RESEARCH ARTICLE

GOVERNING DUAL-MANDATE SOVEREIGN WEALTH FUNDS: A PRINCIPAL-AGENT THEORY WITH EVIDENCE FROM SAUDI ARABIA AND NORWAY

Ahmad Al-Harbi, Ph.D.
Assistant Professor of Finance
Vice Dean of Quality and Development
Alasala Colleges, Dammam, Saudi Arabia

Abstract : This paper develops a multi-task principal-agent framework to analyze sovereign wealth funds (SWFs) operating under dual mandates—balancing financial returns with developmental objectives. Departing from traditional models, it introduces a risk-channel mechanism with separable effort costs to identify how financial and developmental incentives are jointly determined. The model is applied through a comparative analysis of two contrasting cases: Norway's single-mandate Government Pension Fund Global (GPFG) and Saudi Arabia's dual-mandate Public Investment Fund (PIF). Theoretical propositions indicate that high outcome uncertainty causes developmental incentives to collapse. Empirically, the study finds that between 55% and 89% of PIF's $913 billion assets under management (AUM) stems from government capital injections rather than investment returns. Furthermore, GPFG's geometric mean return (7.87%) modestly outperformed PIF's leverage-adjusted return (6.76–7.03%) from 2017 to 2024. To improve dual-mandate governance, the paper proposes independent performance verification, transparent management cost disclosure, and charter-level leverage discipline.

Keyword: Sovereign wealth funds, dual mandate, principal-agent theory, governance, fiscal rules, Norway GPFG, Saudi PIF

DOI : **10.5281/zenodo.20477769**

Introduction

The global universe of Sovereign Wealth Funds now exceeds $11 trillion in assets under management according to the Sovereign Wealth Fund Institute, yet this aggregate figure masks significant institutional diversity. Norway's Government Pension Fund Global stands as the preeminent international benchmark for the savings fund model. Established in 1990 and first funded in 1996, the fund has grown to approximately $1.74 trillion. Its institutional success is built upon a single-mandate architecture where it is legally required to invest Norway's petroleum revenues exclusively in international markets for the benefit of future generations.

The governance of the fund is defined by several pillars designed to minimize principal-agent risks. Central to this is the Fiscal Rule, a binding parliamentary guideline known as the Handlingsregelen, which ensures that the government only spends the expected real return of the fund—currently set at 3%—while leaving the principal untouched. This mechanism is supported by operational independence; while the Ministry of Finance sets high-level strategy, day-to-day management is handled at arm's length by Norges Bank Investment Management.

Furthermore, the fund maintains a commitment to transparency and ethics, holding a perfect 100/100 score on the Global Pension Transparency Benchmark. It is overseen by an independent Council on Ethics that can recommend the exclusion of companies based on environmental or human rights violations. Over 27 years, this disciplined approach has demonstrated significant efficacy, producing an annualized geometric mean return of 6.36%. This performance is achieved with exceptional cost efficiency, as management costs are kept as low as 0.037% of total assets. Ultimately, this model represents a pure financial mandate where success is measured by long-term risk-adjusted returns relative to a global benchmark, entirely free from the complexity of domestic developmental goals.

At the other extreme sits Saudi Arabia's Public Investment Fund (PIF): restructured in 2015–2016 under Vision 2030, it has grown from approximately $59 billion in 2017 to $913 billion by the end of 2024. A central finding of this paper is that, depending on the timing of capital injections, between 55% and 89% of PIF's $913 billion end-2024 AUM reflects government capital injections rather than investment returns—a result robust to all plausible return assumptions (5–12%). The PIF operates under an explicit dual mandate—financial returns and economic diversification—with concentrated political oversight, limited public disclosure, no fiscal rule, and no independent performance verification mechanism.

These two funds maximise institutional contrast while sharing the common characteristic of being resource-based sovereign investors, making them a natural laboratory for questions of SWF governance design. This paper integrates formal theory with multi-year empirical evidence to answer three questions: What institutional design is optimal for a given resource management problem? How do mandate complexity and governance quality interact to determine investment performance and developmental impact? And what does a formal principal-agent model predict about the conditions under which dual-mandate governance succeeds or fails?

We make three contributions. First, we extend the multi-task agency framework to the SWF context with a separable cost structure, identifying a risk-channel rather than cost-channel interaction between financial and developmental incentives. Second, we produce the first systematic performance comparison of PIF and GPFG over 2017–2024 using correct geometric mean methodology and leverage adjustment, documenting the 55–89% capital-injection range. Third, we derive concrete institutional recommendations grounded directly in the model's propositions.

The paper proceeds as follows. Section 2 reviews the literature. Section 3 describes the data and methodology. Section 4 develops the theoretical model. Sections 5 through 9 conduct the systematic empirical comparison. Section 10 maps theoretical results to empirical findings. Sections 11 and 12 present policy implications and conclusions.

2. Literature Review

2.1 SWF Typology and the Governance-Performance Nexus

Sovereign Wealth Funds are distinguished from conventional central bank reserves by their long investment horizon, risk-bearing capacity, and separation from daily monetary operations (IMF, 2008). The foundational taxonomy of Rozanov (2005) and the IFSWF Santiago Principles (2008) identifies five fund types – stabilisation, savings, reserve investment, development finance, and pension reserve – which Al-Hassan et al. (2013) extend to show that governance architecture is endogenous to mandate type. Truman (2010) produces the first systematic SWF governance scorecard, placing GPFG at the top on transparency and accountability. Clark and Monk (2012) argue that governance quality is a first-order determinant of long-run returns, confirmed empirically by Bernstein, Lerner and Schoar (2013): politically motivated SWF investments generate negative abnormal returns while commercially motivated investments generate positive returns. Bortolotti, Fotak and Megginson (2015) confirm this finding in a larger cross-sectional study. Dyck and Morse (2011) document that transparency is positively associated with SWF investment performance, and Fernandes (2014) shows that SWF equity holdings improve target-firm value when governance is strong.

2.2 The Resource Curse and Sovereign Saving

The resource curse hypothesis – that resource abundance tends to impair institutional quality and growth (Sachs and Warner, 1995; Mehlum, Moene and Torvik, 2006) – motivates SWF design. Dutch disease (Gylfason,

2001) and the political economy of resource rents (van der Ploeg and Venables, 2011) provide the mechanisms. Norway represents the canonical success story: Holden (2013) documents how the GPFG combined with a parliamentary fiscal rule allowed Norway to escape the resource curse despite major petroleum revenues. Saudi Arabia's challenge is structurally harder: a younger and larger population, a tighter political timeline for transformation, and oil revenues that must simultaneously fund current expenditure and long-term diversification (El-Katiri, 2016; Hvidt, 2013).

