* rule that adopts member 1’s belief, and I evaluate the metrics of [Section 3](#) both as *perceived* (under the committee’s own belief) and as *realized* (under the true returns $\bar{\boldsymbol{\mu}} = \boldsymbol{\mu}_b + \bar{\boldsymbol{\mu}}^*$), averaging over 100,000 Monte Carlo draws of the members’ signals.

One feature of this baseline design organizes the results. Because every committee’s effective view is an influence-weighted average of unbiased signals, and the influence weights sum to one, *every* committee is an unbiased estimator of the truth: $\mathbb{E}[\boldsymbol{\mu}^*] = \bar{\boldsymbol{\mu}}^*$ for all seven structures. By the closed form $\alpha = \frac{1}{\gamma}(s - r\rho/\sigma_b^2)$ of [Section 3.1](#), the *expected* Jensen’s alpha is therefore identical across committees: the communication structure does not change how much true return the institution earns on average. What it changes is the *risk* taken to earn it. A committee that overweights one member’s signal inherits that member’s idiosyncratic error as tracking error—active risk that carries no compensating return—so the committees separate not on alpha but on the *information ratio*, the return earned per unit of active risk. I report two cases, distinguished by whether the benchmark is the truth.

4.2.1 Rational expectations

In the first case the market is correct: the true view is zero, $\bar{\boldsymbol{\mu}}^* = \mathbf{0}$, so the benchmark is the truth and there is no active return to be had. [Table 1](#) reports realized performance. As [Section 3](#) predicts, the realized alpha and information ratio are zero for each committee under either aggregator—tilting away from a correct market cannot add return. The committees nonetheless differ sharply in the active *risk* they take. The diffuse nonpolar committee, whose pooled belief averages members’ spurious signals toward zero, runs a tracking error of 0.08; the concentrated unipolar and holdout committees, which ride one member’s noise, run more than twice that, 0.19. Because all earn the same benchmark risk premium, the committee that wastes the least active risk has the highest *total* Sharpe ratio—0.45 for the nonpolar committee against 0.33 for the holdout. Together these show that when there is nothing to find, a committee that concentrates on a strong voice merely takes larger uncompensated bets.

The committees are also *overconfident*, and most so exactly where they are worst. Under their own beliefs the concentrated committees perceive large information ratios—0.50 for the unipolar committee and 0.57 for the holdout—and positive alphas, because one unaveraged, high-conviction view looks impressive *ex ante*; under the truth these collapse to zero. The diffuse committees, whose pooled view is muted, are far better calibrated, perceiving an information ratio of 0.25. Deliberation makes matters worse: [Fig. 1](#) shows that as the number of discussion rounds grows, equal-weight pooling drives the unipolar and holdout committees’ tracking error up—from 0.08 at $k = 0$ toward 0.17 and 0.19—as influence concentrates on the dominant member and the committee takes ever-larger uncompensated positions, while the diffuse committees hold flat. A committee that deliberates briefly wastes less risk than the same committee that talks to consensus.

Table 1: Realized performance under rational expectations

Committee	Aggregator	Sharpe	Track. err.	Info. ratio	Alpha
Unipolar	Pooling	0.3550	0.1677	0.0000	0.0000
	Chair	0.3550	0.1677	0.0000	0.0000
Bipolar	Pooling	0.4022	0.1232	0.0000	0.0000
	Chair	0.4022	0.1232	0.0000	0.0000
Nonpolar	Pooling	0.4452	0.0842	0.0000	0.0000
	Chair	0.4452	0.0842	0.0000	0.0000
Faction	Pooling	0.4452	0.0842	0.0000	0.0000
	Chair	0.4183	0.1088	0.0000	0.0000
Dictator	Pooling	0.4452	0.0842	0.0000	0.0000
	Chair	0.3350	0.1884	0.0000	0.0000
Holdout	Pooling	0.3348	0.1885	0.0000	0.0000
	Chair	0.3348	0.1885	0.0000	0.0000
Idealist	Pooling	0.4452	0.0842	0.0000	0.0000
	Chair	0.4344	0.0941	0.0000	0.0000

Notes. Averages over 100,000 Monte Carlo draws at the communication steady state, evaluated against the true returns, in the rational-expectations case ($\bar{\mu}^* = \mathbf{0}$, so the benchmark is the truth). *Sharpe* is the total portfolio Sharpe ratio; the other columns are benchmark-relative. *Pooling* is equal-weight aggregation, *Chair* adopts member 1's belief. Realized alpha and information ratio are zero for all committees, which differ only in the tracking error (active risk) they take, which the total Sharpe ratio reflects in reverse. Beta (omitted) is one throughout, since the institution holds one unit of the benchmark.




