SENTINEL

Structured Evidence-Normalized Transparency and Intelligence for Non-profit Evaluation in Landscapes of Threat

---

Bayesian Sequential Detection, LLM-Augmented Risk Intelligence, and Enterprise Coordination Pattern Analysis in Government-Funded Organizations

---

Brent Figg
PhD in progress via the Global Executive Track pathway, Industrial and Systems Engineering
Wayne State University, Detroit, Michigan

Working Paper — Systems Engineering Applied to
Institutional Integrity and National Security

May 2026

Revised October 2026: author status and rights notice updated; content unchanged.

Classification: Unclassified / Open Academic Distribution

Draws exclusively on publicly accessible data, open-source methodology, and published regulatory frameworks. No classified inputs or outputs.

Rights, status and citation

© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved.

Author status (as of October 2026): M.Sc. in Systems Engineering, Wayne State University; PhD in progress via the Global Executive Track pathway.

Cite as: Figg, B. (2026). SENTINEL: Structured Evidence-Normalized Transparency and Intelligence for Non-profit Evaluation in Landscapes of Threat. Working paper, May 2026 (rev. October 2026). Available at 2ndandcapital.com.

© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com

Abstract

Small local non-profit and non-governmental organizations constitute a structurally underanalyzed risk surface in public accountability and national security frameworks. Their combination of government funding, resource dependency on a small number of funders, community trust, and sub-threshold organizational scale creates structural conditions that are, in organizational theory terms, a convergent vulnerability for co-optation by malfeasant actors. Until now, detecting such co-optation at systematic population scale has been analytically infeasible: no existing instrument possesses sufficient variety, in Ashby's formal sense, to discriminate between genuine governance and strategic impression management across the full spectrum of registered organizations.

This paper introduces **SENTINEL**—*Structured Evidence-Normalized Transparency and Intelligence for Non-profit Evaluation in Landscapes of Threat*— a Bayesian sequential detection architecture that integrates Gaussian Process governance quality modeling, Hidden Markov Model latent state estimation, and Partially Observable Markov Decision Process resource allocation within a Large Language Model interface layer. The result is continuous, calibrated governance risk assessment across all registered organizations, from the smallest local charity to the largest international NGO, updated as public regulatory filings arrive.

Three formal constructs are introduced: the *exploitation calculus* formalizing adversary return-on-investment for organizational co-optation; the *perfect storm index* $\Phi(i)$ characterizing the joint structural conditions that make detection near-impossible under existing mechanisms; and the *enterprise coordination pattern surface*, a network-level analytical layer enabling detection of coordinated governance failures that remain individually below any single-organization detection threshold. A synthetic validation framework specifying detection lift over population-ratio baselines is developed to ground the architecture empirically. Failure modes, feature construction limitations, and boundary conditions are analyzed with the specificity needed to bound the system's operational claims.

Keywords: system of systems; Bayesian optimization; POMDP; hidden Markov model; nonprofit accountability; resource dependency; enterprise coordination patterns; LLM interface; government-funded organizations; program protection; counterintelligence.

© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com

0 Contents

1 Introduction **5**

1.1 The Analytical Gap . . . 5

1.2 The Structural Opportunity for Co-optation . . . 5

1.3 What SENTINEL Is and Is Not . . . 5

2 The Exploitation Calculus **6**

2.1 Adversary Return on Investment . . . 6

2.2 Resource Dependency Theory . . . 6

2.3 The Perfect Storm Condition . . . 7

3 SoS Accountability Foundation **8**

3.1 The Accountability Sink . . . 8

3.2 Why Population Normalization Is Necessary . . . 8

4 SENTINEL POMDP: Formal Specification **9**

5 Mathematical Specification **9**

5.1 GP Governance Quality Model . . . 9

5.2 Hierarchical Population Model . . . 10

5.3 Population-Normalized Anomaly Score . . . 10

5.4 HMM Governance State Estimator . . . 10

5.5 Time-Indexed Detection Reward . . . 10

6 Observable Feature Specification **11**

6.1 The Feature Construction Problem . . . 11

6.2 Missing Data Handling . . . 11

6.3 Feature Reliability and Known Distortions . . . 11

7 Enterprise Coordination Pattern Detection **12**

7.1 From Individual Anomaly to Coordinated Pattern . . . 12

7.2 The Enterprise Network Model . . . 12

© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com

8 LLM-Augmented Risk Detection Interface **13**
8.1 The Access Problem . . . 13
8.2 Representative Queries by Actor Type . . . 14

9 Synthetic Validation Framework **15**
9.1 The Evaluation Problem . . . 15
9.2 Synthetic Injection Design . . . 15
9.3 Evaluation Metrics . . . 16
9.4 Projected Outcome Range and Honest Uncertainty . . . 16

10 SysML Architectural Specification **17**
10.1 Requirements . . . 17
10.2 SysML v2 Textual Specification . . . 17

11 Adjacent System Risk Management **19**
11.1 Program Protection Surface . . . 19
11.2 DCSA FOCI Integration . . . 20

12 Failure Modes and Boundary Conditions **20**
12.1 Adversarial Compliance Optimization . . . 20
12.2 Correlated Sector-Wide Shocks . . . 21
12.3 Sparse Data Collapse . . . 21
12.4 False Positives from Legitimate Organizational Transitions . . . 21
12.5 Computational Scaling . . . 22

13 Legal and Ethical Constraints **22**
13.1 The Public Data Perimeter . . . 22
13.2 Probabilistic Assessment, Not Legal Determination . . . 22
13.3 Analytical Neutrality . . . 23

14 Discussion **23**
14.1 The Structural ROI Effect . . . 23
14.2 What the Validation Would Need to Show . . . 23
14.3 Positioning . . . 24

© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com

15 Conclusion **24**

© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com

1 Introduction

1.1 The Analytical Gap

Government-funded non-profit organizations operate under disclosure obligations that were designed for an era of manual audit at low coverage density. Canada's T3010, the American Form 990, and the UK Annual Return collectively generate substantial compliance data. They generate it annually, in aggregate form, without cross-organizational normalization, and without longitudinal trajectory analysis. The result is an information architecture with insufficient variety, in Ashby's (1956) formal sense, to perform the control function it nominally provides.

This is not an indictment of any regulatory body. It is a structural observation: an oversight mechanism that can distinguish n states in the system it governs cannot correct deviations that require distinguishing $n+k$ states, regardless of the political will or resource allocated to its operation. The aggregate annual financial report of a $\$280,000$-revenue community organization produces a handful of financial ratios. The actual governance state of that organization is a high-variety latent variable—whether its resources flow toward its stated mission, whether its board exercises genuine independence, whether its procurement practices are arm's-length—that the annual filing cannot represent at adequate resolution.

1.2 The Structural Opportunity for Co-optation

The gap between the low-variety representation available to regulators and the high-variety reality of organizational governance creates an exploitable condition. Small local non-profit organizations, in particular, combine four structural properties—resource dependency, near-zero audit coverage, community trust, and sub-threshold organizational scale—that jointly create conditions under which governance failure can accumulate undetected across multiple fiscal cycles.

