Oversight Risk: How Investment Committees Shape Portfolio Performance

Scott Condie* Gabriel Lehnardt† James Tavita‡

ABSTRACT

Most institutional portfolios are overseen by an investment committee that has formal authority over the portfolio manager's decisions. These committee members often bring diverse perspectives and expertise. This paper develops a model of institutional portfolio management in which investment decisions are subject to oversight by a committee with heterogeneous beliefs and risk preferences. Since the committee's effective beliefs and risk aversion governing portfolio choice can reflect changes in committee preferences and power dynamics within a committee, the resulting portfolio of the committee is subject to what we call *oversight risk*. This paper characterizes the consequences of oversight risk for portfolio performance. There are at least two channels through which oversight risk can operate: fluctuations in effective risk aversion and fluctuations in effective beliefs. We show that fluctuations in effective risk aversion affect all standard performance metrics—Sharpe ratio, beta, Jensen's alpha, and the Information Ratio—through a market-timing channel whose direction depends on the correlation between risk aversion and market returns. Fluctuations in effective beliefs erode performance by increasing portfolio churn without necessarily generating compensating active returns. Simulations calibrated to monthly U.S. equity sector data spanning March 1970 to December 2025 confirm these predictions and document a skill–noise tradeoff in which committee skill raises expected performance while oversight noise widens the distribution of outcomes. Implications for investment committees are provided.

Keywords: Investment committees, oversight risk, institutional investment, portfolio performance, committee governance

JEL classification: G11, G23, D71

*Department of Economics, Brigham Young University. Email: ssc@byu.edu.
†Department of Economics, Brigham Young University.
‡Department of Economics, Brigham Young University.

1 Introduction

Many institutional investment decisions are made by committees whose members have diverse backgrounds, objectives, and beliefs about how money should be managed. Endowments, pension funds, foundations, and family offices frequently delegate formal authority over the portfolio to a committee of trustees or board members, while the portfolio manager retains day-to-day responsibility for implementation. In this setting, the manager's ability to generate strong risk-adjusted returns depends not only on market insight and analytical skill, but also on the ability to navigate the preferences and power dynamics of the oversight committee. This paper studies the portfolio performance consequences of this institutional structure.

We define *oversight risk* as the risk that a portfolio's effective investment policy will shift in response to changes in the beliefs, risk tolerances, and relative influence of oversight committee members. We develop a mean-variance framework for measuring and categorizing this risk, and we use simulations calibrated to over five decades of monthly U.S. equity data to quantify its consequences for standard performance metrics.

The core insight is that oversight risk operates through two channels with different signatures in standard performance metrics. When the committee's effective risk tolerance fluctuates—either because different members periodically dominate deliberation and those members have different risk preferences, or because risk preferences of many committee members change—the portfolio is repeatedly scaled away from its original risk position. When this scaling is correlated with market returns, all standard performance metrics are affected: pro-cyclical risk tolerance—scaling up following market downturns—improves the Sharpe ratio and generates positive alpha when negative returns are more commonly followed by high returns as is the case in the sample of data that we use for simulations. In this sample counter-cyclical risk tolerance degrades both the Sharpe ratio and alpha. When the committee's effective beliefs about expected returns fluctuate but are neither wrong, nor right on average the result is excess portfolio tilting that is not compensated by higher expected returns, which reduces the Sharpe ratio monotonically while widening the distributions of both alpha and beta but leaving the Information Ratio largely unaffected. In other words, committee investment skill is necessary to correct for the noise added by belief fluctuations. These distinct signatures allow an observer to diagnose which type of oversight risk is driving performance drag in a specific institutional context.

1.1 Related Literature

The governance literature consistently finds that committee composition and process affect investment outcomes. Binfare et al. (2017) provide direct empirical evidence from university endowments, showing that the professional backgrounds of investment committee members have a substantial impact on fund performance. Endowments whose committees include a higher proportion of members with alternative investment expertise—particularly venture capital and private equity experience—earn significantly higher risk-adjusted returns, with a 10 percent increase in the proportion of board members with venture capital experience associated with 100 to 120 basis points of additional excess return. These composition effects operate through both asset allocation and manager selection: committees with relevant expertise allocate more to alternative assets and select better-performing managers within those asset classes. The Binfare et al. (2017) findings establish that *who* sits on the committee matters for asset allocation, which is the cross-sectional analog of the question we study here. They give an example where the presence of individuals with expertise in a certain area is related to a tilt toward that area. In the language of our model, the Binfare et al. results speak to the tilt of the effective belief vector $\boldsymbol{y}$—committees with more expertise in a particular area have beliefs that are more aligned with investment in those areas. Binfare et al. (2017) focus on the realized performance of this tilt toward alternative investments, as opposed to the structural properties of this tilt that could happen in any market as we do.

Scherer (2024) develops a model of optimal committee design that addresses several of the governance problems we study here. Scherer identifies three pathologies of traditional qualitative consensus-building in investment committees—group shift bias (a form of group polarization in which members' views become more extreme after deliberation), free-rider incentive problems, and aggregation failures that allow dominant voices to crowd out independent signals. His proposed solution is an algorithmic consensus: each committee member independently and anonymously submits a portfolio of active positions, these are rescaled to carry identical tracking error so that no single member's risk appetite dominates, and the committee portfolio is formed by averaging the individual submissions. This design is directly relevant to the oversight risk framework developed here: the rescaling step controls the risk aversion channel by neutralizing differences in members' effective risk tolerances, while the averaging of independent submissions could reduce the belief noise channel by ensuring that idiosyncratic forecast errors cancel rather than compound through group discussion.

Other work on investment committees includes Ellis (2011) which identifies best practices for investment committees, and Yoder (2011) which provides governance guidance specific to educational institutions. Additionally, the formal and real authority framework of Aghion & Tirole (1997) and the oversight and delegation framework of Useem & Zelleke (2006) provide theoretical foundations for understanding when and how managers balance their own expertise against the preferences of overseers. Yet the standard metrics for evaluating institutional portfolio performance—Sharpe ratio, Jensen's alpha, Information Ratio—were not designed with oversight risk in mind. When these metrics deteriorate for a committee-governed portfolio, it is often unclear how much of the underperformance is attributable to poor market insight versus poor governance. Our framework helps disentangle these two sources.

Existing empirical evidence on team-managed investment vehicles provides important context. Bär et al. (2011) study U.S. mutual funds and compare funds managed by teams to those managed by individual portfolio managers. They find that team-managed funds adopt less extreme investment styles, hold more diversified portfolios, and exhibit lower portfolio turnover than individually managed funds. Critically, this moderation does not translate into superior performance: team-managed funds do not outperform their individually managed counterparts, and they display significantly less performance persistence—consistent with the view that team deliberation compresses the signal in investment decisions along with the noise. Prather & Middleton (2002) reach a similar conclusion, finding that team-managed mutual funds do not consistently outperform individually managed funds. These findings are directly relevant to the framework developed here. In our model, the averaging of committee member beliefs through the influence-weighting process produces exactly this kind of moderation when the committee is stable. But when the influence weights *fluctuate*—the dimension that the existing empirical literature does not study—the averaging effect itself becomes time-varying, generating excess portfolio weight variance that imposes additional performance costs beyond those documented in the mutual fund literature. Our contribution is to characterize and quantify these dynamic costs.

This paper relates to several other bodies of literature in the broader social sciences. It connects to the behavioral literature on group decision-making, which documents systematic tendencies—group polarization (Moscovici & Zavalloni 1969, Isenberg 1986), groupthink (Janis 1982), informational and reputational cascades (Sunstein & Hastie 2015), and shared information bias (Stasser & Titus 1985, Reimer et al. 2010)—that can generate exactly the kind of belief and preference fluctuations we model. It draws on expectation states theory (Berger et al. 1972, Correll & Ridgeway 2003) and the organizational politics literature (Eisenhardt & Bourgeois 1988) to explain why influence weights are both unequal and time-varying. It extends the financial economics literature on committee decision-making, where Blinder & Morgan (2005) and Lombardelli et al. (2005) find that committees can outperform individuals in monetary policy settings, while Bär et al. (2011) and Prather & Middleton (2002) find less favorable results for investment management. It connects to the delegated portfolio management literature, where Berk & Green (2004) study how the competitive equilibrium between investors and managers eliminates abnormal returns, and He & Krishnamurthy (2012) model how agency frictions between managers and their investors amplify asset price volatility. The committee setting we study can be viewed as a special case of this broader delegation problem, in which the agency friction arises not from a single investor–manager relationship but from the internal dynamics of a multi-member oversight body whose effective preferences and beliefs are themselves stochastic. It complements the herding literature (Scharfstein & Stein 1990), in which career concerns lead managers to deviate from optimal portfolios—here oversight risk operates through a related channel in which committee pressure, rather than peer imitation, drives the deviation.

2 Behavioral and Organizational Foundations

Before turning to the formal model, we ground the two channels of oversight risk studied here in well-studied behavioral and organizational dynamics that can generate them in practice. The research literature on group decision-making documents several mechanisms that produce the belief and preference fluctuations central to our analysis.

The first potential avenue for committee dynamics to impact investment outcomes is through production blocking. Investment committees form views about expected returns through group discussion of macroeconomic trends, sector developments, valuation signals, and other factors. This deliberation process is analogous to a brainstorming session in which members generate and evaluate hypotheses, with the discussion ideally aggregating individual private signals into a collective forecast. A key source of inefficiency in such deliberation is *production blocking*—the constraint that only one member can speak at a time. Nijstad et al. (2003) show experimentally that the delays imposed by turn-taking interfere with the cognitive process of idea generation in two distinct ways: long delays disrupt the *organization* of idea generation (members lose the thread of a developing line of thought), while unpredictable delays reduce its *flexibility* (members become less able to shift between different categories of ideas). In a committee setting, the practical consequence is that effective beliefs disproportionately reflect the views of early or dominant speakers, whose ideas suffer the least blocking-induced degradation, independently of the informational value of their contributions.

