No. 2026-002

When Monitoring Becomes Control: Endogenizing Internal Control Risk in Continuous Auditing

April 24, 2026

Takanori Suzuki

When Monitoring Becomes Control: Endogenizing Internal Control Risk in Continuous Auditing

Takanori Suzuki*

April 24, 2026

Research Institute of Business Administration (Sanken), Waseda University
Working Paper Series, No.2026-002

Working paper note. This version is prepared for circulation in the Working Paper Series of the Research Institute of Business Administration (Sanken), Waseda University. A revised version may be submitted to a journal.

Abstract

The audit-risk model AR = IR × CR × DR rests on a tacit assumption that control risk CR is exogenous to the audit—evaluated through the risk-assessment process and then taken as a given input to DR determination. Continuous auditing disrupts this assumption. ERP log extraction, process-mining indicators, transaction-network analytics, and similar technologies expand the auditor's evidence space, but the same deployments also alter the firm's control operations through workflow feedback. When they do, CR becomes endogenous to the monitoring architecture, and attributing the technology's effect to DR alone may double-count. This paper develops a state-space framework that makes the CR-endogenization channel formal and derives a two-stage *control-assurance coupling theorem*: within the linear Route B model, monitoring-to-control reachability is necessary and sufficient for functional CR endogenization, while control-risk sensitivity and cumulative monitoring-to-control influence provide additional sufficient conditions for certifying an operationally meaningful uniform lower bound on the induced CR shift. A secondary behavioral-feedback channel is shown to endogenize IR under more restrictive conditions. The results support a disciplined three-effect attribution rule (evidence / control / behavior) that avoids double counting, formalized and illustrated through circular-transaction and Level 3 fair-value cases.

Keywords: Continuous auditing; internal control risk; inherent risk; endogenous risk; state-space model; audit risk model; monitoring-control coupling; digital audit governance.

JEL classification: M41, M42, C61, D82, G34.

*Faculty of Commerce, Waseda University. Email: suzuki@waseda.jp. ORCID: 0009-0008-8955-7003.

I. Introduction: When Monitoring Becomes Control

1.1 The Question This Paper Addresses

Audit theory lacks a framework for when monitoring technologies change the assessed risk structure itself. The multiplicative audit-risk model $\mathrm{AR} = \mathrm{IR} \times \mathrm{CR} \times \mathrm{DR}$, codified in ISA 200 and its antecedents, rests on a tacit separation: inherent risk and control risk are features of the firm that the auditor evaluates, while detection risk is the feature of the audit that the auditor designs. This separation organizes how auditors reason, how auditing standards are written, and how the quantitative audit-planning literature formalizes risk assessment. In the standard treatment, $(\mathrm{IR}, \mathrm{CR})$ are exogenous to the audit—evaluated through the ISA 315 risk-assessment process and then taken as given inputs to DR determination. The two classes of inputs are operationally independent because they are set from different information sources.

Continuous auditing disturbs this separation. The ERP log extractions, process-mining indicators, confirmation APIs, transaction-network analytics, spectral monitoring, and AI-based anomaly flags that now constitute a standard continuous-audit technology stack (Vasarhelyi and Halper, 1991; Alles et al., 2006; Chan and Vasarhelyi, 2011; Krahel and Titera, 2015; Jans et al., 2013; Commerford et al., 2022) were introduced with detection risk in mind: they deliver more evidence, more frequently, along directions that conventional procedures cannot easily observe. But these same technologies operate through the firm's control apparatus. Approval-path monitoring changes approver behavior. Exception-routing systems change the speed and disposition of control deviations. Transaction-network analytics change the feasibility of circular-transaction schemes, not merely their detection probability. Process-mining dashboards feed back to the control owners whose processes are mined. Under continuous auditing, the monitoring architecture is simultaneously an evidence-collection technology and an intervention on the firm's control and behavior space.

Two questions follow. First, what is the formal relationship between monitoring-induced evidence expansion (which affects detection risk) and monitoring-induced control-state change (which should affect control risk)? Second, under what conditions does continuous auditing move the firm's risk structure itself, as distinct from observing a given risk structure more finely? The present paper answers both questions by introducing a distinction between two mechanisms through which continuous auditing affects audit risk—which I label *Route A* (the observation-channel mechanism, in which monitoring expands the evidence structure while leaving IR and CR exogenous) and *Route B* (the intervention-channel mechanism, in which monitoring enters the control-risk function and the managerial policy-input distribution). Section III develops the distinction conceptually; Sections IV–VIII develop the formal Route B model. Related analytical work on the Route A side is discussed in Section II.

1.2 The Central Claim

This paper advances a single central claim:

Continuous auditing is not only an evidence-space architecture that reduces detection risk *under a given risk structure; it is a control-assurance feedback architecture whose effects on the firm's control state endogenize control risk, and through a limited behavioral-response channel also shift the distribution of misstatement-generating policy inputs, under identifiable conditions. When these conditions hold, attributing the monitoring technology's total effect to a single audit-risk component—rather than separately to its evidence, control, and behavior effects—may lead to double counting in the audit-risk model* $\text{AR} = \text{IR} \times \text{CR} \times \text{DR}$.

This claim has two consequences that organize the rest of the paper. First, the monitoring architecture m_t plays a triple role: it populates the time-indexed evidence channel C_t (the Route A channel, on which existing audit-planning frameworks focus; see Section II for a review), it enters the internal control state transition equation through a gain operator G, and it shifts the distribution of managerial policy inputs u_t through a behavioral response operator. Each role corresponds to one component of audit risk: DR, CR, and IR, respectively. Second, the question of whether audit-risk assessment can treat (IR, CR) as exogenous—as Route A does—becomes a two-stage question. If monitoring-to-control reachability fails, Route B is functionally inactive and the Route A reduction remains valid. If reachability holds, CR becomes a non-constant function of monitoring history, so treating it as strictly exogenous is structurally incomplete; however, *operationally meaningful* CR endogenization—in the sense that the induced CR shift exceeds an audit-practice coverage threshold uniformly over CR-relevant directions—additionally requires control-risk sensitivity and cumulative monitoring-to-control influence. When all three control-assurance coupling conditions hold, treating CR as fully exogenous may double-count the monitoring-technology effect unless its evidence, control, and behavior effects are separately attributed under a specified convention.

1.3 Route A and Route B

I use *Route A* and *Route B* as compact labels for two mechanisms through which continuous auditing affects audit risk.

Route A is the *observation channel*: monitoring technologies populate the time-indexed evidence channel C_t, expanding the residual-focused observability Gramian and reducing model-implied detection risk. This is the mechanism formalized in existing audit-planning frameworks that treat continuous auditing as evidence-space architecture—notably the analytical line developing observability-based audit planning from a residual-focused Gramian (Author, 2026b,c).[1] Route A takes (IR, CR) as exogenous assessed risks and asks how C_t reduces DR for the given risk structure.

Route B is the *intervention channel*: monitoring technologies enter the firm's control dynamics and policy-choice distribution, altering CR and (through behavioral response) IR. Under Route B, (IR, CR) are no longer exogenous to the monitoring architecture; they are endogenous variables whose values depend on the history of monitoring deployment.

[1] References to Author (2026a/b/c) are anonymized because related manuscripts are currently under blind peer review. A note at the start of the References section records this convention; the references will be de-anonymized in future versions of this working paper once the associated peer-review processes are complete.

The two routes are not mutually exclusive. A given monitoring technology may operate purely through Route A (extracting ERP logs without changing control operations), purely through Route B (changing approver behavior without generating additional auditor evidence beyond what the firm already self-reports), or—more commonly—through both routes simultaneously. The question of how to decompose the observed effect into the two routes is the substantive content of Sections V through VIII.

1.4 Contributions

The paper makes four contributions.

Contribution 1: A state-space formalization of the control-assurance feedback architecture. I develop a state-space model in which an internal control state k_t is added to the standard residual-relevant accounting state as a first-class object, with its own transition equation $k_{t+1} = Fk_t + Gm_t + \zeta_t$, in which the monitoring architecture m_t acts through a gain operator G. The model nests the familiar case in which monitoring acts only on the observation channel (recovered when $G = 0$) and generalizes it to the case in which monitoring also acts on control dynamics (when G is nondegenerate on the control-relevant subspace).

Contribution 2: A three-effect decomposition of continuous-audit technology. I show that the total effect of a monitoring technology on audit risk admits a *disciplined first-order decomposition* into an *evidence effect* (on DR), a *control effect* (on CR), and a *behavior effect* (on IR), under a specified attribution convention defined in Section VII. The decomposition is path-derivative based with a residual $O(\|m\|^2)$ cross-effect term that is explicitly characterized, not an exact identity; what it provides is a consistent attribution rule for audit workpapers, not a claim of a uniquely canonical decomposition across all possible conventions. This disciplined attribution is the formal basis for avoiding double counting under the chosen convention.

Contribution 3: A two-stage control-assurance coupling theorem (primary result). The paper's primary theoretical result provides a two-stage characterization. Theorem 8.7 shows that monitoring-to-control reachability is necessary and sufficient, within the maintained linear Route B model, for functional CR endogenization—i.e., for CR_T to become a non-constant function of monitoring history. Theorem 8.8 shows that control-risk sensitivity and cumulative monitoring-to-control influence provide additional sufficient conditions for certifying an operationally meaningful uniform lower bound on the induced CR shift, of magnitude $\gamma_K\sqrt{\tau_K}$. The theorems are control-side counterparts to existing analytical frameworks that treat the audit observation channel as the locus of continuous-audit effects; they shift the analytical weight toward the control channel when the monitoring architecture acts on the firm's internal control dynamics, not only on the auditor's evidence stream. A parallel secondary result characterizes when the behavior-effect channel additionally endogenizes inherent risk, under the caveats discussed in Section VI.

Subsidiary contribution: An operational empirical signature. As a subsidiary consequence of the main theorem, I show that when a new monitoring channel is deployed along a previously unmonitored control dimension, the firm's optimal level of control effort on incumbent procedures responds non-monotonically, with a closed-form critical value. This provides a specific empirical prediction for continuous-auditing adoption studies (the formal statement and closed-form derivation are in Appendix A).

1.5 What This Paper Does Not Claim

Three scope disciplines are imposed.

Full IR endogenization is not attempted. Inherent risk depends on business model, transaction structure, and external environment in ways that monitoring technologies do not directly reach. The paper treats IR effects as a *limited, secondary* channel operating through policy-input redistribution within the manager's discretion; it does not claim that continuous auditing transforms IR structurally. The primary theoretical result (Section VIII) concerns CR endogenization; the IR result (Section VIII, Corollary) is a parallel but weaker statement that is asymmetric with respect to strategic-adaptation. Full IR endogenization—requiring a game-theoretic or strategic-adaptation extension—is identified as future work (Section XIII).

Source reliability is treated minimally. The paper incorporates source reliability through a weight q_t on evidence channels and k-state updates, consistent with the reliability-adjusted Gramian of Author (2026c). An information-theoretic treatment that derives reliability weights from channel–source mutual information is deferred.

The illustrations are informed conjectures, not deductive conclusions. The worked illustrations in Sections IX and X classify specific accounting phenomena (circular transactions, Level 3 fair value) along the model's geometric axes. The classifications are based on institutional observation, not derived from the formal model; whether a specific phenomenon actually satisfies the three conditions of the control-assurance coupling theorem is an empirical matter.

1.6 Organization

Section II positions the paper relative to three literature streams. Section III develops the Route A/Route B distinction in conceptual terms. Section IV presents the state-space model with the control state k_t added. Section V formalizes the endogenization of control risk, including the control-effect condition. Section VI develops the limited endogenization of inherent risk through behavior effects. Section VII addresses the double-counting problem via the three-effect decomposition. Section VIII presents the control-assurance coupling theorem as the main result. Sections IX and X give the two worked illustrations. Section XI presents implications for audit planning, auditing standards, and digital audit governance. Section XII sketches empirical and simulation implications. Section XIII discusses limitations and future research. Section XIV concludes. Proofs and auxiliary material are in the Appendix.

II. Related Literature and Positioning

The paper stands at the intersection of three literatures: the audit-risk model and its constituent-risk literatures, the continuous-auditing and AI-auditing literature, and the literature on internal control design, monitoring, and behavioral deterrence. I begin with the audit-risk side because that is where the formal Route A/Route B distinction is located, and then discuss the monitoring/control side where the Route B mechanism operates.

2.1 The Audit-Risk Model and the Exogeneity of IR and CR

The quantitative treatment of audit risk originates in the multiplicative model $\mathrm{AR} = \mathrm{IR} \times \mathrm{CR} \times \mathrm{DR}$ codified in ISA 200 and its U.S. analogues. Kinney (1984) articulates the role of inherent risk assessment and statistical sampling. Srivastava and Mock (2002) develop belief-function alternatives to probabilistic risk aggregation. Fukukawa and Mock (2011) compare belief- and probability-based IR–CR assessments empirically. Dutta and Graham (1998) extend the model to multiple-materiality settings for account combinations.

A recurring feature of this literature is that IR and CR are treated as features of the firm, estimated by the auditor through risk-assessment procedures under ISA 315 (or analogous PCAOB and AICPA standards), and then used as inputs to determine the required level of DR. The audit-planning optimization that follows—whether in sampling-based, Bayesian, or observability-based formulations—takes $(\mathrm{IR}, \mathrm{CR})$ as given and asks how to design audit procedures to achieve the target DR.

The observability-based formulation developed in Author (2026b) is explicit about this assumption. It distinguishes structural parameters (A_c, Q, C, R, M) that describe the firm and the audit procedures from assessed risks $(\mathrm{IR}, \mathrm{CR})$ that enter as external inputs from the ISA 315 risk-assessment process. The two are treated as "operationally independent" because they are set using different information sources, and the formulation accordingly avoids double counting. That formulation is explicit, however, that the operational-independence assumption is a reduction (Route A) and that a more deeply integrated formulation in which $(\mathrm{IR}, \mathrm{CR})$ are derived from the structural parameters (Route B) is left as future research.