Figure 1: Tracking error as a function of the number of discussion rounds k , under equal-weight pooling, in the rational-expectations case. With the benchmark equal to the truth there is no reward for active risk, yet deliberation drives the concentrated committees to take more of it: the unipolar and holdout committees' tracking error rises from the wisdom-of-crowds value at $k = 0$ toward 0.17 and 0.19 as influence concentrates, while the diffuse committees (nonpolar, faction, dictator, idealist) hold flat. Averages over 100,000 Monte Carlo draws.

4.2.2 An informed committee

In the second case the benchmark is wrong: the truth is the benchmark plus a view $\bar{\mu}^*$, calibrated to a standard deviation of 0.02, and the members hold unbiased noisy signals of that view, so their equal-weighted average recovers it. Now there is genuine value to capture. Table 2 reports realized performance. The realized alpha is essentially identical across committees (0.020): because the committees are all unbiased, each earns the same expected active return. But the committees capture it with very different efficiency. The nonpolar committee, by averaging its members’ signals, isolates the true view with little idiosyncratic noise and achieves a realized information ratio of 0.16; the unipolar and holdout committees, relying on one noisy member, achieve only 0.10 and 0.09. The diffuse committee is no more accurate on average but far more *efficient*, earning the same alpha at roughly half the tracking error.

The overconfidence of the first case persists: the holdout perceives an information ratio of 0.61 while realizing 0.09, and the unipolar committee perceives 0.55 while realizing 0.10, so the committees that look most impressive on their own books are the least efficient in fact. Deliberation again sharpens the contrast: Fig. 2 shows that under pooling, additional rounds erode the concentrated committees’ information ratio—from 0.16 at $k = 0$ down toward 0.10 for the unipolar committee and 0.09 for the holdout—as each round concentrates influence and discards the averaging that made the crowd’s estimate efficient, while the diffuse committees hold at the wisdom-of-crowds value. Fig. 3 confirms that the same ordering governs the full distribution: the concentrated committees’ realized information ratios are not only lower on average but more dispersed, with fatter tails of poorly diversified outcomes.

4.3 The efficient committee: matching influence to skill

The seven archetypes above hold every member equally skilled, so equal weighting is best and concentration is unhelpful (Corollary 1). Proposition 2 predicts that this verdict changes once members differ in precision: the efficient committee then weights each member by her precision, $w_i^* \propto p_i$, and concentration matched to skill can beat naive diffusion. To illustrate, I make member 1 an expert (signal standard deviation 0.02), members 2–4 average (0.04), and member 5 low-skill (0.08), so the members’ precisions stand at roughly 12:3:3:3:1 and the efficient weights at $w^* \approx (0.55, 0.14, 0.14, 0.14, 0.03)$. I feed unbiased signals of this heterogeneous quality through four influence allocations—the precision optimum, equal weighting, all weight on the expert, and all weight on the low-skill member—and report realized performance in the informed case over 200,000 draws (Table 3).

The simulation is consistent with Proposition 2. The precision-optimal committee attains

Table 2: Realized performance with an informed committee

Committee	Aggregator	Sharpe	Track. err.	Info. ratio	Alpha
Unipolar	Pooling	0.4069	0.1859	0.0972	0.0204
	Chair	0.4069	0.1859	0.0972	0.0204
Bipolar	Pooling	0.4560	0.1470	0.1239	0.0203
	Chair	0.4560	0.1470	0.1239	0.0203
Nonpolar	Pooling	0.4988	0.1165	0.1584	0.0202
	Chair	0.4988	0.1165	0.1584	0.0202
Faction	Pooling	0.4988	0.1165	0.1584	0.0202
	Chair	0.4722	0.1352	0.1353	0.0203
Dictator	Pooling	0.4988	0.1165	0.1584	0.0202
	Chair	0.3855	0.2047	0.0881	0.0204
Holdout	Pooling	0.3852	0.2048	0.0878	0.0203
	Chair	0.3852	0.2048	0.0878	0.0203
Idealist	Pooling	0.4988	0.1165	0.1584	0.0202
	Chair	0.4882	0.1238	0.1484	0.0202

Notes. Same simulation as [Table 1](#), but the truth is the benchmark plus a view $\bar{\mu}^*$ of standard deviation 0.02, of which the members hold unbiased noisy signals. Realized alpha is essentially equal across committees (all committees are unbiased estimators of the truth), so the committees separate on the information ratio: averaging captures the true view at low active risk, concentration at high active risk. Columns and conventions as in [Table 1](#).




