The Benchmarking Prisoner's Dilemma
Game-Theory Insights for Institutional Investors

Redouane Elkamhi[1] Jacky S.H. Lee[2]

First version: June 8, 2025

This revised version: September 16, 2025

The views and opinions expressed in this study are solely those of the authors and do not necessarily reflect the official policy or position of any other agency, organization, employer, or company, including the authors' employer in the asset management industry. The information provided in this study is for general informational purposes only and should not be construed as legal, tax, investment, financial, or other professional advice.

ACKNOWLEDGMENT

We'd like to thank Michael Wissell for his constant support of research and innovation, and for always pushing us to think differently—looking for simple, practical ways to solve real-world problems. We're also deeply grateful for the many conversations we've had with him about this topic. His ideas, questions, and challenges have shaped our thinking in important ways and contributed a great deal to this paper.

---

[1] Professor of Finance, Rotman School of Management, University of Toronto, Canada
[2] Senior Managing Director, Head of Total Portfolio Group, Healthcare of Ontario Pension Plan (HOOPP), Canada. Editorial Advisory Board Member, *Journal of Portfolio Management* and *Journal of Financial Data Science*, Portfolio Management Research, UK.

ABSTRACT

Benchmarking has long been a cornerstone of institutional investing, offering discipline, structure, and accountability. Yet as portfolios have evolved to include complex, multi-asset strategies, traditional benchmarking frameworks can unintentionally encourage cautious or narrowly focused behavior, limiting the potential for genuine long-term value creation. This paper applies insights from game theory and agency theory to show how such dynamics can arise naturally within well-designed systems, leading to stable but constrained outcomes.

We explore an alternative approach—delegated benchmarking—where high-level objectives remain clearly defined, but the design of mandate-specific benchmarks is entrusted to those closer to the information. This shift preserves the clarity and accountability that benchmarks provide, while better aligning incentives, reducing hidden risks, and supporting stronger talent retention. By repositioning benchmarks as tools for fair measurement rather than day-to-day steering, institutions can foster a more cooperative and adaptive equilibrium—one that strengthens oversight while enabling the pursuit of sustainable outperformance.

INTRODUCTION

Benchmarks were originally created to impose discipline on investment portfolios. For decades this worked well. Portfolios were simpler—mostly public stocks and bonds—and a static Strategic Asset Allocation with asset-class benchmarks provided clear structure. In that legacy equilibrium, a board could set a policy (asset) mix, assign each asset class a standard market index, and evaluate managers against those benchmarks. This approach, rooted in the well-known finding that asset allocation explains most return variability, made sense in a stable environment. If a manager beat their benchmark, they demonstrated skill (alpha) beyond the market's return (beta), and the policy portfolio provided a straightforward yardstick for overall performance. Board members—experienced fiduciaries with broad responsibilities but limited time for detailed market oversight—could therefore supervise professionals by asking a simple question: *"Did this manager do a good job?"*

Over time, however, portfolios have become far more complex. Institutions now span public and private markets, with illiquid assets, dynamic strategies, and interdependent risks. Performance measurement has struggled to keep up. Many funds still rely on a fixed total-fund benchmark (such as a static 60/40 mix or an inflation-plus hurdle). While valuable as a high-level anchor, once that reference point is applied rigidly at the mandate level, it can sometimes create unintended incentives. Each sub-portfolio is judged in isolation against a rigid yardstick, creating what economists call tracking-error sensitivity—where even small deviations from the benchmark become highly visible and career-sensitive. In practice, this dynamic channels managers into two predictable strategies: either benchmark hugging (staying close to the index) or beta gaming (adding extra market exposure to boost apparent results). *In other words, the safest career move is to shadow the index or quietly dial up market exposure, because leverage is the simplest way to appear to outperform in a rising market—even though this kind of outperformance is neither sustainable nor aligned with the fund's true objectives.*

This tension reflects a double-agency problem: the board, acting as principal, delegates to the CIO and CEO (hereafter referred to collectively as the Chief Executives, or CE), who in turn act as principals to individual managers. Each layer introduces familiar agency frictions—moral hazard (gaming the rules), adverse selection (losing top talent when incentives are misaligned), and information asymmetry (boards, by design, cannot access the same level of detail as day-to-day managers). Under a board-controlled benchmarking regime, the dominant strategy for managers is to remain close to the index or economize effort by juicing beta exposure rather than investing energy in finding true orthogonal alpha. The CE's best response, anticipating how performance will be judged, is to conform to rigid optics. In game-theory terms, this setup constitutes a Nash equilibrium: a stable outcome in which no player can improve their own position by unilaterally deviating, even though the collective result is uninspiring. *In other words, the organization settles into a 'safe but constrained' equilibrium—careers are protected and optics are preserved, but genuine value creation is limited.*