This is not a claim about the prevalence of malfeasance in the sector, which the evidence does not support as epidemic. It is a structural claim: the *detection capacity* of existing accountability mechanisms is insufficient relative to the *structural opportunity* for undetected governance failure, particularly in the small-organization tier. SENTINEL addresses this detection gap.

1.3 What SENTINEL Is and Is Not

SENTINEL is a population-level governance anomaly detection system that operates on publicly accessible regulatory filing data, applies Bayesian sequential inference to maintain calibrated belief distributions over organizational governance quality, and delivers those beliefs through a Large Language Model interface accessible to any authorized actor in natural language. It is not an accusation engine, an autonomous investigation system, or a legal determination mechanism. Its

© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com

outputs are probabilistic assessments that support human investigative judgment; they do not replace it.

The system makes three technical contributions beyond the existing literature: a hierarchical GP population model that provides governance assessments for newly registered organizations with no filing history; a network-level enterprise coordination pattern surface that detects governance anomalies that are individually below threshold but collectively significant; and a time-indexed POMDP detection reward that formalizes the opportunity cost of detection delay within closing regulatory windows.

2 The Exploitation Calculus

2.1 Adversary Return on Investment

The selection of organizational co-optation vehicles by malfeasant actors is implicitly an optimization over an expected return function. Define the adversary exploitation ROI for organization i as:

$$\mathrm{ROI}_{\mathrm{adv}}(i) = \frac{V(i)}{\mathrm{Risk}_{\mathrm{detect}}(i) \cdot C_{\mathrm{op}}(i)} \tag{1}$$

where $V(i)$ is the value of organizational access (financial resources, intelligence access, geographic position, influence capacity); $\mathrm{Risk}_{\mathrm{detect}}(i)$ is the probability of detection within the exploitation window; and $C_{\mathrm{op}}(i)$ is the cost of achieving operational influence. For large, professionally governed institutions, both $\mathrm{Risk}_{\mathrm{detect}}$ and C_{op} are materially elevated by audit density, board independence requirements, and management professionalism. For small local organizations, both terms approach their minimum values. The ROI function for organizational co-optation is therefore not uniformly distributed across the organizational population: it is concentrated at the small-organization tail.

2.2 Resource Dependency Theory

Pfeffer and Salancik (1978) established that organizations adapt their behavior in response to the resource dependencies that constrain them. In the government-funded non-profit context, this takes a specific governance form. Define the resource dependency ratio:

$$\mathrm{RD}(i) = \max_j \left(\frac{r_j^{(i)}}{R^{(i)}} \right) \tag{2}$$

where $r_j^{(i)}$ is revenue from funder j and $R^{(i)}$ is total revenue. Organizations with $\mathrm{RD}(i) > 0.6$ face a structural dependency condition: a single funder controls the majority of organizational

© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com

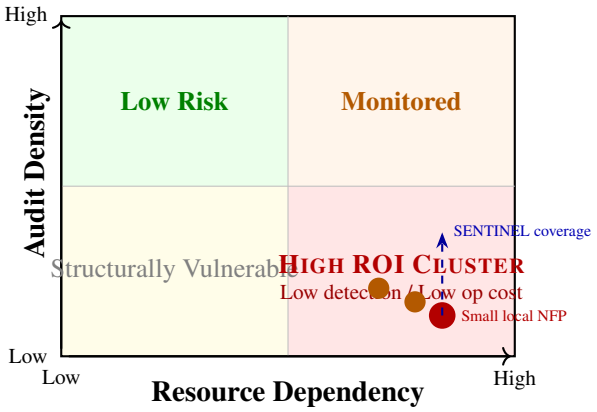

Figure 1: **Exploitation ROI Matrix.** High resource dependency combined with near-zero audit density produces the highest adversary ROI for organizational co-optation. Small local non-profits cluster in the high-ROI quadrant. SENTINEL shifts effective detection coverage upward for the full organizational population, altering the ROI calculation for the entire class.

resources and therefore commands de facto governance influence that formal board structures may be unable to countermand. In this condition, the introduction of a single resource-controlling actor with interests adverse to the organization's stated mission can achieve effective governance capture without any formal governance change.

2.3 The Perfect Storm Condition

Four jointly present structural properties define the *perfect storm condition* for governance vulnerability—the configuration under which existing accountability mechanisms are structurally insufficient to detect governance failure even if it is occurring:

C1 — Resource Dependency: $\mathrm{RD}(i) > \theta_{\mathrm{RD}}$. Governance independence is structurally compromised by funder control.

C2 — Audit Vacuum: $p_{\mathrm{audit}}(i) < \varepsilon$. The probability of independent verification within any five-year window approaches zero for organizations below the full-filing threshold.

C3 — Community Trust: $\mathrm{Trust}(i) > \theta_T$. The organization's community standing provides operational cover that a commercial or government entity could not replicate at equivalent cost.

C4 — Sub-threshold Scale: $R^{(i)} < R_{\mathrm{thresh}}$. Financial scale falls below the thresholds at which enhanced reporting, mandatory audit, and regulatory attention materialize.

The *perfect storm index* for organization i is:

$$\Phi(i) = \mathrm{RD}(i) \cdot \frac{1}{p_{\mathrm{audit}}(i) + \varepsilon} \cdot \mathrm{Trust}(i) \cdot \mathbb{1}\left[R^{(i)} < R_{\mathrm{thresh}}\right] \tag{3}$$

$\Phi(i)$ enters the SENTINEL detection reward as a prior weight: organizations with high $\Phi(i)$ re-

© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com

ceive elevated investigative priority at equivalent anomaly scores, because the social cost of non-detection in structurally unmonitorable organizations is higher than in monitored ones.

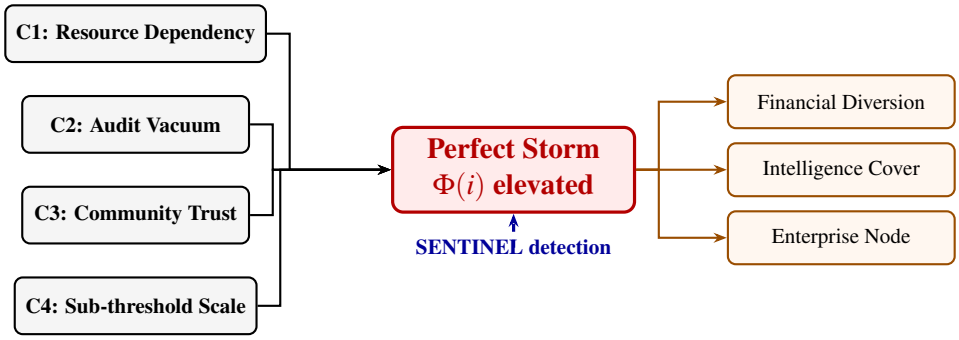

Figure 2: **Perfect Storm Convergence.** The four conditions jointly define the structural configuration under which governance failure is both most likely to occur and least likely to be detected by existing mechanisms. SENTINEL's population-normalized continuous assessment intercepts at the convergence point.