Figure 2: Realized information ratio as a function of the number of discussion rounds k , under equal-weight pooling, in the informed-committee case. All committees begin at the wisdom-of-crowds value at $k = 0$; as deliberation concentrates influence, the information ratios of the unipolar, bipolar, and holdout committees decline (the holdout's only slowly), while the diffuse committees (nonpolar, faction, dictator, idealist) hold. A committee that deliberates briefly outperforms the same committee that talks to consensus. Averages over 100,000 Monte Carlo draws.




Figure 3: Distribution of each realized performance metric across the 100,000 Monte Carlo draws, by committee style, under equal-weight pooling in the informed-committee case. Curves are Gaussian kernel-density estimates; committees whose effective view coincides draw-for-draw (nonpolar, faction, dictator, idealist) are merged into a single curve. The concentrated committees (holdout, unipolar) have lower and more dispersed information ratios and fatter tracking-error tails, whereas the diffuse committees concentrate tightly at a higher information ratio. The realized risk premium is common to all committees—a near-degenerate spike—because each committee is an unbiased estimator of the truth.

the lowest expected squared gap to the truth, hitting the additivity floor $1/\sum_i p_i$ of (38) to within Monte Carlo error, and earns the highest realized information ratio, 0.17. Equal weighting—the wisdom-of-crowds rule that was *optimal* under homogeneous skill—now carries 2.1 times the estimation error and an information ratio of only 0.14, because it dilutes the expert’s signal with three average ones and a poor one. Most pointedly, concentrating *all* influence on the expert beats equal weighting (0.15 against 0.14, at 1.8 versus 2.1 times the optimal gap): when one member is far better informed than the rest, listening only to her wastes less than averaging her away. The mirror image is the catastrophe of concentrating on the low-skill member, whose information ratio collapses to 0.05 at twenty-nine times the optimal gap. Realized alpha is, as always, common across the four committees (≈ 0.020)—all committees are unbiased—so they separate purely on the risk efficiency that [Proposition 2](#) governs. The lesson refines the headline of [Section 4.2](#): diffusion is not a virtue in itself. The goal is not to spread influence but to *align* it with skill—to concentrate exactly as much as the dispersion of member precisions warrants, and no more.

Table 3: The efficient committee under heterogeneous skill

Influence allocation	N_{eff}	$\mathbb{E}\ \boldsymbol{\delta}\ ^2/\text{opt}$	Track. err.	Info. ratio	Alpha
Precision-optimal ($w_i \propto p_i$)	2.76	1.00	0.1069	0.1740	0.0203
Equal-weight ($w_i = 1/I$)	5.00	2.10	0.1293	0.1416	0.0202
Expert-only ($w = \mathbf{e}_1$)	1.00	1.81	0.1239	0.1487	0.0203
Low-skill only ($w = \mathbf{e}_5$)	1.00	29.00	0.3852	0.0461	0.0203

Notes. Realized performance over 200,000 Monte Carlo draws in the informed case ($\bar{\boldsymbol{\mu}}^*$ of standard deviation 0.02), when members hold unbiased signals of *heterogeneous* precision: member 1’s signal has standard deviation 0.02, members 2–4’s 0.04, and member 5’s 0.08, so the precisions are about 12:3:3:3:1 and the efficient weights $w^* \approx (0.55, 0.14, 0.14, 0.14, 0.03)$. The committees differ only in the effective influence vector $\mathbf{w}$. $N_{\text{eff}} = 1/(\mathbf{w}'\mathbf{w})$ is the effective number of members; $\mathbb{E}\|\boldsymbol{\delta}\|^2/\text{opt}$ is the expected squared estimation error relative to the precision optimum $1/\sum_i p_i$ of [Proposition 2](#) (the closed form $\sum_i w_i^2 \sigma_i^2$ matches the simulated gap to within Monte Carlo error). Realized alpha is common across committees, so they separate on the information ratio. Equal weighting, optimal under homogeneous skill, is bettered both by precision weighting and—here—by concentrating on the lone expert, while concentrating on the low-skill member is ruinous.

The same logic governs *deliberation*, and it helps to explain the sign reversal of [Section 4.2](#) from the skill side. [Figure 4](#) traces the realized information ratio against the number of discussion rounds under equal-weight pooling, for committees of this heterogeneous skill whose influence networks are centered on different members. Because deliberation drives effective influence $\mathbf{w}_k = (\mathbf{W}')^k \mathbf{a}$ toward eigenvector centrality, its value has the sign of the movement it induces relative to the precision weights $\mathbf{w}^*$. A network in which the able member is the one everyone heeds—the unipolar structure, centered on the expert—carries $\mathbf{w}_k$ toward skill, and discussion *lifts* the information ratio above the equal-weight starting point. In this case,

most of the gain arrives in the first round, after which further deliberation slightly overshoots the expert and gives a little back. A network centered on the low-skill member—the holdout structure—carries $\mathbf{w}_k$ the other way, and discussion drives the information ratio steeply down. A diffuse network holds it flat. Deliberation is thus neither good nor bad as such: it helps when the committee comes to listen more to the members worth heeding and hurts when it does not.