Delegation changes the rules of the game. When the board empowers the Chief Executives to design fit-for-purpose asset-class benchmarks—while still holding them accountable for outcomes at the total-fund level—the strategic dynamics shift. The CE's dominant strategy is no longer to tolerate benchmarks that are too easy (which would undermine credibility) or too demanding (which would drive talent away). Instead, the rational move is to calibrate fair, role-specific benchmarks that make incentive compatibility credible across the organization. Managers, in turn, face a new dominant response: to pursue orthogonal alpha—genuine, non-overlapping sources of value rather than thinly disguised beta—aligned with the fund's overarching objectives. In this delegated regime, asset-class benchmarks stop steering day-to-day portfolio choices. They become instruments for compensation, evaluation, and accountability. *In other words, benchmarks return to their proper role: a ruler for measuring craft, not a steering wheel for driving the portfolio.*

The outcomes are not guaranteed—no governance model delivers alpha in every market. But in equilibrium, delegation is expected to yield superior results: higher-quality alpha, stronger alignment of incentives, and better talent retention. Unless a board places an unusually high value on direct control—the so-called control premium—delegation is also the board's rational dominant strategy. Even in the worst case, delegation produces outcomes no worse than control, and in the likely case, it delivers better performance and alignment.

Using the language of game theory (dominant strategies, equilibria, payoffs) and agency theory (moral hazard, adverse selection, principal–agent dynamics), we show why delegated benchmarking is the superior governance equilibrium for modern portfolios. We also translate these ideas into plain language and practical tools, so that trustees, CE, and practitioners can apply the lessons directly.

Finally, it is worth noting that many trustees themselves have recognized the limits of rigid, benchmark-driven frameworks and the sense of reassurance that comes from detailed oversight they can create. Directors frequently raise concerns similar to those discussed here—that benchmarks can sometimes constrain innovation, encourage caution, or provide an incomplete picture of success. The intent of this paper is simply to highlight how structural frictions and agency dynamics can allow such patterns to emerge, even in thoughtful and well-intentioned governance settings.

The rest of the paper (i) explains why the legacy benchmarking system can lock funds into predictably average outcomes, (ii) frames the governance trade-off through a game-theory lens, (iii) lays out the delegated model and its expected benefits, and (iv) offers practical guidance for boards considering the shift.

THE LEGACY BENCHMARKING TRAP: MISALIGNED INCENTIVES AND AGENCY PROBLEMS

When boards set detailed benchmarks for each asset class and tie performance evaluation to those benchmarks, they unintentionally create a game of incentives that undermines the total portfolio. What begins as a logical accountability mechanism becomes a classic principal–agent problem: the board seeks total-fund success, but managers are rewarded for outperforming narrow indices—even when that behavior detracts from the whole.

The result is a predictable pattern. Managers focus on their slice of the portfolio, striving to beat the benchmark in front of them rather than considering the portfolio's bigger picture. Each silo may look good in isolation, yet the total fund can drift into concentrations and imbalances. Even the CE is constrained in this environment, because policy weights and benchmark optics leave little room to reallocate capital or adapt mandates to changing conditions. In effect, the CE's flexibility can be constrained by mandate optics, limiting their ability to fully act as a steward of the fund.

Compensation structures reinforce this rigidity. The safest career strategy is to hug the index or, when possible, quietly load up on additional market exposure—beta gaming—to produce apparent outperformance. In game-theory terms, these choices become dominant strategies, yielding an equilibrium of modest, lower-quality alpha. Boards, not involved in day-to-day management, may also face moral hazard as managers push for easier yardsticks or benchmarks that quietly embed market exposure. Over time, what began as neutral reference points drift into padded targets, inflating perceived performance.

Short-termism further compounds the challenge. Frequent scorecard reviews reward quick wins while discouraging contrarian ideas that may need a year or more to prove their value. The situation resembles a "prisoner's dilemma": while everyone recognizes that long-horizon bets can add value, short-term pressures make it difficult for any manager to risk underperforming the benchmark in the interim. As a result, managers often remain close to conventional holdings, which can make it more challenging for innovative ideas to gain traction.

The effects ripple through the organization. Rigid benchmarks prevent capital from moving to its best use: high-conviction ideas that cut across asset classes may never see funding if they do not fit neatly into the established categories. Talent is also at stake. Top investment professionals, who thrive on fair recognition of their skill, are less willing to remain in systems where effort and insight are judged against unrealistic or ill-fitting yardsticks. This produces adverse selection: the strongest agents exit while more complacent ones remain, gradually weakening the fund's human capital.