A related problem is the *shared information bias*, or "hidden profile" effect: groups systematically over-discuss information already held in common by all members and under-surface the unique private signals that only individual members possess, because shared information is statistically more likely to be raised and is socially validating to discuss (Stasser & Titus 1985, Reimer et al. 2010). Together, production blocking and shared information bias introduce two types of distortion into committee beliefs. First, they bias the mean of effective beliefs toward the views of dominant members rather than the informationally efficient aggregation of all members' private signals. Second, because the cognitive disruptions documented by Nijstad et al. (2003) may vary in severity from meeting to meeting—depending on the length and unpredictability of interruptions—they add idiosyncratic across-meeting variance to effective beliefs. Both distortions contribute to the variance of effective beliefs in the formal model below.

Another mechanism that can systematically distort group beliefs is the need for cognitive closure. Kruglanski & Webster (1996) define need for closure as the desire for a firm, definite answer to a question—any answer—over continued ambiguity and uncertainty, and identify two characteristic tendencies: *seizing*, the inclination to converge rapidly on whatever position is available without waiting for additional evidence, and *freezing*, the inclination to maintain that position and resist information that would require reopening the question. Critically, need for closure is not only a stable individual trait but can be situationally elevated by conditions that raise the perceived cost of continued deliberation: time pressure, environmental noise, cognitive fatigue, and ambient uncertainty. In the managing of institutional portfolios, these conditions can be most common during periods of market stress, when the need for closure is likely to be elevated and the consequences of belief distortion most severe.

At the group level, elevated need for closure produces what Kruglanski et al. (2006) call *group-centrism*: social pressure toward opinion uniformity, marginalization of dissenters (Kruglanski & Webster 1991), a preference for homogeneous membership (Kruglanski et al. 2002), and autocratic communication structures in which one or a few members dominate deliberation (Pierro et al. 2003). The cumulative consequence is a committee that has seized on a set of beliefs, frozen in place, and organized its social dynamics to resist correction. Unlike production blocking, which primarily inflates the variance of effective beliefs, need for closure primarily distorts their mean—and does so in a direction that the committee's own social dynamics actively protect from correction.

Additionally, committee deliberation often shifts collective views toward more extreme positions than the average of members' pre-meeting opinions. Moscovici & Zavalloni (1969) provide an early characterization of this *group polarization* phenomenon, and the meta-analysis of Isenberg (1986) documents its robustness across diverse settings. Group polarization is relevant to investment committees because it implies that a committee of investors with heterogeneous risk preferences may not simply average those preferences, but may systematically adopt more extreme collective positions. When committee composition or relative influence changes over time, this mechanism generates the period-to-period belief and preference volatility that is the subject of the present paper. Charness & Sutter (2012) provide an important counterpoint, finding that in strategic settings groups tend to make more rational decisions than individuals—suggesting that whether committee oversight improves or degrades portfolio performance depends critically on the decision environment and committee composition.

Sunstein & Hastie (2015) synthesize these deliberation pathologies into a unified account of how group discussion can amplify rather than correct individual errors. They identify two cascade mechanisms that are especially relevant to investment committees. In *informational cascades*, members who speak later in a discussion rationally discount their own private signals in favor of the views already expressed, so that early speakers exert disproportionate influence on the group's conclusion regardless of the relative quality of their information. In *reputational cascades*, members self-censor views that would put them at odds with the emerging consensus, because the social cost of dissent outweighs the private benefit of contributing an independent signal. Together, these cascades mean that the effective influence weights in our model are not stable structural parameters but are themselves endogenous to the order of discussion, the social standing of early speakers, and the perceived cost of disagreement—all of which can vary from meeting to meeting and generate the time-varying influence dynamics that our framework studies.

A deeper structural explanation for why influence weights are unequal—and can be systematically misaligned with information quality—comes from expectation states theory. Berger et al. (1972) show that in task-oriented groups (see Correll & Ridgeway 2003 for a comprehensive review), members rapidly form performance expectations based on *status characteristics*: observable attributes such as seniority, professional credentials, institutional rank, or demographic categories that carry culturally shared assumptions about competence. These expectations then govern who is given opportunities to contribute, whose contributions are positively evaluated, and who prevails in disagreements—creating a self-reinforcing hierarchy in which high-status members dominate group output regardless of whether their task-relevant expertise warrants it. On an investment committee, status characteristics such as board tenure, professional title, or the size of assets a member personally controls can establish influence hierarchies that persist independently of investment skill. The result is a set of influence weights $\boldsymbol{\phi}$ that reflect social standing rather than information quality. When status hierarchies shift—through board turnover, changes in members' external reputations, or evolving institutional politics—the influence weights shift with them, generating exactly the time-varying oversight dynamics that our model captures.

At the organizational level, the political dynamics of coalition formation provide a further mechanism for time-varying influence. Eisenhardt & Bourgeois (1988) study top management teams in high-velocity environments and find that political behavior—alliance formation, information withholding, and behind-the-scenes lobbying—is both pervasive and consequential for decision quality. Critically, they find that the use of politics is associated with *poorer* firm performance: teams that relied more heavily on political maneuvering made slower decisions and were less likely to adopt high-quality alternatives. In the investment committee context, shifting coalitions among board members—formed around shared views on asset allocation, loyalty to a particular consultant or manager, or institutional power struggles—can abruptly reallocate effective influence from one faction to another. Unlike the gradual drift produced by individual-level status dynamics, coalition realignment can produce discrete jumps in the committee's effective beliefs and risk preferences, corresponding to large jumps in the values of effective beliefs and risk aversion in our model. The performance costs we document in the simulations below can therefore be understood as a quantification of what organizations lose when political dynamics, rather than informational quality, determine whose views prevail.

Finally, it is well documented that risk tolerance can shift at the individual level. Guiso et al. (2018) document that people become more risk averse after experiencing investment losses. This phenomenon, when combined with potential committee dynamics that encourage either "playing it safe" (i.e. an increase in effective aversion to risk following portfolio losses), or "buying the dip" (i.e. a decrease in effective risk aversion following declines in asset returns) can alter the amount of risk in the portfolio held by the institution. Additionally, investment committees face compositional dynamics where shifts in *whose* risk preferences effectively govern the portfolio can lead to these same phenomena occurring. As the relative influence of committee members changes over time, the committee's effective risk aversion changes with it, even when no individual member's preferences have changed. From the portfolio manager's perspective, this manifests as time-varying effective risk aversion: the portfolio that the committee will accept changes from period to period in a way that may or may not be correlated with portfolio performance.

3 A Model of Oversight Risk

Consider a set of overseers $i \in \{1, \dots, I\}$, each with mean-variance preferences characterized by risk aversion γ_i and beliefs $\boldsymbol{\mu}_i$ about expected returns. We assume the N risky assets have true expected returns $\boldsymbol{\mu}$, covariance matrix $\boldsymbol{\Sigma}$, and risk-free rate r_f. The random weight ϕ_i reflects overseer i's current influence, with $\boldsymbol{\phi}$ denoting the vector of influence weights (summing to one).

Portfolio managers face pressure to adhere to the wishes of a divergent set of overseers, even when the fund's charter places sole responsibility for portfolio decisions on the manager. Committee members who hold strong views about particular asset classes may interpret underperformance as management failure, especially since the counterfactual performance under their preferred strategy is typically not considered. We model this by assuming the manager acts as if maximizing a weighted average of overseer utilities, where the weights reflect the current political dynamics of the committee. Although insiders may understand why particular members gain or lose influence, from the portfolio's perspective these fluctuations are an additional source of randomness—and the more they fluctuate, the greater the potential risk to performance.

3.1 Risk aversion volatility

We first isolate the effect of fluctuating risk preferences by assuming all committee members agree on expected returns ($\boldsymbol{\mu}_i = \boldsymbol{\mu}$ for all i) but differ in their risk tolerance. Each overseer i has mean-variance utility $U_i(\boldsymbol{x}) = r_f + \boldsymbol{x}'(\boldsymbol{\mu} - r_f\mathbf{1}) - \frac{\gamma_i}{2}\boldsymbol{x}'\boldsymbol{\Sigma}\boldsymbol{x}$, and the manager solves

$$\max_{\boldsymbol{x}} \sum_{i=1}^{I} \phi_i \left[\boldsymbol{x}'(\boldsymbol{\mu} - r_f\mathbf{1}) - \frac{\gamma_i}{2}\boldsymbol{x}'\boldsymbol{\Sigma}\boldsymbol{x}\right]. \tag{1}$$

The optimal portfolio is

$$\boldsymbol{x} = \frac{1}{\boldsymbol{\phi}'\boldsymbol{\gamma}}\,\boldsymbol{\Sigma}^{-1}(\boldsymbol{\mu} - r_f\mathbf{1}), \tag{2}$$

where $\boldsymbol{\phi}'\boldsymbol{\gamma}$ is the influence-weighted average risk aversion. The portfolio has exactly the same composition as the optimal portfolio of a single manager with correct beliefs—the same relative weights across assets—but is scaled by $(\boldsymbol{\phi}'\boldsymbol{\gamma})^{-1}$ rather than by a fixed γ^{-1}. When the influence weights $\boldsymbol{\phi}$ change over time, this scaling factor changes with them, even if no individual member's preferences have changed.

The portfolio excess return is $r_p - r_f = V\boldsymbol{w}'(\boldsymbol{r} - r_f\mathbf{1})$, where $V = (\boldsymbol{\phi}'\boldsymbol{\gamma})^{-1}$ is the inverse of the influence-weighted risk aversion and $\boldsymbol{w} = \boldsymbol{\Sigma}^{-1}(\boldsymbol{\mu} - r_f\mathbf{1})$ is a fixed vector determined by the investment opportunity set. Since both V (through the fluctuating influence weights) and $\boldsymbol{r}$ (through market randomness) are random, the variance of portfolio returns includes additional terms beyond the standard market risk:

$$\sigma_p^2 = \boldsymbol{w}'\boldsymbol{\Sigma}\boldsymbol{w} + \sigma^2(V) + \mathrm{cov}(\boldsymbol{w}'\boldsymbol{r}, V). \tag{3}$$

In a model without oversight risk the variance would be simply $\boldsymbol{w}'\boldsymbol{\Sigma}\boldsymbol{w}$. The additional terms

$$\sigma_{or}^2 = \sigma^2(V) + \mathrm{cov}(\boldsymbol{w}'\boldsymbol{r}, V) \tag{4}$$

represent the variance contribution of oversight risk. The term $\sigma^2(V)$ captures the direct effect of fluctuating committee preferences: over time, changes in the relative importance of each overseer's perspective cause the desired portfolio riskiness to shift, adding volatility to returns. The covariance term captures the interaction between market returns and committee dynamics. If more risk-averse overseers tend to gain influence when portfolio returns are poor—or equivalently, if most overseers tend to become more averse to risk when returns are low, consistent with the empirical evidence of Guiso et al. (2018)—then this covariance is positive and adds further volatility. The intuitive behavior of becoming more conservative after losses has the counterintuitive effect of *increasing* the volatility of portfolio returns.