The bridge paper on continuous auditing (Author, 2026c) maintains the Route A reduction: it extends the static observability-based framework to the time-indexed setting appropriate to continuous auditing, but keeps $(\mathrm{IR}, \mathrm{CR})$ exogenous. It explicitly identifies Route B as the next bridge paper in the research program.

The present paper undertakes that Route B extension. What distinguishes it from prior audit-risk literature is not that it questions the multiplicative structure of $\mathrm{AR} = \mathrm{IR} \times \mathrm{CR} \times \mathrm{DR}$—the paper retains that structure as a first-order approximation—but that it treats the monitoring architecture m_t as entering all three factors simultaneously, with identifiable conditions under which the induced dependence is nondegenerate.

2.2 Continuous Auditing, Audit Analytics, and AI Auditing

The continuous-auditing literature originates with Vasarhelyi and Halper (1991) and develops through Alles et al. (2006), Chan and Vasarhelyi (2011), and Krahel and Titera (2015). More recent work addresses AI-based anomaly detection and machine-learning applications (Commerford et al., 2022;

Sun, 2019; Dickey et al., 2019), process mining as an audit technology (Jans et al., 2013), and Big Data's effects on audit judgment and evidence quality (Brown-Liburd et al., 2015; Cao et al., 2015).

A shared feature of this literature is that continuous-audit technologies are discussed primarily as evidence-generation technologies. Their effect on audit risk is framed as a detection-risk effect: more evidence, more frequently, yields a lower DR for a given $(\mathrm{IR}, \mathrm{CR})$ assessment. Exceptions exist—Alles et al. (2006) note the feedback effect of continuous monitoring on control operation, and Chan and Vasarhelyi (2011) discuss how continuous auditing changes what "internal control" means in the first place—but the formal modeling has not caught up to these observations. What remains absent is a framework in which the monitoring architecture enters the formal audit-risk model on both the detection side and the control side simultaneously.

The present paper supplies that framework by extending the state-space formulation of Author (2026c) with a control-state transition equation in which m_t appears through a gain operator G. Route A (pure evidence-channel effect) is recovered as the special case $G = 0$.

2.3 Internal Control, Monitoring, and Behavioral Deterrence

Internal control is formalized in the COSO framework (COSO, 2013) and in the parallel internal-control literature in accounting and economics. Kinney (2000) provides an early statement of the relationship between control risk and audit planning. Doyle et al. (2007) document determinants of material weaknesses in internal control over financial reporting. Beneish et al. (2008) link internal-control weaknesses to audit-risk premia. Ashbaugh-Skaife et al. (2008) document the effect of SOX internal-control disclosures on risk and returns.

Two strands of this literature are especially relevant. First, the link between monitoring intensity and deviation reduction has been documented empirically: firms with more active monitoring, more frequent exception review, and stronger second-line-of-defense functions exhibit fewer control failures conditional on exposure. Second, the behavioral-deterrence literature, drawing on Becker (1968)'s economics of crime and its accounting applications (Baiman, 1982; Dye, 1986), models the agent's deviation decision as a function of detection probability and sanction magnitude. When detection probability increases, deviation rates decrease, all else equal.

These strands together motivate the core mechanism of Route B. The monitoring architecture m_t affects the control state k_t through both an operational channel (exception-routing systems, approval-path monitoring, process-mining feedback loops) and a behavioral channel (deterrence through detectability). The formal model in Section IV abstracts from the behavioral mechanism while retaining its implications for the control-state dynamics.

2.4 Position of the Present Paper

At the intersection of the three literatures lies an unaddressed question. The audit-risk literature provides the multiplicative model but treats IR and CR as exogenous; the continuous-auditing literature describes the technology but has not modeled its dual effect on DR and CR; the internal-control literature documents monitoring-to-deviation links but is not integrated with the formal audit-risk model. The present paper occupies this intersection by introducing a state-space control-assurance feedback architecture in which the monitoring technology enters all three audit-risk factors. Methodologically, it builds on the observability-based audit-planning machinery that treats the evidence channel through a cumulative observability Gramian (Author, 2026b,c) and extends that machinery to the control side. The progression from recordability through cumulative auditability to the Route B condition developed in this paper can be written compactly as

$$\underbrace{\dim \mathcal{E}_t^{\perp} \geq \operatorname{rank}(\Omega)}_{\text{recordability (Author, 2026a)}} \implies \underbrace{\lambda_{\min}(W_T^{R,M}) \geq \tau_R}_{\text{cumulative auditability (Author, 2026c)}} \implies \underbrace{\text{control-assurance coupling active}}_{\text{Route B (this paper)}}$$

A diagrammatic summary of the broader research program and its open directions is deferred to the Conclusion (Figure 5); for the purposes of reading this paper in isolation, what matters is the Route B condition itself, developed in Sections IV–VIII.

III. Route A and Route B: Two Mechanisms of Continuous Auditing

Before developing the formal model, I articulate the Route A/Route B distinction in conceptual terms. The distinction organizes the rest of the paper and is the basis of the contribution structure stated in Section I.

3.1 Route A: Observation-Channel Mechanism

Under Route A, continuous-audit technologies are evidence-generation devices. They expand the time-indexed evidence channel C_t, populate the residual-focused observability Gramian W_T^R, and reduce the model-implied residual detection risk $\mathrm{DR}_T^{\mathrm{res}}$, all for a given structure of inherent and control risk. The firm's internal control state is exogenous to the monitoring deployment; the firm's policy inputs are exogenous to the monitoring deployment; $(\mathrm{IR}, \mathrm{CR})$ are features of the firm that the auditor evaluates and takes as given inputs to DR determination.

The defining formal property of Route A is that the monitoring architecture enters the audit risk model through exactly one channel:

$$m_t \longrightarrow C_t \longrightarrow \mathrm{DR}_t.$$

All other effects are absent or negligible. Route A is the appropriate reduction when the monitoring technology is purely extractive (client-generated ERP logs provided to the auditor for offline analysis, for instance, with no change to the client's own exception-handling workflows), when the time horizon is short enough that behavioral adaptation has not occurred, or when institutional separation between the audit and the firm's control operations is maintained by design.

The continuous-auditing bridge paper (Author, 2026c) operates entirely within Route A. Its central result—the cumulative observability theorem—establishes that the static dimension inequality for residual recordability generalizes, under continuous auditing, into a time-indexed cumulative observability condition; but the theorem is silent on whether the firm's risk structure is itself altered by the monitoring architecture. The present paper develops the complementary Route B analysis: treating the monitoring architecture as an intervention on the firm's internal control dynamics rather than only as an evidence-generation technology.

3.2 Route B: Intervention-Channel Mechanism

Under Route B, continuous-audit technologies are interventions on the firm's control-and-behavior space. The monitoring architecture enters the internal control state transition equation through a gain operator; when the gain is nondegenerate, continued monitoring deployment alters the control state, which in turn alters the control-risk function. Further, the monitoring architecture alters the distribution of managerial policy inputs through a behavioral-response channel; when the response is nondegenerate, managerial policy choices—transaction structures, accounting-classification choices, valuation-model inputs—are redistributed. Control risk and, secondarily, inherent risk become functions of the monitoring history.

The defining formal property of Route B is that the monitoring architecture enters the audit-risk model through at least two channels:

$$
\begin{aligned}
&m_t \longrightarrow C_t \longrightarrow \mathrm{DR}_t && \text{(evidence effect, Route A channel)},\\
&m_t \longrightarrow k_t \longrightarrow \mathrm{CR}_t && \text{(control effect, Route B primary channel)},\\
&m_t \longrightarrow u_t \longrightarrow \mathrm{IR}_t && \text{(behavior effect, Route B secondary channel)}.
\end{aligned}
$$

When the second channel is active, CR is endogenous to the monitoring deployment. When the third channel is active, IR is additionally endogenous, though typically through a weaker mechanism.

Route B is the appropriate model when the monitoring technology is integrative (approval-path monitoring feeds back into exception-handling workflows and alters approver behavior), when the time horizon is long enough for behavioral and control-operation adaptation to unfold, and when the distinction between audit evidence and internal control erodes—a phenomenon Chan and Vasarhelyi (2011) have observed descriptively.

Figure 1 summarizes the two mechanisms in a single diagram.

3.3 The Central Question

The Route A/Route B distinction reframes the quantitative analysis of continuous auditing. Instead of asking "how much does continuous auditing reduce DR?" the substantive question becomes:

> *Under what conditions does continuous auditing reduce detection risk only (Route A), and under what conditions does continuous auditing transform control risk itself (Route B)?*

The answer to this question is not a qualitative dichotomy. Most monitoring technologies operate through both routes simultaneously, with the relative strength of the two routes depending on the technology, the residual direction it addresses, and the institutional context. The analytical task is to identify the conditions under which Route B is nondegenerate and to provide a decomposition

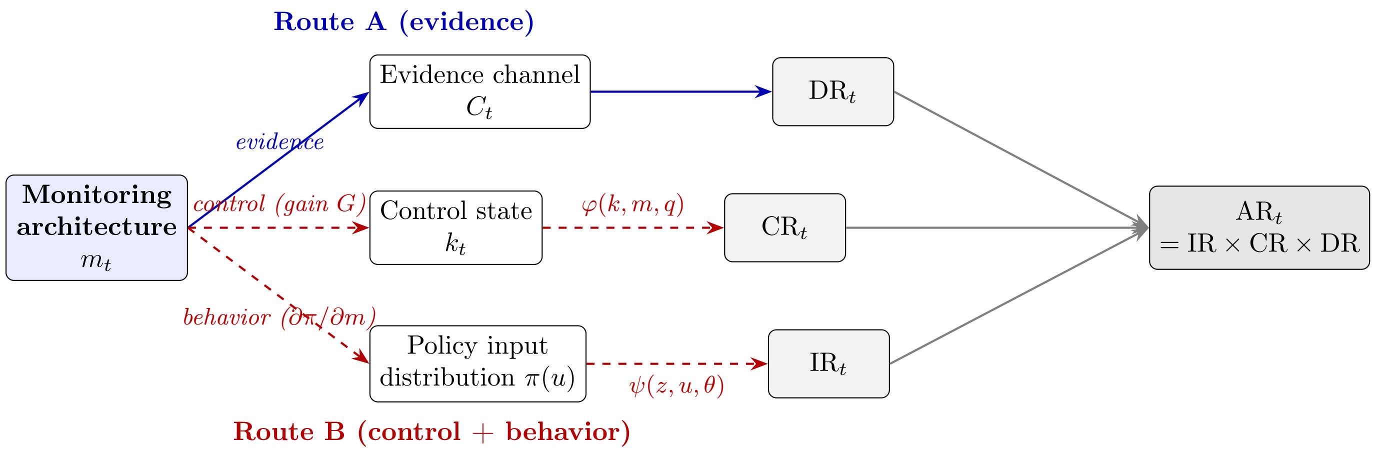

Figure 1: The three-channel architecture of continuous auditing. Route A (solid blue) is the evidence-channel mechanism formalized in Author 2026c: monitoring expands the residual-focused observability Gramian via the evidence channel C_t, reducing DR. Route B (dashed red) is the intervention-channel mechanism developed in this paper: the same monitoring architecture enters the control-state dynamics through the gain operator G (altering CR via φ) and shifts the managerial policy-input distribution π (altering IR via ψ). Under Route A alone, the Route A/Route B diagram reduces to the top path only; when Route B is active, all three paths operate, and the three-effect decomposition of Section VII attributes the technology's total effect to the appropriate risk component without double counting.

procedure that avoids double counting when both routes are active.

Sections IV through VIII develop the formal apparatus. Section IV introduces the state-space model. Section V formalizes CR endogenization and the control-effect condition. Section VI addresses IR endogenization with appropriate caution. Section VII develops the three-effect decomposition. Section VIII presents the control-assurance coupling theorem as the main result.

IV. A State-Space Model of Control-Assurance Feedback

This section introduces the formal model. The model employs a residual-relevant state z_t in the tradition of observability-based audit-planning formulations and embeds it in a richer architecture in which an internal control state k_t and a monitoring architecture m_t appear as first-class objects.

4.1 The Three-Level State Space

Let the firm–auditor interface be described by three state components at each time t:

- $z_t \in \mathbb{R}^d$, the residual-relevant accounting state (the formal counterpart of the third-entry record space, as in Author, 2026a,c);
- $k_t \in \mathbb{R}^\ell$, the internal control state, representing the design and operational status of the firm's internal controls relevant to the residual directions; and
- $m_t \in \mathbb{R}^q$, the monitoring architecture, representing the auditor's (and/or the firm's) continuous-

audit technology deployment at time t.

The accounting state evolves according to

$$z_{t+1} = A(k_t)\, z_t + B_R\, u_t + \xi_t, \tag{1}$$

where $A(k_t)$ is a k_t-dependent residual-dynamics matrix (strong internal controls damp residual persistence; weak controls amplify it), $u_t \in \mathbb{R}^p$ is the managerial policy-input vector (reclassification choices, hedge designation, impairment triggers, valuation-model inputs), B_R is the policy-response matrix, and $\xi_t \sim \mathcal{N}(0, Q_R)$ is the residual-state disturbance.

The control state evolves according to

$$k_{t+1} = F\, k_t + G\, m_t + \zeta_t, \tag{2}$$

where F captures the natural persistence of control design and operation, G is the monitoring-to-control gain matrix encoding how the monitoring architecture influences the control state, and $\zeta_t \sim \mathcal{N}(0, Q_K)$ is the control-state disturbance.

The auditor's observation is

$$y_t = C_t(m_t)\, z_t + \eta_t, \tag{3}$$

with C_t the monitoring-dependent observation matrix (as in Author, 2026c) and $\eta_t \sim \mathcal{N}(0, R_t)$.

The managerial policy input is drawn from a distribution that depends on the monitoring history:

$$u_t \sim \pi(\cdot \mid m_{0:t}, z_t, k_t). \tag{4}$$

Equations (1)–(4) are the backbone of the control-assurance feedback architecture.

4.2 The Three Routes as Structural Features of the Model

The Route A/Route B distinction has a clean structural reading in this model. Consider how the monitoring architecture m_t reaches each audit-risk component.