3 SoS Accountability Foundation

3.1 The Accountability Sink

Figg (2026a) defines the accountability sink as the structural condition in which the separation of decision authority, operational execution, and experienced outcome across organizational boundaries and temporal horizons is sufficiently complete that corrective feedback cannot form. Three conditions jointly constitute the sink:

$$S_{\mathscr{G}} : V(f) < V_{\min}, \quad \text{or} \quad \tau_f > \tau_{\text{correctable}}, \quad \text{or} \quad \text{reach}(f) \cap \mathscr{A}_{\text{response}} = \emptyset \tag{4}$$

The GFNO sector satisfies all three conditions. Compliance data has insufficient variety ($V(f) < V_{\min}$) to discriminate governance states. Regulatory response lags the correctable window ($\tau_f > \tau_{\text{correctable}}$) by the documented multi-year intervals between governance failure onset and detection. And the actors whose resource decisions most immediately influence organizational behavior—dominant funders in high-RD(i) organizations—are structurally excluded from the accountability loop when their interests are adverse.

3.2 Why Population Normalization Is Necessary

Absolute financial ratios—program expense percentages, overhead rates— are the standard currency of non-profit accountability analysis. They are also structurally inadequate for the control function they nominally serve, for two reasons.

© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com

First, they do not control for organizational context. A $\$280,000$-revenue community health organization operates under fundamentally different cost structures than a \$28 million national advocacy organization; comparing their overhead ratios against a common benchmark generates information-free comparisons. Second, they do not provide distributional context. Knowing that an organization's program ratio is 67% conveys no information about whether this is exceptional, typical, or anomalous for organizations of that type, size, and mission.

SENTINEL addresses both deficiencies through matched-cohort population-normalized benchmarking: all assessments are expressed relative to the distribution of organizationally comparable peers, not against absolute thresholds.

4 SENTINEL POMDP: Formal Specification

Definition 4.1 (SENTINEL POMDP)**.** *The SENTINEL detection system is a POMDP* $\mathscr{M} = (S, A, \mathscr{T}, \mathscr{R}, \Omega, \mathscr{O})$*:*

- ***Hidden state*** s_t*: joint governance quality vector* $G_t \in \mathbb{R}^K$*, latent malfeasance state* M_t *per organization, and network enterprise state* $\mathscr{N}_t$*.*
- ***Observation*** o_t*: publicly accessible compliance data across the observable feature space (Section 6).*
- ***Belief state*** $b_t = \big(p(\boldsymbol{g} \mid D_t), \{b_k(t)\}_{k=1}^{W}, p(M_t \mid e_{1:t}), p(\mathscr{N}_t \mid h_{1:t})\big)$*.*
- ***Action*** a_t*: investigative resource allocation—enhanced monitoring, targeted extraction, regulatory referral, program protection escalation.*
- ***Reward*** $\mathscr{R}(s_t, a_t, t)$*: time-indexed detection value weighted by perfect storm index* $\Phi(i)$*.*

5 Mathematical Specification

5.1 GP Governance Quality Model

For each of K governance dimensions:

$$g_k(\boldsymbol{x}) \sim \mathscr{GP}\big(\mu_k(\boldsymbol{x}), k_{\text{M52}}(\boldsymbol{x}, \boldsymbol{x}')\big), \quad k_{\text{M52}} = \sigma_f^2 \left(1 + \frac{\sqrt{5}r}{\ell} + \frac{5r^2}{3\ell^2}\right) e^{-\sqrt{5}r/\ell} \tag{5}$$

The Matérn-5/2 kernel is appropriate for governance surfaces that can exhibit rapid transitions—leadership change, funding crisis—without being discontinuous. The posterior mean $\mu_n(\boldsymbol{x})$ estimates governance quality for configuration $\boldsymbol{x}$; the posterior variance $\sigma_n^2(\boldsymbol{x})$ is elevated for organizations with sparse filing history or internally inconsistent reported configurations, providing explicit uncertainty quantification in all downstream outputs.

© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com

5.2 Hierarchical Population Model

$$g^{(i)} \sim \mathcal{GP}\big(\mu_{\text{pop}}(\boldsymbol{x}), k_{\text{pop}}\big) + \delta^{(i)}, \quad \delta^{(i)} \sim \mathcal{GP}(0, k_{\text{ind}}) \tag{6}$$

A newly registered organization receives its initial governance assessment from the population prior calibrated to its organizational type, budget tier, and mission category. The system explicitly reports elevated uncertainty for cold-start assessments; the posterior converges toward organization-specific inference as filing history accumulates. This mechanism enables SENTINEL to provide meaningful (if uncertain) assessments for the sub-threshold organizations that are, by the exploitation calculus, the highest-risk tier—despite being the tier with the fewest observations.

5.3 Population-Normalized Anomaly Score

$$\bar{A}_k^{(i)}(t) = \frac{y_k^{(i)}(t) - \text{median}\{y_k^{(j)}(t) : j \in \mathcal{P}_i\}}{\text{IQR}\{y_k^{(j)}(t) : j \in \mathcal{P}_i\}} \tag{7}$$

Median and IQR normalization rather than mean and standard deviation is chosen because financial ratio distributions in the non-profit sector are characteristically non-normal—program expenditure ratios exhibit left-skew; overhead ratios exhibit right-skew with heavy tails. Parametric normalization would systematically understate anomaly in the tails that are analytically most relevant.

5.4 HMM Governance State Estimator

Each organization occupies a latent governance state $G_{i,t} \in \{\textsc{Clean}, \textsc{Drifting}, \textsc{Compromised}, \textsc{Captured}\}$ evolving under transition matrix A^{gov}:

$$b_{i,g}(t) \propto P(e_{i,t} \mid G_{i,t} = g) \sum_{g'} P(G_{i,t} = g \mid G_{i,t-1} = g') \cdot b_{i,g'}(t-1) \tag{8}$$

The forward belief runs in $\mathcal{O}(n)$ per observation. Figure 3 shows the state machine. Transition matrix A^{gov} is estimated from the historical record of organizations whose governance states were eventually revealed through regulatory action, whistleblower disclosure, or investigative reporting—the ground truth dataset for model calibration.

5.5 Time-Indexed Detection Reward

$$\mathcal{R}(b_t, i, t) = \Phi(i) \cdot \left[\alpha_{\text{EHVI}}(\boldsymbol{x}_i; D_n) \cdot \prod_j \hat{P}_j(\boldsymbol{x}_i) + \sum_{k=1}^{W} b_k(t) \cdot \eta_k(\boldsymbol{x}_i) - \lambda \cdot c(\boldsymbol{x}_i)\right] \tag{9}$$

The urgency score for organization i under regulatory window k:

$$U_k^{(i)}(t) = b_k(t) \cdot \eta_k(\boldsymbol{x}_i^*) \cdot P(\text{foreclosure before } t + \Delta \mid A^{\text{gov}}, b_{i,g}(t)) \tag{10}$$

© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com

High-urgency organizations—elevated governance state beliefs, open regulatory windows, high foreclosure probability—are surfaced to the LLM interface as priority investigative targets.