2.3 Principal-Agent Theory in Delegated Investment

The governance of SWFs is a multi-layered principal-agent problem (Jensen and Meckling, 1976; Holmström, 1979). Matsen (2017) observes that the Santiago Principles are explicitly derived from principal-agent theory. For single-mandate funds, the problem is tractable – financial benchmarks provide clear performance metrics. For dual-mandate funds, Holmström and Milgrom (1991) demonstrate the multi-task moral hazard: when agents pursue multiple tasks from a shared effort cost pool, incentive provision on observable tasks can undermine effort on unobservable ones. Our model departs from this by specifying separable costs; the incentive interaction is then generated through correlated output risk rather than shared cost. The key result in our Proposition 2 – that developmental incentives collapse under high outcome uncertainty – is an application of Holmström's (1979) informativeness principle to the developmental task in a multi-task environment. The mathematical result itself is not novel; the contribution lies in its application to the institutional black box of sovereign development funds, where the relevant 'noise' is the opacity and complexity of mega-project outcomes rather than the measurement noise of financial returns. Ang (2014) provides the most comprehensive principal-agent treatment of institutional investment relevant to SWF governance.

2.4 ESG and Responsible Sovereign Investment

Richardson (2011) argues that SWFs are universal owners whose returns depend on the health of the global economy, creating structural ESG incentives. Halvorssen (2023) documents the evolution of GPFG's ethical framework as the global benchmark for responsible sovereign investing. Knill, Lee and Mauck (2012) confirm empirically that politically motivated SWF investments underperform commercially motivated ones – consistent with principal-agent predictions when governance mechanisms fail to align incentives.

2.5 Saudi Arabia's PIF in the Literature

Academic literature on the PIF remains sparse, reflecting its limited transparency and recent transformation. El-Katiri (2016) analyses Vision 2030's design; Hvidt (2013) situates Gulf diversification strategies in comparative perspective. Our own assessment using the Truman (2010) framework places PIF at approximately 40% of the GPFG benchmark across transparency, accountability, and ethics dimensions – consistent with the limited available ratings.

3. Data and Methodology

3.1 Comparative Design

We employ a structured most-different-cases comparative design (George and Bennett, 2005; Przeworski and Teune, 1970). Norway and Saudi Arabia share resource-based SWF status but differ maximally across political system, development stage, mandate architecture, and governance model – maximizing the likelihood that observed differences reflect structural rather than context-specific mechanisms. The principal-agent model serves as the unifying analytical lens across five dimensions: mandate architecture, governance structure, portfolio strategy, financial performance, and ESG integration. Crucially, the two-fund design is theory-motivated illustration rather than a formal model test; with n = 2 cases no causal identification of governance effects is possible.

3.2 Data Sources

For GPFG: NBIM Annual Reports 1998–2024 (nbim.no), all audited and GIPS-compliant; Norwegian Ministry of Finance Government Pension Fund White Papers 2001–2024; and the Global Pension Transparency Benchmark 2024 (CEM Benchmarking/Willis Towers Watson). For PIF: PIF Annual Reports 2022–2024 (pif.gov.sa); Fitch Credit Report December 2024; and Saudi Press Agency official releases.

Supplementary data series draw on the following sources: MSCI EM net returns from the MSCI factsheet (msci.com); Tadawul All-Share Index returns from Saudi Exchange annual closing levels; Brent crude prices from EIA annual averages and IMF Primary Commodity Prices; Saudi fiscal break-even oil prices from IMF Regional Economic Outlook; Saudi GDP growth from IMF Article IV Consultation reports.

A critical data asymmetry must be acknowledged. GPFG returns are audited, benchmark-adjusted, GIPS-compliant, and cover 27 complete annual observations spanning multiple market cycles (1998–2024). PIF returns are a single annualized aggregate since 2017, without benchmark comparison, sub-portfolio attribution, or management cost disclosure. The 2023 Annual Report cited 8.7% annualized since inception; the 2024 Annual Report revised this to 7.2%. All performance comparisons account for this asymmetry explicitly.

3.3 Decomposing PIF's AUM Growth: Timing Scenarios

The future value of all injected capital (including returns) is computed as:

$$FV_{\text{injections}} = \$913 - \$59 \times (1.072)^7 = \$817 \text{ billion.}$$

Because PIF does not disclose injection timing, we recover the raw cumulative capital injected (excluding returns) under three polar timing assumptions:

- Back-loaded – all injections occur at end-2024 (earn zero returns). Raw injection = \$817 bn → 89% of 2024 AUM.
- Linear annuity – equal injections at the start of each year 2018–2024. Raw injection ≈ \$657 bn → 72% of 2024 AUM.
- Front-loaded – single injection at start-2017, compounding for 7 years. Raw injection ≈ \$502 bn → 55% of 2024 AUM.

Thus the injection share lies between 55% and 89% depending on the unknown true timing. Table 1 presents this range across different return assumptions.

Table 1. PIF Capital Injection Sensitivity: Timing and Return Scenarios

Return Assumption	Front-loaded (LB)	Linear Annuity (Mid)	Back-loaded (UB)
5.0%	61.2%	75.3%	90.9%
7.2% (reported)	55.0%	67.1%	89.5%
10.0%	47.3%	58.2%	87.4%
12.0%	41.9%	51.5%	85.7%

Notes: LB = lower bound (front-loaded); UB = upper bound (back-loaded). Raw cumulative injected capital as % of end-2024 AUM. Source: authors' calculations.

Figure 1 visualizes the injection share range across return assumptions and timing scenarios.

Figure 1 Injection range line chart (front-loaded, annuity, back-loaded lines; shaded area between bounds; vertical line at 7.2%.)

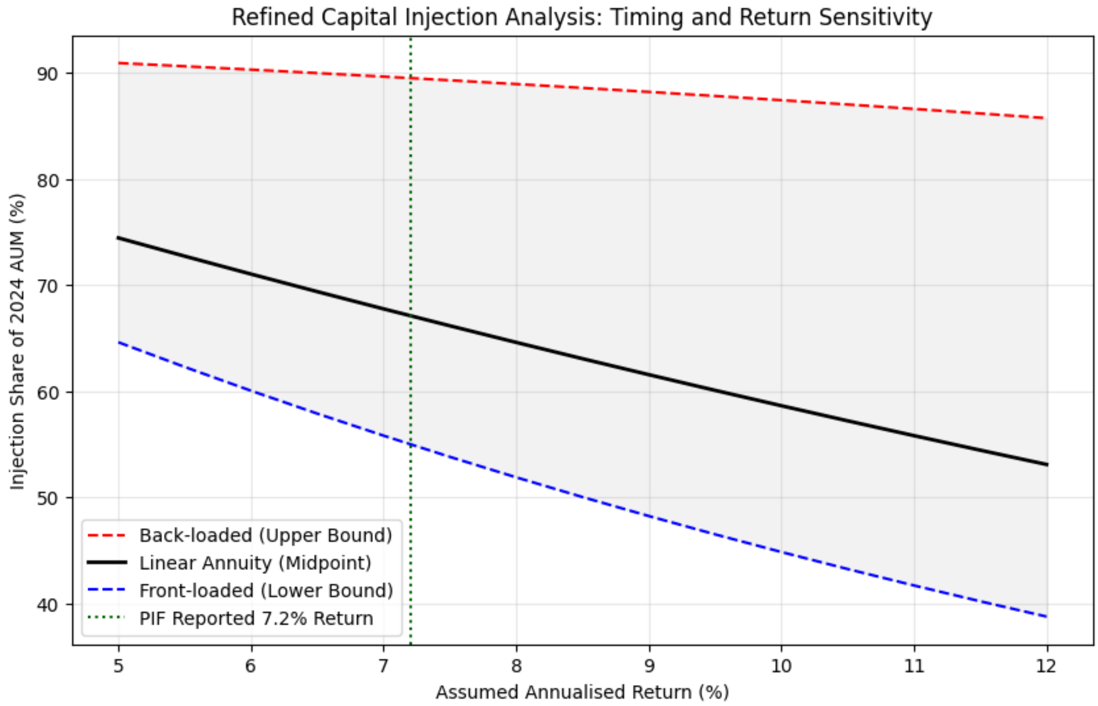

4. The Multi-Task Principal-Agent Model

4.1 Environment

A principal (the sovereign government) delegates investment management to an agent (the SWF manager). The agent exerts effort on two tasks: financial portfolio performance (effort e) and developmental programme effort (effort a). The output technology is:

$$R = e + \varepsilon_r \text{(financial return)} \qquad (1)$$

$$S = ba + \varepsilon_s \text{(developmental outcome)} \qquad (2)$$

where R is the observable financial return, S is the developmental outcome (e.g., non-oil GDP contribution), $b > 0$ is the agent's developmental productivity, and the shocks are jointly normal: $(\varepsilon_r, \varepsilon_s) \sim N(0, \Sigma)$, with covariance matrix $\Sigma = [[\sigma_r^2, \sigma_{rs}], [\sigma_{rs}, \sigma_s^2]]$. The covariance σ_{rs} captures the co-movement of financial and developmental shocks. We estimate σ_{rs} from MSCI EM returns vs Saudi non-oil real GDP growth (2016–2023): $\rho \approx -0.524$ $(p \approx 0.18, n = 8)$, yielding $\sigma_{rs} \approx -0.031$ under the paper's scale normalisation $k = 13.97$. This estimate is directional; the 95% bootstrap confidence interval spans zero ([-0.70,

+0.85]), so the sign of σ_{rs} is not statistically identified. All results that depend on σ_{rs} are therefore presented as conditional theoretical predictions.
Figure 2 illustrates the robustness of the σrs estimate.
Figure 2σrs estimation robustness: (a) data scatter, (b) bootstrap histogram with 95% CI crossing zero, (c) model outputs across ρ, (d) risk premium comparison

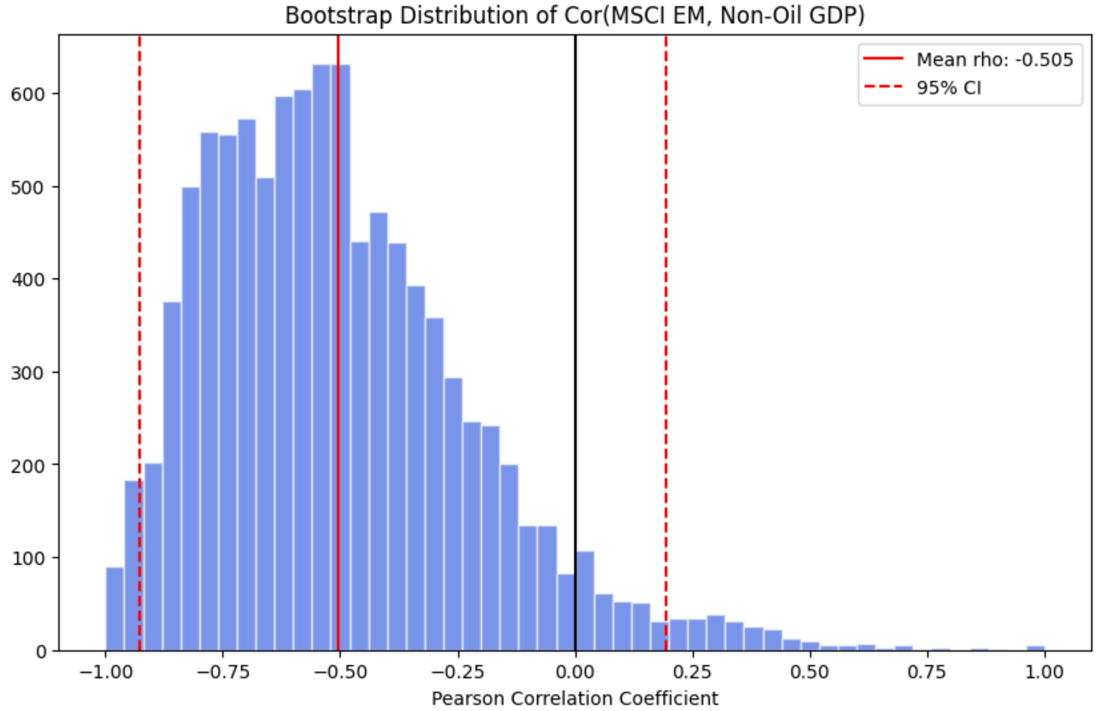

4.2 Contract and Preferences

The agent is risk-averse with CARA coefficient $r > 0$ and has a separable quadratic effort cost:

$$c(e,a) = \frac{1}{2}e^2 + \frac{\delta}{2}a^2 \qquad (3)$$

Separability is the key departure from Holmström–Milgrom (1991): with a shared cost pool, incentivizing one task directly raises the marginal cost of the other. In large SWFs such as PIF, distinct internal teams handle financial portfolio management (CIO office) and developmental program oversight (sector development teams), so effort on developmental due diligence need not crowd out financial analysis effort directly. For PIF at scale – where the 103 newly-created portfolio companies are managed by dedicated sector teams – the assumption is defensible as an approximation. With separable costs, tasks are not cost-substitutes; the incentive interaction arises entirely through the correlated risk premium – a risk-channel mechanism. The linear contract is $w = \alpha + \beta R + \gamma S$, where α is the fixed payment (set by the binding participation constraint), β is the financial incentive weight, and γ is the developmental incentive weight. The principal values developmental outcomes at rate $\lambda \geq 0$. For PIF, λ is politically determined and potentially unstable in the absence of a binding fiscal rule.

4.3 Propositions

Proposition 1 (Optimal Contract). The agent's incentive-compatibility conditions yield $e^* = \beta$ and $a^* = \gamma b/\delta$. The unique optimal incentive weights satisfy the joint linear system:

$$\beta^*(1 + r\sigma_r^2) + r\gamma^*\sigma_{rs} = 1$$
$$\gamma^*\left(\frac{b^2}{\delta} + r\sigma_s^2\right) + r\beta^*\sigma_{rs} = \lambda\frac{b^2}{\delta}$$

Financial and developmental incentives are jointly determined through the correlated risk term $r\sigma_{rs}$: when $\sigma_{rs} < 0$ the incentives are complements; when $\sigma_{rs} > 0$ they are substitutes.
Proposition 2 (Informativeness Principle Applied to the Developmental Task).

$$\frac{\partial\gamma^*}{\partial\sigma_s^2} = -\frac{r\gamma^*}{b^2/\delta + r\sigma_s^2} < 0$$