Figure 4: Realized information ratio versus rounds of discussion k under equal-weight pooling, for committees of heterogeneous skill (member 1 an expert, signal standard deviation 0.02; members 2–4 average, 0.04; member 5 low-skill, 0.08). Deliberation drives effective influence $\mathbf{w}_k = (\mathbf{W}')^k \mathbf{a}$ toward eigenvector centrality: a network centered on the expert (everyone heeds member 1) moves it toward the precision weights and lifts the information ratio above the equal-weight value, with most of the gain in the first round; a network centered on the low-skill member (everyone heeds member 5) moves it away and collapses the ratio; a diffuse network holds it flat. The dashed line is the precision optimum of [Proposition 2](#), which pooling on a centered network approaches but does not reach because it overshoots the optimal concentration. Averages over 200,000 Monte Carlo draws.

4.4 A biased chair

The examples above studied committees whose members hold unbiased beliefs: all seven styles draw unbiased signals, so the committees separate on the *risk* they take and not on

systematic error. [Proposition 3](#) predicts what changes when the chair—member 1, the voice the chair rule adopts—instead holds a *biased* view whose systematic component $\mathbf{b}$ differs from the truth. To test it, I give member 1 a fixed bias $\mathbf{b}$ added to her signal on each draw, leaving the other four members unbiased. The simulation changes the magnitude $\|\mathbf{b}\|$ from zero to four times the per-member signal scale σ_ε . Three configurations differ in the chair’s effective influence $\omega = w_{k,1}$: she dominates because everyone defers to her (the unipolar committee, $\omega \approx 0.89$), she dominates because the rule adopts her view (chair-take-all, $\omega = 1$), or she is one of five equals (the nonpolar committee under pooling, $\omega = 0.20$). The bias is oriented relative to the benchmark-orthogonal true view $\bar{\boldsymbol{\mu}}_\perp^*$: a *favorable* bias points toward it, an *adverse* bias is its negation, and the two have equal magnitude.

The simulations confirm [Proposition 3](#) and [Corollary 3](#) on all counts. First, the effective view inherits a fraction ω of the chair’s bias, and the mean squared gap to the truth tracks the bias–variance prediction $\omega^2\|\mathbf{b}\|^2 + \sigma_\varepsilon^2 H(\boldsymbol{w})$ to within Monte Carlo error: for the dominant chair at $\|\mathbf{b}\| = 4\sigma_\varepsilon$ the simulated gap is 4.58 against a predicted 4.57. Second, the damage scales with influence ([Table 4](#)). As her bias grows the dominant chair’s realized tracking error nearly quadruples (from 0.19 to 0.71) while her information ratio collapses from 0.10 toward zero. The chair-take-all committee behaves almost identically; the diffuse committee, which places only one-fifth of its weight on her, barely moves—with its information ratio slipping from 0.16 to 0.11. A biased chair leads to poor performance only to the extent that the committee listens to her.

Third, the information ratio falls whether the bias is favorable or adverse, even though the alpha moves in opposite directions ([Fig. 5](#)). A favorable bias—one that happens to point toward the truth—raises realized alpha (from 0.020 to 0.035 for the dominant chair), yet the information ratio still falls, because the bias adds tracking error faster than it adds return. An adverse bias lowers alpha (here from 0.020 toward zero) and lowers the information ratio by even more. Conviction in the wrong proportion is costly on a risk-adjusted basis whether or not the committee is lucky in the direction of its mistake. Under rational expectations ($\bar{\boldsymbol{\mu}}^* = \mathbf{0}$) this last point is starkest: realized alpha is then *exactly* zero for any bias and any committee. The chair’s bias still inflates tracking error, so it buys pure uncompensated risk. Finally, the committees remain *overconfident*, and most so where they are worst: the chair’s bias, which she mistakes for insight, drives the *perceived* information ratio up—past 2.0 for the dominant chair at the largest bias—exactly as the realized one falls, widening the overconfidence wedge most for the committee that defers to her most.