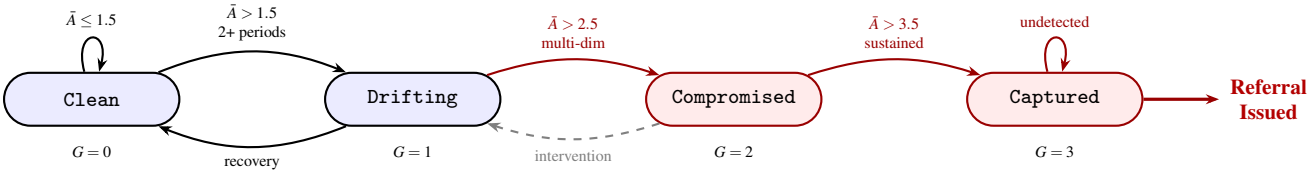

Figure 3: **Governance Quality State Machine.** HMM latent state space. Transition guards are driven by population-normalized anomaly scores. The CAPTURED state triggers priority escalation and referral brief generation. Transition matrix A^{gov} is estimated empirically from the historical ground-truth record.

6 Observable Feature Specification

6.1 The Feature Construction Problem

A technically credible claim about GP posterior quality depends on specifying what signals populate the GP observation vector $\boldsymbol{y}$, how each signal is constructed from raw filing data, what its reliability and completeness characteristics are, and how the system degrades when signals are missing or corrupted. Table 1 provides this specification.

6.2 Missing Data Handling

The GP framework handles missing observations naturally: a missing feature leaves the posterior unchanged in that configuration region, elevating posterior variance and increasing the exploration weight assigned to organizations with incomplete observations. This behavior is operationally appropriate: organizations with limited observable features receive wider uncertainty bands in their governance assessments, driving investigative resources toward data-collection actions rather than premature conclusions. Missing signals are never imputed with point estimates that would falsely reduce posterior uncertainty.

6.3 Feature Reliability and Known Distortions

Several features carry known reliability limitations that bound the system's claims. Related party transactions are self-reported with no external verification below the mandatory audit threshold—making the feature a weak signal for precisely the organizations where it would be most informative. Board composition data from corporate registries carries a structural lag of three to six months in most jurisdictions, limiting its timeliness for rapidly evolving governance situations.

© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com

Vendor concentration data is sparse for small organizations because public procurement databases are populated primarily by organizations large enough to engage in formal tendering processes.

These limitations are not remediated by the SENTINEL architecture; they constrain what the system can reliably assess. The appropriate response to systematic feature sparsity in a specific organizational sub-tier is elevated posterior variance in that sub-tier, which the GP naturally produces, rather than confidence suppression or compensatory imputation.

7 Enterprise Coordination Pattern Detection

7.1 From Individual Anomaly to Coordinated Pattern

Individual organization anomaly scores, however well-calibrated, cannot detect the coordinated governance failures that represent the most consequential threat condition: a network of nominally independent organizations that together constitute a coherent exploitation platform, in which no individual organization exceeds detection thresholds but the network configuration is statistically inconsistent with independent benign operation.

This is a direct emergent property of the SoS structure of the civil society ecosystem: organizations are assessed independently by existing oversight mechanisms, while their emergent collective behavior is analytically unobserved. SENTINEL's network layer addresses this by treating the organizational ecosystem as a graph and detecting subgraph-level anomalies that are invisible at the node level.

7.2 The Enterprise Network Model

Let $\mathscr{H}_t = (V, E_t)$ be the organizational network graph, with edges representing shared board members, common funding sources above threshold, geographic proximity within defined radii, and commercial relationships. For subgraph $\mathscr{S} \subseteq \mathscr{H}_t$:

$$\mathscr{A}^{\mathscr{S}}(t) = \phi\left(\sum_{i \in \mathscr{S}} w_i \bar{A}^{(i)}(t)\right) \cdot \Psi(\mathscr{S}, \mathscr{F}_t) \tag{11}$$

where $\phi(\cdot)$ aggregates individual anomaly scores while detecting coordinated anomaly patterns, and $\Psi(\mathscr{S}, \mathscr{F}_t)$ is the proximity alignment score measuring overlap between the subgraph footprint and the sensitive facility map $\mathscr{F}_t$.

The output of this analysis—when $\mathscr{A}^{\mathscr{S}}$ exceeds the enterprise threshold—constitutes a *pattern detection package* that is structurally consistent with the evidentiary requirements of enterprise coordination frameworks including RICO (18 U.S.C. §§ 1961–1968), Canada Criminal Code Part XII.2, and the UK Serious Crime Act 2007. It is not itself evidence; it is the analytical predi-

© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com

cate that supports focused evidentiary development by appropriately authorized investigators.

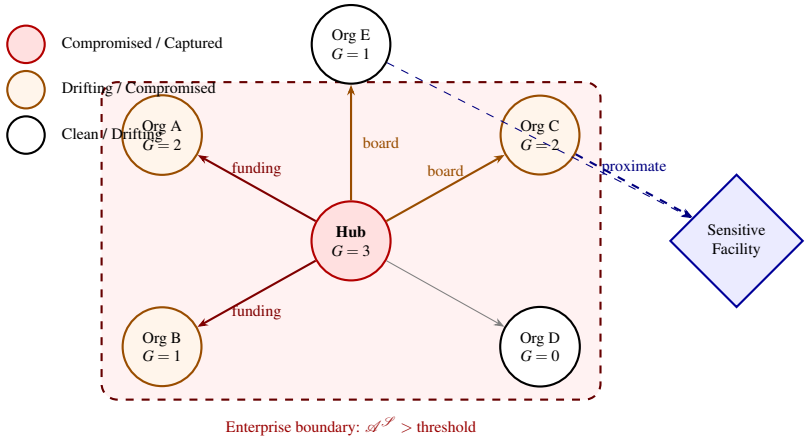

Figure 4: **Enterprise Coordination Network Graph.** The SENTINEL network graph identifies organizational subgraphs whose aggregated anomaly score $\mathscr{AS}$ exceeds the enterprise threshold, establishing the coordination pattern for investigative referral. Individual organization states (Clean/Drifting/Compromised/Captured) are shown. Network edges represent funding flows and board interlocking.

8 LLM-Augmented Risk Detection Interface

8.1 The Access Problem

The analytical architecture of SENTINEL—GP posteriors, HMM belief states, POMDP acquisition functions, network graph anomaly scores—is not directly accessible to most actors with the authority and motivation to act on its outputs. The Model Context Protocol integration layer resolves this by exposing the SENTINEL detection engine as a set of tool-callable functions to a language model inference service, enabling natural language interaction with the full Bayesian belief state.

The LLM functions as an interface to a rigorously specified Bayesian backend, not as the source of analytical judgment. It translates queries into API calls against the GP posterior and HMM state distribution; it translates those distributions into natural language responses calibrated to the querying actor's context. This is the appropriate role for LLMs in decision-critical systems: access layer, not inference engine.

© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com

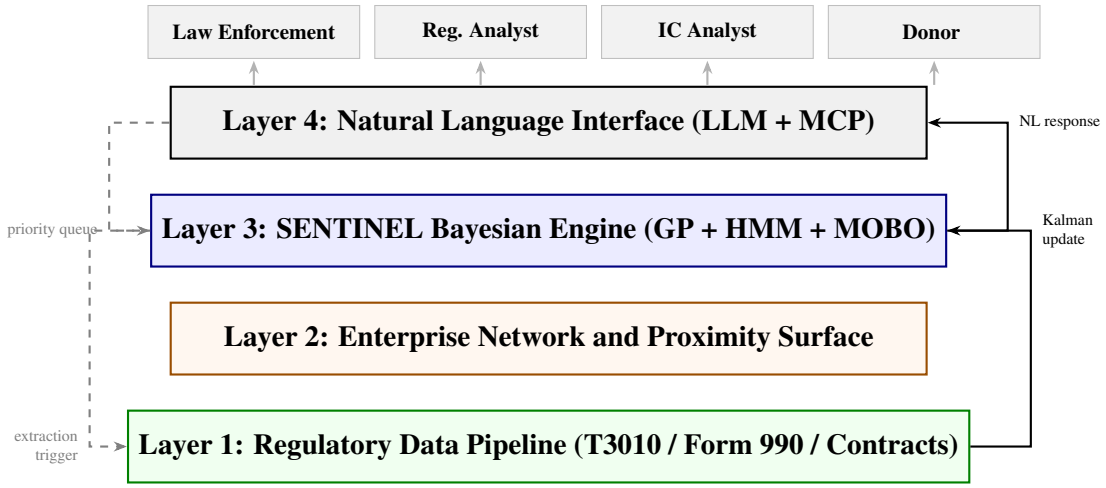

Figure 5: **SENTINEL Four-Layer Architecture.** The LLM / MCP layer provides natural language access to the Bayesian backend for any actor type. Data flows upward as streaming Kalman belief updates; responses flow downward as natural language. The LLM is an interface layer; all analytical inference occurs in Layers 2–3.

8.2 Representative Queries by Actor Type

Law Enforcement — Pattern Detection Query

Query: "Show non-profits within 10 km of [Facility] with resource dependency above 70% and governance anomaly scores above the 85th percentile for peer organizations. Flag shared board members."
Response: Ranked list with individual anomaly scores, matched-cohort percentiles, resource dependency ratios, board interlocking indicators, and enterprise pattern assessment if $\mathscr{A}^{\mathscr{S}}$ exceeds threshold. Uncertainty bounds explicit for each assessment.

CRA / IRS Audit Prioritization

Query: "Which registered organizations in my portfolio have shown the most adverse governance trajectory over three years while remaining below the mandatory full-audit threshold?"
Response: Ranked by trajectory slope of population-normalized anomaly score, weighted by $\Phi(i)$, with regulatory window urgency scores indicating which organizations have open intervention windows approaching closure.

Program Protection Officer

Query: "What is the governance risk profile of civil society organizations within the program protection perimeter? Are there coordinated patterns?"
Response: Proximity heatmap, individual governance state beliefs, enterprise network score for the perimeter subgraph, and recommendation on counterintelligence referral priority. Un-

© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com

certainty bounds and data sparsity flags explicit for each organization.

Individual Donor

Query: "Is [Organization] financially sound relative to comparable organizations, and are there governance concerns I should know about?"
Response: Percentile position on program ratio, overhead ratio, and governance structure indicators relative to matched peer cohort; longitudinal trajectory; plain-language governance state assessment with explicit uncertainty.

9 Synthetic Validation Framework

9.1 The Evaluation Problem

The absence of a validation framework is the central technical gap in the current architecture specification. The claim that SENTINEL provides materially better governance anomaly detection than existing baseline methods is an architectural inference, not a demonstrated empirical result. This section specifies the validation framework required to convert that inference into a defensible empirical claim, and states explicitly what the framework would need to show.

9.2 Synthetic Injection Design

The proposed validation uses a synthetic injection experiment. An organizational population of N_{org} simulated organizations is generated from the hierarchical GP prior, with true governance states assigned and governance state trajectories simulated forward under the estimated HMM transition matrix A^{gov}. Filing sequences consistent with each organization's governance state and HMM emission model B^{gov} are then generated as the synthetic observation stream.

Anomaly injection: M_{inject} organizations are selected for synthetic governance deterioration: true state transitions Clean $\rightarrow$ Drifting at period t_1, then Drifting $\rightarrow$ Compromised at $t_1 + \Delta t$, with Δt drawn from a plausible deterioration timescale distribution. Filing sequences are modified to be consistent with the injected true state while maintaining surface-level compliance with filing requirements—representing the strategic impression management behavior that sophisticated organizations would exhibit.

Detection comparators:

Baseline A: Single-period program ratio threshold. Flag if program expense ratio < 0.65 (the conventional charity watchdog threshold).

Baseline B: Single-period z-score on financial ratios. Flag if $|z| > 2.0$ on any feature, using the full-population mean and standard deviation (non-normalized).

© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com

SENTINEL: GP posterior anomaly score $\bar{A}^{(i)}(t)$ combined with HMM governance state belief $b_{i,g}(t)$, using matched-cohort population normalization.

9.3 Evaluation Metrics

Detection lag Δ_{lag}: Number of filing periods between the injected true state transition and the system's flag, defined as the first period in which the detection criterion exceeds threshold. Lower is better. The hypothesis is that SENTINEL achieves $\Delta_{\text{lag}}^{\text{SENTINEL}} < \Delta_{\text{lag}}^{\text{Baseline}}$ by at least two filing periods on average.

Precision at threshold: $P(\text{true positive} \mid \text{flag})$ at the operating threshold. Population normalization is expected to reduce false positives arising from structural differences between organizations (e.g., a legitimate rural health clinic with high overhead due to geographic cost structure being flagged by a non-normalized threshold).

Recall at threshold: $P(\text{flag} \mid \text{anomalous state})$. The hypothesis is that SENTINEL achieves materially higher recall for the small-organization tier specifically, where structural features (high $\Phi(i)$) suppress the noise that creates false negatives in non-normalized comparators.

Network detection lift: For injected coordinated enterprise scenarios (multiple organizations deteriorating under common control), the proportion of enterprise patterns detected within five periods of the first node transition, across all comparators. Baselines A and B have no network layer; this metric quantifies the added value of the enterprise coordination surface.

9.4 Projected Outcome Range and Honest Uncertainty

Based on theoretical analysis of the architecture, the following outcome ranges are projected for the full validation experiment. These are projections from the model structure, not empirical results; the validation experiment is required to confirm or refute them.

The projection range is wide. This reflects genuine uncertainty about three parameters that cannot be specified without empirical data: the HMM emission distributions B^{gov} (how well filing patterns discriminate governance states in practice); the GP lengthscale and signal variance that best fit the organizational population; and the rate at which sophisticated organizations produce surface-level-clean filings despite deteriorating governance states. The validation experiment is designed to estimate these parameters from the historical ground-truth record and propagate them into the detection performance projections.

© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com

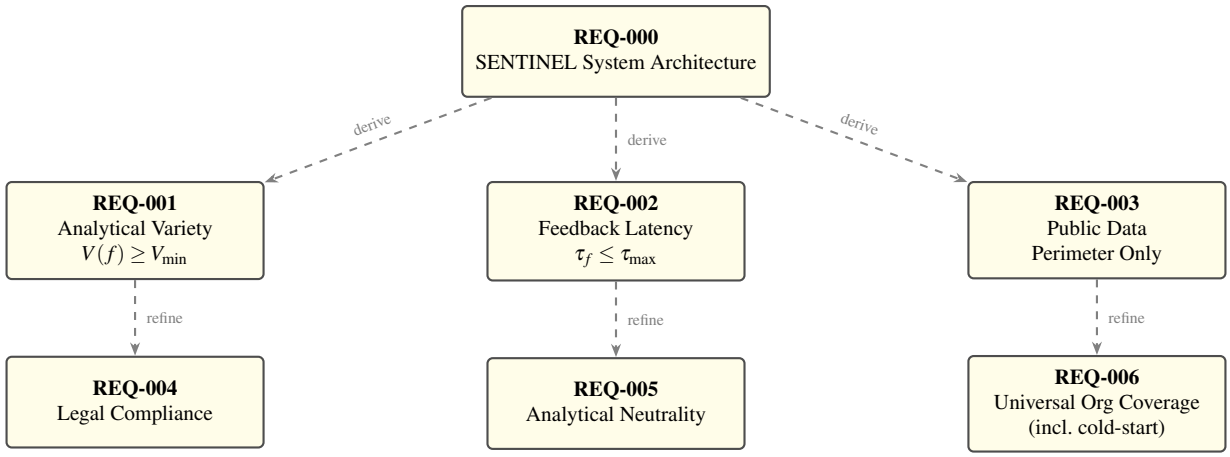

Figure 6: **SysML Requirements Diagram.** REQ-006 (Universal Organizational Coverage) explicitly includes the cold-start mechanism for small organizations with no prior filing history.

10 SysML Architectural Specification

10.1 Requirements

10.2 SysML v2 Textual Specification

```
package SENTINEL {

  // -- Requirements ----------------------------------------------------
  requirement <'REQ-001'> AnalyticalVarietyAdequacy {
    doc /* GP posterior representations shall have V(f) >= V_min,
           sufficient to discriminate governance patterns invisible
           in compliance-data representations. Explicitly verified
           by synthetic injection validation (Section 9). */
    assume constraint { V_signal >= V_min }
  }

  requirement <'REQ-006'> UniversalOrganizationalCoverage {
    doc /* System shall provide governance assessments for all
           registered organizations including those below full-filing
           thresholds, via hierarchical population prior cold-start.
           Cold-start assessments shall carry explicit uncertainty
           quantification; GP posterior variance shall not be
           suppressed to create false confidence. */
  }

  // -- Core Blocks -----------------------------------------------------
  block def GPGovernanceSurrogate {
    value dimension       : Integer;
    value posteriorMean   : Real[*];
```


© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com

```
    value posteriorVar    : Real[*];
    value populationPrior : HierarchicalGPPrior;
    perform action coldStartAssessment {
      in  orgProfile   : OrganizationConfig;
      out assessment   : GovernanceEstimate;
      // Returns high posterior variance for sparse orgs.
      // Assessment calibrated to Eq. (2) population prior.
    }
    perform action updatePosterior {
      in  newFiling  : FilingRecord;
      out updatedGP  : GPGovernanceSurrogate;
      // O(n) Kalman-equivalent Matern-5/2 update (Figg 2026b)
    }
  }

  block def HMMGovernanceEstimator {
    value transitionMatrix : Real[4,4]; // A^{gov}, empirically
        estimated
    value emissionModel    : Real[4,*]; // B^{gov}, from ground-truth
        record
    value stateBeliefs     : Real[4];   // [Clean, Drifting, Comp.,
        Captured]
    state machine GovernanceStateMachine {
      state Clean;
      state Drifting;
      state Compromised;
      state Captured;
      initial transition -> Clean;
      transition Clean -> Drifting {
        guard { anomalyScore > 1.5 and consecutivePeriods >= 2 }
      }
      transition Drifting -> Compromised {
        guard { anomalyScore > 2.5 and multiDimensional }
      }
      transition Compromised -> Captured {
        guard { anomalyScore > 3.5 and sustained }
      }
    }
  }

  block def EnterprisePatternLayer {
    value orgGraph        : OrganizationGraph;
    value sensitiveAssets : GeographicPolygon[*];
    value enterpriseThreshold : Real;
    perform action detectCoordinatedPattern {
      in  subgraph    : OrganizationGraph;
      out patternScore: Real;
```


© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com

```
      out referralPkg : PatternDetectionPackage;
      // Output is analytical predicate, not legal determination.
      // Consistent with evidentiary frameworks: RICO, CCC Pt XII.2,
      // UK Serious Crime Act 2007.
    }
  }

  block def LLMInterfaceLayer {
    satisfy REQ-001, REQ-002, REQ-006;
    value mcpEndpoint  : URI;
    // LLM is interface layer only -- all inference in GP/HMM/MOBO.
    perform action processNLQuery {
      in  query       : NaturalLanguageQuery;
      in  actorType   : ActorTypeEnum;
      in  beliefState : CompoundBeliefState;
      out response    : NaturalLanguageResponse;
      // Response includes explicit uncertainty quantification.
      // Probabilistic language only; no determinations of fact.
    }
  }

  block def SENTINEL_System {
    satisfy REQ-000;
    part gpSurrogate    : GPGovernanceSurrogate[K];
    part govHMM         : HMMGovernanceEstimator;
    part enterpriseLayer: EnterprisePatternLayer;
    part moboLoop       : MOBODetectionLoop;
    part llmInterface   : LLMInterfaceLayer;
    part dataLayer      : DataIngestionLayer;
    value perfectStormIdx  : Real[*]; // Phi(i) per org, Eq. (3)
    value beliefState      : CompoundBeliefState;
  }
}
```

Listing 1: SysML v2 – SENTINEL Detection System (core blocks)

11 Adjacent System Risk Management

11.1 Program Protection Surface

For a sensitive program $\mathscr{P}$ with perimeter $\Pi_{\mathscr{P}}$, the program protection risk surface at time t is:

$$\mathscr{R}_{\mathscr{P}}(t) = \sum_{i:d(i,\mathscr{P})\leq\delta} \left[\sum_{\theta} b_{i,\theta}(t)\cdot v_{\theta}\right] \cdot \Psi(i,\mathscr{P}) \cdot \Phi(i) \tag{12}$$

© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com

where $b_{i,\theta}(t)$ is the threat modality belief (collection, personnel, technology transfer, infrastructure, influence); ν_θ is the consequence weight; $\Psi(i, \mathscr{P})$ is the proximity alignment; and $\Phi(i)$ is the perfect storm index. Output is formatted for direct integration into Program Protection Plans under DoDI 5200.39.

11.2 DCSA FOCI Integration

SENTINEL provides DCSA analysts with a continuously updated view of government-funded organizations proximate to contractors under FOCI review, with threat modality belief distributions indicating whether nominally independent organizations constitute undisclosed FOCI vectors. The LLM interface enables analysts to query this view in natural language without data science support requirements.

12 Failure Modes and Boundary Conditions

The following failure modes are not edge cases requiring extraordinary conditions. They are predictable behaviors under realistic operational conditions. Each is specified with the mechanism, detection indicator, and mitigation.

12.1 Adversarial Compliance Optimization

Mechanism: A sophisticated organization with legal and accounting resources designs its filings to produce GP anomaly scores below detection threshold while continuing governance misconduct. The related-party transaction disclosure, for example, requires disclosure only of the existence and aggregate amount of transactions with related parties, not their terms or individual amounts—leaving substantial optimization space for strategic presentation.