As $\sigma_s^2 \to \infty$, $\gamma^* \to 0$: the optimal developmental incentive vanishes regardless of the social weight λ placed on development. This is an application of Holmström's (1979) informativeness principle to the developmental task

in our multi-task setting. With separable costs, the developmental incentive depends on σ_s^2 exactly as in the single-task problem; the σ_{rs} cross-term enters only as a second-order correction. The financial incentive converges to $\beta^* = 1/(1 + r\sigma_r^2) > 0$.
Proposition 3 (Governance Structure Preference). Denote

$$RP_{\text{Int}} = \frac{1}{2} r(\beta^{*2}\sigma_r^2 + \gamma^{*2}\sigma_s^2 + 2\beta^*\gamma^*\sigma_{rs})$$

and

$$RP_{\text{Sep}} = \frac{1}{2} r(\beta^{*2}\sigma_r^2 + \gamma^{*2}\sigma_s^2)$$

as the agent's risk premia under integrated and separated governance respectively. Then $RP_{\text{Int}} < RP_{\text{Sep}}$ if and only if $\sigma_{rs} < 0$. When mandate risks are negatively correlated, integrated governance is preferred; when positively correlated, separation is preferred. For $\sigma_{rs} = -0.031$ (Scenario A), integration saves approximately 18% of the risk premium ($RP_{\text{Int}} = 0.118$ vs $RP_{\text{Sep}} = 0.144$). This result holds conditional on the sign of σ_{rs}; since the 95% bootstrap confidence interval spans zero, the sign is not empirically identified. The 18% saving is illustrative of the mechanism's magnitude under the calibrated parameters, not an estimate of the actual saving realised by PIF.

4.4 Adverse Selection Extension

When the agent's developmental productivity $b \in \{b_L, b_H\}$ is private information, standard mechanism design (Mirrlees, 1971) yields a closed-form separating contract. The optimal γ_L^{AS} for the low type satisfies:

$$\gamma_L^{\text{AS}} = \frac{\lambda b_L^2/\delta}{\frac{b_L^2}{\delta} + r\sigma_s^2 + \frac{\pi}{1-\pi} \cdot \frac{b_H^2 - b_L^2}{\delta}} \qquad (4)$$

where π is the principal's prior probability of the high type. This expression shows that $\gamma_L^{\text{AS}} < \gamma^*(b_L)$ (downward distortion from first-best), with the distortion increasing in π and in the type gap $b_H - b_L$. Corner solutions ($\gamma_L^{\text{AS}} = 0$) emerge at high π and large type gaps. Applied to PIF: of the 225 total portfolio companies, 103 are PIF-created since 2017 (agents with initially unknown developmental productivity b). The adverse selection model predicts graduated contracts for these companies – lower initial developmental incentives with upward revision as sector productivity is revealed through operational performance.

4.5 Calibration

Parameters are estimated from publicly available data or set as calibration assumptions. Financial return volatility: $\sigma_r^2 = 0.50$ (normalised from MSCI EM 2016–2023 using scale factor $k = 13.97$; see Appendix A, Table A2 for k-sensitivity). Developmental outcome volatility: $\sigma_s^2 \in \{0.0087, 1.759, 3.50\}$ across scenarios (IMF Saudi non-oil GDP variance; KPI variance; extreme-uncertainty mega-project). Covariance: $\sigma_{rs} = -0.031$ (from MSCI EM × non-oil GDP, $n = 8, p = 0.18$; directional). Risk aversion: $r = 2.0$ (Hall and Murphy, 2002). Social weight: $\lambda = 1.5$ (calibration assumption). Effort-cost ratio: $b = 1.2, \delta = 1.5$. Table 2 presents results for six scenarios.

Table 2. Calibration Scenarios and Model Outputs

Scenario	σ_s^2	β^* (Fin)	γ^* (Dev)	Integration Saving
A: Baseline (non-oil GDP)	0.0087	0.547	1.508	18.0%
B: Mega-project / KPI	1.759	0.510	0.329	3.2%
C: Extreme-uncertainty mega-project	3.500	0.506	0.185	1.1%
D: Fiscal Deficit (low λ)	0.0087	0.516	0.524	12.4%
E: Fiscal Surplus (high λ)	0.0087	0.542	1.533	35.8%

Note: The shift from D to E represents a 192.6% variation in γ^* over the oil cycle. The saving is relative to separated governance under same parameters.

Figure 4 demonstrates the collapse of developmental incentives under extreme uncertainty (Proposition 2)

Figure 4 Proposition 2 stress test: γ vs σ_s^2, showing collapse at σ_s^2=3.5

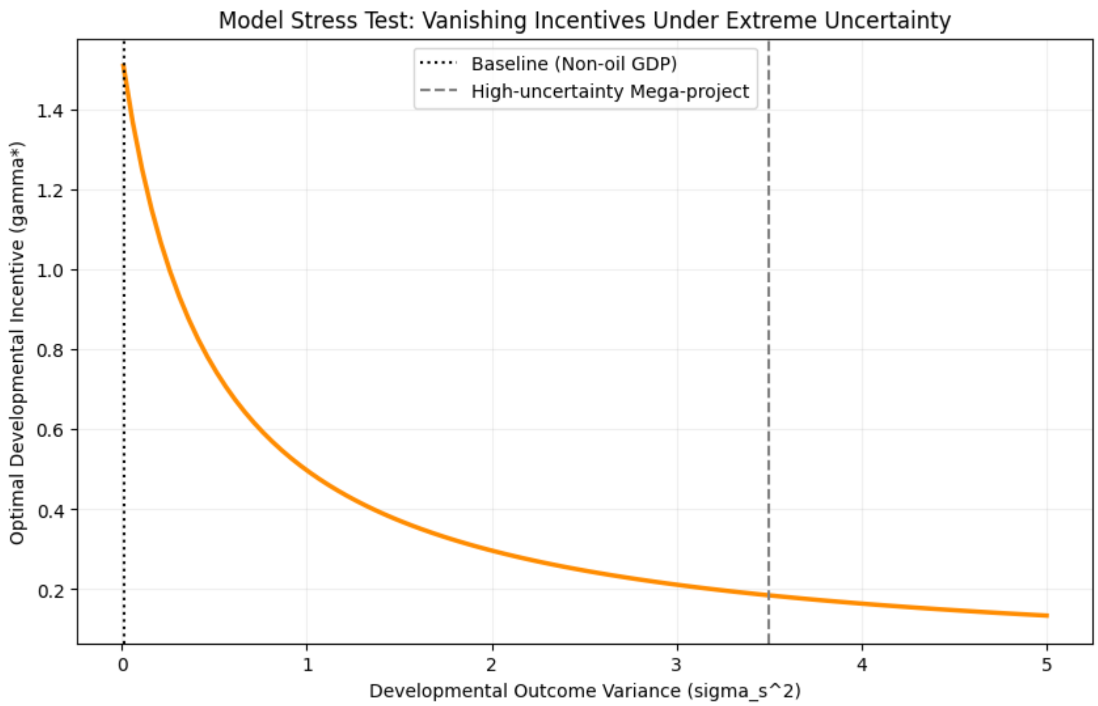

The sensitivity of β^* and γ^* to each parameter is shown in Appendix B, Figures 1–2.

5. Mandate Architecture

5.1 GPFG – The Single-Mandate Savings Model

The GPFG was established under the Government Petroleum Fund Act of 1990, with its mandate singular and legally codified: invest Norway's petroleum revenues internationally, with domestic investment prohibited by statute. This prohibition collapses the principal-agent problem to the single-mandate case (effectively $\lambda = 0$), making the governance problem fully tractable. The 3% fiscal rule (introduced 2001, reduced from 4% in 2017) locks λ at zero institutionally.