Table 4: Realized performance under an adverse-biased chair

Committee	$\ \mathbf{b}\ /\sigma_\varepsilon$	Track. err.	Info. ratio	Alpha	Perceived IR
Dominant chair ($\omega = 0.89$)	0	0.1858	0.0968	0.0202	0.5505
	1	0.2477	0.0868	0.0166	0.7639
	2	0.3848	0.0636	0.0129	1.1848
	4	0.7079	0.0425	0.0056	2.1589
Chair-take-all ($\omega = 1.00$)	0	0.2047	0.0881	0.0203	0.6075
	1	0.2762	0.0790	0.0161	0.8501
	2	0.4319	0.0586	0.0121	1.3267
	4	0.7960	0.0395	0.0038	2.4231
Diffuse chair ($\omega = 0.20$)	0	0.1166	0.1583	0.0203	0.3380
	1	0.1207	0.1587	0.0194	0.3628
	2	0.1360	0.1463	0.0186	0.4196
	4	0.1882	0.1133	0.0169	0.5886

Notes. Realized performance in the informed case ($\bar{\mu}^*$ of standard deviation 0.02), averaged over 100,000 Monte Carlo draws, when the chair (member 1) carries an *adverse* systematic bias $\mathbf{b}$ swept in units of the per-member signal scale σ_ε ; $\omega = w_{k,1}$ is the chair’s effective influence. The three committees differ only in how much weight the deliberation and aggregator place on the chair: the dominant (unipolar) and chair-take-all committees collapse as her bias grows, while the diffuse (nonpolar, pooling) committee is largely insulated. *Perceived IR* is the information ratio under the committee’s own belief; the gap from the realized value is the overconfidence the bias produces. Under a favorable bias the alpha rises instead of falling, but the information ratio still declines (Fig. 5).




Figure 5: Realized tracking error, information ratio, and Jensen’s alpha for the dominant chair (unipolar, pooling) as her bias $\|\mathbf{b}\|$ grows, under a *favorable* bias (aligned with the benchmark-orthogonal true view, raising alpha) and an *adverse* bias (its negation, lowering alpha), in the informed case. Both directions raise tracking error and lower the information ratio; only the alpha moves in opposite directions (the dotted line marks the unbiased alpha). Averages over 100,000 Monte Carlo draws.

4.5 Consensus and the role of the aggregator

When does a committee agree? Whether deliberation ends in a single shared belief is governed entirely by the structure of the influence matrix $\mathbf{W}$, and the conditions under which repeated averaging of this kind reaches a consensus have been studied in other settings. DeGroot (1974) gave a sufficient condition, Berger (1981) the necessary and sufficient one, and DeMarzo et al. (2003) and Golub and Jackson (2010) study the same dynamics in economic settings, characterizing the limiting consensus as a network-weighted average of members’ initial views (see Golub and Sadler 2016, for a survey). In the present notation, a committee reaches consensus—every member converging on the same belief—if and only if its communication network contains a single closed group that the rest of the committee can reach: formally, one closed communicating class toward which influence flows, directly or indirectly, from every member. Intuitively, the committee agrees when it holds “one conversation,” so that everyone’s view is ultimately traceable to a common set of voices. The unipolar, bipolar, and nonpolar committees are connected in this way and so converge. The holdout also converges: although member 5 listens to no one, every other member listens to her, so she is the single closed group toward which all influence flows, and the committee converges on her belief. A stubborn member, by herself, does not prevent agreement.

What prevents agreement is *mutual deafness*. The faction, dictator, and idealist committees each split into two groups that neither listen to nor are heard by one another—two factions that ignore each other, or a stubborn member and a bloc that place no weight on one another. Each group is closed, so the network has two recurrent classes and no single belief can emerge: the subgroups converge internally but never to each other. The contrast between the holdout and the idealist is interesting. Both seat one member who listens only to herself, but the holdout is heeded by the rest and so unifies the committee, whereas the idealist is ignored and so fractures it. Consensus, in short, requires not that the committee be free of stubborn members but that no subgroup be closed to the others’ influence.

Consensus makes the aggregator a formality; disagreement makes it decisive.

This distinction determines how large a role the portfolio manager—the aggregator f —plays in the outcome. When the committee reaches consensus, all members hold the identical belief $\bar{\mu}$, and any aggregation rule that respects unanimity returns it: equal-weight pooling, a chair rule, and a seniority-weighted average all yield $\mu^* = \bar{\mu}$. The manager is a bystander, because deliberation has determined the portfolio. This is why, in Tables 1 and 2, the pooling and chair rules deliver *identical* metrics for the four committees that reach consensus. When the committee does not agree, members hold different beliefs and the aggregator must adjudicate among them. In this case μ^* , and hence the portfolio, is determined by the decision rule

f . The decision then rests with the manager rather than with the committee. Aggregation therefore matters most when the committee is internally divided, which is also when the institution’s choice of governance rule has the largest effect on the portfolio it ends up holding.