In short, board-led benchmarking tends to produce a stable but suboptimal equilibrium. Each agent's rational choice—protect career risk, shadow the index, or game the system—adds up to stable outcomes with limited room for innovation. *In other words, the rules of the game create safety at the cost of genuine success: the fund avoids major disasters, but also misses the chance for meaningful outperformance.*

A GAME-THEORETICAL PERSPECTIVE: WHY THE CURRENT EQUILIBRIUM PERSISTS

It is useful to frame the interaction through the lens of game theory. Strictly speaking, the setup is a double-agency problem: the board, as principal, delegates to the CE, who themselves act as principals to individual managers. Each actor faces distinct incentives—the board seeks fiduciary security, the CE is judged on total-fund outcomes, and managers are evaluated on mandate performance. In principle, this creates a three-player game—a "cube" of strategies. For clarity, however, we collapse the CE and managers into a single player in the main text. The logic is that the CE has better information and monitoring capacity, and can internalize manager behavior (catching issues such as index-hugging or beta gaming). This allows the system to be represented as a 2×2 matrix between the board and the CE, preserving the essence of the trade-offs without losing tractability. Readers interested in the full three-actor version can refer to the "strategy cube" in the appendix A. The strategic payoff matrix is summarized in Figure A and elaborated below

In this simplified framing, the board has two choices: retain direct control of benchmarks (the status quo), or delegate benchmark design to the CE. The CE, in turn, has two responses: play it safe (staying close to benchmarks and minimizing active risk), or pursue alpha (taking high-conviction positions when warranted).

When the board controls and the CE plays it safe, the outcome is familiar: benchmarks are tight, so the CE and managers rationally shadow indices to avoid visible underperformance. The payoff is moderate for the board (predictable but unexciting results) and low for the CE (little upside, low career risk). In game-theory terms, this is a Nash equilibrium: neither side has an incentive to deviate unilaterally. *In other words, it is a stable stalemate—an uneasy truce where the fund avoids disaster but rarely achieves genuine outperformance.*

If the board controls while the CE pursues alpha, the CE takes on significant career risk. Any short-term underperformance against rigid benchmarks becomes highly visible, inviting scrutiny or even termination. The board, meanwhile, sees volatility or tracking error that was not expected, and may regret the loss of control. This is not a stable strategy pair; it is punished by the rules of the game. *Like a prisoner's dilemma, unilateral cooperation leads to worse outcomes for the one who takes the risk.*

If the board delegates but the CE continues to play it safe, trust is is extended but may not be fully utilized at first. The outcomes may look similar to the status quo, though with less emphasis on direct oversight. Eventually, the board may revert to control, questioning why authority was granted in the first place. Again, this is not a stable outcome.

The most attractive case emerges when the board delegates and the CE pursues alpha. Here, the board defines the total-fund benchmark and guardrails, while the CE designs fit-for-purpose benchmarks and allocates flexibly. Both sides win: the board sees better risk-adjusted returns and a more agile portfolio; the CE and managers are rewarded for genuine skill. Importantly, this combination can become self-reinforcing—a Pareto-superior equilibrium. If the board delegates, the CE's best response is to pursue alpha. If the CE delivers, the board's best response is to continue delegating. Trust builds on itself: empowerment begets effort and results, which beget further empowerment.

In this framing, the status quo reflects a cautious equilibrium, where coordination across roles can be difficult to achieve. Each party's cautious move makes sense individually, but collectively it locks the fund into low trust and low excess returns. Delegation, by contrast, unlocks a cooperative equilibrium where incentives align and both sides are better off. For years, many institutions have remained stuck in the former equilibrium because it feels safer—boards value the assurance that comes from maintaining direct oversight, while CEs hesitate to take risk without explicit backing. These agency frictions, amplified by information asymmetry, explain why the inefficient outcome persists. *In other words, the system stays in place not because anyone chooses poorly, but because the existing rules make it difficult for either side to move first on their own.*

The lesson is straightforward: by redesigning the rules of the game—specifically, who sets benchmarks and how success is measured—the institution can move from a low-level equilibrium of caution to a high-level equilibrium of alignment. In this new regime, the CE acts as a trusted agent with authority to act in the board's best interest, and evaluation metrics shift so that managers are rewarded for total-fund success rather than narrow outperformance. *In other words, what was once an adversarial oversight game becomes a cooperative one: the board and management on the same team, competing not against each other, but against the market.*