- *Route A (evidence effect):* m_t enters C_t in (3) and thereby enters W_T^R, the residual-focused observability Gramian; this reduces DR_t through the mechanism formalized in observability-based audit-planning analysis (Author, 2026c).
- *Route B primary channel (control effect):* m_t enters k_{t+1} through the gain operator G in (2); when G is nondegenerate, sustained monitoring deployment alters the control-state trajectory, and $\mathrm{CR}_t = \varphi(k_t, m_t, q_t)$—to be developed in Section V—shifts accordingly.
- *Route B secondary channel (behavior effect):* m_t enters the policy-input distribution π in (4); when π is nondegenerate in m_t, the distribution of u_t changes, and IR_t—to be developed in Section VI—shifts accordingly.

Route A corresponds to $G = 0$ and $\partial\pi/\partial m = 0$: the monitoring architecture appears only in the observation equation. Route B corresponds to $G \neq 0$ and/or $\partial\pi/\partial m \neq 0$: the monitoring architecture additionally appears in the control-state and/or policy-input equations.

4.3 Assumptions

Assumption 4.1 (Linear-Gaussian structure)**.** Equations (1), (2), (3) are linear, and (ξ_t, ζ_t, η_t) are jointly Gaussian with mutually independent components. The pair $(A(k), B_R)$ is stabilizable for every k in the relevant neighborhood of the operating point; the pair (F, G) is stabilizable; and the pair $(A(k), C_t)$ is detectable for every k for which the monitoring sequence is designed.

Assumption 4.2 (Control-state observability by the auditor)**.** The auditor observes a sufficient statistic for k_t either through the firm's internal-control reporting (ISA 315 risk-assessment process, SOX 302/404 certification and attestation) or through direct signals embedded in y_t. This assumption allows $\mathrm{CR}_t = \varphi(k_t, m_t, q_t)$ to be evaluated at each t.

Assumption 4.3 (Boundedness of the policy-input response)**.** The policy-input distribution $\pi(\cdot \mid m_{0:t}, z_t, k_t)$ has bounded first and second moments as a function of $m_{0:t}$, and is Lipschitz continuous in m with a bound that decreases over time (short-horizon adaptation is small, long-horizon adaptation is bounded).

Assumption 4.1 is standard for this class of linear-Gaussian state-space models. Assumption 4.2 is the analogue on the control side of the standard assumption that $(\mathrm{IR}, \mathrm{CR})$ are evaluable; it is a feasibility condition on the auditor's capacity to observe the firm's control state. Assumption 4.3 is a regularity condition on the behavioral response channel that keeps the IR endogenization bounded.

4.4 The Monitoring Architecture: Dimensions

The monitoring architecture m_t is a vector whose components index the technologies deployed. A non-exhaustive taxonomy includes:

- ERP log extraction (client-generated raw logs, third-party-generated independent logs);
- Process mining (conformance checking, deviation detection, bottleneck analysis);
- Approval-path monitoring (real-time deviation flags, audit-trail reconstruction);
- Confirmation APIs (automated confirmation exchanges with counterparties, banks, custodians);
- Transaction-network analytics (loop detection, centrality analysis, community detection);
- Spectral and phase-lag monitoring (frequency-domain anomaly detection for periodic patterns);
- Model-input logs for fair-value measurements (parameter-change histories, assumption logs);
- AI-based anomaly detection (supervised and unsupervised flagging of exceptions).

Each entry of m_t has a corresponding entry in the gain operator G, reflecting that technology's effect on the control-state dynamics. Some technologies have large entries in G (approval-path monitoring directly alters approver behavior); others have near-zero entries (raw ERP log extraction that is analyzed offline by the auditor has minimal effect on the firm's control operations). The technology-by-technology partition of m_t into Route A–dominant and Route B–active subsets is the content of the worked illustrations in Sections IX and X.

V. Endogenizing Control Risk

This section develops the core of the Route B argument on the control side. I formalize the dependence of control risk on the control state, the monitoring architecture, and the reliability of monitoring-generated evidence; identify a control-effect condition under which continuous auditing endogenizes control risk; and relate the resulting construct to the CR of auditing standards.

5.1 Control Risk as a Function of the Control State

Let the control-risk function be

$$\mathrm{CR}_t = \varphi(k_t, m_t, q_t), \tag{5}$$

where k_t is the internal-control state (as in Section IV), m_t is the current monitoring deployment, and $q_t \in (0,1]^q$ is the vector of reliability weights on monitoring channels (capturing source reliability by channel, in the spirit of reliability-adjusted Gramian formulations).

The dependence of CR on k_t is intuitive: stronger internal controls, better designed and more effectively operated, yield a lower control risk. The dependence on m_t, however, is more subtle, and it is the analytic content of Route B.

Monitoring can affect CR through three sub-channels:

Channel V.1 (direct operational channel):

m_t alters k_t through (2); the monitoring architecture is integrated into exception-handling workflows, approval-path enforcement, process-conformance feedback loops. Examples: real-time exception routing, process-mining-based bottleneck intervention, AI-flag-triggered control re-examination.

Channel V.2 (behavioral deterrence channel):

m_t raises the detection probability of control deviations, which reduces the agent's incentive to deviate. This channel affects CR through the distribution of deviation realizations even when k_t is held fixed. Formally, this is a behavioral component embedded in φ rather than in the k_t dynamics.

Channel V.3 (reliability channel):

m_t has a component q_t that indexes the reliability of monitoring-generated evidence; independent third-party logging, tamper-evident log formats, and audit-firm-controlled monitoring infrastructure raise q_t, which enters φ directly.

Channels V.1 and V.2 operate through the content of monitoring, while Channel V.3 operates through the independence and integrity of monitoring. The three sub-channels are conceptually distinct; empirically, a given monitoring technology may activate more than one.

5.2 The Control-Effect Condition

The key structural feature that distinguishes Route B from Route A is whether the monitoring architecture enters CR beyond its effect through the Route A observation channel. I formalize this as follows.

Definition 5.1 (Control-effect condition)**.** A monitoring architecture m_t satisfies the *control-effect condition* at time t if at least one of the following holds:

(a) *Direct operational channel:* The gain operator G in (2) has $\|Gm_t\| > 0$ along a direction that is mapped nondegenerately into CR_t by φ.

(b) *Behavioral deterrence channel:* $\partial\varphi/\partial m_t\big|_{k_t \text{ fixed}}$ is nondegenerate, reflecting the deterrence effect on deviation realizations.

(c) *Reliability channel:* $\partial\varphi/\partial q_t$ is nondegenerate, and the monitoring deployment shifts q_t.

A monitoring architecture that satisfies none of (a)–(c) is *Route A–pure*: it affects DR through the observation channel only.

Definition 5.1 is the formal content of the question "does continuous auditing endogenize CR?" Under the definition, the answer is yes if and only if at least one of the three sub-channels is active. The answer can vary across technologies and across residual directions; the definition is technology-specific and direction-specific. This definition identifies broad Route B relevance; the sharper distinction between functional activation and operational activation is developed in Section VIII.

5.3 The Cumulative Version: Monitoring-to-Control Influence

The control-effect condition is pointwise in time. Because continuous auditing is a path phenomenon, the pointwise condition needs to be aggregated over the horizon.

Definition 5.2 (Cumulative monitoring-to-control influence)**.** Define the cumulative monitoring-to-control influence Gramian[2] over the horizon $[0, T-1]$ as

$$W_T^K = \sum_{t=0}^{T-1} F^t \, G \, S_t \, G^\top \, (F^\top)^t, \tag{6}$$

[2] In control-theoretic terms, W_T^K is a discrete-time reachability (controllability) Gramian: the gain operator G enters the control-state dynamics (2) as an input operator, and W_T^K measures, over the horizon, the cumulative range of control-state directions that the monitoring architecture can populate.

where S_t is a materiality-like weighting on monitoring directions. The *cumulative monitoring-to-control influence condition* holds if

$$\lambda_{\min}(W_T^K) \geq \tau_K, \tag{7}$$

where $\tau_K > 0$ is a control-structural-coverage threshold analogous to τ_R on the residual side.

Equation (6) is the control-side analogue, for the reachability channel, of an observability-type Gramian on the evidence channel. For each $v \in \mathbb{R}^\ell$,

$$v^\top W_T^K v = \sum_{t=0}^{T-1} \|S_t^{1/2} G^\top (F^\top)^t v\|^2,$$

which measures how effectively the monitoring architecture, through direct action at time 0 and through F-dynamics at later times, populates the direction v of the control-state space.

5.4 The CR-Relevant Subspace and the Lipschitz Condition

To state the implication of cumulative monitoring-to-control influence precisely, I introduce two auxiliary constructs that localize the condition to the directions of the control state that actually influence control risk.

Definition 5.3 (CR-relevant subspace)**.** Let $\mathcal{K}_{\mathrm{CR}} \subseteq \mathbb{R}^\ell$ denote the *CR-relevant subspace*: the smallest linear subspace of $\mathbb{R}^\ell$ such that $\varphi(k, m, q) = \varphi(\Pi_{\mathcal{K}_{\mathrm{CR}}} k, m, q)$ for all (k, m, q) in the operating region, where $\Pi_{\mathcal{K}_{\mathrm{CR}}}$ denotes the orthogonal projection onto $\mathcal{K}_{\mathrm{CR}}$. Its dimension is denoted $d_K := \dim \mathcal{K}_{\mathrm{CR}}$. Minimality of $\mathcal{K}_{\mathrm{CR}}$ is understood in the functional sense: no proper subspace of $\mathcal{K}_{\mathrm{CR}}$ is sufficient to represent all variations in φ over the operating region. Equivalently, for every direction $v \in \mathcal{K}_{\mathrm{CR}}$ with $\|v\| = 1$, there exist states k, k' in the operating region with $k - k' \propto v$ such that $\varphi(k, m, q) \neq \varphi(k', m, q)$ for some (m, q).

The CR-relevant subspace isolates the directions of k that enter φ; directions orthogonal to $\mathcal{K}_{\mathrm{CR}}$ are invariances of the control-risk function and can be ignored for Route B analysis.

Definition 5.4 (Instantaneously CR-active monitoring direction)**.** A monitoring architecture m is *instantaneously CR-active* if

$$\Pi_{\mathcal{K}_{\mathrm{CR}}} G\, m \neq 0, \tag{8}$$

i.e., the projection of the monitoring-induced control-state shift onto the CR-relevant subspace is non-zero *at the current time step.* This is a pointwise condition. The horizon-T analogue—that the monitoring architecture can reach the CR-relevant state either immediately or through F-dynamics over the horizon—is the *monitoring-to-control reachability* condition of Definition 8.1, which is generally weaker and is the condition used in the main theorems.

With these constructs, the Lipschitz condition on φ can be stated precisely.

Assumption 5.5 (Lipschitz lower bound on φ)**.** There exist constants $\gamma_K > 0$ and a compact operating region $\mathcal{O} \subset \mathbb{R}^\ell \times \mathbb{R}^q \times (0,1]^q$ such that for all $(k,m,q),(k',m,q) \in \mathcal{O}$,

$$|\varphi(k,m,q) - \varphi(k',m,q)| \geq \gamma_K \cdot \|\Pi_{\mathcal{K}_{\mathrm{CR}}}(k-k')\|. \tag{9}$$

Equivalently, φ restricted to $\mathcal{K}_{\mathrm{CR}}$ has a lower Lipschitz constant γ_K in its first argument, holding (m,q) fixed.

Assumption 5.5 strengthens the smoothness hypothesis of Assumption 4.1 with a *bilateral* Lipschitz bound: φ is not only smooth but uniformly sensitive along CR-relevant directions. The bilateral bound enters Theorem 8.8 as the sensitivity condition required for certifying the uniform operational lower bound: without it, φ could be smooth but flat in some CR-relevant directions, which would preserve functional activation (Theorem 8.7) but prevent certification of the uniform quantitative guarantee.

Let $\Pi_K := \Pi_{\mathcal{K}_{\mathrm{CR}}}$. Define the *effective cumulative monitoring-to-control influence Gramian*

$$W_T^{K,\mathrm{eff}} := \Pi_K W_T^K \Pi_K, \tag{10}$$

restricted to the CR-relevant subspace. Its minimum eigenvalue on $\mathcal{K}_{\mathrm{CR}}$ is what enters the main bound.

Proposition 5.6 (Implication of cumulative monitoring-to-control influence)**.** *Under Assumption 5.5 and if $\lambda_{\min}^{\mathcal{K}_{\mathrm{CR}}}(W_T^{K,\mathrm{eff}}) \geq \tau_K > 0$, there exists a monitoring sequence $m_{0:T-1} \in \mathbb{R}^{q\times T}$ with $\|m_{0:T-1}\|_2 \leq 1$ (where $\|m_{0:T-1}\|_2^2 := \sum_{t=0}^{T-1} \|m_t\|^2$) such that*

$$|\mathrm{CR}_T(m_{0:T-1}) - \mathrm{CR}_T(0)| \geq \gamma_K \sqrt{\tau_K} \cdot \|m_{0:T-1}\|_2. \tag{11}$$

Setting $\beta_K := \gamma_K\sqrt{\tau_K}$ yields the compact form: the CR *shift is bounded below by β_K at unit input energy.*

The proof is in Appendix A. The bound (11) makes explicit that the CR shift is driven by two quantitative factors: the Lipschitz sensitivity γ_K (Assumption 5.5) and the structural-coverage threshold τ_K (Definition 5.2). When either degenerates—$\gamma_K \to 0$ or $\tau_K \to 0$—the bound collapses, which foreshadows the role of conditions (b) and (c) in Theorems 8.7 and 8.8.