Performance: consensus fixes it, disagreement leaves it to the manager. The performance consequences follow this logic, and the informed case (Table 2) shows them most clearly. Under equal-weight pooling, the three committees that fail to reach consensus—faction, dictator, and idealist—match the *best* consensus committee, the nonpolar, each realizing an information ratio of 0.16. Pooling re-mixes the fragments’ separate consensus back into the grand mean, so a divided committee whose views are combined by an averaging manager recovers the full wisdom of the crowd. However, under a chair rule the same divided committees vary with whichever subgroup the chair belongs to, from near-best—the idealist, whose chair sits inside the four-member bloc—to worst—the dictator, whose chair is the isolated stubborn member herself. Among the committees that *do* reach consensus, by contrast, the aggregator has no effect: once deliberation has concentrated everyone on one or two members’ views, as in the unipolar and holdout committees, the portfolio is committed to a poorly diversified position that no aggregation rule can undo, and realized information ratios remain low (0.10 and 0.09) under either rule.

Fig. 3 shows that the same logic governs the whole distribution, not merely its mean. The dispersion of realized performance across draws is driven by the effective number of members whose signals enter μ^* . The more members, the tighter and more right-shifted the information-ratio distribution and the thinner the right tails of variance and tracking error. The concentrated committees—holdout and unipolar—place weight on only one or two members and so display the widest, most left-shifted information-ratio distributions and the fattest tails of occasional large, poorly diversified positions, whereas the well-diversified nonpolar consensus and the pooled fragments cluster tightly around a higher information ratio. Consensus fixes this effective number through deliberation, beyond the manager’s reach. Without consensus, the manager sets it through the choice of f .

5 Testable Implications

The model’s comparative statics are stated in observables—the concentration of influence within a committee, the rule by which its views are aggregated, the length of deliberation, and the risk-adjusted performance of the resulting portfolio. This leads to some falsifiable predictions. The empirical analogue of the influence matrix $\mathbf{W}$ is whatever captures who carries weight (relative speaking time, voting power, seniority, or the share of decisions

traceable to a member); the analogue of the effective view μ^* is the portfolio’s active tilt, recoverable from disclosed holdings. Six predictions follow, with the simulated magnitudes of [Section 4.2](#) as effect-size priors:

- H1.** *Concentration raises active risk without active return.* Committees with more concentrated influence—a high Herfindahl index of estimated centralities, or a dominant CIO/chair—run higher tracking error and lower information ratios than otherwise-comparable diffuse committees, with statistically indistinguishable alphas. This joint pattern (higher active risk, lower information ratio, unchanged alpha) is the model’s distinguishing signature, which no skill, size, or fee story reproduces; a positive concentration effect on alpha, or none on tracking error, would reject it ([Corollaries 1 and 2](#)).
- H2.** *Concentrated committees are overconfident.* The wedge between perceived performance (under the committee’s own forecasts, where recorded in minutes or return targets) and realized performance widens with concentration ([Corollary 3](#)).
- H3.** *Governance matters only under disagreement.* The aggregation rule moves the portfolio if and only if the committee is divided ([Section 4.5](#)), so a governance reform should move divided committees’ portfolios but not cohesive ones—a built-in placebo.
- H4.** *Deliberation length has an aggregator-dependent sign.* The information ratio falls with deliberation length under pooling and rises under a chair rule—a sign reversal difficult to rationalize without the influence-concentration channel.
- H5.** *Sway tracks centrality, not volume.* A member’s implied weight in the collective decision loads on her eigenvector centrality rather than her raw participation ([Section 3.3](#)).
- H6.** *Performance tracks the alignment of influence to skill.* Risk-adjusted performance declines in the precision-weighted misallocation $\sum_i \sigma_i^2 (w_i - w_i^*)^2$ of [\(39\)](#) between realized and inverse-variance influence ([Proposition 2](#)). This refines H1: holding skill fixed concentration is pure cost, but once members differ in skill its sign flips with its target—concentration on an able member can beat diffusion—so an unconditional benefit of diffusion would reject precision weighting.

[Table 5](#) maps the model’s primitives to observable proxies in the two leading settings—team- versus single-managed mutual funds and public pension boards—and [Table 6](#) states the predictions as estimating equations with signs. The transcript-free core (H1–H3 and H6) turns only on a concentration proxy, disclosed holdings, and the one belief institutions report directly: a public pension’s assumed actuarial return, which makes the overconfidence wedge

of H2 measurable. H4 and H5, which need deliberation length and members’ individual views, await richer records such as those of monetary-policy committees. Identification is the central challenge, as influence is chosen rather than assigned: staggered governance reforms and chair or CIO turnover supply plausibly exogenous variation in $\mathbf{W}$ and f for a difference-in-differences design, with the consensual committees as a built-in placebo.