If the committee's effective risk aversion were constant at some value γ, the expected excess return would be $\gamma^{-1}(\boldsymbol{\mu} - r_f\mathbf{1})'\boldsymbol{\Sigma}^{-1}(\boldsymbol{\mu} - r_f\mathbf{1})$. With fluctuating influence weights, the expected excess return instead depends on $E[V] = E[(\boldsymbol{\phi}'\boldsymbol{\gamma})^{-1}]$. Whether oversight risk raises or lowers expected returns relative to this benchmark depends on whether $E[V]$ is greater or less than γ^{-1}—that is, on the shape of the distribution of effective risk aversion. But regardless of the effect on expected returns, the Sharpe ratio is directly impacted because oversight volatility inflates the denominator:

$$S = \frac{E[V]\,\boldsymbol{w}'(\boldsymbol{\mu} - r_f\mathbf{1})}{\sqrt{\boldsymbol{w}'\boldsymbol{\Sigma}\boldsymbol{w} + \sigma^2(V) + \mathrm{cov}(\boldsymbol{w}'\boldsymbol{r}, V)}}. \tag{5}$$

The denominator grows with oversight volatility, so the Sharpe ratio falls as oversight risk increases. Although risk aversion scales the portfolio without altering relative asset weights at any given moment, the correlation between the scaling factor V and market returns creates a market-timing effect that propagates through all standard performance metrics. When the committee's risk tolerance moves pro-cyclically—scaling the portfolio up in rising markets and down in falling markets—the resulting convex payoff pattern manifests in a linear CAPM regression as slightly lower beta paired with positive alpha, because the curvature of the payoff is absorbed into the intercept rather than the slope. The reverse concave pattern under counter-cyclical risk tolerance produces slightly higher beta and negative alpha. The Information Ratio is similarly affected, because the systematic timing of portfolio scaling generates active return patterns relative to the benchmark. The magnitude of these effects is quantified in the simulations below.

3.2 Belief volatility

Now suppose committee members share a common risk aversion γ but differ in their beliefs about expected returns. Such differences can arise because overseers have different areas of expertise, different exposure to information sources, or different priors about macroeconomic conditions. Let $\boldsymbol{M} = [\boldsymbol{\mu}_1, \ldots, \boldsymbol{\mu}_I]$ be the $N \times I$ matrix of member-specific expected return vectors. The committee's *effective beliefs* are the influence-weighted average $\boldsymbol{y} = \boldsymbol{M}\boldsymbol{\phi}$, and the manager solves

$$\max_{\boldsymbol{x}} \sum_{i=1}^{I} \phi_i \left[\boldsymbol{x}'(\boldsymbol{\mu}_i - r_f\mathbf{1}) - \frac{\gamma}{2}\boldsymbol{x}'\boldsymbol{\Sigma}\boldsymbol{x}\right] = \max_{\boldsymbol{x}} \left[\boldsymbol{x}'(\boldsymbol{y} - r_f\mathbf{1}) - \frac{\gamma}{2}\boldsymbol{x}'\boldsymbol{\Sigma}\boldsymbol{x}\right]. \tag{6}$$

The optimal portfolio is

$$\boldsymbol{x} = \frac{1}{\gamma}\boldsymbol{\Sigma}^{-1}(\boldsymbol{y} - r_f\mathbf{1}). \tag{7}$$

The effective belief vector $\boldsymbol{y}$ has a random component because both the beliefs of individual overseers and the influence weights $\boldsymbol{\phi}$ can change over time. From the portfolio manager's perspective, this randomness has the potential to affect the nature of the portfolio's returns.

Unlike the risk aversion channel, where fluctuations scale the entire portfolio up or down, fluctuations in effective beliefs alter the *composition* of the portfolio—shifting weights across asset classes in ways that may or may not reflect genuine information about future returns.

Assume that the random vector $\boldsymbol{y}$ has mean $\boldsymbol{\mu}_y$ and covariance matrix $\boldsymbol{\Sigma}_y$. The variation in portfolio weights is $\boldsymbol{\Sigma}_{\boldsymbol{x}} = \gamma^{-2}\boldsymbol{\Sigma}^{-1}\boldsymbol{\Sigma}_y\boldsymbol{\Sigma}^{-1}$. This variation in portfolio weights sits on top of any variation due to changing market conditions, and it directly inflates the variance of portfolio returns.

The interaction between effective beliefs and realized returns is critical for understanding whether oversight adds or destroys value. Let $\boldsymbol{\Sigma}_{ry} = \mathrm{cov}(\boldsymbol{r}, \boldsymbol{y})$ capture this relationship. The expected excess return is

$$E[r_p - r_f] = \frac{1}{\gamma}(\boldsymbol{\mu}_y - r_f\mathbf{1})'\boldsymbol{\Sigma}^{-1}(\boldsymbol{\mu} - r_f\mathbf{1}) + \frac{1}{\gamma}\,\mathrm{tr}(\boldsymbol{\Sigma}^{-1}\boldsymbol{\Sigma}_{ry}). \tag{8}$$

The first term is the expected return that would obtain if effective beliefs were constant at $\boldsymbol{\mu}_y$. When beliefs are correct on average ($\boldsymbol{\mu}_y = \boldsymbol{\mu}$), this term equals the expected return of a portfolio that mimics the market portfolio. The second term captures the relationship between beliefs and returns. When the committee tends to be optimistic about assets that subsequently perform well—that is, when the entries of $\boldsymbol{\Sigma}_{ry}$ are positive—this term adds to expected return, reflecting genuine committee skill. When the committee's deliberation process tends to produce beliefs that are negatively correlated with future returns, this term subtracts from expected return.

The next section presents simulations that quantify the performance consequences of belief and risk aversion volatility under different assumptions about the skill and noise in committee beliefs.

4 Simulations

We calibrate the model to monthly U.S. equity sector data and simulate three forms of oversight risk. Monthly sector returns and market-capitalization weights are constructed from CRSP data spanning 18 NAICS sectors from March 1970 to December 2025, and the monthly risk-free rate is taken from the Fama-French five-factor data file.[1] The covariance matrix of excess returns is estimated at each month using a 60-month rolling window.

The market-implied expected excess returns are computed at each month t using the Black-Litterman reverse-optimization formula:

$$\boldsymbol{\mu}_t - r_{f,t} = \gamma\, \boldsymbol{\Sigma}_t\, \boldsymbol{x}_{m,t}, \tag{9}$$

where $\boldsymbol{x}_{m,t}$ is the vector of market-capitalization weights for the 18 sectors, $\boldsymbol{\Sigma}_t$ is the rolling-window covariance matrix, $r_{f,t}$ is the monthly risk-free rate, and $\gamma = 1$ is the market risk-aversion parameter. The resulting vector $\boldsymbol{\mu}_t$ serves as the benchmark "true" expected returns against which committee beliefs are measured.

All three simulation exercises follow the same Monte Carlo protocol: for each parameter configuration, 10,000 independent paths through the sample period are drawn. In each draw, optimal portfolio weights are computed at every month in the sample using $\boldsymbol{x}_t = \gamma_t^{-1}\boldsymbol{\Sigma}_t^{-1}\boldsymbol{y}_t$, and the resulting time series of portfolio returns is used to calculate the Sharpe ratio, Jensen's alpha, and Information Ratio relative to the value-weighted market portfolio.

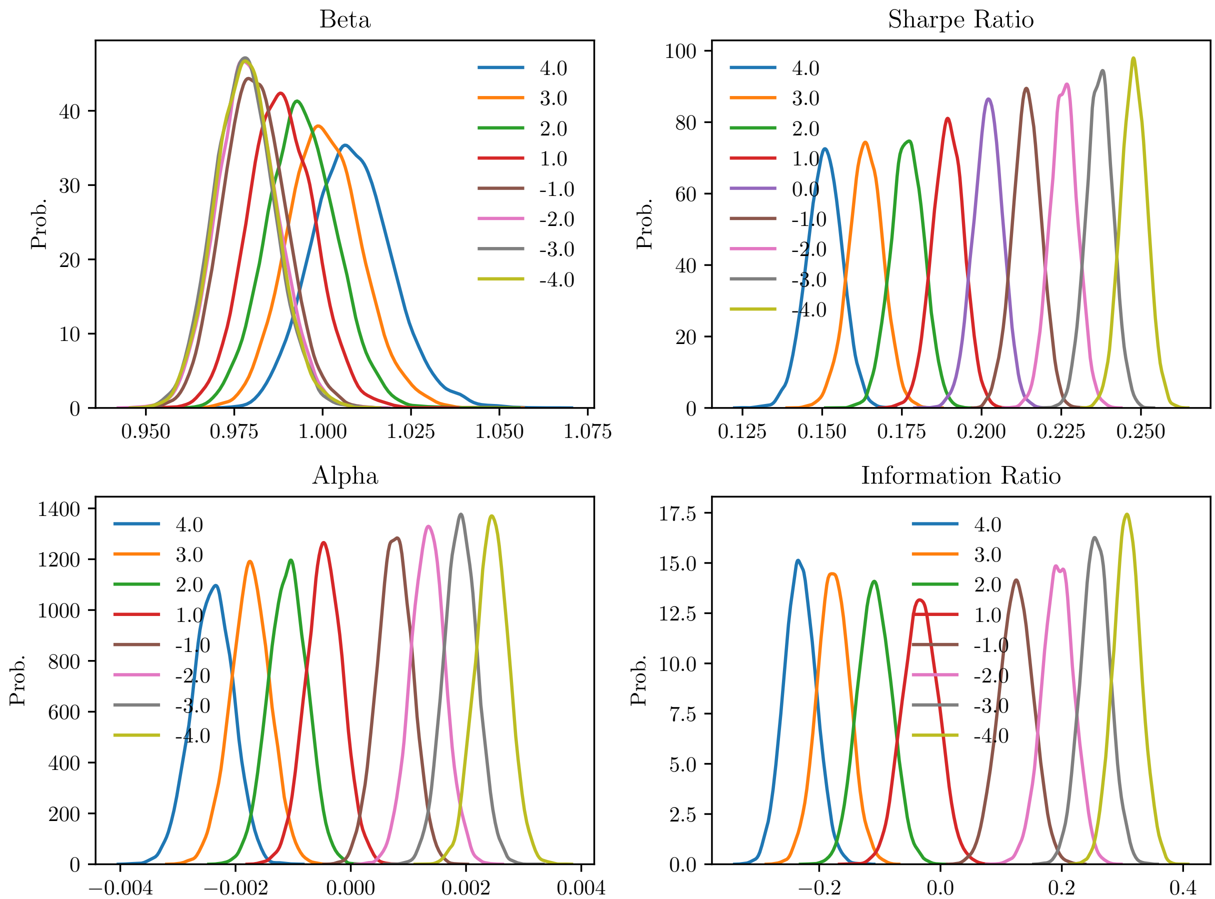

Figure 1: Portfolio performance across nine levels of risk aversion sensitivity (k). Each curve is the distribution of per-draw outcomes under a fixed sensitivity of risk aversion to market returns, ranging from strongly pro-cyclical risk tolerance ($k = -4$, meaning risk aversion falls sharply when markets rise) to strongly counter-cyclical ($k = +4$). All four metrics vary monotonically with k: pro-cyclical risk tolerance is associated with higher Sharpe ratios, positive alpha, positive Information Ratios, and beta slightly below one, while counter-cyclical risk tolerance produces the opposite pattern.