Figure A: Strategic Payoff Matrix for Benchmark Control

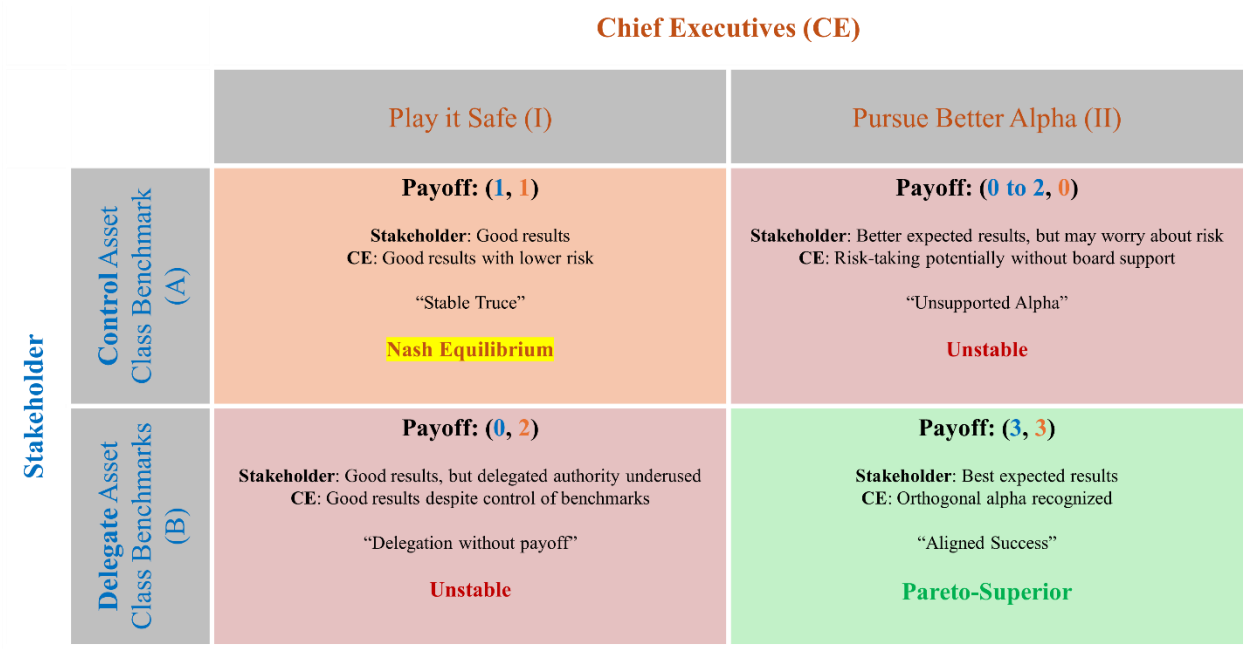


Figure: This matrix illustrates the strategic interaction between the stakeholder (board) and the Chief Executives (CIO/CEO). The board can either retain control of asset-class benchmarks or delegate their design, while the CE can either play it safe by hugging indices or pursue alpha through high-conviction positions. The top-left cell (Control + Safe) represents the status quo: predictable but constrained outcomes, a Nash equilibrium that is stable but limited in upside. The top-right cell (Control + Alpha) reflects efforts to generate alpha that are not recognized under board-set benchmarks, producing an unstable outcome where genuine skill appears as excess risk rather than added value. The bottom-left cell (Delegate + Safe) shows authority granted but underused, another unstable scenario. The bottom-right cell (Delegate + Alpha) delivers the superior outcome: the board defines total-fund objectives, the CE designs fit-for-purpose benchmarks, and managers pursue orthogonal alpha, creating a Pareto-superior but not stable without a change in the rules of the game (e.g., repeated play, delegation, incentive redesign, trust-building). In terms of best responses, when the board controls, the CE rationally plays it safe; when the board delegates, the CE's best move is to pursue alpha; if the CE plays safe, the board prefers to retain control; and if the CE pursues alpha, the board's rational choice is to delegate—demonstrating why institutions tend to remain stuck in the low-return "stable truce" and how they can move toward the high-alignment "outperformance equilibrium."

DELEGATED BENCHMARKING: SOLVING THE DOUBLE-AGENCY PROBLEM

The proposed solution is to delegate authority over asset-class benchmarks to the CE and investment team, while the board defines overall objectives and constraints at the total-fund level. This creates a governance structure best described as a double-agency model: the board (as principal) holds the CE accountable for total results, and the CE in turn acts as principal to individual managers, setting their benchmarks and evaluating their performance. At first glance, adding another agency layer might seem to worsen the problem. In practice, however, the CE is not simply another agent but also a principal—one with better information, better monitoring capacity, and stronger alignment with the fund's ultimate goals than a distant board could provide. Paradoxically, this two-tier structure reduces, rather than compounds, classic agency frictions.

First, information asymmetry narrows. The CE and their team have deeper insight into asset classes, current market conditions, and each manager's strategy than the board can reasonably possess. By delegating benchmark design, benchmarks can be set in ways that are realistic and context-specific rather than generic. For example, instead of a blunt "public equity + X%" target for private equity, the CE can tailor a benchmark to reflect the actual opportunity set or peer group. Because these benchmarks are more adaptive—they can evolve as markets or strategies change—they challenge managers while still treating them fairly. *In other words, decisions are being made by the people closest to the information, which makes benchmarks both tougher to game and more relevant to reality.*

Second, moral hazard is contained. Under board-set benchmarks, managers can quietly juice returns by adding hidden beta, lobbying for softer yardsticks, or otherwise gaming the rules. In a delegated structure, however, those tricks no longer pay off. Since the CE is accountable for total-fund outcomes, any sleeve-level maneuver that doesn't truly add value is immediately visible at the aggregate level. The CE also has the authority to set and adjust benchmarks directly, which means ongoing monitoring can catch unintended risks far sooner than an infrequently meeting board. Importantly, delegation does not remove oversight—it redistributes it. The board still evaluates the CE on total-fund outcomes, while the CE monitors managers day to day. Each layer keeps the next in check, but each focuses on what it does best: the board defines the mission, while the CE ensures fair and effective implementation.