5.2 PIF – The Dual-Mandate Transformation Model

The PIF was established in 1971 as a domestic development finance institution and transformed under Vision 2030 into a dual-mandate vehicle simultaneously tasked with competitive financial returns (targeted 7–9% annualized) and serving as 'the engine of economic diversification' for Saudi Arabia. The dual mandate is operationalized through approximately 80% domestic allocation, creation of 103 companies since 2017, and a cumulative non-oil GDP contribution of $243 billion (2021–2024). In the model's terms, $\lambda = 1.5$ reflects the social premium placed on non-oil GDP beyond its direct financial value.

Table 3 compares the mandate architectures of the two funds.

Table 3. Mandate Architecture Comparison

Dimension	GPFG	PIF
Legal basis	Parliamentary statute	Royal decree
Mandate type	Single (savings)	Dual (financial + developmental)
Domestic investment	Prohibited	~80%
Fiscal rule	Binding 3%	Absent
Developmental weight λ	0	~1.5 (politically determined)

6. Governance Structure

6.1 GPFG – Parliamentary Anchoring

The GPFG's governance hierarchy runs from the Storting (Parliament) through the Ministry of Finance to NBIM as day-to-day manager. Key features: the 3% fiscal rule; the Council on Ethics (independent; binding exclusion powers; established 2004; public list of 200+ excluded companies); a tracking error constraint of 1.25%; and annual parliamentary white paper review. Management cost amounts to 0.037% of AUM – among the lowest for any large institutional investor globally. The information ratio averages approximately +0.39 over 1998–2024, computed as mean active return (0.315pp) divided by NBIM's reported ex-ante tracking error (~80bp annualized). Both measures confirm consistent positive active returns over the majority of years.

6.2 PIF – Concentrated Governance

PIF's governance concentrates authority in the Crown Prince as chair of the Supreme Authority, with the Board comprising key ministers. This structure enables strategic agility but creates the principal-agent risks identified in Section 4: concentrated authority without independent measurement. PIF holds Fitch A+ (stable, long-term issuer default) and Moody's A1 (positive, senior unsecured) – approximately equivalent credit quality on different scales – reflecting implicit sovereign support rather than standalone institutional strength.

A note on governance endogeneity is warranted: PIF's concentrated structure and absent fiscal rule could represent either (a) political capture that reduces institutional quality, or (b) a deliberately agile design to meet the urgent and politically binding diversification timeline of Vision 2030, where the slower deliberative processes of the GPFG model would be strategically counterproductive. This paper does not adjudicate between these interpretations; it provides a formal framework within which either can be evaluated given additional evidence on the principal's true objective function.

Table 4 presents the governance scorecard.

Table 4. Governance Scorecard

Dimension	GPFG	PIF
Political insulation	Arm's-length (MoF → NBIM)	Crown Prince chairs Supreme Authority
Independent ethics	Council on Ethics (binding)	None
Transparency	100/100 benchmark, full holdings	Annual report only
Cost disclosure	0.037% of AUM	Not disclosed
Composite score	~95/100 (Truman)	~40/100 (authors' assessment)

7. Portfolio Strategy

7.1 GPFG – Global Passive Diversification

The GPFG owns an average of 1.5% of every listed company in the world – 8,659 companies across 71 countries at end-2024. Asset allocation (71.4% equities, 26.6% fixed income, 1.8% real estate, 0.1% renewable infrastructure) is set by the Ministry of Finance through a statutory benchmark revised via parliamentary white papers. Management is primarily passive (FTSE Global All Cap benchmark) with a tracking error limit of 1.25% constraining active discretion.

7.2 PIF – Concentrated Domestic Active Ownership

PIF's portfolio inverts GPFG's across almost every dimension: approximately 80% domestic, active controlling-stake ownership, new company creation from scratch, and balance-sheet leverage. Portfolio allocation at end-2024 comprises Saudi equity holdings (36%), Saudi sector development (30%), Saudi real estate and infrastructure (7%), giga-projects (6%), international strategic investments (8%), and treasury/other (13%).

Table 5 contrasts the portfolio strategies.

Table 5. Portfolio Strategy Comparison

Dimension	GPFG	PIF
AUM (end-2024)	\$1.744 tn	\$913 bn
AUM at 2017 base	\$1,034 bn	\$59 bn
Geographic allocation	100% international, 71 countries	~80% Saudi, ~17% international
Holdings	8,659 companies, 6,934 bonds	225 companies (103 PIF-created)
Leverage	None	Bonds + sukuk (A+/A1)

8. Financial Performance

8.1 GPFG – Twenty-Seven Years of Audited Returns

The GPFG's return history from 1998 to 2024 is the most comprehensive and credibly benchmarked long-run performance record available for any SWF. Over 27 years, the fund generated a 6.36% geometric mean return (CAGR) with an average active return over benchmark of approximately +0.315 percentage points.The lower full-period CAGR (6.36%) reflects early-stage scaling and the 2008/2022 drawdowns, while the 2017–2024 window aligns with PIF's reporting horizon for comparability. The record spans four major crises: the dot-com collapse (−4.7% in 2002), the Global Financial Crisis (−23.3% in 2008, the worst single year), the eurozone crisis (−2.5% in 2011), and the 2022 inflation shock (−14.1%, the first year of simultaneous equity-bond loss in fund history). The full annual return table is provided in Appendix A, Table A1 (online supplement). For illustration: in 1998 the fund returned 10.0% (active +0.8pp), in 2008 –23.3% (active 0.0pp), in 2017 13.7% (active +0.5pp), and in 2024 13.1% (active –0.4pp).

8.2 PIF – Capital Injection Finding

As shown in Table 1, the injection share ranges from 55% to 89% depending on timing. Table 6 shows PIF's AUM milestones.

Table 6. PIF AUM Growth and Milestones, 2017–2024

Year	AUM	Reported Return	Milestones
2017	~$59bn	—	Restructuring; Vision 2030 launch
2018	~$200bn	—	NEOM announced
2019	~$320bn	—	Aramco IPO anchor
2020	~$400bn	—	COVID purchases
2021	~$620bn	—	Domestic surge
2022	~$700bn	—	International correction
2023	$765bn	8.7% (since 2017, revised)	29 new companies
2024	$913bn	7.2% (since 2017)	225 companies, 103 created

8.3 Comparative Performance Assessment

GPFG's 2017–2024 geometric mean = 7.87%. PIF's reported 7.20% is levered. Applying WACC unlevering (cost of debt 4.14%, D/E 20–30%) yields PIF unlevered return of 6.50–6.69%. Table 7 summarizes the comparison.