4.1 Risk Aversion Volatility

To simulate time-varying effective risk aversion, the committee's effective risk aversion $\gamma_t \equiv \boldsymbol{\phi}_t' \boldsymbol{\gamma}$ is modeled as a logistic function of the current month's excess market return:

$$\gamma_t = 0.5 + \frac{1.0}{1 + \exp(-k\,(r_{m,t} - \bar{r}_m))} + \varepsilon_t, \qquad \varepsilon_t \sim N(0,\, 0.01), \tag{10}$$

where $r_{m,t}$ is the excess market return, $\bar{r}_m$ is its sample mean, k is the sensitivity parameter, and ε_t is idiosyncratic noise. Positive k implies that risk aversion rises with market returns (counter-cyclical risk tolerance), while negative k implies pro-cyclical risk tolerance. In all periods the committee's beliefs about expected returns are consistent with market beliefs and are equal to $\boldsymbol{\mu}_t$; only the risk aversion fluctuates. The simulation is run for $k \in \{-4, -3, -2, -1, 0, 1, 2, 3, 4\}$.

Figure 1 displays the results. The Sharpe ratio panel tells the central story. The mean Sharpe ratio decreases monotonically with k: it is highest under strongly pro-cyclical risk tolerance ($k = -4$, mean Sharpe ≈ 0.248) and lowest under strongly counter-cyclical risk tolerance ($k = +4$, mean Sharpe ≈ 0.151), with the no-sensitivity baseline ($k = 0$) at approximately 0.202. Pro-cyclical risk tolerance—scaling the portfolio up in rising markets and down in falling markets—captures more of the equity risk premium when it is being realized, while counter-cyclical risk tolerance forgoes upside by deleveraging in advancing markets.

Table 1 presents the corresponding summary statistics. All four performance metrics vary systematically with k, confirming that the correlation between effective risk aversion and market returns propagates through every standard diagnostic. The pattern is consistent with the convex-versus-concave payoff structure predicted by the model. Under pro-cyclical

[1] The NAICS sectors "Public Administration" and "Other Services (except Public Administration)" are excluded because there are spells during the sample period where no data is available for publicly traded firms in these sectors. The sample is restricted to months for which all 18 sectors have complete return data.

risk tolerance (negative k), the portfolio creates a convex payoff by scaling up in rising markets and down in falling markets; in a linear CAPM regression this convexity appears as beta slightly below one paired with positive alpha and a strongly positive Information Ratio. Under counter-cyclical risk tolerance (positive k), the concave payoff yields beta near or slightly above one, negative alpha, and a negative Information Ratio. The monotonic decrease in Sharpe ratio, alpha, and Information Ratio with k reflects the performance advantage of timing that is aligned with the equity risk premium.

Table 1: Portfolio performance metrics by effective risk aversion sensitivity (k). Mean and standard deviation across 10,000 Monte Carlo draws, March 1970–December 2025.

k	statistic	Beta	Sharpe Ratio	Alpha	Information Ratio
-4.0	mean	0.9783	0.2478	0.0025	0.3077
	std	0.0085	0.0041	0.0003	0.0229
-3.0	mean	0.9779	0.2370	0.0019	0.2554
	std	0.0085	0.0042	0.0003	0.0241
-2.0	mean	0.9788	0.2257	0.0013	0.1937
	std	0.0086	0.0044	0.0003	0.0262
-1.0	mean	0.9807	0.2142	0.0008	0.1242
	std	0.0087	0.0045	0.0003	0.0282
0.0	mean	1.0000	0.2020	0.0000	0.0000
	std	0.0000	0.0046	0.0000	0.0000
1.0	mean	0.9884	0.1894	-0.0005	-0.0329
	std	0.0093	0.0049	0.0003	0.0296
2.0	mean	0.9942	0.1767	-0.0011	-0.1084
	std	0.0099	0.0052	0.0003	0.0282
3.0	mean	1.0007	0.1636	-0.0017	-0.1765
	std	0.0105	0.0054	0.0003	0.0269
4.0	mean	1.0084	0.1506	-0.0024	-0.2317
	std	0.0114	0.0055	0.0004	0.0263

These results carry a notable practical implication. The direction of the correlation between effective risk aversion and market returns matters: pro-cyclical risk tolerance improves performance relative to the no-sensitivity baseline, while counter-cyclical risk tolerance degrades it. A committee that becomes more conservative in rising markets—for example, by taking profits or reducing equity exposure after strong performance—forgoes upside and earns a lower Sharpe ratio, negative alpha, and a negative Information Ratio. The performance gap is substantial: the Sharpe ratio under strongly counter-cyclical risk tolerance ($k = +4$, mean ≈ 0.151) is roughly 25% below the no-sensitivity baseline ($k = 0$, mean ≈ 0.202), and the Information Ratio swings from $+0.31$ at $k = -4$ to -0.23 at $k = +4$. The four metrics together provide a rich diagnostic: the sign and magnitude of alpha, beta's deviation from one, and the Information Ratio jointly reveal not only whether risk aversion volatility is present but also whether the committee's effective risk tolerance moves pro-cyclically or counter-cyclically with market returns.

4.2 Belief Volatility Without Skill

The next experiment isolates the performance cost of belief noise in the absence of directional skill. The committee's effective beliefs at month t are modeled as

$$\boldsymbol{y}_t = \boldsymbol{\mu}_t + \boldsymbol{\varepsilon}_t, \qquad \boldsymbol{\varepsilon}_t \sim N(\mathbf{0}, \sigma^2 \boldsymbol{I}), \tag{11}$$

where σ is the belief noise standard deviation. Noise is drawn independently across sectors and time periods, so beliefs are correct on average but fluctuate randomly around the market-implied expected returns each month. The effective risk aversion is fixed at $\gamma_t = 1$ throughout. The simulation is run for $\sigma \in \{0.02, 0.04, 0.06, 0.08, 0.10, 0.20\}$.

Figure 2 shows the results. The Sharpe ratio panel is the most revealing. The six distributions shift steadily leftward as belief noise increases. At the lowest noise level ($\sigma = 0.02$), the committee holds a portfolio close to the market optimum and achieves a mean Sharpe ratio of approximately 0.23. As noise rises to $\sigma = 0.20$, the mean Sharpe ratio falls to approximately 0.05—a decline of about 80% from the low-noise benchmark. This monotone decline quantifies the cost of governance churn: a committee whose views are correct on

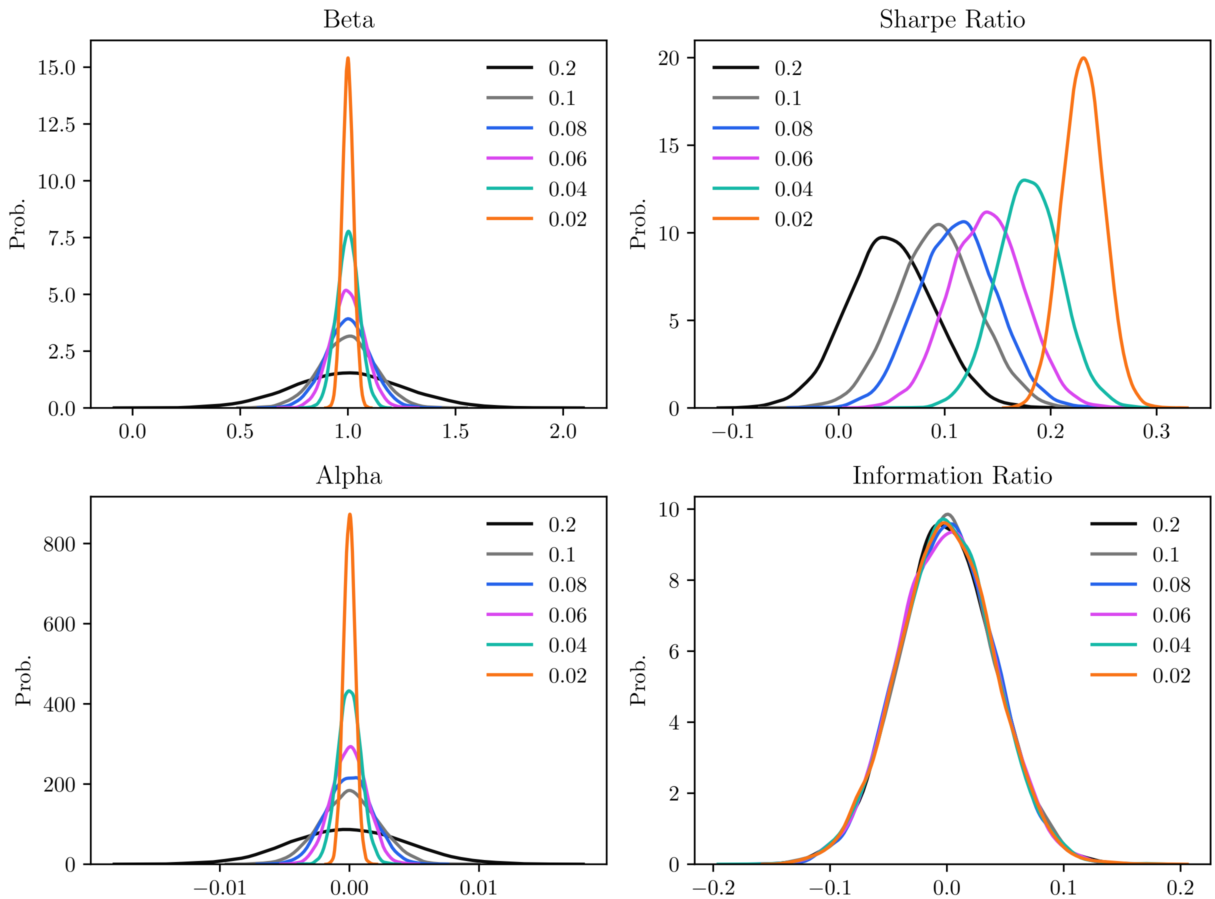

Figure 2: Portfolio performance across six levels of belief noise (σ). The committee has no directional skill: effective beliefs are correct on average but fluctuate randomly around the true market-implied expected returns. As belief noise increases, the Sharpe ratio declines monotonically, the distributions of Jensen's alpha and beta widen, and the Information Ratio distribution is largely unchanged across noise levels.