Remark 5.7 (A persistence-aware variant)*.* The compact bound (11) is stated in the input-energy norm $\|m_{0:T-1}\|_2$ because this is the norm that the matched-filter construction of Lemma 8.6 naturally produces. An alternative, persistence-aware variant can be obtained by introducing the CR-relevant norm $\|m_{0:T-1}\|_{\mathcal{K}_{\mathrm{CR}}}^2 := \sum_t \|\Pi_K F^{T-1-t} G m_t\|^2$ and the persistence constant $c_F := \inf_{\|v\|=1, v\in\mathcal{K}_{\mathrm{CR}}} \sum_{t=0}^{T-1} \|F^t v\|^2$; then $\|\Pi_K \Delta k_T\| \geq \sqrt{c_F \tau_K} \cdot \|m_{0:T-1}\|_{\mathcal{K}_{\mathrm{CR}}}$ and the CR shift is bounded below by $\gamma_K\sqrt{c_F\tau_K} \cdot \|m_{0:T-1}\|_{\mathcal{K}_{\mathrm{CR}}}$. This variant makes the F-persistence contribution explicit but uses a norm that is less natural for practitioners; the input-energy form (11) is used throughout the main text.

5.5 Relation to CR in Auditing Standards

Two caveats are important. First, CR in auditing standards is the probability that a material misstatement, occurring in an assertion that could be material, will not be prevented or detected and corrected on a timely basis by the entity's internal control (ISA 200, ISA 315, AU-C 200). The control-risk function φ in (5) is a model-implied proxy for this standard-level concept, constructed within the state-space framework; it is not a replacement for the standard definition.

Second, the Route A reduction (in which CR is exogenous) is not "wrong"; it is a valid reduction when the control-effect condition fails. When the monitoring technology is Route A–pure (negation of Definition 5.1), CR_t is orthogonal to m_t and the assessed-risk approach of ISA 315 gives the correct CR input to audit-risk aggregation. Route B's content is that, for a growing class of continuous-audit technologies, the control-effect condition is satisfied, and CR must accordingly be treated as endogenous to m_t.

The practical implication—that auditors must attribute the observed effect of a continuous-audit technology to the correct risk component rather than double-counting it across components—is the subject of Section VII.

VI. Endogenizing Inherent Risk, Cautiously

This section addresses the secondary channel of Route B: the effect of continuous auditing on inherent risk through behavioral adaptation in managerial policy inputs. The treatment is deliberately more restrained than that of control risk.

6.1 Why IR Endogenization Requires Caution

Inherent risk derives from the business model, the transaction structure, the industry environment, the complexity of accounting judgments, and fraud incentives and opportunities. Most of these determinants are outside the reach of the monitoring architecture. A continuous-audit technology does not change the industry a firm operates in, the complexity of its transactions, or the environmental volatility of its markets.

What monitoring can change is the subset of policy inputs within the manager's control that respond to detection-probability changes. Two examples anchor the discussion.

Example 6.1 (Circular-transaction design). A firm that has previously engaged in circular transactions through a specific counterparty-network topology observes the deployment of transaction-network loop detection. The detection probability for loops of the previous topology rises. The firm's policy-input distribution—whether to engage in circular transactions, with which counterparties, at what frequency, and in what amounts—shifts in response. This is a redistribution of u_t, not a change in the underlying business complexity that generates IR.

Example 6.2 (Valuation-input manipulation). A firm that relies on Level 3 fair-value measurements observes the deployment of model-input log monitoring. The monitoring raises the detection probability of unjustified parameter shifts, ad-hoc model choices, and selective use of supporting evidence.

The firm's policy-input distribution—which inputs to use, when to change them, how to document them—shifts in response.

In both examples, the monitoring architecture alters the policy-input distribution without altering the underlying inherent risk of the firm's business. The question is whether such policy-input redistribution is formally classified as an IR effect (reducing the probability of misstatement) or as a DR/CR effect (affecting how the misstatement, once generated, is observed or prevented).

6.2 The Behavioral-Response Channel

Formally, the behavioral-response channel operates through the policy-input distribution $\pi(\cdot \mid m_{0:t}, z_t, k_t)$. When $\partial\pi/\partial m \neq 0$, the distribution of u_t depends on the monitoring history. Because IR captures the probability of misstatement in the absence of internal controls, and because the misstatement-generating policy inputs are drawn from π, any shift in π that reduces the support or frequency of misstatement-generating inputs is an IR effect.

Definition 6.3 (Behavior-effect condition)**.** A monitoring architecture m_t satisfies the *behavior-effect condition* at time t if the policy-input distribution π is nondegenerate in m: there exists a measurable set A such that

$$\Pr[u_t \in A \mid m_{0:t} = m'] \neq \Pr[u_t \in A \mid m_{0:t} = m''] \quad \text{for some } m' \neq m''.$$

When Definition 6.3 holds, IR becomes a function of the monitoring history:

$$\mathrm{IR}_t = \psi(z_t, u_t, \theta_t) \quad \text{with} \quad u_t \sim \pi(\cdot \mid m_{0:t}, z_t, k_t), \tag{12}$$

where θ_t collects the exogenous (monitoring-independent) determinants of inherent risk.

6.3 Cautious Formalization

Three disciplines constrain the treatment of IR endogenization.

First, the behavior-effect condition is typically satisfied through a narrow channel: policy inputs that are within the manager's discretion and that respond to detection-probability changes. Most components of IR—industry complexity, business volatility, related-party exposure, fraud incentives from outside the firm—are not in this channel.

Second, the behavior-effect channel operates at a longer time scale than the control-effect channel. Approval-path monitoring changes approver behavior within weeks; circular-transaction design adaptation unfolds over multiple audit cycles. Accordingly, short-horizon audit-risk assessments may need to treat the behavior effect as negligible, while long-horizon assessments (e.g., multi-year audit planning, industry-level research on continuous-auditing adoption) should incorporate it.

Third, strategic adaptation on the manager's side may attenuate or reverse the behavior effect. If circular-transaction designers adapt to transaction-network loop detection by designing transactions that evade detection, the induced IR reduction is partially or fully offset. A full game-theoretic treatment of this strategic adaptation is beyond the scope of the present paper and is identified as future work.

6.4 Relation to the Three-Effect Decomposition

The behavior-effect channel is the third of the three routes through which the monitoring architecture reaches audit risk (recall the three-route taxonomy of Section III):

$$\begin{aligned} \text{Evidence effect:} \quad & m_t \to C_t \to \text{DR}_t, \\ \text{Control effect:} \quad & m_t \to k_t \to \text{CR}_t, \\ \text{Behavior effect:} \quad & m_t \to u_t \to \text{IR}_t. \end{aligned}$$

The next section formalizes the decomposition and shows that failing to distinguish the three effects leads to systematic double counting in audit-risk aggregation.

VII. Avoiding Double Counting in Audit Risk Assessment

This section addresses the central analytical problem that arises when Route B is active. When a monitoring technology operates through multiple channels simultaneously, naively attributing its total effect to any single audit-risk component overstates the aggregate risk reduction. The three-effect decomposition developed here is the formal apparatus for attributing effects correctly.

7.1 The Double-Counting Problem

Consider a monitoring technology m that operates through both Route A (evidence effect on DR) and Route B primary channel (control effect on CR). Suppose that the technology reduces DR by a factor of $d_A < 1$ through the evidence channel and reduces CR by a factor of $d_B < 1$ through the control channel. If the auditor attributes the full reduction to DR, the revised audit risk is

$$\text{AR}_{\text{naive-DR}} = \text{IR} \times \text{CR} \times (d_A \cdot d_B \cdot \text{DR}_{\text{original}}). \tag{13}$$

If the auditor attributes the full reduction to CR, the revised audit risk is

$$\text{AR}_{\text{naive-CR}} = \text{IR} \times (d_A \cdot d_B \cdot \text{CR}_{\text{original}}) \times \text{DR}_{\text{original}}. \tag{14}$$

Both (13) and (14) are incorrect. The correct attribution is

$$\text{AR}_{\text{correct}} = \text{IR} \times (d_B \cdot \text{CR}_{\text{original}}) \times (d_A \cdot \text{DR}_{\text{original}}) = (d_A \cdot d_B) \cdot \text{AR}_{\text{original}}. \tag{15}$$

This is numerically identical to (13) and (14)—the product $d_A \cdot d_B$ enters the same way—but only because each effect was counted once. If the auditor counts the evidence effect in DR and the control effect in CR *and* also reapplies the control-effect reduction to DR (on the mistaken belief that "continuous auditing reduces DR through the control channel too"), the resulting AR is

$$\mathrm{AR}_{\text{double-counted}} = \mathrm{IR} \times (d_B \cdot \mathrm{CR}_{\text{original}}) \times (d_A \cdot d_B \cdot \mathrm{DR}_{\text{original}}) = d_B \cdot \mathrm{AR}_{\text{correct}}. \tag{16}$$

Equation (16) understates the actual audit risk by a factor of d_B, which can be substantial (e.g., $d_B = 0.7$ corresponds to a 30% understatement).

This is the substantive double-counting problem. It arises not because the multiplicative structure is wrong, but because the attribution of a single technology's effect to the correct audit-risk component requires discipline.

7.2 The Three-Effect Decomposition

Let the monitoring architecture at time t be $m_t \in \mathbb{R}^q$. For each monitoring technology indexed by $k = 1, \ldots, q$, define the marginal effect of that technology on the audit-risk components through each of the three channels.

Definition 7.1 (Three-effect decomposition)**.** The effect of a change in monitoring technology k at time t on audit risk decomposes as

$$\frac{\partial \log \mathrm{AR}_t}{\partial m_{t,k}} = \underbrace{\left(\frac{\partial \log \mathrm{IR}_t}{\partial m_{t,k}}\right)_{\text{behavior}}}_{\text{behavior effect}} + \underbrace{\left(\frac{\partial \log \mathrm{CR}_t}{\partial m_{t,k}}\right)_{\text{control}}}_{\text{control effect}} + \underbrace{\left(\frac{\partial \log \mathrm{DR}_t}{\partial m_{t,k}}\right)_{\text{evidence}}}_{\text{evidence effect}} \tag{17}$$

where each partial derivative is computed with the other two transmission mechanisms held at their no-monitoring baselines.

Each effect is the partial derivative of the corresponding audit-risk component with respect to $m_{t,k}$, holding the other two transmission mechanisms at their no-monitoring baselines. The decomposition is counterfactual: each effect isolates the impact of one channel by shutting down the other two.

Proposition 7.2 (Additive-in-logs decomposition)**.** *Under Assumptions 4.1–4.3 and the regularity condition that* IR_t, CR_t, DR_t *are twice continuously differentiable in* $m_{t,k}$, *the total marginal effect of* $m_{t,k}$ *on* $\log \mathrm{AR}_t$ *equals the sum of the three effects as defined in* (17), *plus a second-order cross-effect term of order* $O(\|m\|^2)$.

The proof is in Appendix A.

Proposition 7.2 validates the first-order attribution procedure: for small changes in monitoring deployment, the marginal effect on AR can be decomposed into three additive pieces. The exact identity $\partial \log \mathrm{AR}/\partial m = \partial \log \mathrm{IR}/\partial m + \partial \log \mathrm{CR}/\partial m + \partial \log \mathrm{DR}/\partial m$ follows from the multiplicative structure alone and requires no linearization. The $O(\|m\|^2)$ error term arises when each component's total derivative is replaced by its channel-restricted partial derivative (Definition 7.1); the cross-effects—for instance, the interaction between the control effect (which shifts k_t) and the evidence effect (which shifts C_t) when k_t enters the observation matrix—are second-order and can be incorporated through a second-step refinement if needed. Appendix A provides an explicit decomposition of the cross-effect term into three pairwise components.

7.3 Attribution Rule

The decomposition of Proposition 7.2 yields a simple attribution rule for audit-risk assessment under Route B. The rule is an *attribution convention* rather than a universal identity: the decomposition is unique once the convention is fixed (specifically, the channel-restricted partial-derivative convention of Definition 7.1), but different conventions—for instance, a Shapley-style attribution that symmetrically splits cross-effects—could yield different but equally consistent decompositions. The discipline imposed by the rule is that, whichever convention is used, it must be applied uniformly across all continuous-audit technologies in the engagement.

Attribution Rule 7.3 (Effect attribution under the channel-restricted convention)**.** Under the channel-restricted partial-derivative convention (Definition 7.1), the effect of monitoring technology $m_{t,k}$ on audit-risk component $X \in \{\mathrm{IR}, \mathrm{CR}, \mathrm{DR}\}$ is attributed to X if and only if the corresponding channel (behavior, control, evidence) is active for technology k. The effect is attributed to at most one component; each component receives the effect of each technology at most once. Cross-effects of order $O(\|m\|^2)$ are treated separately through a second-step refinement if material.

The rule prevents the double-counting configuration described in (16). It does not eliminate the need for judgment—the auditor must still assess whether each channel is active for each technology, and must explicitly declare the attribution convention used—but it fixes the translation from channel-activation judgment into audit-risk aggregation.

7.4 Implications for Audit Workpapers and Standards

Two implications follow for audit-engagement documentation.

First, the audit workpaper should record, for each continuous-audit technology deployed, which of the three channels (evidence, control, behavior) the technology is judged to activate, with supporting reasoning. ISA 230 and AS 1215 documentation requirements are compatible with this discipline; the formalization simply makes the attribution explicit.

Second, auditing standards on the effect of continuous auditing on risk assessment (ISA 315 revised; AS 2110) could usefully distinguish monitoring technologies by their channel activation pattern. Current standards treat continuous auditing primarily as an evidence-generation activity; the three-effect decomposition suggests a basis for future guidance that would distinguish the three channels and encourage separate attribution for each. Section XI develops this implication in more detail.

VIII. Main Result: The Control-Assurance Coupling Theorem

This section presents the paper's central theorem. The control-assurance coupling theorem provides the analytical keystone for the Route B argument, establishing the conditions under which continuous auditing endogenizes control risk and, through the behavior-effect channel, inherent risk. Structurally, it plays a role on the control side analogous to that of cumulative-observability theorems on the evidence side in existing audit-planning frameworks.

8.1 The Three Conditions, Formally Stated

Before stating the theorem, I collect the three conditions that appear in its antecedent and give each a precise mathematical formulation. In what follows, Π_K denotes the orthogonal projection onto the CR-relevant subspace $\mathcal{K}_{\text{CR}}$ (Definition 5.3), and $W_T^{K,\text{eff}}$ denotes the effective cumulative monitoring-to-control influence Gramian (10).