Third, adverse selection is reversed. Delegation signals to talented managers that their skill will be judged on meaningful results, not arbitrary hurdles. This makes the organization far more attractive to high performers who want autonomy and fair recognition. A private equity manager, for instance, is more likely to stay if benchmarks reflect illiquidity and J-curve dynamics rather than being naively compared to roaring public markets. Over time, this environment compounds performance potential: better people stay, make better decisions, attract more strong talent, and build a virtuous cycle. *In other words, instead of driving innovators out, the governance model keeps them in and gives them room to add value.*

Taken together, these shifts restore a unified objective: everyone aims at total-fund success, just from different roles. The board defines what success looks like—say, outperforming a reference portfolio or achieving CPI + 4% within risk limits—and holds the CE accountable for delivering it. The CE translates those goals into fit-for-purpose benchmarks for managers, who then execute their mandates. Unlike the siloed approach, there is no misalignment between what a manager is asked to do and what the fund actually needs. In game-theory terms, delegation realigns the payoff structure so that each agent's dominant strategy is also the best outcome for the principal. The CE's best move is to maximize total-fund value, while each manager's best move is to maximize their mandate's contribution to that value. As a result, delegation becomes a dominant strategy for the institution itself: given the choice between governance regimes, the delegated model consistently yields stronger performance, clearer alignment, and better retention of talent, effectively dominating the board-controlled regime on every meaningful dimension. *In other words, delegation transforms a system of competing silos into a coordinated team where everyone's incentives point toward the same goal.*

A common concern with delegated benchmarking is the possibility that a CIO might deliberately set benchmarks that are too easy in order to show outperformance. In practice, however, this behavior is neither rational nor sustainable. First, the board continues to define the total-fund benchmark or reference portfolio, which serves as the ultimate yardstick of success. Even if mandate-level benchmarks were softened, the overall results would quickly reveal whether genuine value was being added. Second, credibility is central to the CIO's role. Any attempt to lower the bar artificially would, over time, undermine trust and risk the CIO's own career prospects, since long-term outcomes cannot be disguised.

Delegation also works because incentives are aligned. A CIO benefits most by setting benchmarks that are challenging but fair, ensuring talented managers remain motivated and engaged. If benchmarks were consistently too soft, strong professionals would recognize this misalignment and disengage, or worse, leave the organization—an outcome no CIO would willingly create. Finally, safeguards remain firmly in place: boards can establish clear total-fund hurdles, require transparent reporting that links mandate results back to fund objectives, and review benchmark calibration on a regular basis. In this sense, delegation does not mean benchmarks are easier—it means they are more realistic, more relevant, and ultimately more effective at driving long-term results.

EXPECTED OUTCOMES UNDER DELEGATION

A theoretical discussion is only as valuable as the outcomes it implies in practice. While no governance model guarantees higher returns, thinking through the incentives of managers, the CE, and the board suggests that delegation is likely to produce meaningful improvements across several dimensions. In effect, delegation represents a different equilibrium—one tilted toward alignment rather than caution.

One clear benefit is **improved total-portfolio efficiency**. With flexibility to rebalance across silos, the CE can cut exposures that exist merely to "match the benchmark" and reallocate to strategies that truly diversify risk. *In other words, capital is no longer trapped in yesterday's allocations, but can flow to where it adds the most value.*

Delegation also **decouples benchmarks from investment decisions**. In this regime, benchmarks cease to drive day-to-day choices and instead function primarily as tools for compensation, accountability, and oversight. Because the CE can adjust benchmarks as markets evolve, managers are free to focus on generating orthogonal alpha. The result is less gaming, more fairness, and benchmarks that serve governance rather than constrain initiative.

This shift leads to **higher-quality alpha**. Under board control, some of what passes as alpha can be disguised beta or small benchmark tilts. Delegation encourages managers to find genuinely uncorrelated opportunities, improving both the quality and the diversification of active returns.

Another outcome is **better talent alignment**. Benchmarks that are too easy or too punitive distort incentives and risk driving strong professionals out of the organization. Delegation enables fit-for-purpose benchmarks that challenge but also reward fairly, encouraging skilled managers to stay and innovate. *In other words, people no longer leave because the system feels unfair; they stay because the system recognizes genuine skill.*

Delegation also promotes **more honest reporting**. With less room for benchmark gaming, reported results provide a truer reflection of skill, giving boards clearer oversight and reinforcing trust. Finally, the combination of flexibility and better incentives supports **higher risk-adjusted returns**. By freeing the CE to shift capital where it adds the most value, and by rewarding managers for genuine contributions rather than narrow benchmark wins, delegation captures returns that would otherwise be lost to inefficiencies.