average but inconsistent from month to month will substantially underperform a committee with stable beliefs, even if neither committee has any directional insight into future returns.

The beta panel shows how this churn transmits into portfolio holdings. At low noise levels, beta is tightly concentrated near 1.0, reflecting a portfolio that closely tracks the market. At high noise levels, the beta distribution widens substantially—a standard deviation of 0.26 at $\sigma = 0.20$ versus 0.03 at $\sigma = 0.02$—even though the mean beta stays near 1.0. The large transient tilts induced by noisy beliefs add return volatility without adding expected return.

Table 2 confirms the key pattern numerically. Jensen's alpha is centered at zero for every noise level—the committee generates no systematic active return when beliefs match market beliefs on average—but the distribution of alpha widens substantially as noise increases (standard deviation rising from near zero at $\sigma = 0.02$ to 0.005 at $\sigma = 0.20$), reflecting the growing dispersion of realized active returns across draws. The Information Ratio distributions are virtually indistinguishable across noise levels, with means near zero and standard deviations around 0.041 in every case. The insensitivity of the Information Ratio to belief noise reflects the fact that noise inflates both active returns and tracking error in proportion, leaving their ratio unchanged.

Table 2: Portfolio performance metrics by belief noise standard deviation (σ). Mean and standard deviation across 10,000 Monte Carlo draws, March 1970–December 2025.

σ	Statistic	Beta	Sharpe Ratio	Alpha	Information Ratio
0.02	mean	1.000	0.231	−0.000	−0.000
	std	0.026	0.020	0.000	0.041
0.04	mean	1.000	0.180	−0.000	−0.001
	std	0.051	0.030	0.001	0.041
0.06	mean	1.000	0.140	−0.000	−0.000
	std	0.077	0.035	0.001	0.041
0.08	mean	1.000	0.112	−0.000	−0.000
	std	0.102	0.037	0.002	0.041
0.10	mean	0.999	0.093	0.000	0.000
	std	0.127	0.039	0.002	0.041
0.20	mean	1.000	0.049	−0.000	−0.000
	std	0.259	0.041	0.005	0.041

These results sharply distinguish the effect of belief noise from that of systematic bias. A committee whose beliefs are correct on average but unstable over time will nonetheless underperform a committee with stable beliefs, and the performance gap grows with the magnitude of the instability. This finding qualifies the common view that diverse perspectives and active debate improve investment outcomes. Charness & Sutter (2012) find that groups can outperform individuals in strategic settings, and the production-gain literature surveyed by Kerr & Tindale (2004) documents conditions under which group deliberation adds genuine value. But these process gains require that the deliberation process aggregate information efficiently. The behavioral foundations in Section 2 identify several reasons why investment committees may fail to do so: informational and reputational cascades cause members to defer to early or high-status speakers rather than contributing independent signals (Sunstein & Hastie 2015), status hierarchies allocate influence based on perceived standing rather than information quality (Berger et al. 1972), and coalition dynamics can abruptly reallocate effective control when internal alliances shift (Eisenhardt & Bourgeois 1988). When these forces cause beliefs to fluctuate from meeting to meeting without being anchored to genuine information about future returns, the resulting portfolio weight variance imposes a cost that grows monotonically with the magnitude of the instability.

4.3 The Skill–Noise Tradeoff

The preceding simulation varies beliefs around market-derived beliefs in order to isolate the extent to which effective belief noise can alter portfolio performance metrics. However, one might argue that committees bring diverse sets of information to the investment process that, when properly assimilated, can increase potential returns to the portfolio. That is, if the investment committee effectively increases the skill of the portfolio manager, then it can mitigate the negative impact of belief noise, reducing the overall impact of oversight risk. The following simulation seeks to understand this interaction.

In this simulation, the committee's beliefs at month t are constructed to have some ability to predict next period's market-implied expected returns. Managers whose beliefs are better aligned with future risk premia can tilt portfolio weights toward sectors that will subsequently outperform, earning higher risk-adjusted returns.

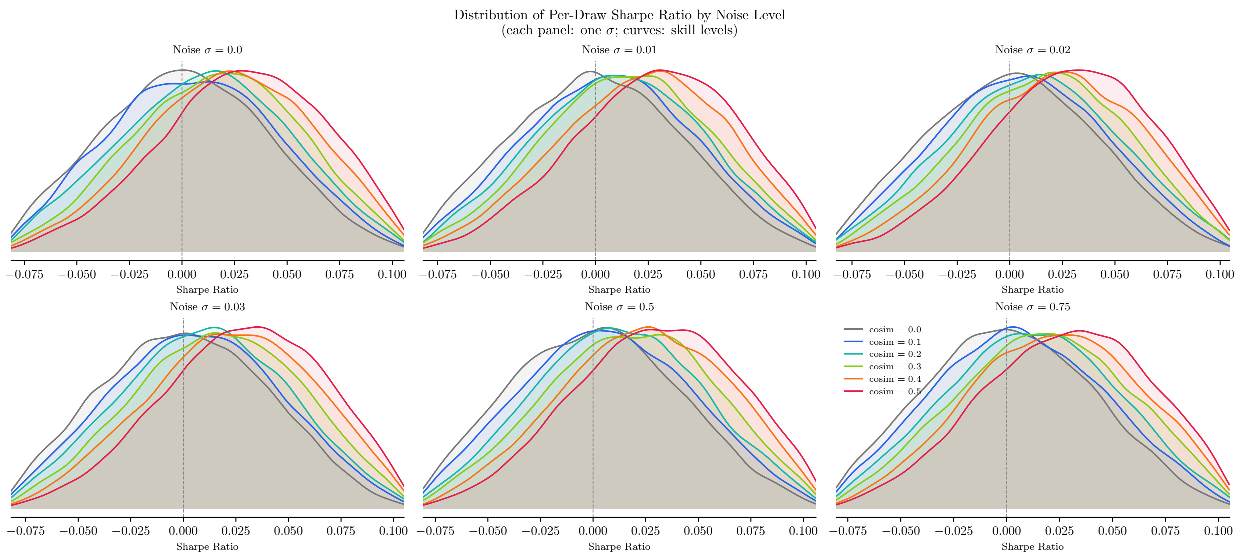

Figure 3: The skill–noise tradeoff in Sharpe ratio performance. Each panel shows the distribution of per-draw Sharpe ratios across six levels of committee skill (cosine similarity between effective beliefs and next-period market-implied returns), for a given level of belief noise. Higher skill (warmer colors) consistently shifts the distribution rightward within every panel, while higher noise (moving from left to right across panels) widens the distribution without shifting its center.

Skill is modeled in these simulations by the similarity between the committee's effective beliefs $\boldsymbol{y}_t$ and the next period's market-implied expected returns $\boldsymbol{\mu}_{t+1}$. The cosine similarity between these two vectors is a natural measure of skill, as it captures the degree to which the committee's beliefs are directionally aligned with future returns. The simulation constructs effective beliefs that have a specified target cosine similarity $\bar{\rho}$ with next period's market-implied returns, while also incorporating belief noise.

This is done by starting with the next period's market-implied expected returns $\boldsymbol{\mu}_{t+1}$ and adding a noise component that is orthogonal to it. The resulting effective belief vector is a weighted average of the true future returns and the noise, with the weights determined by the target cosine similarity and a random shock.

The simulation is run for average skill (cosine similarities) $\bar{\rho} \in \{0.0, 0.1, 0.2, 0.3, 0.4, 0.5\}$ and noise around that skill of $\sigma_\rho \in \{0.00, 0.01, 0.02, 0.03, 0.50, 0.75\}$.

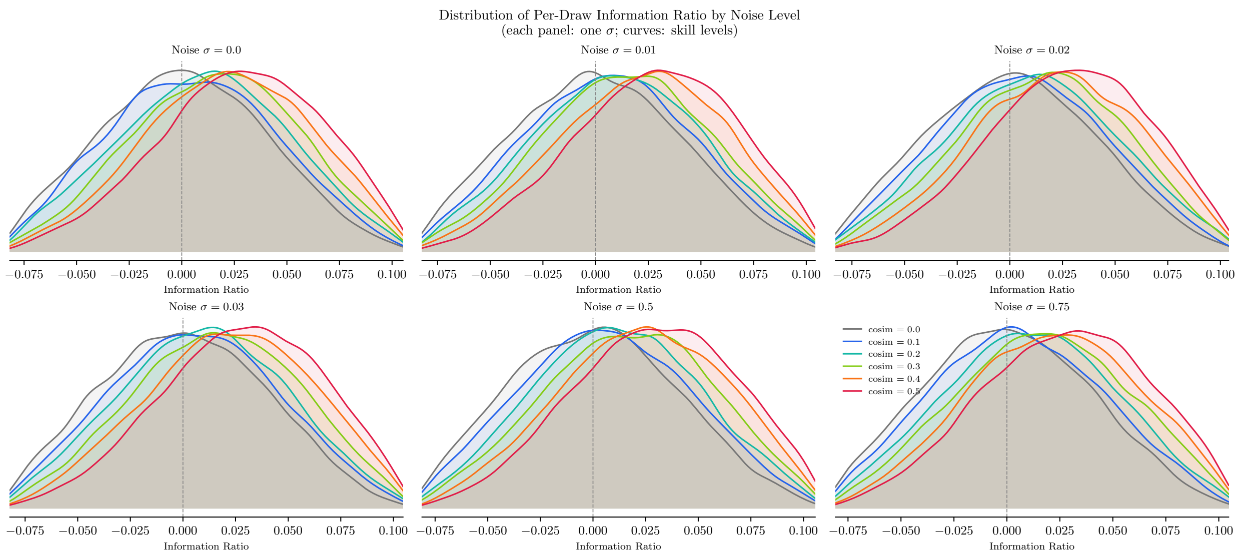

Figure 4: The skill–noise tradeoff in Information Ratio performance. The pattern mirrors the Sharpe ratio results: skill raises the center of the performance distribution while noise broadens it. The parallel behavior of Sharpe ratio and Information Ratio confirms that the performance benefit of skill arises from genuine active return rather than from mechanical leverage on market exposure.