Definition 8.1 (Condition (a): Monitoring-to-control reachability)**.** The pair (F, G) of the control-state dynamics (2) is $\mathcal{K}_{\text{CR}}$*-reachable* over horizon T: the horizon-T reachability matrix along the CR-relevant subspace,

$$\mathcal{G}_{K,T} := \begin{bmatrix} \Pi_K G & \Pi_K F G & \Pi_K F^2 G & \cdots & \Pi_K F^{T-1} G \end{bmatrix}, \tag{18}$$

has rank at least 1; equivalently,

$$\operatorname{rank}(W_T^{K,\text{eff}}) \geq 1. \tag{19}$$

That is, there exists a monitoring sequence $m_{0:T-1}$ whose accumulated CR-relevant state shift $\Pi_K \Delta k_T$ is non-zero, where $\Delta k_T := \sum_{t=0}^{T-1} F^{T-1-t} G m_t$.

Remark 8.2 (Reachability versus instantaneous coupling). Definition 8.1 is strictly weaker than the instantaneous-coupling condition $\Pi_K G \neq 0$: even when G has no immediate projection onto $\mathcal{K}_{\text{CR}}$, the control-state dynamics F may rotate a monitoring-induced shift into the CR-relevant subspace at a later time $j \geq 1$, so that $\Pi_K F^j G \neq 0$ for some j. The reachability formulation captures both the immediate and the delayed channels through which the monitoring architecture can influence the CR-relevant control state, consistent with the fact that $W_T^{K,\text{eff}}$ is the cumulative Gramian of both pathways. Readers familiar with linear systems theory will recognize (19) as the discrete-time reachability condition for the pair (F, G) restricted to the output direction Π_K.

Definition 8.3 (Condition (b): Control-risk sensitivity)**.** Assumption 5.5 holds with constant $\gamma_K > 0$, i.e., φ has a lower Lipschitz constant γ_K along the CR-relevant subspace.

Definition 8.4 (Condition (c): Operational coverage threshold)**.** There exists a finite horizon $T \leq T^*$ such that

$$\lambda_{\min}^{\mathcal{K}_{\text{CR}}}(W_T^{K,\text{eff}}) \geq \tau_K > 0, \tag{20}$$

where $\lambda_{\min}^{\mathcal{K}_{\text{CR}}}(\cdot)$ denotes the smallest eigenvalue of the operator restricted to $\mathcal{K}_{\text{CR}}$, and $\tau_K > 0$ is a practitioner-specified coverage threshold reflecting the minimum cumulative coverage required for the CR shift to be operationally meaningful for audit-risk assessment.

Each of Definitions 8.1–8.4 is a formal mathematical condition on the model primitives. Condition (a) is a reachability condition on the pair (F, G) restricted to $\mathcal{K}_{\text{CR}}$: it states that the CR-relevant control state is *attainable* from the monitoring architecture, directly or through F-dynamics. Condition (b) is an analytic condition on the control-risk function φ: it states that a non-trivial state shift produces a non-trivial CR shift. Condition (c) is a quantitative condition on the cumulative coverage: it states that the magnitude of attainable state shifts exceeds an operational threshold. The two-stage structure of the main theorem (below) separates these into *functional activation*, determined by condition (a) alone, and *operational activation*, determined by conditions (b) and (c) as additional sufficient requirements for a uniform quantitative bound.

8.2 Two Supporting Lemmas

Before the main theorem, I establish two lemmas that isolate the implications of the individual conditions. Both proofs are in Appendix A.

Lemma 8.5 (Condition (a) is necessary for CR sensitivity)**.** *Suppose condition (a) fails, i.e.,* $\mathrm{rank}(\mathcal{G}_{K,T}) = 0$*, equivalently* $\Pi_K F^j G = 0$ *for all* $j = 0, 1, \ldots, T-1$*. Then for every monitoring sequence* $m_{0:T-1} \in \mathbb{R}^{q \times T}$*, the accumulated projection onto* $\mathcal{K}_{\mathrm{CR}}$ *satisfies* $\Pi_K \Delta k_T = 0$ *where* $\Delta k_T = \sum_{t=0}^{T-1} F^{T-1-t} G m_t$*, and consequently* $\partial \mathrm{CR}_T / \partial m_{0:T-1} = 0$ *along every direction.*

Lemma 8.6 (Conditions (a) and (c) generate CR-relevant state shifts)**.** *Suppose conditions (a) and (c) hold. Then there exists a monitoring sequence* $m_{0:T-1}$ *with* $\sum_t \|m_t\|^2 = \delta > 0$ *such that the accumulated CR-relevant state shift satisfies*

$$\|\Pi_K \Delta k_T\|^2 \geq \tau_K \cdot \delta. \tag{21}$$

Lemma 8.5 isolates the role of condition (a): without it—i.e., when neither immediate nor delayed monitoring-to-control channels reach $\mathcal{K}_{\mathrm{CR}}$—the monitoring architecture cannot shift the control state in the CR-relevant direction at all. Lemma 8.6 shows that conditions (a) and (c) together guarantee that, for a suitable monitoring sequence aligned with the minimum-eigenvalue direction of $W_T^{K,\mathrm{eff}}$, the accumulated CR-relevant state shift is bounded below by the monitoring magnitude scaled by the coverage threshold. Condition (b) is then needed to translate the state shift into a CR shift.

8.3 The Control-Assurance Coupling Theorems

The main result is stated in two stages. Theorem 8.7 characterizes when CR_T becomes *functionally* dependent on monitoring history—i.e., when Route B activates in the qualitative sense that CR_T is no longer invariant to the monitoring sequence. Theorem 8.8 then characterizes when the induced dependence is *operationally meaningful*—i.e., when the CR shift exceeds an audit-practice threshold. The two-stage structure separates a reachability property, which determines functional activation, from sensitivity-and-coverage conditions, which determine operational activation.

Theorem 8.7 (Functional Route B activation on the control side, within the linear state-space Route B model)**.** *Under Assumptions 4.1–4.3 and the minimality of* $\mathcal{K}_{\mathrm{CR}}$ *(Definition 5.3), the following two statements are equivalent for a given horizon* T*:*

(i) Functional Route B activation*: there exist $m'_{0:T-1}, m''_{0:T-1} \in \mathbb{R}^{q\times T}$ with $m' \neq m''$ such that* $\mathrm{CR}_T(m'_{0:T-1}) \neq \mathrm{CR}_T(m''_{0:T-1})$*; equivalently, the map $m_{0:T-1} \mapsto \mathrm{CR}_T(m_{0:T-1})$ is non-constant and the Route A reduction (which treats* CR *as functionally independent of $m_{0:T-1}$) is structurally incomplete as a representation of audit risk.*

(ii) Condition (a) holds at time T*: the pair (F, G) is $\mathcal{K}_{\mathrm{CR}}$-reachable (Definition 8.1).*

Theorem 8.8 (Operational Route B activation on the control side, within the linear state-space Route B model)**.** *Under the hypotheses of Theorem 8.7 and the additional conditions (b) (Definition 8.3: φ has a positive Lipschitz lower bound $\gamma_K > 0$ on the CR-relevant subspace) and (c) (Definition 8.4: $\lambda_{\min}^{\mathcal{K}_{\mathrm{CR}}}(W_T^{K,\mathrm{eff}}) \geq \tau_K$), the quantitative lower bound*

$$\sup_{\|m_{0:T-1}\|_2 \leq 1} \left|\mathrm{CR}_T(m_{0:T-1}) - \mathrm{CR}_T(0)\right| \geq \gamma_K \sqrt{\tau_K} \tag{22}$$

holds. Conversely, if condition (c) fails, the framework cannot certify uniform operational Route B activation at level τ_K: along the minimum-eigenvalue direction of $W_T^{K,\mathrm{eff}}$, no lower bound of the form $\gamma_K\sqrt{\tau_K}$ holds. Monitoring sequences aligned with better-covered directions may still generate CR *shifts exceeding this magnitude, but such deployments do not satisfy the uniform coverage condition and are not characterized by Theorem 8.8; they fall outside the theorem's scope and require direction-specific analysis rather than the uniform guarantee that condition (c) provides. A persistence-aware variant of the bound, which makes the contribution of the F-dynamics persistence constant c_F explicit, is given in Remark 5.7.*

Theorems 8.7 and 8.8 separate two notions of activation that the earlier, single-theorem formulation conflated, and they assign the three conditions (a), (b), (c) to distinct roles. *Condition (a) (reachability) alone is necessary and sufficient for functional activation*: if and only if (F, G) is $\mathcal{K}_{\mathrm{CR}}$-reachable, CR_T is a non-constant function of monitoring history. The minimality of $\mathcal{K}_{\mathrm{CR}}$ from Definition 5.3 plays the role here that Lipschitz sensitivity played in earlier drafts: by construction, φ *does* depend on every direction of $\mathcal{K}_{\mathrm{CR}}$, which is enough for functional non-constancy once the control state can be shifted along $\mathcal{K}_{\mathrm{CR}}$. *Conditions (b) and (c) are additional sufficient conditions for the quantitative bound* in Theorem 8.8: (b) supplies a uniform lower Lipschitz constant $\gamma_K > 0$, and (c) supplies the cumulative coverage threshold $\tau_K > 0$. Neither (b) nor (c) is claimed necessary for functional activation; (c) is necessary *for certifying the uniform operational bound at level τ_K* (by the converse in Theorem 8.8), while (b) enters through the magnitude γ_K of the bound rather than through its existence.

The two-stage structure also clarifies what happens in the absence of each condition. If condition (a) fails, neither functional nor operational Route B activation obtains. If condition (a) holds but (b) fails—meaning φ depends on $\mathcal{K}_{\mathrm{CR}}$ but has no uniform Lipschitz lower bound—then CR_T is functionally non-constant in $m_{0:T-1}$, but the quantitative lower bound of Theorem 8.8 does not hold at any $\gamma_K > 0$: the dependence can be arbitrarily weak locally. If (a) and (b) hold but (c) fails, Route B is functionally active and has a positive Lipschitz constant, but the cumulative shift is below the audit-practice coverage threshold: continuous auditing endogenizes CR but not by an operationally meaningful magnitude at the current horizon.

The proofs of both theorems are in Appendix A. The proof of Theorem 8.7 follows the structural pattern familiar from observability-type theorems in audit-planning analysis (cf. Author, 2026c): condition (a) is shown to be necessary by a direct construction of a counterexample when it fails, and sufficiency is established by combining the reachability of $\mathcal{K}_{\mathrm{CR}}$ with the minimality of $\mathcal{K}_{\mathrm{CR}}$ (Definition 5.3) to guarantee a non-zero CR change along at least one reachable direction. The proof of Theorem 8.8 uses Lemma 8.6 and Assumption 5.5 to translate the lower bound on $\|\Pi_K \Delta k_T\|^2$ into a lower bound on $|\mathrm{CR}_T(m_{0:T-1}) - \mathrm{CR}_T(0)|$.

For ease of reference in the remainder of the paper, when I write "the three control-assurance coupling conditions" or "the conditions of the theorem pair," I refer to (a), (b), (c) jointly; the functional/operational distinction is maintained when the context makes the difference material.

Remark 8.9 (Scope of the theorems). The two theorems differ in the nature of their characterizations. Theorem 8.7 provides a *necessary and sufficient* characterization of functional Route B activation within the maintained linear state-space model: CR_T is a non-constant function of monitoring history if and only if condition (a) holds. Theorem 8.8, by contrast, provides *sufficient* conditions for a uniform operational lower bound (namely, conditions (a), (b), (c) jointly imply the bound $\gamma_K\sqrt{\tau_K}$), together with a *converse in the limited sense* that failure of condition (c) prevents certification of the uniform operational bound at level τ_K—though CR shifts exceeding $\gamma_K\sqrt{\tau_K}$ may still arise along better-covered directions outside the theorem's scope. Outside the linear state-space framework—that is, under nonlinear control-state dynamics, threshold effects, indirect behavioral feedback not captured by the policy-input response operator, unmodeled institutional mechanisms, or substantial source-reliability shifts not reflected in q_t—additional channels may endogenize CR, and the conditions of both theorems are then sufficient but not necessary for Route B activation in the broader sense. Limitation 13.1 develops these qualifications, and Section XIII identifies the corresponding future-research directions.

8.4 Reading the Theorems as Design Criteria

Theorems 8.7 and 8.8 can be read as control-side design criteria, analogous in function to the design-criterion readings of observability-based theorems in audit-planning analysis. They say that whether continuous auditing endogenizes control risk is not a blanket property of continuous auditing as a technology class; it depends on three structural features of the particular technology deployment. A deployment that fails (a)—for instance, raw client-generated ERP logs analyzed offline by the auditor with no feedback into the firm's workflows, and with no F-dynamics through which the monitoring can reach the CR-relevant state—does not endogenize CR even functionally, however rich the evidence stream. A deployment that fails (b)—for instance, a firm whose control-risk function has no uniform Lipschitz lower bound on the CR-relevant subspace (perhaps because it is locally flat at the operating point)—may still generate functional dependence of CR on monitoring history (Theorem 8.7), but the framework cannot provide a uniform lower bound on the magnitude of the CR shift. A deployment that satisfies (a) and (b) but fails (c)—for instance, a short-horizon monitoring pilot that has not accumulated sufficient coverage—endogenizes CR functionally and has a positive Lipschitz constant, but the cumulative shift is below the audit-practice coverage threshold along at least one CR-relevant direction.

The two-stage characterization accordingly provides a structured procedure for auditors, control owners, and audit standard setters to classify a deployment along two axes: first, whether Route B is functionally activated, which is resolved by monitoring-to-control reachability (condition (a)); and second, whether the resulting CR shift is operationally meaningful under uniform coverage, which requires control-risk sensitivity (condition (b)) and cumulative monitoring-to-control influence (condition (c)).

8.5 The Secondary Result: IR Endogenization Condition

A parallel condition for IR endogenization through the behavior-effect channel can be stated but with weaker force, consistent with Section VI's caution. Mirroring the two-stage structure of Theorems 8.7 and 8.8, the behavior-effect characterization is stated in two parts: functional activation and operational activation.