Equally important, delegation strengthens risk oversight: it reduces hidden risks from beta-gaming, improves transparency in reporting, and allows boards to focus on systemic exposures and long-term resilience rather than line-item optics.

Taken together, these outcomes form an **expected Pareto improvement**: the board gains clearer accountability, the CE gains flexibility to manage the total fund, and managers gain fairer evaluation. The trade-off is a shift from detailed line-item control toward broader, outcome-focused oversight. *In other words, delegation trades the comfort of detailed control for the progress of broader, outcome-focused oversight.*

THE TRADE-OFF AND THE "CONTROL PREMIUM" – OVERCOMING RELUCTANCE

If delegating benchmark design to the CE is so effective, why don't all boards adopt it immediately? The answer often relates to the natural value boards place on **control**. Trustees carry a fiduciary duty and feel a deep responsibility to safeguard the fund. For many boards, the instinct to retain control reflects not stubbornness but genuine concern for accountability, especially in an era of heightened scrutiny from stakeholders and regulators. With that responsibility comes an understandable comfort in setting the rules directly and being able to see that they are followed. Economists might describe this assurance as a **control premium**—an intangible benefit that arises from retaining authority, even if it sometimes comes at the cost of incremental performance.

Suppose, for illustration, that maintaining close control provides a sense of assurance equivalent to 0.4% per year in return. Delegation, by contrast, may offer the potential for significantly higher long-term results. If the "comfort value" of control is less than the return foregone by holding onto it, the net effect for beneficiaries is negative. Phrased differently, the question becomes: *what is the right balance between the reassurance of direct oversight and the potential benefits of empowerment?*

For many boards, once the expected benefits of delegation are made explicit, the case for adjusting the balance becomes clear. Importantly, delegation does not mean giving up oversight. Boards retain authority over the most important dimensions: defining objectives, setting risk parameters, and holding the CE accountable for results. What changes is that the detailed mechanics—the "how"—are entrusted to those closer to the information. This division of roles reflects comparative advantage: boards focus on the *what*, while CEs focus on the *how*. In practice, this elevates the board's role to defining mission, risk appetite, and long-term purpose—the areas where trustees' independence and broad expertise create the greatest value.

Of course, moving to a new equilibrium requires trust, which can feel uncomfortable at first. External pressures—consultants, regulators, or stakeholders accustomed to detailed benchmarks—can also reinforce the old model. Yet maintaining an overly detailed framework can carry hidden costs, not unlike an avoidable fee drag. Framed this way, delegation is less about "letting go" and more about modernizing governance so that board time and attention are spent where they create the greatest value.

Concerns about worst-case scenarios are natural: *what if delegation leads to mistakes or misalignment?* Here, strong safeguards remain in place. Delegation does not mean abdication. Instead, it reframes the board's role: not as pilots micromanaging every control, but as navigators who set the course, define the risk boundaries, equip the CE with the right instruments, and ensure the ship arrives safely at its destination. Boards still monitor diligently, but the focus shifts to outcomes and processes rather than line-by-line inputs. Effective checks include setting a total-fund benchmark or hurdle, establishing risk limits, and requiring regular reporting that links mandate performance back to the total fund.

Viewed this way, delegation strengthens rather than weakens governance. By concentrating oversight on total-fund health, risk management, and long-term strategy, boards free themselves

from the minutiae of line-item benchmarks and elevate discussions to a more strategic level. Many leading funds that have embraced delegation report that board meetings have become less about oversight mechanics and more about insight into the big picture. *In other words, delegation allows boards to preserve what matters most—their fiduciary responsibility and oversight—while enabling management to execute with the flexibility needed to deliver stronger results.*

CONCLUSION: TOWARD AN EQUILIBRIUM OF TRUST AND PERFORMANCE

In summary, the traditional approach of board-controlled benchmarking was effective in a simpler era, but today's complex environment calls for modernization. Left unchanged, it can lead boards and managers into a low-risk, lower-reward equilibrium—akin to a multi-player prisoner's dilemma. A game-theory lens shows that while each side's cautious choices are individually rational, the collective outcome falls short of the institution's true potential. To escape this trap, the rules of the game need to change—chiefly by delegating detailed benchmark setting to those closer to the action, aligning evaluation with overall fund success, and encouraging cooperative, value-creating behavior.

The delegated benchmarking model transforms the investment organization into a more cohesive and adaptive system. It addresses principal–agent issues by placing a knowledgeable agent—the CE—in charge of mandate-level design and oversight, under the board's high-level direction. This double-agency structure, when supported by transparency and accountability, consistently outperforms the single-tier control structure on key dimensions: higher returns, stronger alpha, better risk management, improved talent retention, and clearer information flow. It creates a new equilibrium of trust, where each individual's dominant strategy also advances the collective payoff. In this setting, moral hazard is constrained not by rigid external rules but by internal alignment, and adverse selection is reversed as talented professionals are drawn to an environment that rewards genuine skill.