Table 5: Mapping the model to observable data

Model object	Meaning	Mutual funds (#1)	Public pensions (#2)
$H(\mathbf{w}) = 1/N_{\text{eff}}$	influence concentration	manager Herfindahl; $1/(\#\text{managers})$; single-manager dummy	inverse board size; board-composition Herfindahl
ω	dominant member’s effective influence	lead-manager dummy; longest-tenure weight share	sole-trustee dummy ; ex-officio/appointed share
σ_i^{-2}	member precision (skill)	manager’s historical forecast accuracy	trustee financial expertise; staff delegation
$\ \boldsymbol{\mu}^*\ $	conviction $\rightarrow$ tracking error	tracking error; active share	allocation distance from peer policy book
$s(k)$	accuracy / realized active return	realized alpha; information ratio	realized return – policy benchmark
μ_p (perceived)	committee’s own expected return	— (not disclosed)	assumed actuarial return
overconf. wedge	perceived – realized	—	assumed – realized (or – CMA-implied)
$\mathbf{b}$	dominant member’s bias	dominant manager’s past error	political-cycle pressure on a sole trustee

Notes. $N_{\text{eff}} = 1/H(\mathbf{w})$ is the effective number of members (Corollary 1). The concentration and dominance proxies are observable without committee transcripts; the active tilt and its tracking error come from disclosed holdings; the assumed actuarial return is the rare directly disclosed “perceived” belief that makes the overconfidence wedge (Corollary 3) measurable.

6 Conclusion

Most institutional capital is managed by committees, yet the standard theory of portfolio choice treats the investor as a single mind. This paper has taken the gap between the two seriously by embedding the DeGroot (1974) model of opinion formation into the mean-variance framework of Markowitz (1952): committee members begin with heterogeneous, noisy views about expected returns, revise them through a communication network, and a portfolio manager aggregates the post-deliberation beliefs into the active view that drives the institution’s portfolio. The communication structure $\mathbf{W}$, the aggregation rule f , and the length of deliberation k thereby map directly into the portfolio’s tracking error, information ratio, and alpha, making the internal dynamics of the committee—ordinarily invisible in the

Table 6: Hypotheses, specifications, and predicted signs

	Prediction	Dependent var. $\sim$ regressor	Sign	Source
H1a	concentration raises active risk	tracking error (alloc. distance) $\sim$ conc.	+	Corollary 2
H1b	but not risk-adjusted return	information ratio $\sim$ conc.	−	Corollary 1
H1c	and not raw return (discriminating null)	realized alpha $\sim$ conc.	0	Corollary 1
H2	concentration widens overconfidence	(assumed − realized) wedge $\sim$ conc.	+	Corollary 3
H2′	a biased dominant voice imports systematic error	wedge $\sim$ sole-trustee $\times$ political pressure	+	Proposition 3
H6	influence should align with skill	info. ratio $\sim$ misallocation $\sum_i \sigma_i^2 (w_i - w_i^*)^2$	−	Proposition 2
H3	governance matters only when divided	outcome $\sim$ conc. $\times$ disagreement	$\neq 0^a$	Section 4.5
DiD	deconcentration narrows wedge, cuts active risk	post sole $\rightarrow$ board $\sim$ wedge, active risk	−, −	Proposition 3

Notes. “Conc.” is any concentration proxy of [Table 5](#) (sole-trustee dummy, inverse board size, manager Herfindahl); all specifications carry year and entity controls. The sole $\rightarrow$ board difference-in-differences is identified off staggered governance reforms.

^a The interaction is nonzero while the concentration effect among consensual committees is zero—a built-in falsification test.

data—a measurable determinant of performance.

Three findings stand out. First, whether a committee reaches a shared view is a property of the structure of attention rather than of the stubbornness of individuals: consensus fails not because some member is immovable but because the committee fractures into subgroups that neither hear nor are heard by one another. Second, the manager’s aggregation rule is immaterial when the committee agrees and decisive when it does not, so the institution’s choice of governance matters most precisely when its committee is internally divided. Third, and most consequentially for performance, concentrating influence on a few members imports their idiosyncratic estimation error as uncompensated active risk: in the simulations, the most concentrated committees run more than double the tracking error of the diffuse ones (roughly 0.19 against 0.08) for no gain in realized alpha, and they are systematically overconfident, perceiving information ratios near 0.6 while realizing under 0.10. Two corollaries sharpen the picture. Under equal-weight pooling, whether a committee reaches consensus is by itself a poor predictor of performance, because averaging re-mixes a divided committee’s fragments back into the wisdom of the crowd; and deliberation is not neutral—additional rounds can *erode* risk-adjusted performance under pooling by concentrating influence, while improving it under a chair rule as the chair learns from her colleagues.