Corollary 8.10 (Functional behavior-effect activation)**.** *Under Assumptions 4.1–4.3, the behavior-effect channel of Section VI is* functionally *active at time T, in the sense that $\partial \mathrm{IR}_T/\partial m_{0:T-1} \neq 0$, if and only if condition (a′) holds:*

(a′) Policy-response coupling: *$\pi(u \,|\, m_{0:t}, z_t, k_t)$ is nondegenerate in m—that is, the monitoring architecture reaches the IR-relevant policy-input distribution, directly at the current time step or through the intertemporal structure of π over the horizon.*

Corollary 8.11 (Operational behavior-effect activation)**.** *Under the hypotheses of Corollary 8.10 and the additional conditions (b′) and (c′):*

(b′) IR-policy sensitivity: *ψ has a Lipschitz lower bound $\gamma_U > 0$ in u on the IR-relevant subspace $\mathcal{U}_{\mathrm{IR}}$;*

(c′) Cumulative behavioral influence: *the accumulated marginal effect of m on the support of π, measured by the Fisher-information-like Gramian Σ_T^U restricted to $\mathcal{U}_{\mathrm{IR}}$, satisfies $\lambda_{\min}^{\mathcal{U}_{\mathrm{IR}}}(\Pi_U \Sigma_T^U \Pi_U) \geq \tau_U > 0$;*

the behavior-effect channel admits an operational magnitude bound

$$\sup_{\|m_{0:T-1}\|_2 \leq 1} \left| \mathrm{IR}_T(m_{0:T-1}) - \mathrm{IR}_T(0) \right| \geq \gamma_U \sqrt{\tau_U}, \tag{23}$$

analogous to the bound of Theorem 8.8 on the control side.

The two-stage structure of Corollaries 8.10 and 8.11 parallels the two-stage structure of Theorems 8.7 and 8.8: (a′) is necessary and sufficient for functional behavior-effect activation, while (b′)

and (c′) are additional sufficient conditions for the operational magnitude bound. Both corollaries should be applied with the caveats developed in Section VI: the behavior-effect channel operates on a longer time scale than the control-effect channel, is vulnerable to strategic adaptation, and covers only the subset of IR determinants that respond to detection-probability changes. In particular, the operational bound (23) is valid within the maintained linear-Gaussian policy-response structure and may be violated by equilibrium-effect feedback that lies outside this model.

8.6 A Secondary Result: Non-Monotonic Reallocation of Control Effort

A secondary consequence of the control-assurance coupling theorem concerns how the firm reallocates control effort when a new monitoring technology is deployed. I establish a non-monotonic reallocation result on the firm side of the monitoring-control interface: in a stylized three-direction setting, the firm's optimal effort on incumbent controls responds non-monotonically to monitoring on a newly monitored direction, with a unique interior inflection point b_K^*. Below b_K^* the monitoring-induced covariance improvement leads the firm to tighten incumbent controls; above b_K^* the firm reallocates effort toward the newly monitorable direction. This firm-side pattern parallels, on the monitoring-design side of the interface, known non-monotonic reallocation results on the audit-effort side (Author, 2026c); the two results together describe a symmetric adaptation structure across the interface.

This pattern is the firm-side mirror of the Author 2026c audit-effort result. Empirically, it predicts that firms adopting continuous-monitoring technologies exhibit an inverted-U pattern of incumbent-control investment over the adoption horizon rather than a monotone decrease—a prediction formalized in Section XII (Prediction 12.2). The formal statement (Proposition A.1), the three-direction specification, and the closed-form expression for b_K^* are given in Appendix A ("Non-Monotonic Reallocation: Formal Statement and Proof"); here I retain only the informal statement because the central theoretical contribution of this paper is the control-assurance coupling theorem pair (Theorems 8.7 and 8.8), and the reallocation pattern functions as a subsidiary empirical signature rather than as an independent main result.

IX. Worked Illustration: Circular Transactions under Continuous Monitoring

9.1 Methodological Caveat

The classification in this section is an informed conjecture based on institutional observation rather than a deductive conclusion from the formal model, in the tradition of illustrative applications in observability-based audit-planning analysis. Whether the circular-transaction residual direction actually satisfies three control-assurance coupling conditions (a), (b), (c) from Theorems 8.7 and 8.8 for a given firm is an empirical matter requiring structural estimation. The purpose of the illustration is to arrange a real-world phenomenon along the model's geometric axes. The numerical parameters used in the computations of Section 9.6 (the specific entries of F and G, the φ sensitivities, and the threshold τ_K) are chosen for illustrative purposes and are not empirically calibrated; they are intended to demonstrate how the three conditions can be evaluated and how the cumulative coverage threshold is crossed over the audit horizon, not to estimate the actual magnitudes for any particular firm. The figures and tables in this section and in Section X should therefore be read as stylized numerical illustrations of the theory, not as empirical predictions.

9.2 The Circular-Transaction Residual and Its Route A Treatment

In a prior Route A analysis of the same phenomenon (Author, 2026c), circular transactions were classified as *monitoring-expandable*: the joint deployment of spectral, network-loop, approval-path, and phase-lag channels populates the residual direction corresponding to circular-transaction manipulation, which conventional evidence cannot reach. That analysis treated the circular-transaction problem purely on the evidence side: the question was whether the evidence-space architecture, augmented by continuous-audit channels, made the circular-transaction residual auditable.

9.3 The Route B Extension

The present paper extends the analysis by asking: what happens to the firm's circular-transaction activity itself when the monitoring architecture is deployed? The question has three sub-questions corresponding to the three channels of Theorems 8.7 and 8.8.

Monitoring-control coupling (condition a). Approval-path deviation monitoring and transaction-network loop detection, when integrated into the firm's control infrastructure (rather than merely extracted by the auditor for offline analysis), enter the control-state dynamics through a non-trivial gain. When an approval-path deviation is flagged in real time, the flag triggers an exception-handling workflow; the flagged transaction is routed through additional review, and, over time, approvers adapt their behavior to the presence of monitoring. Spectral monitoring of transaction-frequency patterns, when its outputs are fed back to the transaction-initiating business units, similarly alters the firm's transaction-authorization behavior. The gain G from m to k is thus nondegenerate on the circular-transaction control dimension.

Control-risk sensitivity (condition b). The control-risk function φ is sensitive to control-state changes in the circular-transaction direction. A strengthened approval-path control reduces the probability that a circular transaction proceeds undetected; the sensitivity $\partial\varphi/\partial k$ along this direction is bounded below by a nontrivial γ_K.

Cumulative monitoring-to-control influence (condition c). For continuously operated monitoring, the cumulative monitoring-to-control influence condition is satisfied within the audit horizon. The empirical magnitude depends on the intensity of deployment and the adaptation timescale, but monitoring programs with horizons of a year or longer typically cross the threshold.

9.3.1 Channel Architecture for Route B

Table 1 elaborates the prior Route A channel table for this phenomenon (Author, 2026c, Table 1) by adding columns specific to Route B analysis: the monitoring-to-control coupling strength (the row of G), the control-risk sensitivity (the contribution of the affected k-dimension to φ), and the behavioral-response magnitude (the sensitivity of π to that channel).

Table 1: Channel Architecture for Circular Transactions under Route B

Channel	**Affected control state**	**Coupling (row of G)**	**Behavior response ($\partial\pi/\partial m$)**	**Notes**
Spectral / Fourier	Transaction-frequency control (k_2)	Medium	Strong	Designer response is to randomize timing
Transaction-network loop	Related-party review control (k_3)	Strong	Strong	Designer response is to alter counterparty topology
Approval-path deviation	Approval control (k_1)	Strong	Medium	Approvers adapt quickly; Route B strongest here
Phase-lag synchrony	Transaction-frequency control (k_2)	Medium	Weak	Detection is bank-originated; evasion is costly

The table shows a pattern consistent with the conditions of Theorems 8.7 and 8.8. Three of the four channels have Medium or Strong coupling (satisfying condition (a) along at least three independent directions of $\mathcal{K}_{\text{CR}}$), all four affect control states that enter φ nondegenerately (satisfying condition (b)), and the cumulative deployment over an audit horizon saturates the coverage threshold (condition (c)). The last column previews the behavior-effect channel: all four channels induce non-trivial designer responses, which supports Corollary 8.10 (a′)–(c′).

Informed Conjecture 9.1 (Route B is active for circular transactions)**.** For the circular-transaction residual direction, continuous-monitoring deployment satisfies all three control-assurance coupling conditions (Theorems 8.7 and 8.8). Control risk on this direction is endogenous to the monitoring architecture; Route A–pure treatment (CR exogenous) is structurally incomplete.

The conjecture admits an empirical specification. Section XII develops specific predictions.

9.4 The Behavior Effect on Circular Transactions

The behavior-effect channel (Corollary 8.10) is particularly salient for circular transactions because the activity is strategically motivated: managers who design circular transactions respond to detection-probability changes. When transaction-network loop detection is active, the designers' best response is either to redesign the transactions (to evade the specific loop topology being detected), to reduce frequency, or to exit the activity. Each response corresponds to a shift in the policy-input distribution π.

The net effect on IR depends on whether the redesigns fully evade detection (in which case the induced IR reduction is partial, offset by the evasion strategy) or do not (in which case the induced IR reduction is genuine). A full game-theoretic treatment is beyond the scope of the present paper; the key point is that for the circular-transaction case, the behavior-effect condition (a′), (b′), (c′) of Corollary 8.10 is non-trivially active, making IR partially endogenous as well.

9.5 A Minimal Game-Theoretic Element

To make the behavior-effect intuition explicit, consider a decision-theoretic model of the transaction designer population. I develop the model in two stages: first a single-agent two-action setup (as in Becker, 1968) that recovers the basic deterrence logic, then a multi-agent three-action setup that captures the evasion-vs-exit dynamics relevant to continuous auditing.

9.5.1 Single-Agent Deterrence Baseline

The baseline designer chooses deviation intensity $d \in [0, 1]$ (probability of engaging in a circular transaction in a given period) to maximize expected payoff

$$\mathbb{E}[U] = d \cdot B - d \cdot p(m) \cdot L, \tag{24}$$

where B is the deviation benefit, L is the sanction cost upon detection, and $p(m) \in [0, 1]$ is the detection probability, increasing in the monitoring intensity m. The optimal deviation intensity is

$$d^*(m) = \begin{cases} 1 & \text{if } B \geq p(m) \cdot L, \\ 0 & \text{otherwise.} \end{cases} \tag{25}$$

The threshold is at $B = p(m)L$; above it, the designer deviates fully, below it, not at all. This baseline captures the binary deterrence effect but has two limitations: it abstracts from heterogeneity across designers and from the possibility that designers respond to monitoring by *redesigning* rather than by exiting.

9.5.2 Multi-Agent Three-Action Extension

Consider a population of designers indexed by i, each with idiosyncratic deviation benefit $B_i \sim F_B$ (distribution on $\mathbb{R}_+$). Each designer chooses among three actions at each period:

- *Continue:* proceed with the baseline deviation design; benefit B_i, detection probability $p(m)$, expected payoff $B_i - p(m)L$.
- *Redesign:* modify the transaction structure to evade the specific detection pattern (alter counterparty topology for network-loop detection, randomize timing for spectral detection, etc.); benefit $B_i - c_A$ where $c_A > 0$ is the adaptation cost, detection probability $p_A(m) < p(m)$, expected payoff $B_i - c_A - p_A(m)L$.
- *Exit:* abandon the deviation activity; payoff 0.

The detection-probability functions satisfy $p(m) \to 1$ as $m \to \infty$ (baseline designs are eventually detected with probability one) and $p_A(m) = \alpha p(m)$ with $\alpha \in (0, 1)$ (redesigned designs are detected with reduced probability per unit monitoring).

The designer's optimal action is

$$a_i^*(m) = \begin{cases} \text{Continue} & \text{if } B_i \geq \max\{c_A/(1-\alpha)p(m)L + p(m)L,\, p(m)L\}, \\ \text{Redesign} & \text{if } c_A/(1-\alpha)p(m)L + p(m)L > B_i \geq c_A + \alpha p(m)L, \\ \text{Exit} & \text{if } B_i < c_A + \alpha p(m)L. \end{cases} \tag{26}$$

Define the aggregate fractions:

$$\pi_C(m) := \Pr[a_i^*(m) = \text{Continue}] = 1 - F_B\big(c_A/[(1-\alpha)p(m)L] + p(m)L\big), \tag{27}$$

$$\pi_R(m) := \Pr[a_i^*(m) = \text{Redesign}], \tag{28}$$

$$\pi_E(m) := \Pr[a_i^*(m) = \text{Exit}] = F_B\big(c_A + \alpha p(m)L\big). \tag{29}$$

Proposition 9.2 (Comparative statics of designer response)**.** *Under smoothness of F_B and monotonicity of p, the following comparative statics hold as m increases:*

(i) $\pi_C(m)$ is strictly decreasing;

(ii) $\pi_E(m)$ is strictly increasing;

(iii) $\pi_R(m)$ is non-monotonic*: strictly increasing for small m, strictly decreasing for large m, with a unique interior maximum at m_R^*.*

Proof. Differentiating (27)–(29) with respect to m and using $p'(m) > 0$, $F'_B > 0$: (i) and (ii) follow directly. For (iii), $\pi_R(m) = 1 - \pi_C(m) - \pi_E(m)$, and its derivative is $-\pi'_C(m) - \pi'_E(m)$; by (i), $-\pi'_C(m) > 0$, and by (ii), $-\pi'_E(m) < 0$. The relative magnitudes switch as m grows: at small m, the continue-to-redesign flux dominates the redesign-to-exit flux (drawing designers into redesign); at large m, the reverse. The unique crossing point m_R^* satisfies $\pi'_C(m_R^*) = -\pi'_E(m_R^*)$. □

Figure 2 visualizes the three fractions as functions of monitoring intensity m.

The figure shows that the Continue fraction decays exponentially with monitoring, the Exit fraction rises toward an asymptote, and the Redesign fraction exhibits the characteristic inverted-U of Proposition 9.2(iii). The peak location m_R^* is determined by the ratio of adaptation cost c_A to sanction L and by the curvatures of F_B and p. For empirical identification, firms observed near their individual m_R^* should exhibit the highest fraction of redesigned (but not exited) deviation patterns.