For CEs and investment leaders, delegation brings greater empowerment to deliver on the fund's mandate, but also greater responsibility to design fair benchmarks and incentives for their teams. For boards, delegation calls for a higher-level mindset: defining success in terms of mission, risk appetite, and total-fund objectives, and then holding the CE accountable for results. It means resisting the natural pull to oversee every sub-portfolio, and instead embracing a role of *trust but verify*. Ultimately, that is a more strategic, higher-impact role for the board to play.

In closing, the evolution of institutional investing requires an evolution in governance. Embracing delegated benchmark design is not about relinquishing oversight—it is about modernizing it. Much like the classic prisoner's dilemma, the legacy system locks each player into choosing the "safe" strategy of control and conformity, even though everyone would be better off if they cooperated. By applying the principles of economics and game theory to realign incentives, institutions can escape this low-trust stalemate and move to a new paradigm where benchmarks serve the portfolio rather than steer it. The leaders in this shift will be those boards that recognize the value of comparative advantage in information and decision-making, and those CEs who rise to the responsibility of greater authority. Such organizations will not only deliver better performance, but also foster a more harmonious and purpose-driven approach to managing assets—an outcome where everyone, from the CE to the ultimate beneficiaries, can confidently say the team did "a

good job" for the right reasons. In this sense, the board remains firmly at the helm of purpose and oversight, acting as the navigator that ensures the fund stays true to its mission, while empowered executives handle the mechanics of the voyage. In this way, boards can strengthen their oversight impact by focusing more on strategy, purpose, and long-term resilience, while management ensures benchmarks are calibrated and implemented effectively.

APPENDIX A: FULL TRI-PARTY STRATEGY CUBE

Under a **board-controlled benchmarking regime**, authority is centralized at the board level. The board sets the strategic asset allocation and most sleeve benchmarks, approves active risk budgets, and judges value-add at the sleeve level against board-set indices. These results are then rolled up to the policy portfolio. The incentive signal this sends is clear: managers and executives are encouraged to minimize tracking error and preserve predictable optics.

Within this framework, the CE has limited room to maneuver. Implementation is confined to the benchmarks and weights defined by the board, with little ability to reshape incentives or adjust mandates dynamically. Evaluation becomes a mixture of sleeve-level value-add rollups and total-fund contribution. As a result, the CE's rational incentive is to avoid tracking-error breaches and minimize career risk, which in practice leads to conformity rather than innovation.

For managers, the dynamic is equally restrictive. Portfolios are run against rigid index yardsticks, and performance is judged strictly on value-add versus those indices, typically with narrow risk bands. A natural tendency is to stay close to benchmarks, given how performance is commonly evaluated. Cross-sleeve collaboration is limited, because each manager is evaluated on their own isolated target rather than the contribution to the total fund.

By contrast, under a **delegated benchmark authority regime**, the distribution of roles changes fundamentally. The board retains responsibility for setting the purpose of the fund, defining the total-fund reference or hurdle and establishing risk and liquidity guardrails. Crucially, however, evaluation now focuses on the CE's ability to deliver outcomes against that total-fund benchmark within the agreed limits. The incentive signal is one of adaptability and coherence: boards prefer total-fund execution that is responsive to changing conditions rather than rigid adherence to static benchmarks.

In this delegated framework, the CE designs fit-for-purpose benchmarks for each mandate, defines roles and incentives, and allocates capital dynamically across sleeves. Evaluation is based on total-fund results, not sleeve-by-sleeve rollups. The incentive structure therefore encourages the CE to maximize total-fund outcomes, pursue orthogonal alpha, and redeploy capital to its best use.

Managers also experience a shift in their evaluation framework. Instead of being judged solely on rigid index comparisons, they are evaluated primarily against mandate-specific benchmarks tailored by the CE. A minority component of their evaluation is linked to total-fund outcomes, which strengthens alignment. The incentive signal changes accordingly: managers are encouraged to deliver idiosyncratic or orthogonal alpha that fits the overall design of the fund, while also collaborating across sleeves rather than working in silos.