Detection indicator: Abrupt improvement in anomaly scores following a period of elevated scores is itself a weak signal; natural recovery trajectories are gradual. Statistical consistency testing across the temporal sequence of anomaly scores can distinguish organic improvement from threshold gaming. Network-layer signals (board composition, funding flows) are harder to optimize simultaneously with financial ratios, creating consistency pressure.

Mitigation bound: For highly sophisticated organizations with dedicated legal and accounting resources, SENTINEL's detection capacity is materially reduced. This is an honest limitation. The network layer provides the primary residual signal; single-organization financial analysis is of limited value against sustained, sophisticated evasion.

© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com

12.2 Correlated Sector-Wide Shocks

Mechanism: A macroeconomic shock, pandemic, or sector-wide funding change causes legitimate correlated governance stress across many organizations simultaneously. The GP posterior for affected organizations shifts together, potentially generating false-positive anomaly signals across the sector.

Detection indicator: A sector-wide anomaly spike that is correlated with a documented external event and that reverts within 1–3 filing periods is characteristic of genuine shock rather than governance failure.

Mitigation: The matched-cohort population normalization absorbs common shocks partially—if all organizations of a given type experience a similar ratio change, the peer-relative anomaly score remains near zero. Post-hoc temporal controls for documented macro events should be applied before generating regulatory referral packages.

12.3 Sparse Data Collapse

Mechanism: An organization with fewer than three years of filing history provides insufficient data to separate individual posterior from the population prior. The GP posterior variance remains elevated; anomaly scores are dominated by prior uncertainty rather than genuine signal. A real governance anomaly may be indistinguishable from a cold-start organization with an unusual-but-legitimate mission profile.

Detection indicator: Posterior variance above the cold-start threshold for an organization's mission-type category indicates insufficient data for reliable individual assessment.

Mitigation bound: For organizations with fewer than three complete filing periods, SENTINEL's assessment is explicitly a population-prior estimate with wide uncertainty. The system is not designed to produce high-confidence individual assessments for new organizations; it is designed to place them in a distribution and flag when their limited observable signals are inconsistent with that distribution. The appropriate response to cold-start uncertainty is monitoring, not referral.

12.4 False Positives from Legitimate Organizational Transitions

Mechanism: A merger, mission pivot, leadership transition, or major program expansion generates filing patterns that are anomalous relative to the organization's prior trajectory without reflecting governance failure.

Detection indicator: A single-period anomaly spike followed by reversion is characteristic of transition rather than systematic deterioration. The HMM state machine requires sustained anomaly (multiple periods above threshold) before transitioning to COMPROMISED; this design

© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com

choice directly mitigates false positives from legitimate transition events.

Mitigation bound: Transition-related false positives are expected to constitute the majority of flagged cases at moderate anomaly thresholds. The VADM-structured investigative interface explicitly labels the proportion of historical flags in the organization's peer cohort that resolved to legitimate transitions versus confirmed governance failures, providing base-rate context for investigative priority decisions.

12.5 Computational Scaling

Mechanism: Exact GP inference is $\mathcal{O}(n^3)$ in the number of observations. At full regulatory population scale—approximately 85,000 registered charities in Canada, 1.5 million Form 990 filers in the United States, 170,000 in the UK—exact GP inference is computationally intractable without approximation.

Mitigation: The Matérn-5/2 / Kalman filter equivalence of Figg (2026b, Proposition 1) reduces the temporal streaming component to $\mathcal{O}(n)$ per observation. The inducing-point approximation (Titsias, 2009) reduces the spatial component to $\mathcal{O}(M^2 n)$ where $M \ll n$ is the number of inducing points. The combination makes full-population inference tractable on standard cloud infrastructure, though not in real time on a single workstation. The computational architecture requires explicit acknowledgment in any deployment specification.

13 Legal and Ethical Constraints

13.1 The Public Data Perimeter

SENTINEL operates exclusively within the public data perimeter defined by regulatory filing obligations. This is both an ethical constraint—the constitutional rights of civil society organizations are not subordinated to analytical convenience—and a practical one: an architecture that requires organizational cooperation can be disabled by organizational resistance. Organizations with governance anomaly to conceal are precisely those most likely to resist cooperative access.

13.2 Probabilistic Assessment, Not Legal Determination

All SENTINEL outputs are probabilistic assessments derived from publicly accessible data. A governance state belief of 0.73 on COMPROMISED means that organizations with this observable pattern have historically been found to have genuine governance failures at high rates; it does not constitute evidence that this specific organization has committed any offense. The LLM interface is constrained by design to use probabilistic language in all responses and to include explicit uncertainty bounds. It cannot generate determinative claims about any organization.

© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com

13.3 Analytical Neutrality

Population normalization is the technical implementation of analytical neutrality: assessments are always relative to comparable organizations, and mission content has no bearing on governance anomaly scoring. An organization whose financial patterns, governance structure, and programmatic claims are consistent with sector norms receives a correspondingly low anomaly score regardless of its advocacy positions.

14 Discussion

14.1 The Structural ROI Effect

The most durable systemic impact of SENTINEL may be deterrence rather than detection. The exploitation calculus of Equation 1 is currently favorable to co-optation of small organizations because $\text{Risk}_{\text{detect}}$ approaches zero. When continuous population-level monitoring raises effective detection probability materially, the ROI calculation changes for the entire organizational class, not merely for those organizations directly flagged. The architecture's existence changes the behavior it monitors—a property of accountability systems identified empirically by Fung, Graham, and Weil (2007) across multiple transparency policy contexts.

14.2 What the Validation Would Need to Show

For the SENTINEL architecture to cross from "coherent theoretical contribution" to "defensible empirical contribution," the synthetic validation of Section 9 needs to demonstrate three things with statistical rigor:

1. Detection lag reduction of at least two filing periods on average relative to baseline ratio methods, for the small-organization tier specifically.
2. False positive rate reduction attributable to population normalization, quantified separately from detection lag to isolate the normalization contribution from the sequential inference contribution.
3. Non-trivial network-layer detection lift for coordinated enterprise scenarios, demonstrating that the subgraph anomaly score $\mathscr{A}^{\mathscr{S}}$ captures patterns invisible to single-organization analysis.

These are the three claims that the architecture's design most directly supports. They are also the three claims most easily challenged by a reviewer who doubts the practical value of the additional complexity relative to simpler approaches.

© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com

14.3 Positioning

SENTINEL is positioned as the skeleton of a new accountability intelligence capability, not a deployed system. The theoretical architecture is complete. The mathematical specification is internally consistent. The validation framework is specified but not yet executed. The next stage of development is empirical: executing the synthetic injection experiment, calibrating the HMM transition and emission matrices against historical ground-truth records, and reporting detection performance against the projected ranges of Table 2. The architecture is designed to support this validation without modification.

15 Conclusion

Small local non-profit organizations have been, until now, structurally positioned as the highest-ROI co-optation vehicle in the malfeasant actor's toolkit—not despite their legitimacy but because of it. The combination of resource dependency, near-zero audit coverage, community trust, and sub-threshold organizational scale produces structural conditions in which existing accountability mechanisms cannot form the corrective feedback that Ashby's Law requires for purposive governance.