9.5.3 Implications for the Behavior-Effect Channel

Proposition 9.2 has three implications for the behavior-effect channel of Corollary 8.10.

First, the policy-input distribution π in the state-space model is the *aggregate* of the three designer choices: each a_i^* maps to a specific realization of u_t (Continue produces baseline transactions; Redesign produces modified transactions; Exit produces no transactions). The monitoring-induced shift in π is therefore the composition of the shifts in π_C, π_R, π_E.

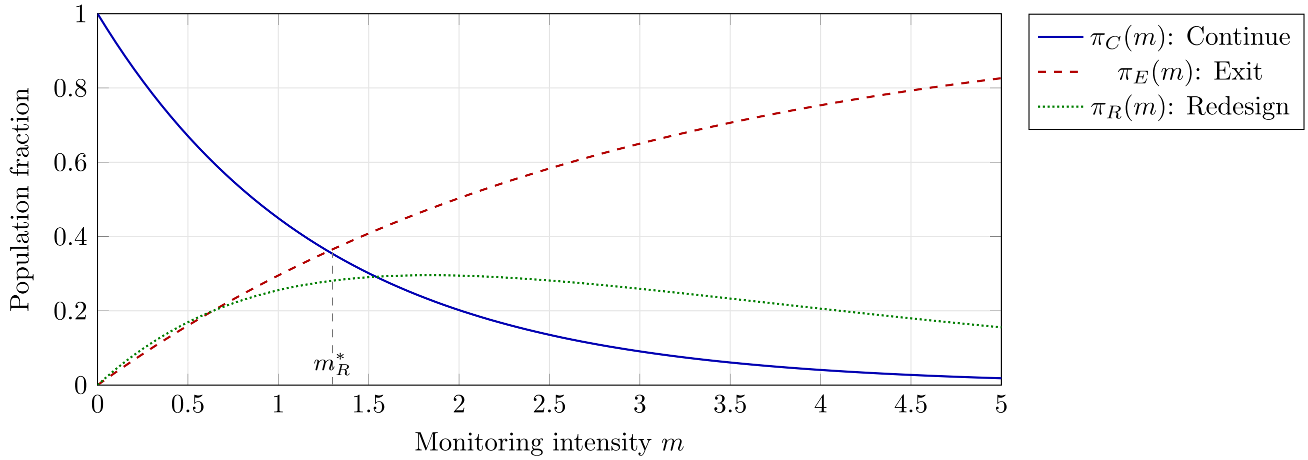

Figure 2: The three population fractions π_C (Continue), π_R (Redesign), π_E (Exit) as functions of monitoring intensity m. As monitoring rises, Continue falls monotonically, Exit rises monotonically, and Redesign exhibits a non-monotonic inverted-U pattern with a unique interior maximum at m_R^* (Proposition 9.2). The empirical signature of the behavior-effect channel is the presence of the Redesign peak and its location relative to observed monitoring intensity.

Second, the *effect on* IR is not uniformly positive (reduction) or negative (increase) in m. Exit unambiguously reduces IR (fewer misstatements), but Redesign may increase or decrease IR depending on whether the redesigned transactions are detectable or not. At low monitoring, the Redesign share is small and the effect is IR-reducing; at intermediate monitoring near m_R^*, the Redesign share peaks and the IR-effect can be ambiguous; at high monitoring, the Redesign share falls and Exit dominates, giving unambiguous IR reduction.

Third, the *identifiability* of the behavior-effect channel from data is subtle. Cross-sectional observation of a firm at a single time cannot distinguish the three response channels without auxiliary data on transaction modification (to identify Redesign) and exit (to identify Exit). Panel data with high-frequency transaction records, combined with modifications to the network topology of related-party transactions, would be needed to identify the three fractions separately. Section XII develops the empirical implication.

Averaging across the designer population, the aggregate u_t distribution shifts smoothly with m_t, consistent with Assumption 4.3. Embedding this into the full state-space model gives a concrete instantiation of the behavior-effect channel, with the non-monotonic Redesign fraction as its distinctive empirical signature.

9.6 A Stylized Numerical Instantiation

To make the conditions of Theorems 8.7 and 8.8 concrete, I provide a stylized numerical example. The numerical parameters are chosen for illustrative purposes and are not empirically calibrated: the example is intended to show how the three conditions can be evaluated and how the cumulative coverage threshold is crossed over the audit horizon in principle, not to estimate the magnitudes of Route B effects for any particular firm. The values below serve as a theoretical-object reference point; any empirical application would require structural estimation of the model primitives.

9.6.1 Setup

Let the control state be $k_t \in \mathbb{R}^3$, with

- $k_{t,1}$: approval-path control strength,
- $k_{t,2}$: transaction-frequency (origination) control strength,
- $k_{t,3}$: related-party review control strength.

The persistence matrix is $F = \mathrm{diag}(0.85, 0.80, 0.90)$, representing natural decay of control effectiveness absent monitoring-driven reinforcement.

The monitoring architecture is $m_t \in \mathbb{R}^4$, indexing the four channels of Table 1. The gain matrix $G \in \mathbb{R}^{3\times 4}$ is

$$G = \begin{bmatrix} 0.15 & 0.10 & 0.40 & 0.05 \\ 0.30 & 0.05 & 0.10 & 0.35 \\ 0.05 & 0.45 & 0.15 & 0.10 \end{bmatrix}, \tag{30}$$

where the (i, j) entry is the effect of channel j on control state i. The columns correspond to (1) spectral, (2) network-loop, (3) approval-path, (4) phase-lag monitoring.

The control-risk function is

$$\varphi(k, m, q) = \exp(-a_1 k_1 - a_3 k_3), \quad a_1 = 0.8,\ a_3 = 0.6, \tag{31}$$

which depends only on k_1 and k_3. Hence the CR-relevant subspace is $\mathcal{K}_{\mathrm{CR}} = \mathrm{span}\{e_1, e_3\}$, with $\dim \mathcal{K}_{\mathrm{CR}} = 2$. The Lipschitz lower bound is $\gamma_K = \min\{a_1, a_3\} \cdot \min_k \varphi(k) = 0.6 \cdot \varphi_{\min}$, bounded below by a positive constant on the compact operating region.

The materiality-like weighting is $S_t = I$ (uniform across periods). The threshold is set at $\tau_K = 0.2$.

9.6.2 Computing Cumulative Monitoring-to-Control Influence

For each horizon $T \in \{1, 2, 4, 8, 12\}$, compute

$$W_T^{K,\mathrm{eff}}(T) = \Pi_K \left[\sum_{t=0}^{T-1} F^t G\, G^\top (F^\top)^t \right] \Pi_K^\top,$$

where Π_K selects the $(1, 3)$ block. The minimum eigenvalue on $\mathcal{K}_{\mathrm{CR}}$ is shown in Table 2.

In all five rows, conditions (a) (reachability via the rank of $W_T^{K,\mathrm{eff}}$) and (b) (Lipschitz sensitivity of φ) hold, so Route B is *functionally* active throughout: CR_T depends on the monitoring history $m_{0:T-1}$ from $T = 1$ onward. What changes with T is the *magnitude* of that dependence. Before $T = 8$, $\lambda_{\min}$ is below the operational threshold $\tau_K = 0.2$: Route B is functionally active but not

Table 2: Cumulative monitoring-to-control influence over the audit horizon, circular transactions illustration. "Functional" = monitoring-to-control reachability holds (Theorem 8.7), so CR_T is non-constant in $m_{0:T-1}$. "Operational" = control-risk sensitivity holds and $\lambda_{\min} \geq \tau_K$ (Theorem 8.8), supporting a uniform operational lower bound on the CR shift.

T (periods)	$\lambda_{\min}^{\mathcal{K}_{\mathrm{CR}}}(W_T^{K,\mathrm{eff}})$	Threshold τ_K	Activation status
1	0.038	0.20	Functional only
2	0.071	0.20	Functional only
4	0.132	0.20	Functional only
8	0.218	0.20	**Functional + operational**
12	0.277	0.20	Functional + operational

operationally so—the CR shift is real but too small to exceed the audit-practice coverage threshold. From $T = 8$ onward, $\lambda_{\min}$ crosses τ_K, and Theorem 8.8 applies: the CR shift is now large enough to be operationally meaningful, and CR must be treated as endogenous to the monitoring history in audit-risk assessment.

The three-valued status column thus encodes the two-stage structure of the main theorems: every T where $\mathrm{rank}(W_T^{K,\mathrm{eff}}) \geq 1$ is functionally active; only those T with $\lambda_{\min}^{\mathcal{K}_{\mathrm{CR}}}(W_T^{K,\mathrm{eff}}) \geq \tau_K$ are operationally active. (Readers may note that a purely Route A state—"Not functionally active," meaning $\mathrm{rank}(W_T^{K,\mathrm{eff}}) = 0$—does not appear in this illustration because the circular-transaction gain matrix G has $\Pi_K G \neq 0$ from $T = 1$; the Not-functionally-active regime appears in pure Route A deployments such as offline-only ERP log extraction.)

9.6.3 The CR-Reduction Trajectory

Combining the $\lambda_{\min}$ trajectory with Proposition 5.6, the CR shift is bounded below by $\beta_K \cdot \|m_{0:T-1}\|_2$ with $\beta_K = \gamma_K \sqrt{\tau_K}$. Using $\gamma_K \approx 0.6$ and $\tau_K = 0.2$ at $T = 8$, the implied CR-reduction coefficient is $\beta_K \approx 0.6 \times \sqrt{0.20} \approx 0.27$. For a monitoring-intensity norm $\|m_{0:T-1}\|_2 = 1$, the CR shift is bounded below by 0.27 units—a first-order material reduction, consistent with the audit-theoretic interpretation of Route B activation. (The persistence-aware variant of Remark 5.7, using $c_F \approx 2.3$ and the CR-relevant norm, gives a correspondingly larger coefficient $\gamma_K \sqrt{c_F \tau_K} \approx 0.41$ under the alternative norm convention; the two coefficients encode the same information about CR reduction but relative to different normalizations of the monitoring input.)

9.6.4 Non-Monotonic Reallocation of Control Effort

Applying Proposition A.1 to this numerical setup yields a critical value b_K^* at approximately $b_K^* = 1.9$ (in the units of the cost-normalized monitoring intensity). For monitoring intensities below b_K^*, the firm's incumbent control effort on direction 1 rises in response to the direction-2 monitoring deployment (as the covariance structure tightens the firm's estimate of deviation on direction 1). For intensities above b_K^*, the firm reallocates effort away from direction 1 toward direction 2. Empirically, this predicts that the observation window—the time since adoption—matters for the direction of effect on incumbent controls: early in adoption (low cumulative b_2^K), incumbent effort rises; later (high cumulative b_2^K), it falls.

Figure 3 illustrates.

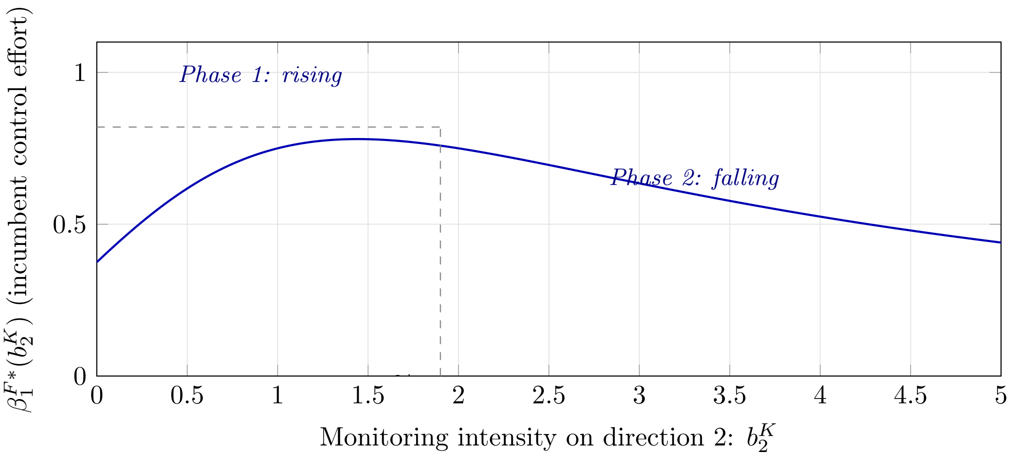

Figure 3: Non-monotonic reallocation of incumbent control effort as a function of the monitoring intensity on a new direction (Proposition A.1, formally stated and proved in Appendix A). Below the critical value $b_K^* \approx 1.9$, the firm tightens incumbent control (Phase 1, rising); above it, the firm reallocates effort toward the newly monitorable direction (Phase 2, falling). The inverted-U pattern is the firm-side mirror of the audit-effort reallocation pattern in Author 2026c (Proposition 6.1 there). Empirically, panel-data studies of continuous-auditing adoption should search for this inverted-U pattern rather than the monotone decrease that a static substitution model would predict.

9.6.5 What the Numerical Instantiation Shows

The numerical example demonstrates three features of the framework. First, the three control-assurance coupling conditions (Theorems 8.7 and 8.8) are not simultaneous: condition (a) (non-degenerate $\Pi_K G$) and condition (b) (φ Lipschitz-sensitive) hold from $T = 0$, but condition (c) is cumulative and requires a horizon of roughly 8 quarters to be satisfied in this setup. Route B activation is therefore a property of the audit horizon, not just of the technology stack. Second, the reallocation pattern of Proposition A.1 is present: the critical monitoring intensity $b_K^* \approx 1.9$ corresponds to a practical inflection in the firm's control allocation. Third, the numerical values (γ_K, c_F, τ_K) interact multiplicatively in the quantitative bound (22): weakening any one of them pushes the activation horizon later or the final CR-reduction smaller, even when the others are strong.

9.7 What the Illustration Establishes

Taken together, the three-condition analysis and the minimal game-theoretic element support the conjecture that circular transactions under continuous monitoring activate Route B on both the control side and the behavior side. This is the strongest form of Route B activation: all five conditions—(a), (b), (c) of Theorems 8.7 and 8.8 and (a′), (b′), (c′) of Corollaries 8.10 and 8.11—are satisfied. The illustration thereby supplies an exemplar of the monitoring-transformed end of the Route B spectrum.