In short, under board control, all actors tend to default toward protecting optics: the board over-specifies, the CE conforms, and managers hug benchmarks. Under delegation, authority is redistributed: the board defines the "what" (purpose and guardrails), while the CE defines the "how" (benchmark design and capital allocation). This shift realigns incentives so that each actor's best response supports the total-fund outcome, rather than sleeve-level optics. Figure A-I illustrates the $2 \times 2 \times 2$ Strategic Payoff Cube:

Figure A-I: Strategic Payoff Cube for the Tri-Party Game

Face 1 – Board **controls** asset class benchmarks:

Portfolio Managers (PM)	Chief Executives (CE): Play it Safe (I)	Chief Executives (CE): Pursue Better Alpha (II)
Play it Safe (A)	**Payoff: (2, 2, 2)** Both CE and PM choose caution, producing predictable but unremarkable results. This forms the likely Nash equilibrium: stable, but locked in less than exceptional outcomes. **(Likely) Nash Equilibrium**	**Payoff: (2, 0, 0)** CE attempts to pursue alpha under rigid benchmarks, but without PM support. The result is misalignment: risks rise, benefits are unclear, and the outcome is unstable. **Unstable**
Pursue Alpha (B)	**Payoff: (2, 0, 0)** PM pursues alpha despite rigid benchmarks, but without CE backing. Misalignment prevents fund-wide benefits, leaving the strategy unstable. **Unstable**	**Payoff: (1 to 3, 1, 1)** Both CE and PM pursue alpha, and results could improve, but rigid benchmarks create tension. Without board recognition, risks may be penalized rather than rewarded. The equilibrium is unstable. **Unstable**

Face 2 – Board **delegates** asset class benchmarks:

Portfolio Managers (PM)	Chief Executives (CE): Play it Safe (I)	Chief Executives (CE): Pursue Better Alpha (II)
Play it Safe (A)	**Payoff: (2, 2, 2)** Both CE and PM play it safe, producing only average results. The board may question whether delegation is being fully used, making this outcome unstable. **Unstable**	**Payoff: (2, 2, 1)** CE seeks alpha by redesigning benchmarks, but PM remains cautious. The result is mixed and misaligned, so the equilibrium is unstable. **Unstable**
Pursue Alpha (B)	**Payoff: (2, 1, 2)** PM pursues alpha, but CE has not fully exercised benchmark authority, leaving old board benchmarks in place. This mismatch creates friction, making the outcome unstable. **Unstable**	**Payoff: (3, 3, 3)** CE sets fit-for-purpose benchmarks and PM responds by pursuing orthogonal alpha. The result is aligned outperformance: higher total fund returns, fair rewards for CE and PM, and a Pareto-superior equilibrium. **Pareto-Superior**

Figure: This illustrates the tri-party game between the board, CE, and portfolio managers (PM). Payoffs are shown in the order (Board, CE, PM) on a qualitative scale of 0–3. The board's choice sets the "face" of the cube: either retaining control of benchmarks (Face 1) or delegating benchmark authority to the CE (Face 2). Within each face, CE and PM choose between "Play it Safe" (staying close to benchmarks) or "Pursue Alpha". Under Board Control,

the likely Nash equilibrium is a "safe–safe" outcome with predictable but limited in upside. Attempts at alpha in this regime are unstable because initiative is not recognized under rigid benchmarks. Under Delegation, by contrast, the best responses converge on the "delegate–design–active" configuration: CEs set fit-for-purpose benchmarks, PMs pursue orthogonal alpha, and the board realizes stronger total-fund outcomes. This equilibrium is Pareto-superior, delivering higher payoffs for all parties, while other cells represent unstable or transitional outcomes.

Dominant-Strategy Logic (Best-Response Flow)

When viewed through the lens of best responses, the logic of delegation becomes clear. For the board, if the objective is to maximize beneficiary value and the "control premium" is not unusually high, delegation weakly dominates control. In the worst case, if delegation is underused and managers remain passive, the outcome is essentially the same as the control regime: a baseline, medium-level result. In the best case, however, delegation enables the active–orthogonal equilibrium, yielding high payoffs for all parties. By contrast, control caps the upside at "stable but limited outcomes" while leaving the system vulnerable to benchmark gaming.

From the CE's perspective, the dominant strategy depends directly on the board's choice. Under control, the rational move is to conform to board-set benchmarks and minimize tracking error, since alpha attempts are penalized under rigid evaluation. Under delegation, however, the dominant strategy shifts to designing fit-for-purpose benchmarks. This unlocks the ability to optimize the total fund and align team incentives with overarching goals.

For portfolio managers, the same logic applies. Under control—or when the CE plays it safe because benchmarks have not been reformed—the dominant strategy is to remain passive. Benchmark-hugging protects evaluation and career prospects, even if it adds little to the fund. But under genuine delegation, where benchmarks are tailored by the CE and tied partly to total-fund results, the dominant strategy becomes active pursuit of orthogonal alpha. In this setting, managers are fairly rewarded and protected for taking idiosyncratic risks that contribute to the total-fund design.

Taken together, these best-response dynamics show how delegation weakly dominates control. The worst-case outcome under delegation is no worse than the control regime, while the best-case outcome is strictly superior. For the board, CE, and managers alike, delegation aligns dominant strategies toward cooperation and value creation. The game therefore shifts from the (Control, Conform, Passive) low-alpha equilibrium to the (Delegate, Design, Active) high-alignment equilibrium.