Figures 3 and 4 display the results for the Sharpe ratio and Information Ratio respectively. Two patterns emerge consistently.

Skill raises the level of performance. Across every noise level, higher cosine similarity shifts the performance distribution to the right. Each 0.1 increase in cosine similarity raises the mean Sharpe ratio by approximately 0.006. Moving from a committee with no directional skill (cosine similarity = 0.0) to one with significant alignment (cosine similarity = 0.5) raises the mean Sharpe ratio by approximately 0.030, with the Information Ratio showing a nearly identical gradient. This is the "alpha" of genuine committee skill: committees whose effective beliefs are better aligned with future returns systematically earn higher risk-adjusted returns. This simulation result provides a structural interpretation of the empirical finding of Binfare et al. (2017) that endowment committees with greater investment expertise earn significantly higher returns—in our framework, expertise raises the cosine similarity between effective beliefs and future returns, shifting the entire performance distribution to the right.

Noise widens the distribution without shifting its center. Higher belief noise does not lower the average level of performance across skill levels—the ranking of skill curves is preserved at every noise level. But it substantially widens the distribution of outcomes around whatever average is implied by the committee's skill level. A high-skill committee that also has high belief noise will earn high returns on average, but will also experience much greater variability in realized performance from one evaluation period to the next.

This distinction between skill and noise effects is a direct simulation counterpart to the theoretical result that oversight volatility enters the variance of portfolio returns but not the mean, absent correlation between effective beliefs and realized returns. Together, the two patterns highlight a key structural feature of investment committee quality: skill and noise operate on different moments of the performance distribution. A committee with moderate skill and low noise is a fundamentally different governance proposition from a committee with the same average skill level but high noise. The first will reliably produce above-benchmark performance; the second will produce the same average but with substantially larger deviations. For institutions subject to drawdown constraints, investment policy restrictions, or peer-relative performance benchmarking, this "second-moment" effect of oversight risk is of first-order concern, as it is the tail outcomes that trigger even larger governance consequences like manager termination, strategy change, and reduced discretion. These changes have the potential to add even more noise to the process.

5 Implications for Practice

The preceding analysis identifies specific mechanisms through which committee dynamics degrade portfolio performance and provides metric-level signatures that distinguish different sources of degradation. This section translates those findings into actionable guidance for those who design committee structures—boards, chief investment officers, consultants—and for those who chair or facilitate committee meetings.

5.1 Detecting belief noise

When portfolio performance disappoints, the natural response is to question the quality of investment ideas. But a significant share of institutional performance drag may arise from governance instability rather than a fundamental lack of market insight. The belief volatility simulations show that when a committee's effective beliefs fluctuate from meeting to meeting—even around a correct average—the Sharpe ratio declines monotonically while the Information Ratio remains essentially unchanged (Table 2). This pattern arises because belief noise inflates both active returns and tracking error in proportion, leaving their ratio constant. A portfolio whose Sharpe ratio has deteriorated relative to peers or to its own history, but whose Information Ratio is stable, is exhibiting the signature of governance-driven churn rather than a skill deficit. The appropriate response is not to replace the manager or alter the investment strategy, but to stabilize the committee's deliberation process.

5.2 Detecting risk aversion volatility

The risk aversion simulations show that changes in the committee's effective risk tolerance erode the Sharpe ratio (Table 1), and that all four standard metrics—beta, alpha, Sharpe ratio, and Information Ratio—vary systematically with the correlation between risk aversion and market returns. Unlike the belief volatility channel, where metric distributions widen symmetrically, risk aversion volatility shifts the *mean* of beta and other metrics in a direction that depends on whether the committee's risk tolerance is pro-cyclical or counter-cyclical. A declining Sharpe ratio accompanied by a directional shift in beta—rather than a symmetric widening of its distribution—therefore points toward risk aversion volatility as the primary source of performance drag. The relevant governance intervention is to anchor the committee's risk tolerance through explicit policy rather than allowing it to drift with committee composition or market sentiment.

5.3 Distinguishing skill from noise

The skill–noise tradeoff simulations (Figures 3 and 4) reveal that skill and noise operate on different moments of the performance distribution: skill shifts the distribution's center while noise widens it. A committee that generates highly variable performance from one evaluation period to the next—alternating between strong and weak results—may have genuine investment skill that is being obscured by governance noise. Before concluding that such a committee lacks skill, one should investigate whether process instability is inflating the variance of outcomes around what may be a favorable mean. This distinction is especially consequential for institutions subject to drawdown constraints, investment policy restrictions, or peer-relative performance benchmarking, because the tail outcomes generated by high governance noise are what trigger more substantive governance consequences like manager termination, strategy change or reduced discretion. These large changes have the potential to add even more noise to the process.

5.4 Committee composition and stability

The belief noise results demonstrate that a committee with consistent views outperforms one that is correct on average but inconsistent from period to period—even when both committees have the same average accuracy. Frequent turnover in committee membership is itself a source of oversight risk, because each change in composition alters the vector of influence weights and therefore the committee's effective beliefs and risk preferences. The expectation states framework of Berger et al. (1972) explains why these changes can be so disruptive: when a new member joins, the committee's status hierarchy reorganizes around the new member's external status, redistributing effective influence in ways that may have little to do with investment skill. Similarly, the coalition dynamics documented by Eisenhardt & Bourgeois (1988) suggest that turnover can trigger realignment of internal alliances, producing discrete shifts in whose views prevail. Consistency has investment value: predictable biases can be managed or offset, while unpredictable fluctuations cannot. Where turnover is necessary, staggered terms—in which only a fraction of membership changes at any time—can dampen the resulting fluctuations in effective beliefs. Scherer (2024) develops a model of optimal committee design that can address some of these composition questions: his algorithmic-consensus approach—in which members independently submit portfolios that are rescaled to equal tracking error and then averaged—effectively eliminates the time-varying influence weights that generate oversight risk in our framework. By enforcing equal weighting mechanically rather than through deliberation, the Scherer design can suppress both the risk aversion channel (through tracking-error rescaling) and the belief noise channel (through averaging of independent forecasts rather than group discussion).

5.5 Explicit risk tolerance governance

Investment committees often devote substantial deliberation to expected returns and comparatively little to explicitly governing the portfolio's risk posture. Yet the simulations show that the Sharpe ratio cost of risk aversion misalignment can rival or exceed the cost of moderate belief errors. The cascade mechanisms identified by Sunstein & Hastie (2015) make this channel especially difficult to control through informal deliberation: when a senior member expresses heightened concern about risk, informational and reputational cascades can rapidly propagate that concern through the committee, shifting the group's effective risk tolerance without any member explicitly voting to change it. Committees should establish an explicit, quantified risk tolerance—expressed, for example, as a target volatility range or a maximum drawdown threshold—that is reviewed and reaffirmed on a fixed schedule rather than allowed to drift with the emotional tenor of the most recent market environment. The Swensen (2009) model of disciplined policy portfolio management can be understood in part as a governance mechanism for controlling exactly this channel. Guiso et al. (2018) document that individual risk aversion increases after experiencing investment losses, suggesting that committees are especially vulnerable to pro-cyclical tightening of risk tolerance during drawdowns—precisely when the opportunity cost of excessive conservatism is highest.

Although our model focuses on portfolio construction, the risk aversion channel can also propagate through payout policy—a margin that is formally outside the portfolio optimization problem but is often governed by the same committee. When a high-status member becomes concerned about portfolio risk, the resulting cascade may not stop at portfolio weights: the committee may also reduce the fund's spending rate or payout ratio as an additional expression of caution. For endowments, foundations, and other mission-driven institutions, this response sacrifices the organization's primary objectives—the programs, scholarships, grants, or operations that the fund exists to support—in order to preserve corpus value. The cost of such a payout reduction is borne not by the portfolio but by the organization's beneficiaries, and it is therefore invisible in standard performance metrics. Moreover, payout reductions triggered by market stress are pro-cyclical in the same sense as the risk aversion channel studied here: they curtail spending precisely when the organization's operating environment may also be under strain (e.g., increased demand for services during economic downturns). An explicit payout policy—specifying, for example, a fixed percentage of a multi-year moving average of portfolio value—serves the same governance function for spending that a target volatility range serves for portfolio construction: it insulates the payout decision from the committee's time-varying risk sentiment and ensures that the organization's mission is not held hostage to the same status-driven cascades and coalition dynamics that can distort portfolio choice.

5.6 Separating belief formation from portfolio construction

A structural remedy for the deliberation failures discussed in Section 2—production blocking, shared information bias, informational and reputational cascades, and the status-driven influence hierarchies documented by Berger et al. (1972)—is to separate the stage at which beliefs are formed from the stage at which portfolio weights are chosen. If the committee's role is explicitly limited to approving a set of expected return assumptions—elicited through a structured process—while portfolio construction is delegated to the manager subject to those assumptions, the committee's influence is channeled through a dimension (beliefs) that is observable and auditable, rather than through ad hoc pressure on individual position sizes. This separation also makes it easier to detect belief noise: the committee's approved assumptions create a written record that can be compared across meetings to assess consistency.

