

SERIES INTRODUCTION

The Foundations of Predictive Governance

A Guide to the Series for Practitioners and Decision-Makers

Rolf Sartorius | Saipion

Foundations of Predictive Governance Series | March 2026

About this document

This introduction provides an overview of the Foundations of Predictive Governance Series produced by Saipion. The series is intended to frame with evidence the concept for an Integrated Outcomes Management System—**Civryn**. The series consists of ten papers written for practitioners in government, international development, civil society, and philanthropy who want to improve how programs and policies are designed, implemented, and managed. This guide describes what the series covers, how the papers connect, and what readers can expect to find in each one. It is intended both as a standalone orientation and as a companion to the individual papers.

The Problem This Series Addresses

Governments and development institutions invest enormous resources in designing policies and programs that then fail to deliver. This is one of the best-documented and least-addressed problems in public administration. The analysis is often sound. The intentions are genuine. But somewhere between approval and implementation, things go wrong—budgets are consumed, political capital is spent, and the people the programs were meant to reach are left waiting.

The causes are predictable. Program design does not adequately take into account the views of program clients, enhance client agency or reflect local context. Political support that existed at approval quietly erodes as implementation progresses. Programs are too difficult for the people they serve to actually access. The agencies responsible don't have the capacity to deliver what they've been asked to do. Frontline workers adapt in ways that weren't intended. Multi-agency delivery falls apart without anyone managing the whole. These aren't random failures—they follow patterns that four decades of research have described with increasing precision.

At the same time, something significant has changed. Advances in AI and data systems have created capabilities that didn't previously exist: synthesizing large bodies of evidence rapidly, modelling what a policy will actually do before it's implemented, monitoring how delivery is going in something close to real time, and detecting early warning signals before they become crises. These tools don't fix the underlying political and organizational challenges of governance—but used well, they can make those challenges substantially more manageable.

This series is about how to use those tools well. It covers the full cycle from the initial question of whether a program is worth doing, through design, implementation, evaluation, and portfolio management—addressing at each stage what good practice looks like and what AI can genuinely contribute. The series is written for practitioners, not academics. It draws on research, but the goal throughout is to give decision-makers and program managers tools they can actually use.

The Series Argument

Effective governance means building the sustained capacity to solve problems at scale—not just running programs. That capacity rests on five things: knowing what works; managing the political conditions under which it can work; learning from what's actually happening and adapting; ensuring that what works reaches those who need it most; and closing the gap between what policy intends and what delivery produces. AI helps with all five. It replaces none of them. The work remains human.

The Ten Papers

The series is organized around the program and policy cycle—from foundational concepts and design, through delivery and learning, to integrity, strategic management and evaluation. Each paper is self-contained and can be read independently, but they are more useful together: each builds on the preceding papers and connects forward to the ones that follow.

THE FOUNDATIONS OF PREDICTIVE GOVERNANCE SERIES — PAPER OVERVIEW

Paper	What It Covers
1. Predictive Governance Foundations, Evolution, and a Framework for Outcome-Driven Public Systems	Sets out the central argument: governance systems weren't built for the complexity and pace of the problems governments now face, and AI-enabled predictive tools offer a response—but only if the institutional foundations are in place. Covers the history of predictive analytics in government, the promises and documented failures, and a practical framework for what predictive governance requires. The essential starting point for the series.
2. Client Agency in the Program Cycle	Shifts the frame from beneficiary participation to client and user agency: programs do not grant agency, they either respect, constrain, or expand what the people they serve already possess—the capacity for informed choice, voice, accountability, and co-governance. Synthesizes several agency frameworks into a five-mode spectrum mapped to each stage of the program cycle, and addresses how predictive analytics should amplify client signals while preserving standing to contest them.
3. Program Design Building Programs That Can Actually Adapt	Addresses quality at entry: the design decisions that determine whether a program can be managed and adapted once implementation begins. Covers theory of change development, assumption-testing, and the design of feedback loops that make adaptive management possible. A weak design constrains everything that follows; this paper is about preventing that constraint.
4. ERR and SROI Analysis Making the Case for Investment—and Knowing When to Stop	Covers the economic and social value frameworks governments and funders use to assess whether a program is worth doing: Economic Rate of Return (ERR), Social Return on Investment (SROI), and cost-benefit analysis. Addresses how to use these tools rigorously, what their limits are, and how AI lowers the cost of design-stage value assessment without changing its underlying disciplines.
5. Adaptive Management Learning from What's Actually Happening	Addresses how organizations can tell whether what they're doing is working—and actually change course based on that evidence, rather than just reporting it. Covers the conditions that distinguish genuine adaptation from routine compliance, and the specific contributions AI makes to real-time monitoring, early warning detection, and evidence-based course correction.
6. Program Evaluation Building the Evidence That Makes Management Credible	Covers the evaluation infrastructure that makes adaptive management credible to stakeholders, funders, and political actors. Addresses evaluation design, the trade-offs between rigor and timeliness, and how AI tools accelerate evidence synthesis without replacing the methodological standards that make evaluation findings trustworthy.
7. Fraud and Corruption Protecting Integrity When It Matters Most	Identifies the organizational and political conditions under which programs leak resources and distort results—and the monitoring, accountability, and deterrence mechanisms that limit this. Covers AI-enabled anomaly detection and financial monitoring, while being clear about the governance requirements that make these tools effective rather than merely visible.
8. Strategic Management Governing Organizations That Deliver	Addresses the leadership and management challenges of organizations responsible for implementing complex programs and policies: how to align incentives with outcomes, manage performance in ways that don't create perverse incentives, and build the organizational cultures that adaptive governance requires.