These results carry a common practical message. For institutions, the lesson is that good

committee outcomes are less about recruiting a single strong voice than about preserving the diversity of views that averaging exploits: governance that concentrates influence—through a domineering chair, a deferential culture, or simply talking to consensus—buys conviction at the cost of diversification, and conviction earns nothing when it is not also accuracy. The model makes this prescription precise: the efficient committee weights each member by the precision of her information—the Bayesian aggregate—so the goal is not to diffuse influence indiscriminately but to *match* it to skill. Equal weighting is optimal only when members are equally able; otherwise the cost of any committee is the distance between the attention it pays and the attention members deserve, and naive deliberation underperforms precisely when it lets influence accrue to the central or voluble rather than the well-informed. For those who oversee such institutions, the model suggests that the observable features of a committee—how concentrated its influence is, whether it reaches consensus, how its views are aggregated—carry information about the risk its portfolio is taking that its headline holdings alone do not reveal, and that a committee’s own perceived performance is least trustworthy exactly where its influence is most concentrated.

The framework is deliberately stylized, and its limitations mark the natural next steps. Prices are held fixed and the committee’s deliberation feeds a single institution’s portfolio; members update naively in the [DeGroot \(1974\)](#) manner rather than communicating strategically; the influence matrix $\mathbf{W}$ and aggregator f are taken as primitives rather than chosen; and members’ signals are unbiased, so committees separate on the risk they take rather than on systematic error. Relaxing any of these—endogenizing the influence structure as the outcome of members’ incentives, allowing biased or correlated signals, embedding many such institutions in a market that clears, or letting members anticipate how their stated views will be aggregated—would extend the analysis in a natural direction. The most valuable next step, however, is empirical: the predictions collected in [Section 5](#) are stated in terms of measurable proxies for influence concentration, governance, and deliberation, and confronting them with data from team-managed funds, pension and endowment committees, or other deliberative bodies would tell us how much of measured institutional performance is attributable, as the model contends it should be, to the way a committee listens to itself.

Data availability statement

The results, tables, and figures in this paper are fully reproducible from the accompanying code, `committee_simulations.py`; no external or proprietary data are used.

References

- Berger, Roger L.**, “A Necessary and Sufficient Condition for Reaching a Consensus Using DeGroot’s Method,” *Journal of the American Statistical Association*, 1981, 76 (374), 415–418.
- Black, Fischer and Robert Litterman**, “Global Portfolio Optimization,” *Financial Analysts Journal*, 1992, 48 (5), 28–43.
- DeGroot, Morris H.**, “Reaching a Consensus,” *Journal of the American Statistical Association*, 1974, 69 (345), 118–121.
- DeMarzo, Peter M., Dimitri Vayanos, and Jeffrey Zwiebel**, “Persuasion Bias, Social Influence, and Unidimensional Opinions,” *Quarterly Journal of Economics*, 2003, 118 (3), 909–968.
- Evans, Richard B., Melissa Porras Prado, A. Emanuele Rizzo, and Rafael Zambrana**, “Identity, Diversity, and Team Performance: Evidence from U.S. Mutual Funds,” *Management Science*, 2025, 71 (4), 3026–3051.
- Garlappi, Lorenzo, Ron Giammarino, and Ali Lazrak**, “Ambiguity and the Corporation: Group Disagreement and Underinvestment,” *Journal of Financial Economics*, 2017, 125 (3), 417–433.
- Golub, Benjamin and Evan Sadler**, “Learning in Social Networks,” in Yann Bramoullé, Andrea Galeotti, and Brian Rogers, eds., *The Oxford Handbook of the Economics of Networks*, Oxford University Press, 2016, pp. 504–542.
- **and Matthew O. Jackson**, “Naïve Learning in Social Networks and the Wisdom of Crowds,” *American Economic Journal: Microeconomics*, 2010, 2 (1), 112–149.
- Hatcher, Michael and Tim Hellmann**, “Networks, Beliefs, and Asset Prices,” *Journal of Economic Dynamics and Control*, 2025, 173, 105059.
- Kuchler, Theresa and Johannes Stroebel**, “Social Finance,” *Annual Review of Financial Economics*, 2021, 13, 37–55.
- Markowitz, Harry**, “Portfolio Selection,” *Journal of Finance*, 1952, 7 (1), 77–91.
- Pedersen, Lasse Heje**, “Game On: Social Networks and Markets,” *Journal of Financial Economics*, 2022, 146 (3), 1097–1119.