5.7 Reducing production blocking through pre-meeting elicitation

Nijstad et al. (2003) show that the cognitive costs of production blocking—disrupted organization of thought and reduced flexibility—are driven by the delays and unpredictability inherent in turn-taking. A committee chair can directly mitigate these effects by collecting members' views in writing before the meeting begins. When each member submits an independent assessment of the key questions on the agenda—expected returns by asset class, risk concerns, proposed tilts—the chair can identify areas of agreement and disagreement before discussion starts, ensuring that the meeting focuses on genuinely contested issues rather than on the views that happen to be voiced first. This approach is closely related to the Delphi method (Dalkey & Helmer 1963), which was designed specifically to aggregate expert judgment while minimizing the distortions of face-to-face interaction.

5.8 Surfacing private information

The shared information bias documented by Stasser & Titus (1985) and Reimer et al. (2010)—the tendency for groups to over-discuss commonly held information and neglect the unique signals that only individual members possess—is particularly costly in an investment committee, where the value of the committee lies in part in the diversity of its members' information sets. Chairs can counteract this bias by explicitly structuring the agenda to allocate time for each member to present information that others are unlikely to have encountered. Gigone & Hastie (1993) show that group judgments are disproportionately driven by information that members share in common, suggesting that deliberate procedural effort is required to ensure that private signals are voiced and weighted appropriately. The pre-meeting written elicitation described above also helps, because members record their private signals before learning what others think, reducing the social incentive to converge on shared information.

5.9 Managing need for closure

The need-for-closure dynamics described by Kruglanski & Webster (1996) and Kruglanski et al. (2006)—seizing on an early position and freezing against disconfirming evidence—are most likely to arise under time pressure, fatigue, and ambient uncertainty. Committee chairs can manage these triggers directly. Scheduling sufficient meeting time reduces the pressure to converge prematurely. Placing the most consequential agenda items early in the meeting, when cognitive resources are freshest, reduces the fatigue channel. Explicitly assigning a dissent role—a member whose job is to articulate the strongest case against the emerging consensus—provides a structural counterweight to the social pressure toward uniformity that Kruglanski & Webster (1991) document. Sunstein & Hastie (2015) provide a comprehensive treatment of procedural reforms designed to counteract premature closure and other group deliberation failures.

5.10 Dampening group polarization

Moscovici & Zavalloni (1969) and Isenberg (1986) document that group deliberation tends to shift collective positions toward more extreme versions of the pre-discussion average. In an investment committee, this means that a group of members who are individually mildly bullish on an asset class may emerge from discussion strongly bullish—an amplification that introduces the kind of belief variance the model shows to be costly. Pre-meeting elicitation serves as a partial remedy: when members record their views before discussion, the chair can anchor post-discussion decisions to the distribution of pre-discussion views and flag cases where the group has moved substantially beyond the pre-discussion average. Larrick & Soll (2006) show that simple averaging of independent judgments is a surprisingly effective aggregation rule—one that eliminates the polarization channel entirely—yet decision-makers systematically underestimate the value of this approach. A chair who computes and shares the simple average of members' pre-meeting views provides a powerful anchor against polarization-driven drift.

5.11 Mitigating informational and reputational cascades

Sunstein & Hastie (2015) show that group deliberation is vulnerable to two cascade mechanisms that can propagate errors through a committee. In informational cascades, members who speak later rationally discount their own private signals in favor of views already expressed, so that the committee's conclusion is disproportionately shaped by whoever speaks first. In reputational cascades, members withhold dissenting views to avoid the social cost of opposing the emerging consensus. Both mechanisms cause the committee's effective influence weights to shift endogenously toward early and high-status speakers, generating exactly the kind of time-varying influence that our model shows to be costly.

Several procedural interventions can attenuate these cascades. Randomizing the order in which members speak—or, equivalently, rotating who presents first across meetings—prevents any single member from systematically anchoring the discussion. The pre-meeting written elicitation described above is especially effective against informational cascades, because members commit to views before learning what others think, preserving the independence of their signals. To counteract reputational cascades, chairs can establish norms that explicitly reward dissent: publicly acknowledging the value of contrarian views, tracking whether minority opinions subsequently prove correct, and ensuring that disagreement carries no implicit career or social penalty. Anonymous polling on key questions—for example, each member's preferred risk tolerance or sector allocation—can surface the true distribution of views without requiring anyone to publicly oppose a dominant voice.

5.12 Managing status-driven influence hierarchies

Expectation states theory (Berger et al. 1972) predicts that committee members rapidly form influence hierarchies based on status characteristics—seniority, professional credentials, institutional rank, and demographic attributes—that may have little relationship to investment skill. These hierarchies are self-reinforcing: high-status members receive more opportunities to contribute, their contributions are evaluated more favorably, and they prevail more often in disagreements, regardless of the quality of their views. The result is a set of influence weights that reflects social standing rather than information quality.

Chairs can counteract status-driven hierarchies through deliberate process design. Structured turn-taking that allocates equal speaking time to each member prevents high-status members from dominating by volume. Requiring members to justify views with specific evidence—data, analysis, or identified signals—shifts evaluation criteria from "who said it" to "what supports it." The practice of eliciting independent views first and then using the discussion period to probe disagreements, with the chair responsible for ensuring that low-status members' views receive the same scrutiny as high-status members' views, can help mitigate the impact of status-driven influence.

5.13 Managing coalition dynamics

Eisenhardt & Bourgeois (1988) find that political behavior in top management teams—alliance formation, information withholding, and behind-the-scenes lobbying—is associated with poorer decision quality and slower decision-making. In investment committees, coalitions can form around shared views on asset allocation, loyalty to a particular consultant or investment manager, or broader institutional power struggles. When coalitions shift—through changes in external relationships, institutional reorganizations, or simply the arrival of a persuasive new member—the committee's effective beliefs and risk preferences can change abruptly, corresponding to the discrete jumps in influence weights that our model shows to be especially costly.

Several governance practices can reduce the influence of coalition dynamics on portfolio decisions. Transparency in decision rationale—requiring that every material portfolio decision be accompanied by a written justification referencing the committee's stated investment beliefs—makes it harder for coalitions to steer decisions through informal channels. Conflict-of-interest policies that require members to disclose relationships with external managers, consultants, or other parties whose interests may be affected by committee decisions can surface hidden alignments. Periodic external reviews of committee process—conducted by an independent governance consultant or the organization's audit function—can identify patterns of coalition behavior that insiders may be too close to recognize. Finally, the separation of belief formation from portfolio construction described above limits the scope of political influence: even if a coalition succeeds in shaping the committee's return assumptions, the translation of those assumptions into portfolio weights remains with the manager, providing a buffer against politically motivated position-level interference.

5.14 Monitoring governance consistency

The model's finding that consistency itself has investment value suggests that committees should track the stability of their own decision processes over time. Recording pre-meeting and post-meeting views, tracking the evolution of approved risk tolerance and return assumptions, and periodically reviewing the dispersion of these quantities across meetings provides a direct measure of governance noise. A committee whose views are stable from meeting to meeting, even if modestly biased, will tend to produce better risk-adjusted performance than a committee that oscillates. This is a direct application of the noise-versus-bias framework of Kahneman et al. (2021) to the specific context of investment governance.

6 Conclusion

This paper develops a model of institutional portfolio management in which investment decisions are subject to oversight by a committee with heterogeneous beliefs and risk preferences, and characterizes the performance consequences of the resulting *oversight risk*. The analysis identifies two distinct channels through which oversight dynamics degrade portfolio performance, each with a different signature in standard performance metrics.

The first channel—risk aversion volatility—affects all standard performance metrics through a market-timing channel. Pro-cyclical risk tolerance improves the Sharpe ratio and generates positive alpha, in periods when low returns tend to be mean-reverting. In these periods, counter-cyclical risk tolerance degrades both the Sharpe ratio and alpha.

The second channel—belief volatility—also erodes the Sharpe ratio, but does so through inflated portfolio weight variance rather than through risk tolerance variation. Even when beliefs are correct on average, period-to-period fluctuations reduce the Sharpe ratio monotonically while leaving the Information Ratio unchanged and widening the alpha and beta distributions.

These differential fingerprints provide a diagnostic framework for institutional investors seeking to identify and remediate the sources of governance-related performance drag. Section 2 identifies specific behavioral and organizational mechanisms that generate the belief and preference fluctuations we model: production blocking and shared information bias distort the aggregation of committee members' private signals; need for closure and group polarization push collective views toward premature or extreme positions; informational and reputational cascades cause influence to concentrate on early and high-status speakers; status-driven hierarchies allocate influence based on perceived standing rather than information quality; and coalition dynamics produce discrete shifts in effective control when internal alliances realign. Section 5 translates these mechanisms into actionable governance guidance—from pre-meeting elicitation and structured turn-taking to conflict-of-interest disclosure and prediction record keeping.

Several avenues for future research include the following. The formal model does not endogenize the deliberation failures discussed here or quantify their individual contributions to oversight risk. Empirical work linking observable committee characteristics—turnover rates, status composition, coalition structure, deliberation procedures—to the skill and noise parameters of the model would provide direct guidance for institutional governance design. Field experiments testing whether the procedural interventions proposed in Section 5 reduce measurable oversight risk in actual committee settings would be especially valuable.

A Technical Details

This appendix outlines the mathematical structure underlying the simulation results for readers who wish to engage with the formal model.