Table 7. Comparative Performance Summary (2017–2024)

Metric	GPFG (Norway)	PIF (Saudi Arabia)
Starting AUM (2017)	$1,034 billion	$59 billion
Ending AUM (2024)	$1,744 billion	$913 billion
Nominal Return (geom. mean)	7.87%	7.20% (self-reported)
Unlevered Return	7.87%	6.50% – 6.69%
Management Cost	0.037% of AUM	Not disclosed

Figure 3 compares GPFG's annual returns with PIF's reported return and shows leverage-adjusted PIF returns. Figure 3 GPFG geometric mean computed from NBIM annual reports (see Appendix A, Table A1 for the full series)

Figure 3. GPFG vs PIF Financial Performance Comparison (2017–2024)

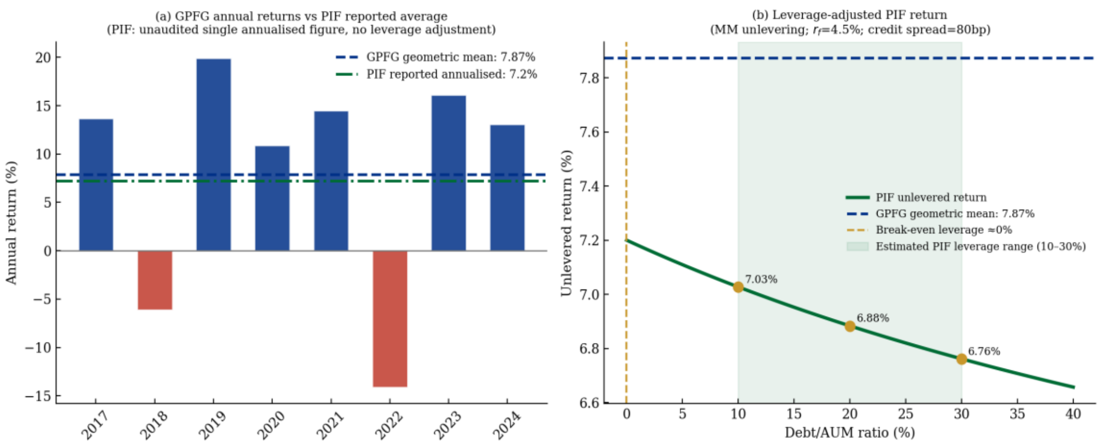

Notes: see Appendix A, Table A1.

9. ESG Integration and Responsible Investment

9.1 GPFG – The Institutional Benchmark

The GPFG's ESG framework – Council on Ethics with binding exclusion powers, full voting record disclosure (~12,000 AGMs annually, published within 24 hours), TCFD-aligned climate reporting, and a 2025 Climate Action Plan – represents the most institutionally robust responsible investment architecture among sovereign investors globally (Halvorssen, 2023). The framework requires ongoing institutional maintenance; it is not a permanent achievement but rather an evolving governance structure subject to parliamentary scrutiny.

9.2 PIF – Commitments Without Architecture

PIF has published ESG commitments including net-zero by 2060 (aligned with Saudi Arabia's NDC) and investments in renewable energy through ACWA Power. The 2023 green bond ($5.5 billion) signals ESG intent. However, the institutional architecture remains weak: no independent ethics body, no public exclusion list, no voting record disclosure, and no independent verification. Saudi Aramco (~16% of portfolio) represents a structural tension with the net-zero commitment that the fund has not publicly addressed. The extreme uncertainty surrounding mega-project outcomes (e.g., projects with outcome variance σ_s^2 = 3.5) is used as a stress-test calibration for Proposition 2, not as an empirical claim about any specific project.

10. Theory-to-Evidence Mapping

Table 8 maps each theoretical result to empirical observations.

Table 8. Theory-to-Evidence Mapping

Theoretical Result	Model Value	Empirical Illustration
P1: β* and γ* jointly determined via σrs	β*=0.547, γ*=1.508 (Scenario A)	GPFG: pure β incentive; PIF: no credible γ-linked accountability
P2: High σs² drives γ*→0	γ*=0.185 at σs²=3.5	Extreme-uncertainty mega-project – near-zero developmental incentive
P3: σrs< 0 favours integrated governance	RP_Int=0.118, RP_Sep=0.144; saving 18%	ρ̂=−0.524, 95% CI spans zero → conditional theoretical result
λ instability	∂γ*/∂λ≈0.984	No fiscal rule: γ* varies 192.6% over oil cycle (0.524→1.533)
Adverse selection	γ^AS_L< γ*(bL)	103 PIF-created companies – graduated contracts predicted

10.1 The Extreme-Uncertainty Prediction (Proposition 2)

With $\sigma_s^2 = 3.5, \gamma^* \approx 0.185$ – near zero. This is a structural prediction about incentive collapse when outcomes are unmeasurable, not a judgement about any specific project.

10.2 λ-Instability and the Capital Injection Finding

The absence of a fiscal rule causes λ to vary with oil prices. γ* varies by 192.6% from fiscal deficit years (0.524) to surplus years (1.533). This instability undermines both financial discipline and developmental commitment.

Figure 5 shows the λ-instability simulation over the oil price cycle, illustrating the 192.6% variation in γ*

Figure 5 λ-instability simulation: γ over 2016–2023, showing 192.6% variation.

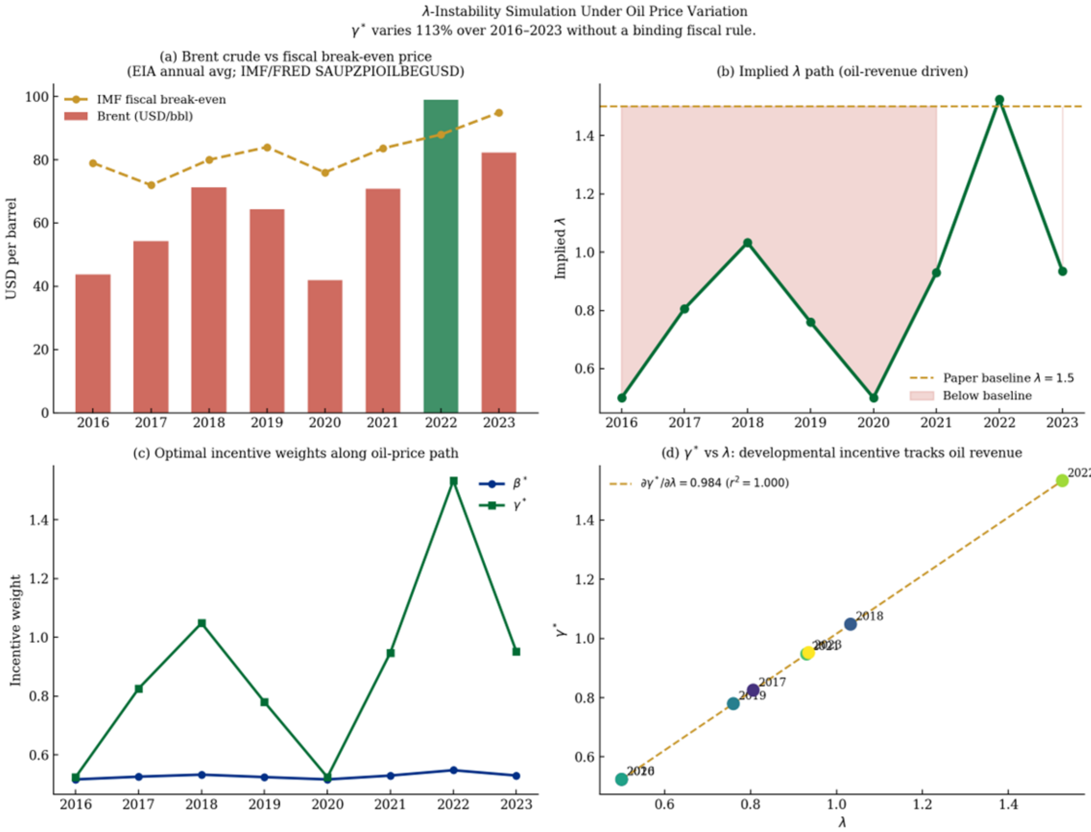

10.3 Integrated Governance and the σrs Result

The sign of σrs is not identified (95% CI includes zero). Proposition 3 is therefore a conditional theoretical principle: if σrs< 0, integration is preferred. It is not an empirical validation of PIF's current structure.Thus, Proposition 3 provides a normative benchmark rather than a diagnostic of PIF's actual risk correlation.

