

Paper X

The Precognitive Commons

Why Siloed Precognitive Tools Fail and a Unified Measurement Architecture Can Eliminate Mass Shootings

Robert Thomas Fair Jr.

Independent Global Systems Policy Research
Great-Nephew of Bill Fair — Co-Founder of FICO
May 2026

Abstract

The Fair Framework has mapped seven planetary commons and established the prerequisite that governance must precede AGI. This paper applies the same five-step diagnostic to a new, critical commons: the precognitive commons — the shared signals of harmful intent that exist across behavioral, digital, psychological, and physiological data streams.

Current precognitive tools (threat assessment algorithms, behavioral analytics, digital signal monitoring) exist in silos. When synthesized through a single, lean measurement architecture, they can detect and interrupt mass-shooting pathways before the act. The Receiver Problem (Paper IX) provides the mechanistic foundation: if intent manifests as a patterned, multi-domain broadcast, then properly calibrated receivers can surface it earlier than today's fragmented systems allow.

This is not pre-crime fiction. It is the logical extension of the Fair Framework: make the unseen measurable, protect civil liberties through constitutional architecture, and govern the commons before preventable harm becomes irreversible.

I. Identify the System — The Precognitive Commons

The precognitive commons is the shared, cross-domain pool of signals that precede targeted violence: anomalous behavioral patterns, digital activity clusters, psychological markers, and physiological anomalies. These signals are real, observable, and increasingly quantifiable. They are not thought crime; they are measurable precursors — analogous to the chokeholds mapped in Papers I–VII.

The empirical foundation is well-established. The U.S. Secret Service National Threat Assessment Center (NTAC) has documented for two decades that targeted violence is rarely impulsive: attackers follow identifiable pathways of ideation, planning, and preparation before acting. The FBI's Behavioral Threat Assessment Center (BTAC) corroborates this, finding that in the vast majority of studied mass casualty events, warning signs were present and observable in advance — across behavioral, digital, and social domains — but were never synthesized into a unified signal.

The commons exists. It is simply ungoverned.

II. Map the Chokehold — Siloed Actors

The chokehold is extreme fragmentation across four principal domains:

- Law enforcement holds behavioral records, criminal histories, and prior contact data.
- Technology platforms hold digital and social signals: search histories, communication patterns, content engagement, anomalous behavioral metadata.
- Mental health systems hold psychological markers, crisis contact records, and clinical risk assessments.
- Intelligence agencies hold classified threat streams and foreign-influence indicators.

No single entity sees the cumulative picture. Each silo optimizes for its own mandate, its own legal authority, and its own institutional risk tolerance. The result is a set of blind spots that mass shooters repeatedly exploit — not because the signals were absent, but because no authorized body was tasked with reading them together.

This is the structural definition of a commons failure: individually rational actors, collectively catastrophic outcomes.

III. Expose the Governance Vacuum

There is no institution, protocol, or cross-domain synthesis mandate that treats precognitive signals as a governed commons. No equivalent to the proposed GFSOB, PEBO, or GOAMN exists for intent-risk monitoring. Current efforts are reactive, post-event, or narrowly scoped within a single agency's statutory authority.

The vacuum is not incidental — it is the product of legitimate but competing constraints. Privacy law, agency jurisdiction, Fourth Amendment doctrine, and HIPAA all operate as vertical barriers preventing horizontal synthesis. These barriers serve vital purposes.

The governance failure is not that they exist, but that no architecture has been designed to synthesize signals across them in a constitutionally compliant manner.

The system currently waits for the act instead of measuring the precursors. Every incident leaves a retrospective trail — across law enforcement contacts, online platforms, school or workplace interactions, and health systems — that, when integrated after the fact, forms a clearly detectable pattern. The governance challenge is to build the architecture that reads this pattern before, not after.

IV. Name the Cascading Risk

The cascading risk is clear, measurable, and preventable. Mass shootings and other forms of targeted violence that could be interrupted earlier impose costs across three dimensions:

- Human cost: lives lost and communities permanently altered by preventable violence.
- Institutional cost: eroded public trust in the systems — law enforcement, schools, mental health — that were present but failed to act on available signals.
- Policy cost: reactive legislative cycles that produce surveillance overreach or ineffective restrictions precisely because no measurement-first alternative has been institutionalized.

The empirical evidence supports a strong claim: integrated multi-domain threat assessment can interrupt a substantial proportion of mass-casualty pathways. The NTAC's 2021 study of targeted school violence found that in every case examined, the attacker communicated intent to others before acting. The FBI's 2018 study of active shooter incidents found that over 60% of incidents ended before law enforcement arrived — meaning earlier interruption, not armed response, is the primary leverage point.

The claim here is not that every act of targeted violence is detectable in advance with certainty. It is that the current system detects far fewer than it could, because cumulative signal synthesis is architecturally impossible under the present governance structure. A measurement-first architecture can close that gap substantially.