Portfolio choice under committee oversight. The portfolio manager is modeled as choosing portfolio weights $\boldsymbol{x}$ to maximize the influence-weighted average of committee member utilities. When committee members share beliefs but differ in risk aversion, the effective risk aversion is $\gamma_t = \boldsymbol{\phi}'\boldsymbol{\gamma}$, where $\boldsymbol{\phi}$ is the vector of influence weights and $\boldsymbol{\gamma}$ is the vector of member risk aversions. The optimal portfolio is:

$$\boldsymbol{x} = \frac{1}{\boldsymbol{\phi}'\boldsymbol{\gamma}}\,\boldsymbol{\Sigma}^{-1}(\boldsymbol{\mu} - r_f\mathbf{1}),$$

where $\boldsymbol{\mu}$ is the vector of expected excess returns, $\boldsymbol{\Sigma}$ is the covariance matrix, and r_f is the risk-free rate. When members share risk tolerance but differ in beliefs, the effective beliefs are $\boldsymbol{y} = \boldsymbol{M}\boldsymbol{\phi}$, where $\boldsymbol{M}$ is the matrix of member-specific expected return vectors, and the optimal portfolio is:

$$\boldsymbol{x} = \frac{1}{\gamma}\,\boldsymbol{\Sigma}^{-1}(\boldsymbol{y} - r_f\mathbf{1}).$$

Performance metrics. For a portfolio with weights $\boldsymbol{x}$, true expected returns $\boldsymbol{\mu}$, covariance matrix $\boldsymbol{\Sigma}$, risk-free rate r_f, and market portfolio weights $\boldsymbol{x}_m$, the three performance metrics are:

$$\text{Sharpe ratio} = \frac{\boldsymbol{x}'(\boldsymbol{\mu} - r_f\mathbf{1})}{\sqrt{\boldsymbol{x}'\boldsymbol{\Sigma}\boldsymbol{x}}}, \qquad \alpha = \boldsymbol{x}'(\boldsymbol{\mu} - r_f\mathbf{1}) - \frac{\boldsymbol{x}'\boldsymbol{\Sigma}\boldsymbol{x}_m}{\boldsymbol{x}_m'\boldsymbol{\Sigma}\boldsymbol{x}_m}\,\boldsymbol{x}_m'(\boldsymbol{\mu} - r_f\mathbf{1}),$$

$$\text{IR} = \frac{(\boldsymbol{x} - \boldsymbol{x}_m)'\boldsymbol{\mu}}{\sqrt{\boldsymbol{x}'\boldsymbol{\Sigma}\boldsymbol{x} - 2\boldsymbol{x}'\boldsymbol{\Sigma}\boldsymbol{x}_m + \boldsymbol{x}_m'\boldsymbol{\Sigma}\boldsymbol{x}_m}}.$$

Variance of portfolio returns under belief volatility. When the committee's effective beliefs $\boldsymbol{y}$ are random with mean $\boldsymbol{\mu}_y$ and covariance $\boldsymbol{\Sigma}_y$, the variance of portfolio returns is:

$$\sigma_p^2 = \frac{1}{\gamma^2}(\boldsymbol{\mu}_y - r_f\mathbf{1})'\boldsymbol{\Sigma}^{-1}(\boldsymbol{\mu} - r_f\mathbf{1}) + \frac{1}{\gamma^2}\,\mathrm{tr}(\boldsymbol{\Sigma}^{-1}\boldsymbol{\Sigma}_y) + \frac{1}{\gamma^2}\,\mathrm{tr}(\boldsymbol{\Sigma}^{-1}\boldsymbol{\Sigma}_{ry}\boldsymbol{\Sigma}^{-1}\boldsymbol{\Sigma}_{ry}'),$$

where $\boldsymbol{\Sigma}_{ry} = \mathrm{cov}(\boldsymbol{r}, \boldsymbol{y})$ captures the relationship between realized returns and effective beliefs. The middle term, $\mathrm{tr}(\boldsymbol{\Sigma}^{-1}\boldsymbol{\Sigma}_y)/\gamma^2$, is the direct contribution of belief variance to portfolio return variance—the formal expression of the second channel of oversight risk.

Simulation design. The market-implied expected excess returns are computed using the Black-Litterman reverse-optimization formula: $\boldsymbol{\mu}_t - r_{f,t} = \gamma\boldsymbol{\Sigma}_t\boldsymbol{x}_{m,t}$, where $\gamma = 1$. For the risk aversion simulation, the effective risk aversion follows a logistic function of the excess market return: $\gamma_t = 0.5 + 1.0/(1 + \exp(-k(r_{m,t} - \bar{r}_m))) + \varepsilon_t$. For the belief noise simulation, effective beliefs are $\boldsymbol{y}_t = \boldsymbol{\mu}_t + \boldsymbol{\varepsilon}_t$ with $\boldsymbol{\varepsilon}_t \sim N(\mathbf{0}, \sigma^2\boldsymbol{I})$. For the skill–noise simulation, beliefs are constructed to achieve a target cosine similarity $\bar{\rho}$ with the next period's market-implied returns while preserving their magnitude.

References

Aghion, P. & Tirole, J. (1997), 'Formal and real authority in organizations', *Journal of political economy* **105**(1), 1–29.

Bär, M., Kempf, A. & Ruenzi, S. (2011), 'Is a team different from the sum of its parts? evidence from mutual fund managers', *Review of Finance* **15**(2), 359–396.

Berger, J., Cohen, B. P. & Zelditch, Jr., M. (1972), 'Status characteristics and social interaction', *American Sociological Review* **37**(3), 241–255.

Berk, J. B. & Green, R. C. (2004), 'Mutual fund flows and performance in rational markets', *Journal of Political Economy* **112**(6), 1269–1295.

Binfare, M., Harris, R., Brown, G. & Lundblad, C. T. (2017), How important are investment committees? evidence from university endowments, *in* 'the 2017 UNC Alternative Investments Conference'.

Blinder, A. S. & Morgan, J. (2005), 'Are two heads better than one? Monetary policy by committee', *Journal of Money, Credit and Banking* **37**(5), 789–811.

Charness, G. & Sutter, M. (2012), 'Groups make better self-interested decisions', *Journal of Economic Perspectives* **26**(3), 157–176.

Correll, S. J. & Ridgeway, C. L. (2003), Expectation states theory, *in* J. Delamater, ed., 'Handbook of Social Psychology', Kluwer Academic/Plenum Publishers, New York, pp. 29–51.

Dalkey, N. & Helmer, O. (1963), 'An experimental application of the DELPHI method to the use of experts', *Management Science* **9**(3), 458–467.

Eisenhardt, K. M. & Bourgeois, III, L. J. (1988), 'Politics of strategic decision making in high-velocity environments: Toward a midrange theory', *Academy of Management Journal* **31**(4), 737–770.

Ellis, C. D. (2011), 'Best practice investment committees', *Journal of Portfolio Management* **37**(2), 139.

Gigone, D. & Hastie, R. (1993), 'The common knowledge effect: Information sharing and group judgment', *Journal of Personality and Social Psychology* **65**(5), 959–974.

Guiso, L., Sapienza, P. & Zingales, L. (2018), 'Time varying risk aversion', *Journal of Financial Economics* **128**(3), 403–421.

He, Z. & Krishnamurthy, A. (2012), 'A model of capital and crises', *Review of Economic Studies* **79**(2), 735–777.

Isenberg, D. J. (1986), 'Group polarization: A critical review and meta-analysis', *Journal of Personality and Social Psychology* **50**(6), 1141–1151.

Janis, I. L. (1982), *Groupthink: Psychological Studies of Policy Decisions and Fiascos*, 2nd edn, Houghton Mifflin, Boston.

Kahneman, D., Sibony, O. & Sunstein, C. R. (2021), *Noise: A flaw in human judgment*, Hachette UK.

Kerr, N. L. & Tindale, R. S. (2004), 'Group performance and decision making', *Annual Review of Psychology* **55**, 623–655.

Kruglanski, A. W., Pierro, A., Mannetti, L. & De Grada, E. (2006), 'Groups as epistemic providers: Need for closure and the unfolding of group-centrism', *Psychological Review* **113**(1), 84–100.

Kruglanski, A. W., Shah, J. Y., Pierro, A. & Mannetti, L. (2002), 'When similarity breeds content: Need for closure and the allure of homogeneous and self-resembling groups', *Journal of Personality and Social Psychology* **83**(3), 648–662.

Kruglanski, A. W. & Webster, D. M. (1991), 'Group members' reactions to opinion deviates and conformists at varying degrees of proximity to decision deadline and of environmental noise', *Journal of Personality and Social Psychology* **61**(2), 215–225.

Kruglanski, A. W. & Webster, D. M. (1996), 'Motivated closing of the mind: "Seizing" and "Freezing"', *Psychological Review* **103**(2), 263–283.

Larrick, R. P. & Soll, J. B. (2006), 'Intuitions about combining opinions: Misappreciation of the averaging principle', *Management Science* **52**(1), 111–127.

Lombardelli, C., Proudman, J. & Talbot, J. (2005), 'Committees versus individuals: An experimental analysis of monetary policy decision-making', *International Journal of Central Banking* **1**(1), 181–205.

Moscovici, S. & Zavalloni, M. (1969), 'The group as a polarizer of attitudes', *Journal of Personality and Social Psychology* **12**(2), 125–135.

Nijstad, B. A., Stroebe, W. & Lodewijkx, H. F. M. (2003), 'Production blocking and idea generation: Does blocking interfere with cognitive processes?', *Journal of Experimental Social Psychology* **39**(6), 531–548.

Pierro, A., Mannetti, L., De Grada, E., Livi, S. & Kruglanski, A. W. (2003), 'Autocracy bias in informal groups under need for closure', *Personality and Social Psychology Bulletin* **29**(3), 405–417.

Prather, L. J. & Middleton, K. L. (2002), 'Are $N + 1$ heads better than one? the case of mutual fund managers', *Journal of Economic Behavior and Organization* **47**(1), 103–120.

Reimer, T., Reimer, A. & Czienskowski, U. (2010), 'Decision-making groups attenuate the discussion bias in favor of shared information: A meta-analysis', *Communication Monographs* **77**(1), 121–142.

Scharfstein, D. S. & Stein, J. C. (1990), 'Herd behavior and investment', *American Economic Review* **80**(3), 465–479.

Scherer, B. (2024), 'Optimal design of investment committees', *Journal of Asset Management* **25**(2), 129–135.

Stasser, G. & Titus, W. (1985), 'Pooling of unshared information in group decision making: Biased information sampling during discussion', *Journal of Personality and Social Psychology* **48**(6), 1467–1478.

Sunstein, C. R. & Hastie, R. (2015), *Wiser: Getting Beyond Groupthink to Make Groups Smarter*, Harvard Business Review Press, Boston, MA.

Swensen, D. F. (2009), *Pioneering portfolio management: An unconventional approach to institutional investment, fully revised and updated*, Simon and Schuster.

Useem, M. & Zelleke, A. (2006), 'Oversight and delegation in corporate governance: Deciding what the board should decide', *Corporate Governance: An International Review* **14**(1), 2–12.

Yoder, J. A. (2011), 'The investment committee. agb effective committee series.', *Association of Governing Boards of Universities and Colleges* .