11. Policy Implications

11.1 Three Model-Grounded Improvements for PIF

1. **Independent developmental performance verification** – An external body (like GPFG's Council on Ethics) to verify non-oil GDP claims and project cost-efficiency. This directly addresses the γ-channel failure (Proposition 2).

2. **Management cost disclosure** – At 913 billion AUM, even a 0.1913 million in annual drag on returns – a compounding performance penalty that is currently invisible to principals and external observers. GPFG's 0.037% cost provides a benchmark.
3. **Charter-level leverage discipline** – A graduated counter-cyclical rule (higher leverage in surpluses, binding floor in deficits) to stabilize λ and protect against pro-cyclical cuts.

11.2 General Design Taxonomy

Four-dimensional SWF design space: (i) mandate breadth determining governance tractability; (ii) fiscal rule strength (λ stability); (iii) governance integration vs separation (optimised by sign of σrs); (iv) outcome measurement architecture determining whether developmental incentives are credible. GPFG occupies the low-complexity corner; the optimal dual-mandate SWF would require all four dimensions aligned.

12. Conclusion

This paper has developed a unified framework for evaluating dual-mandate SWFs. The central robust empirical contribution – independent of model assumptions and governance debates – is that depending on injection timing, between 55% and 89% of PIF's $913 billion AUM reflects government capital injections rather than investment returns (mid-point ~72%). Fund-building must be distinguished from fund-performance.

GPFG solved the intergenerational savings problem with optimal single-mandate design. PIF faces a harder structural transformation problem, with critical gaps: λ instability (γ* varies 192.6% over oil cycle), γ ineffectiveness (Proposition 2), and opacity from undisclosed costs. Three model-grounded improvements are proposed.

Future research should extend the framework to a broader panel of dual-mandate SWFs (Temasek, Khazanah, Mubadala) and, as disclosure improves, test the governance-performance relationship with formal panel econometrics.

References

Al-Hassan, A., Papaioannou, M., Skancke, M. and Sung, C.C. (2013), "Sovereign wealth funds: aspects of governance structures and investment management", *IMF Working Paper*, No. WP/13/231.

Ang, A. (2014), *Asset Management: A Systematic Approach to Factor Investing*, Oxford University Press, Oxford.

Bernstein, S., Lerner, J. and Schoar, A. (2013), "The investment strategies of sovereign wealth funds", *Journal of Economic Perspectives*, Vol. 27 No. 2, pp. 219-238.

Bortolotti, B., Fotak, V. and Megginson, W.L. (2015), "The sovereign wealth fund discount: evidence from public equity investments", *Review of Financial Studies*, Vol. 28 No. 11, pp. 2993-3035.

CEM Benchmarking/Willis Towers Watson (2024), *Global Pension Transparency Benchmark 2024*, CEM Benchmarking, Toronto.

Clark, G.L. and Monk, A.H. (2012), "Sovereign wealth funds: form and function in the 21st century", in Sauvant, K. et al. (Eds), *Sovereign Investment: Concerns and Policy Reactions*, Oxford University Press, Oxford.

Dyck, A. and Morse, A. (2011), "Sovereign wealth fund portfolios", *Chicago Booth Research Paper* No. 11-15.

El-Katiri, L. (2016), "Saudi Arabia's Vision 2030 and the challenge of economic transformation", *Oxford Institute for Energy Studies Insight*, Oxford.

Fernandes, N. (2014), "The impact of sovereign wealth funds on corporate value and performance", *Journal of Applied Corporate Finance*, Vol. 26 No. 1, pp. 76-84.

Fitch Ratings (2024), *Rating Report: Public Investment Fund (PIF)*, December.

George, A.L. and Bennett, A. (2005), *Case Studies and Theory Development in the Social Sciences*, MIT Press, Cambridge, MA.

Gylfason, T. (2001), "Natural resources, education, and economic development", *European Economic Review*, Vol. 45 Nos 4-6, pp. 847-859.

Hall, B.J. and Murphy, K.J. (2002), "Stock options for undiversified executives", *Journal of Accounting and Economics*, Vol. 33 No. 1, pp. 3-42.

Halvorssen, A.M. (2023), "How the Norwegian SWF balances ethics, ESG risks, and returns", in *Pension Funds and Sustainable Investment*, Oxford University Press, Oxford.

Holden, S. (2013), "Avoiding the resource curse: the case Norway", *Energy Policy*, Vol. 63, pp. 870-876.

Holmström, B. (1979), "Moral hazard and observability", *Bell Journal of Economics*, Vol. 10 No. 1, pp. 74-91.

Holmström, B. and Milgrom, P. (1991), "Multitask principal-agent analyses: incentive contracts, asset ownership, and job design", *Journal of Law, Economics, and Organization*, Vol. 7, Special Issue, pp. 24-52.

Hvidt, M. (2013), "Economic diversification in GCC countries: past record and future trends", *LSE Kuwait Programme Research Paper Series*, No. 27.

IFSWF (2008), *Sovereign Wealth Funds: Generally Accepted Principles and Practices (Santiago Principles)*, IFSWF, Washington, DC.

IMF (2008), *Sovereign Wealth Funds: A Work Agenda*, International Monetary Fund, Washington, DC.
Jensen, M.C. and Meckling, W.H. (1976), "Theory of the firm: managerial behavior, agency costs and ownership structure", *Journal of Financial Economics*, Vol. 3 No. 4, pp. 305-360.
Knill, A., Lee, B.S. and Mauck, N. (2012), "Bilateral political relations and sovereign wealth fund investment", *Journal of Corporate Finance*, Vol. 18 No. 1, pp. 108-123.
Matsen, E. (2017), "Government Pension Fund Global: experiences and organisation", *BIS Review*, No. R171030A, Norges Bank, Oslo.
Mehlum, H., Moene, K. and Torvik, R. (2006), "Institutions and the resource curse", *Economic Journal*, Vol. 116 No. 508, pp. 1-20.
Mirrlees, J.A. (1971), "An exploration in the theory of optimum income taxation", *Review of Economic Studies*, Vol. 38 No. 2, pp. 175-208.
Norges Bank Investment Management (NBIM) (1999-2024), *Annual Reports 1998-2024*, NBIM, Oslo.
PIF (2022-2024), *Annual Reports 2022, 2023, 2024*, Public Investment Fund, Riyadh.
Richardson, B.J. (2011), "From fiduciary duties to fiduciary relationships for socially responsible investing", *Investment Management and Financial Innovations*, Vol. 8 No. 3, pp. 196-207.
Rozanov, A. (2005), "Who holds the wealth of nations?", *Central Banking Journal*, Vol. 15 No. 4, pp. 52-57.
Sachs, J.D. and Warner, A.M. (1995), "Natural resource abundance and economic growth", *NBER Working Paper*, No. 5398.
SWFI (2023), *Fund Rankings*, Sovereign Wealth Fund Institute.
Truman, E.M. (2010), *Sovereign Wealth Funds: Threat or Salvation?*, Peterson Institute for International Economics, Washington, DC.
van der Ploeg, F. and Venables, A.J. (2011), "Harnessing windfall revenues: optimal policies for resource-rich developing economies", *Economic Journal*, Vol. 121 No. 551, pp. 1-30.