V. Propose Lean Institutional Architecture

The Fair Framework's standard response applies: identify the minimum viable governance structure that closes the commons failure without replicating the pathologies of the system it replaces. In this case, that requires confronting the civil liberties surface area directly — not with a gesture, but with constitutional architecture.

A. The National Precognitive Monitoring Body (NPMB)

A lean, single-mandate federal body tasked solely with cumulative signal synthesis. The NPMB does not investigate, arrest, prosecute, or surveil. Its sole function is to receive synthesized risk scores from participating agencies and flag cases that cross a judicially established threshold for multi-agency review. Institutional design constraints:

- No direct data collection authority — the NPMB synthesizes scores, not raw data.
- Mandatory sunset review every five years by an independent oversight board with civil liberties, legal, and technical representation.
- All algorithmic models subject to public audit by the Government Accountability Office and annual third-party bias review.
- Explicit prohibition on use of race, religion, national origin, political affiliation, or First Amendment-protected activity as input variables.

B. The Cumulative Precognitive Impact Assessment Protocol (CPIAP)

A cross-domain data-sharing framework built on the following constitutional and statutory guardrails:

- Judicial authorization required before any individual-level synthesis — a warrant standard modeled on the Foreign Intelligence Surveillance Court but with broader civil liberties representation and mandatory defense-side advocacy.
- Strict data minimization: agencies share risk-relevant signal scores, not underlying records. Raw psychological, medical, or communications data never enters the NPMB layer.
- Automatic expungement: synthesis scores older than 18 months without escalation are destroyed, with mandatory audit trail confirmation.
- Independent civil liberties monitor embedded within NPMB operations with full access and public reporting authority.

C. The Receiver Problem Integration

Paper IX established that intent, under certain conditions, manifests as a patterned, multi-domain signal detectable by properly calibrated receivers. The CPIAP operationalizes this by treating behavioral, digital, psychological, and physiological signals as a composite broadcast — no single channel sufficient, cumulative synthesis necessary.

This mirrors the original FICO insight: no single credit variable predicts risk reliably; the synthesis of variables across domains produces a score that is meaningfully more

predictive than any component. The precognitive application follows identical measurement logic, with constitutional constraints appropriate to its higher stakes.

D. What This Architecture Is Not

Precision on exclusions is as important as precision on inclusions:

- Not a surveillance state. The NPMB has no persistent monitoring authority and no data collection power. It synthesizes agency-provided scores on a warrant-authorized basis only.
- Not pre-crime punishment. Flagging generates a multi-agency review and, where warranted, targeted intervention (mental health outreach, law enforcement contact) — not pre-emptive detention or prosecution.
- Not a replacement for social investment. Threat assessment architecture is a last-layer intervention. It operates downstream of the community, educational, and mental health investments that reduce the pipeline of grievance and isolation that produces attackers.

Conclusion

The Fair Framework has shown that ungoverned commons become catastrophic when optimizing forces — whether market actors or AGI systems — arrive before governance does. The precognitive commons is no different. The signals are real. The silo structure is the failure. The governance vacuum is total.

A unified measurement architecture — built on the same first-principles logic that produced FICO and the rest of this series — can close the gap. Not by watching everyone, but by synthesizing what agencies already observe, across domain barriers that currently make cumulative reading impossible.

The civil liberties risk is real and must be governed with the same rigor applied to the violence risk. The architecture proposed here — warrant-based synthesis, score-not-data sharing, mandatory expungement, embedded civil liberties monitoring, and explicit prohibitions on protected-class inputs — is designed to hold both constraints simultaneously.

The pieces are there. The vacuum is clear. The governance solution is ready.

Govern the commons first.

Evidentiary Basis

U.S. Secret Service National Threat Assessment Center (NTAC). Averting Targeted School Violence: A U.S. Secret Service Analysis of Plots Against Schools (2021). Findings: every attacker in the study communicated intent before acting; warning behaviors were observable across multiple domains in advance.

FBI Behavioral Threat Assessment Center (BTAC). A Study of the Pre-Attack Behaviors of Active Shooters in the United States Between 2000 and 2013 (2018). Finding: 63% of active shooter incidents ended before law enforcement arrived, establishing pre-event interruption as the primary leverage point.

Meloy, J.R. & Hoffmann, J. (Eds.). International Handbook of Threat Assessment (2nd ed., 2021). Oxford University Press. Provides empirical grounding for multi-domain behavioral threat assessment as a clinical and institutional practice.

Fair Framework cross-references: Paper IX (The Receiver Problem — mechanistic foundation for multi-domain signal synthesis); Papers I–VII (planetary commons diagnostic methodology); Paper VIII (governance-before-AGI prerequisite).