9. Portfolio Management Managing Programs as a System, Not a Collection	Addresses the challenge of governing not a single program but a portfolio of interconnected interventions: allocating finite capacity across competing demands, managing the dependencies and synergies between programs, and making resource decisions that reflect evidence about relative impact rather than organizational habit.
10. From Programs to Policy: Implementing Systems Level Change Evidence, Insights, and AI Augmentation for Adaptive Governance	The final paper. Addresses the question that runs under all the others: what does it actually take to get from a good design to real change on the ground? Covers the five dimensions along which implementation most commonly fails—coalition dynamics, access barriers, delivery capacity, frontline discretion, and multi-agency coordination—and what AI can and cannot contribute to managing each one.

Five Findings That Run Across All Ten Papers

Certain lessons recur with enough consistency across the series to be worth stating directly. They are not always comfortable—some cut against common assumptions about what AI can do and what governance requires. But they are grounded in the evidence, and they are the things this series most wants readers to take away.

1. The gap between what programs and policies are designed to do and what they deliver is an organizational and political problem—not an analytical one

Every paper in the series encounters some version of this gap. The binding constraint on better governance is rarely a lack of data or analysis—it is the absence of the feedback channels, adaptive authority, and institutional cultures that turn evidence into action. Investing in better analysis without investing in those organizational conditions produces better-documented failures, not better outcomes.

2. AI amplifies whatever institutional capacity already exists—it doesn't create capacity where there is none

This is the most important thing to understand about AI in governance. The tools are genuinely powerful: they can synthesize evidence at scales that weren't previously feasible, monitor delivery in something close to real time, and detect early warning signals that human oversight would miss. But they work by amplifying the capacity of institutions that are already functioning. An organization without the systems, skills, and culture to act on evidence will not be transformed by better evidence. The hard institutional work has to happen alongside the technology—not after it.

3. Equity is a diagnostic signal, not a separate workstream

When access barriers fall hardest on the most vulnerable, when services are rationed away from the most complex cases, when fraud and leakage divert resources from intended beneficiaries—something in the system is failing. The equity lens, applied consistently across the design, delivery, and evaluation of programs, is one of the earliest and most reliable indicators that something has gone wrong. Treating equity as a separate compliance requirement rather than a diagnostic tool means missing the signal until it's too late.

4. Political management is a learnable skill, not a background condition

The political dimensions of governance—building coalitions, managing opposition, navigating the dynamics of blame and credit—are consistently treated in practice as contextual factors to be noted and set aside, rather than as core management disciplines. The evidence across the series suggests this is a mistake. Coalition management, stakeholder strategy, and political mapping are skills that can be developed, tools that can be applied, and investments that pay off in implementation performance.

5. The kind of change that produces durably better outcomes takes years, not project cycles

The governance improvements documented across the series—data systems that generate reliable evidence, analytical skills that can interpret and act on it, learning cultures that treat course correction as competence, trust relationships that enable genuine collaboration—all accumulate slowly. They cannot be shortcut by better technology or larger budgets. This is not a counsel of despair: the series documents many cases where genuine progress was made. But it was made by institutions that invested consistently over time, not by those that deployed the right tools and expected transformation.

How to Use the Series

The papers are designed to be useful both together and independently. Readers who are new to predictive governance or to AI in public policy should start with Paper 1, which establishes the foundational concepts and the framework the rest of the series builds on. Readers with a specific question—how to evaluate a program, how to reduce access barriers, how to manage a portfolio—can go directly to the relevant paper.

Each paper includes a brief executive summary written for decision-makers who need the key findings quickly, followed by detailed sections that practitioners can use to apply the material in their own contexts. All papers include endnotes with full citations to the research and practitioner literature they draw on.

Two papers are particularly foundational for readers who want to understand the series as a whole: Paper 1, which sets out the case for predictive governance and the conditions under which it can work; and Paper 10, which addresses what policy implementation actually requires and brings together the practical tools and evidence from across the series. Taken together, they frame the argument the series makes from beginning to end.

A Note on AI Capabilities

The series was written in 2025–26, and the AI capabilities described throughout reflect conditions as of early 2026. The field is developing rapidly. Readers should treat specific capability claims as reflecting the state of the art at the time of writing rather than as settled facts—and should expect both the tools and their governance requirements to continue evolving. The underlying analytical and management frameworks the series develops are intended to remain useful regardless of how the technology changes.