SENTINEL addresses this not by increasing audit density—a resource-constrained intervention—but by raising the analytical variety of the feedback mechanism across the full organizational population. The GP governance quality surface, HMM state estimator, POMDP detection reward, enterprise coordination pattern layer, and LLM natural language interface together constitute an accountability architecture whose variety is adequate to discriminate the governance states that compliance data cannot distinguish.

Three structural claims are made. First, the hierarchical population prior provides meaningful governance assessments for organizations at every scale, including those below current detection thresholds, with uncertainty quantification that is explicit rather than suppressed. Second, the network-level enterprise coordination surface detects coordinated governance failures that are individually invisible, addressing the emergent threat condition that single-organization oversight cannot reach. Third, the LLM interface makes this analytical capability accessible to any actor with natural language competency and appropriate authorization—from a local investigator to an Intelligence Community analyst—without requiring data science expertise.

Each claim is supported by theoretical analysis and bounded by explicit failure modes. All three require empirical validation through the synthetic injection framework specified in Section 9. The architecture is designed to support that validation; the research program's next stage is to execute it.

© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com

15 References

Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall.

Beer, S. (1972). *Brain of the Firm*. Allen Lane.

Canada Revenue Agency (2023). *Registered Charities: Statistical Overview*. Government of Canada.

Chelst, K. and Canbolat, Y. B. (2011). *Value-Added Decision Making for Managers*. CRC Press.

Defense Counterintelligence and Security Agency (2022). *Trends in Foreign Targeting of Cleared Industry*. DCSA.

Daulton, S., Balandat, M., and Bakshy, E. (2020). Differentiable expected hypervolume improvement. *NeurIPS*, 33.

Europol (2021). *Exploiting Isolation: Offences Against Older Members of Society*. Europol.

Federal Bureau of Investigation (2023). *Foreign Influence Operations: An Overview*. FBI.

Figg, B. (2026a). Accountability as a control mechanism in System-of-Systems development. Working Paper, Wayne State University.

Figg, B. (2026b). TEMPO: Time-indexed evidence-based marginal gains under partial observability. Working Paper, Wayne State University.

Frazier, P. I. (2018). A tutorial on Bayesian optimization. *arXiv:1807.02811*.

Fung, A., Graham, M., and Weil, D. (2007). *Full Disclosure*. Cambridge University Press.

U.S. Government Accountability Office (2022). *Tax-Exempt Organizations: IRS Should Strengthen Oversight*. GAO-22-104423.

Kaelbling, L. P., Littman, M. L., and Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1–2):99–134.

Maier, M. W. (1998). Architecting principles for systems-of-systems. *Systems Engineering*, 1(4):267–284.

O'Neill, O. (2002). *A Question of Trust*. Cambridge University Press.

Perrow, C. (1984). *Normal Accidents*. Basic Books.

Pfeffer, J. and Salancik, G. R. (1978). *The External Control of Organizations*. Harper & Row.

© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com

Power, M. (1997). *The Audit Society*. Oxford University Press.

Puterman, M. L. (1994). *Markov Decision Processes*. Wiley.

Rabiner, L. R. (1989). A tutorial on hidden Markov models. *Proceedings of the IEEE*, 77(2):257–286.

Romzek, B. S. and Dubnick, M. J. (1987). Accountability in the public sector. *Public Administration Review*, 47(3):227–238.

U.S. Senate Select Committee on Intelligence (2020). *Report on Russian Active Measures, Volume 5*. 116th Congress.

Srinivas, N., Krause, A., Kakade, S. M., and Seeger, M. W. (2012). Information-theoretic regret bounds for GP optimization. *IEEE Trans. Information Theory*, 58(5):3250–3265.

Sterman, J. D. (2000). *Business Dynamics*. McGraw-Hill.

Titsias, M. K. (2009). Variational learning of inducing variables in sparse Gaussian processes. *AISTATS*, 5:567–574.

UK Charity Commission (2023). *Annual Report 2022–23*. HMSO.

Weick, K. E. (1995). *Sensemaking in Organizations*. SAGE.

Wiener, N. (1948). *Cybernetics*. MIT Press.

© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com

Table 1: **SENTINEL Observable Feature Space.** Each row specifies an observable governance signal: its source, construction from raw filing data, known reliability characteristics, handling of missing data, and the governance dimension it updates. Reliability ratings (H/M/L) reflect both filing completeness and self-reporting accuracy concerns.

Feature	Source	Construction	Rel.	Missing	GP Dimension
Program expense ratio	T3010 / 990	Prog. exp / Total exp	H	Posterior var elevated	g_1 Financial integrity
Overhead ratio	T3010 / 990	Mgmt&Gen / Total exp	H	Posterior var elevated	g_1 Financial integrity
Revenue concentration RD(i)	T3010 / 990	$\max_j r_j/R$	M	Category-imputed	g_1, g_3
Exec compensation ratio	T3010 / 990	Top exec pay / Total exp	M	Often absent $<$ \$500K	g_1 Financial integrity
Filing timeliness	CRA / IRS records	Days past deadline	H	Complete from registry	g_3 Governance structure
Net asset trajectory	T3010 / 990	ΔNet assets (3yr)	H	Requires 3+ years	g_6 Longitudinal consistency
Related party flag	T3010 / 990	Disclosure + magnitude	L	Self-reported only	g_4 Procurement integrity
Audit qualification	Filed audit report	Binary + qualifier type	M	Only above threshold	g_1, g_3
Board size	Corporate registry	Director count at year-end	M	Registry lag 3–6 mo	g_3 Governance structure
Board turnover rate	Registry time-series	Annual director change / N	M	Requires 2+ years	g_3 Governance structure
Vendor concentration	Procurement DB	HHI of contract awards	L	Sparse for small orgs	g_4 Procurement integrity
Grant concentration	Open data portals	HHI of grant sources	M	Multi-source join	g_2, g_5
Board interlocking degree	Registry cross-ref	Shared directors / board size	M	Partially complete	$\mathscr{H}_t$ Network layer
Geographic proximity	GIS + registered address	Distance to sensitive sites	H	Complete	$\Psi(\cdot)$ Threat surface
Foreign funder indicator	T3010 / 990	Non-domestic source flag	M	Aggregate disclosure	g_1, g_3, Ψ

© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com

Table 2: **Projected vs. Baseline Detection Performance (Synthetic).** Ranges represent expected variation across organizational tier (small / large) and deterioration speed (rapid / gradual). All values are theoretical projections pending empirical validation. Enterprise detection is undefined for baselines lacking a network layer.

Metric	**Baseline A**	**Baseline B**	**SENTINEL**
Detection lag (filing periods)	4–8 (if detected)	3–6	1–3
Precision at threshold	0.25–0.45	0.35–0.55	0.60–0.80
Recall at threshold	0.30–0.50	0.40–0.60	0.65–0.85
Small-org recall specifically	0.10–0.25	0.20–0.35	0.55–0.75
Enterprise pattern detection	N/A	N/A	0.70–0.85

© 2026 Brent Figg. Free to read, download, and cite with attribution; all other rights reserved. brent@2ndandcapital.com