Appendix A – Robustness Tables

Table A1. Parameter Sensitivity Analysis

Parameter	Value	σ_r^2	σ_s^2	σ_{rs}	β^*	γ^*
λ	0.5	0.50	0.0087	–0.031	0.516	0.524
	1.0	0.50	0.0087	–0.031	0.531	1.016
	1.5 ★	0.50	0.0087	–0.031	**0.547**	**1.508**
	2.0	0.50	0.0087	–0.031	0.562	2.000
	2.5	0.50	0.0087	–0.031	0.577	2.492
	3.0	0.50	0.0087	–0.031	0.593	2.984
b	0.6	0.50	0.0087	–0.031	0.547	1.530
	0.8	0.50	0.0087	–0.031	0.547	1.518
	1.0	0.50	0.0087	–0.031	0.547	1.511
	1.2 ★	0.50	0.0087	–0.031	**0.547**	**1.508**
	1.5	0.50	0.0087	–0.031	0.547	1.505
	2.0	0.50	0.0087	–0.031	0.547	1.503
δ	0.5	0.50	0.0087	–0.031	0.547	1.503
	1.0	0.50	0.0087	–0.031	0.547	1.505
	1.5 ★	0.50	0.0087	–0.031	**0.547**	**1.508**
	2.0	0.50	0.0087	–0.031	0.547	1.511
	3.0	0.50	0.0087	–0.031	0.547	1.516
	4.0	0.50	0.0087	–0.031	0.547	1.521
σ_s^2	0.009	0.50	0.009	–0.031	0.547	1.507
	0.10	0.50	0.10	–0.031	0.539	1.270
	0.50	0.50	0.50	–0.031	0.523	0.751
	1.0	0.50	1.0	–0.031	0.515	0.497
	1.759	0.50	1.759	–0.031	0.510	0.329
	3.5	0.50	3.5	–0.031	0.506	0.185

Notes: ★ indicates baseline value. All other parameters held at baseline: $\sigma_r^2 = 0.50$, $\sigma_{rs} = -0.031$, $r = 2.0$, $b = 1.2$, $\delta = 1.5$, $\lambda = 1.5$.

Table A2. Scale Factor k Sensitivity ($\rho = -0.47$ fixed)

k	σ_r^2	σ_s^2	σ_{rs}	β^*	γ^*	Note
1.0	0.036	0.001	–0.002	0.940	1.502	
5.0	0.179	0.003	–0.011	0.762	1.508	

10.0	0.357	0.006	–0.022	0.622	1.509	
13.97 ★	**0.499**	**0.009**	**–0.031**	**0.547**	**1.508**	**Paper baseline**
20.0	0.714	0.013	–0.044	0.467	1.504	
30.0	1.072	0.019	–0.067	0.382	1.495	
50.0	1.786	0.031	–0.111	0.290	1.471	

Notes: ★ indicates baseline value. All qualitative results (sign patterns in Propositions 1–3) are invariant to k. Quantitative outputs should be interpreted relative to the chosen scale normalisation.

Appendix B – Supplementary Figures

Figure 1. Sensitivity of Optimal Incentive Weights to Key Parameters. (Baseline: σr²=0.50, σs²=0.0087, σrs=−0.031, r=2.0, b=1.2, δ=1.5, λ=1.5.)

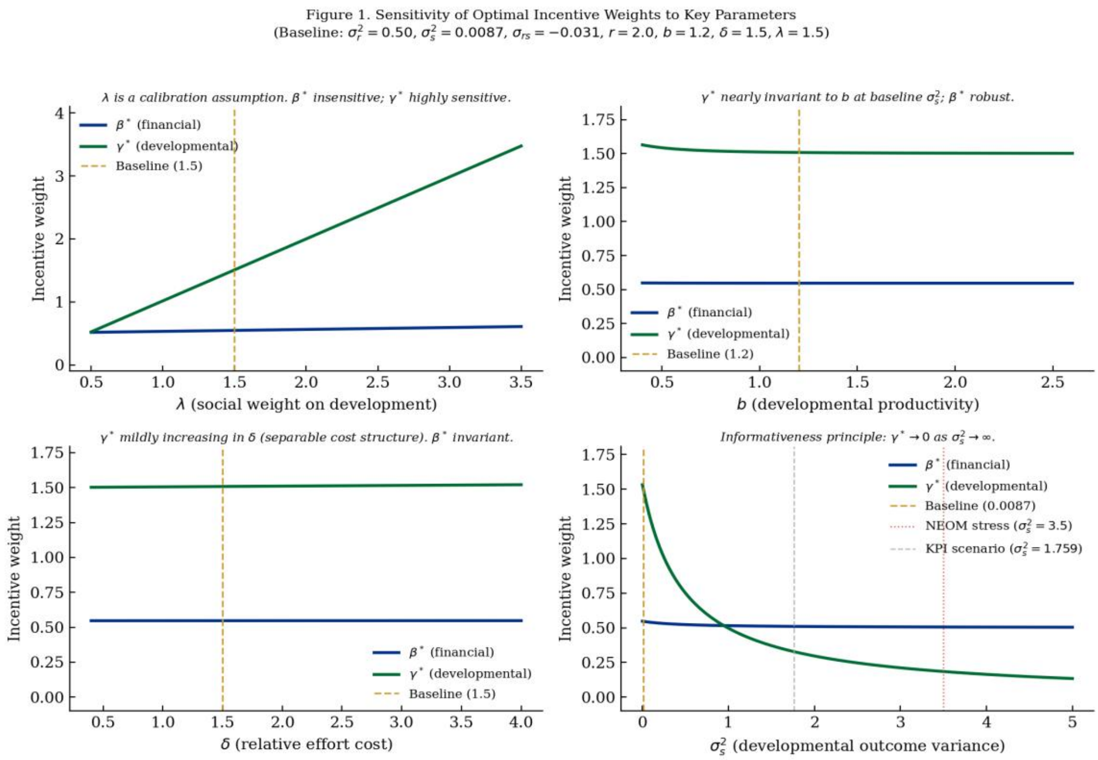

Notes: Each panel varies one parameter while holding all others at baseline. Dashed vertical line marks the baseline value.

Figure 2. Joint Sensitivity of γ* to λ and b. Dashed contours: γ* = 0.185, 0.329, 1.0, 1.508, 2.0.

Figure 2. Joint Sensitivity of γ^* to λ and b
Dashed contours: $\gamma^* = 0.185$ [NEOM], 0.329 [KPI], 1.0, 1.508 [Scenario A], 2.0