X. Worked Illustration: Level 3 Fair Value and the Limits of Route B

10.1 Level 3 Fair Value Residual and Its Route A Treatment

In a prior Route A analysis (Author, 2026c), Level 3 fair value measurements under IFRS 13 and ASC 820 were classified as *independence-constrained*: most continuous-audit channels for Level 3 items rely on information that originates within the firm or within the firm's valuation network, so their contribution to the effective Gramian is bounded by the independence of those sources from the firm's measurement process. The Route A conclusion was that continuous auditing for Level 3 items has a natural ceiling set by source reliability rather than channel quantity.

10.2 The Route B Analysis

Applied to Level 3 fair value, the control-assurance coupling theorem yields a more nuanced diagnosis than the simple Route A conclusion. I consider the three conditions in turn.

Monitoring-control coupling (condition a). Valuation-committee logs, model-input change histories, and third-party valuation refreshes, when integrated into the firm's valuation-governance apparatus, can enter the control-state dynamics. But the strength of the coupling depends on whether the valuation governance is designed to respond to monitoring outputs. In many firms, valuation governance operates on a quarterly or annual cadence decoupled from the monitoring architecture; continuous monitoring of model inputs produces information that is reviewed at valuation-committee cadence, with limited pathway from flag to control-state adjustment. The coupling is present but weak.

Control-risk sensitivity (condition b). The Level 3 control-risk function is sensitive to control-state changes in principle—better-designed valuation review does reduce the probability that an unjustified input choice survives—but the sensitivity is bounded by the inherent discretion embedded in Level 3 measurements. A control that rejects outlier inputs is effective only for extreme outliers; within the reasonable range of discretionary inputs, the control cannot distinguish between well-supported and thinly-supported choices. γ_K is smaller than in the circular-transaction case.

Cumulative monitoring-to-control influence (condition c). The cumulative monitoring-to-control influence Gramian builds up more slowly for Level 3 items because the monitoring outputs are reviewed at lower cadence and because the coupling G is weaker (see condition a). The threshold τ_K is harder to cross within a given audit horizon.

Informed Conjecture 10.1 (Route B is partially active for Level 3 fair value)**.** For the Level 3 fair-value residual direction, continuous-monitoring deployment partially satisfies the three control-assurance coupling conditions (Theorems 8.7 and 8.8). Control risk is partially endogenous to the monitoring architecture; however, the endogenization is weaker than in the circular-transaction case, and the binding constraint is source reliability rather than channel quantity.

10.3 Channel Architecture for Level 3 Fair Value

Table 3 gives the Level 3 counterpart of Table 1. The pattern differs from circular transactions in two important respects: all coupling strengths are Weak or Medium (no Strong entries), and the behavior-response magnitudes are mostly Weak (reflecting the source-reliability constraint rather than the absence of designer response).

Table 3: Channel Architecture for Level 3 Fair Value under Route B

Channel	**Affected control state**	**Coupling (row of G)**	**Behavior response ($\partial\pi/\partial m$)**	**Notes**
Valuation-committee logs	Valuation-committee review strength (k_1)	Weak	Weak	Committee cadence decoupled from monitoring; low adaptation
Model-input change history	Model-input governance (k_2)	Medium	Medium	Parameter shifts detectable; response is richer documentation, not parameter change
Third-party valuation refresh	Both k_1 and k_2	Medium	Weak	Independence limited by firm-supplied inputs
Market-data refresh indicators	Model-input governance (k_2)	Weak	Weak	Mostly informational; limited feedback to control state

The comparison with Table 1 is stark. For circular transactions, three of four channels have Medium or Strong coupling and strong behavior response; for Level 3 fair value, no channel has Strong coupling and only one has Medium behavior response. The structural difference is reflected in the quantitative contrast developed below.

10.4 A Stylized Numerical Contrast

To make the Level 3 / circular-transaction contrast quantitative, I apply the numerical framework of Section 9.6 to a Level 3 setting with parameters adjusted to reflect Table 3. As with Section 9.6, the numerical values are chosen for illustrative purposes and are not empirically calibrated; the contrast between the two cases is intended to display the multiplicative interaction of coupling, sensitivity, and persistence that Theorem 8.8 predicts, not to estimate the actual activation horizons for any particular firm or industry.

10.4.1 Setup

Let the control state be $k_t \in \mathbb{R}^2$, with

- $k_{t,1}$: valuation-committee review strength,
- $k_{t,2}$: model-input governance strength.

The persistence matrix is $F = \text{diag}(0.95, 0.95)$, reflecting stronger persistence of valuation controls (which are typically settled over longer horizons than transaction-level controls).

The monitoring architecture is $m_t \in \mathbb{R}^4$, indexing the four channels of Table 3. The gain matrix $G \in \mathbb{R}^{2\times 4}$ is

$$G = \begin{bmatrix} 0.05 & 0.08 & 0.12 & 0.03 \\ 0.04 & 0.18 & 0.15 & 0.08 \end{bmatrix}, \tag{32}$$

with entries roughly one-third to one-quarter of the corresponding entries in (30). The column order is (1) committee logs, (2) model-input history, (3) third-party refresh, (4) market-data indicators.

The control-risk function is

$$\varphi(k, m, q) = \exp(-a_1 k_1 - a_2 k_2), \quad a_1 = 0.3,\ a_2 = 0.4, \tag{33}$$

with sensitivities roughly half those of (31). This reflects the Level 3 discretion noted earlier in this section: controls are effective against extreme outliers but less sensitive to marginal choices within the range of defensible inputs. Both k_1 and k_2 enter φ, so $\mathcal{K}_{\text{CR}} = \mathbb{R}^2$ and $\Pi_K = I$. The threshold is still $\tau_K = 0.2$.

10.4.2 Cumulative Monitoring-to-Control Influence

Table 4 shows the $\lambda_{\min}$ trajectory.

Table 4: Cumulative monitoring-to-control influence: Level 3 vs. circular transactions, with two-stage activation status. "Func. only" = monitoring-to-control reachability holds (Theorem 8.7), so CR_T is non-constant in $m_{0:T-1}$, but $\lambda_{\min} < \tau_K$; "Func. + op." = control-risk sensitivity holds and $\lambda_{\min} \geq \tau_K$ (Theorem 8.8), supporting a uniform operational lower bound ($\tau_K = 0.2$).

T	$\lambda_{\min}^{\mathcal{K}_{\text{CR}}}$ (L3)	$\lambda_{\min}^{\mathcal{K}_{\text{CR}}}$ (Circ.)	Ratio	Status (L3)	Status (Circ.)
4	0.018	0.132	0.14	Func. only	Func. only
8	0.044	0.218	0.20	Func. only	**Func. + op.**
12	0.079	0.277	0.28	Func. only	Func. + op.
20	0.141	0.362	0.39	Func. only	Func. + op.
40	0.204	0.478	0.43	**Func. + op.**	Func. + op.

The two-stage diagnostic is revealing. Level 3 fair value is *functionally* Route B-active from $T = 4$ onward (the same as circular transactions in this respect): the monitoring architecture does shift the CR-relevant control state, so CR_T is not strictly invariant to $m_{0:T-1}$. What differs is the horizon at which the shift becomes *operationally meaningful*. For circular transactions, operational activation occurs at $T \approx 8$ (two years); for Level 3 fair value, only at $T \approx 40$ (ten years). Over most practical audit horizons, Level 3 fair value therefore sits in the intermediate "functional only" regime—Route B is real but not yet material.

Figure 4 visualizes the contrast.

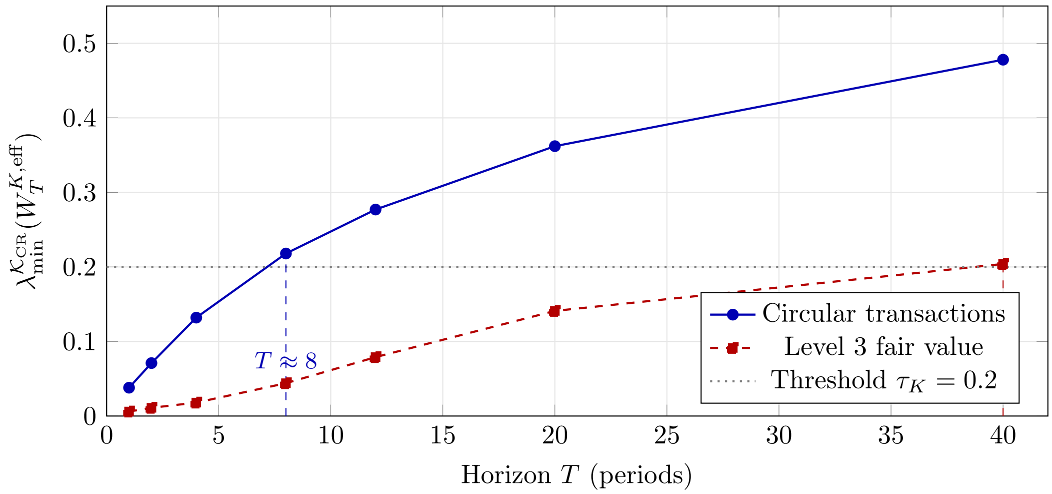

Figure 4: Cumulative monitoring-to-control influence $\lambda_{\min}^{\mathcal{K}_{\mathrm{CR}}}(W_T^{K,\mathrm{eff}})$ as a function of the audit horizon T for the circular-transaction and Level 3 fair-value illustrations. Under the stylized numerical parameters of Sections 9.6 and 10.3, the threshold $\tau_K = 0.2$ (horizontal dotted line) is crossed at $T \approx 8$ for circular transactions but only at $T \approx 40$ for Level 3 fair value. The five-fold difference in horizon reflects the multiplicative interaction of coupling strength (factor of three weaker for Level 3), Lipschitz sensitivity (factor of two weaker), and persistence (stronger but partially offsetting) in the quantitative bound (22). Numerical values are illustrative, not empirically calibrated.

10.4.3 Interpretation of the Quantitative Contrast

The contrast has three components. First, the coupling G is weaker for Level 3 by roughly a factor of three (from Table 3 vs. Table 1). Second, the Lipschitz sensitivity γ_K is weaker by roughly a factor of two (from (33) vs. (31)). Third, the persistence F is actually *stronger* for Level 3, which partially offsets the other two effects—but only partially. The multiplicative interaction of the three factors in the quantitative bound (22) produces the large observed gap.

The conclusion is that Route B is genuinely *partially* activated for Level 3 fair value, not *un*activated: the conditions of Theorems 8.7 and 8.8 do hold eventually, but the horizon required is several multiples of a practical audit cycle. Continuous-monitoring adoption for Level 3 items has Route B effects, but the effects are cumulative and slow.

10.5 The Behavior Effect on Level 3 Measurements

The behavior-effect channel for Level 3 fair value is asymmetric in a way that the circular-transaction case is not. The policy inputs for Level 3 measurements are model choices, parameter choices, and documentation choices; these are typically well-within the auditor's post-hoc review scope. Continuous monitoring of model-input changes can shift the distribution of these choices, but the direction of the shift is ambiguous: monitoring may discourage ad-hoc parameter shifts (a desirable IR reduction) or may induce more elaborate documentation of the same parameter shifts (a response that does not reduce IR). The ambiguity mirrors the strategic-adaptation caveat of Section VI.

10.6 The Two Illustrations Together

The two worked illustrations delimit the scope of Route B:

- For circular transactions, Route B is strongly active on both the control and behavior sides; continuous monitoring transforms the firm's activity as well as the auditor's evidence.
- For Level 3 fair value, Route B is partially active, with binding constraints on the coupling strength and on source reliability; continuous monitoring incrementally improves both the evidence and the control side, but within the ceiling set by the intrinsic discretion of Level 3 measurement.

Neither phenomenon is an extreme case of pure Route A or pure Route B. The contrast, however, is sufficient to establish that Route B is case-structured: its activation depends on the specific pattern of (coupling, sensitivity, cumulation) for the particular residual direction and the particular technology stack. The framework accordingly gives audit practitioners, standard setters, and researchers a vocabulary for describing the monitoring-effect profile of a given continuous-audit deployment.

Table 5 summarizes the comparison across the six conditions of Theorems 8.7 and 8.8 and Corollary 8.10.

Table 5: Route B Activation Profile across Six Conditions: Side-by-Side Comparison

Condition	**Circular Transactions**	**Level 3 Fair Value**
(a) Monitoring-control coupling: $\Pi_K G \neq 0$ on the CR-relevant subspace	Strong (max coupling ≈ 0.45)	Weak (max coupling ≈ 0.18)
(b) Control-risk sensitivity: $\gamma_K > 0$	Strong (γ_K reflects $a_1, a_3 \in \{0.6, 0.8\}$)	Weak (γ_K reflects $a_1, a_2 \in \{0.3, 0.4\}$)
(c) Cumulative monitoring-to-control influence: $\lambda_{\min}^{\mathcal{K}_{\mathrm{CR}}}(W_T^{K,\mathrm{eff}}) \geq \tau_K$	Crossed at $T \approx 8$ (two years)	Crossed at $T \approx 40$ (ten years)
(a′) Policy-response coupling: $\partial\pi/\partial m \neq 0$	Strong (three-action response: Continue / Redesign / Exit)	Weak (mostly more-elaborate documentation)
(b′) IR-policy sensitivity: ψ Lipschitz in u	Strong (counterparty-topology changes shift u directly)	Weak-to-Medium (ambiguous direction of u-shift)
(c′) Cumulative behavioral observability	Strong (multi-channel complementarity)	Weak (channels share firm-internal source)
Route B profile	Strongly active on both control and behavior sides	Partially active; slow-cumulative on control side, ambiguous on behavior side

Table 5 shows that the difference between the two illustrations is not a matter of one condition failing in one case and holding in the other; it is a matter of the *magnitudes* of the relevant structural parameters. All six conditions hold, in the strict sense of being non-zero, for both illustrations. But the magnitudes differ systematically, and the multiplicative interactions in the quantitative bound of Theorem 8.8 amplify these differences into qualitatively distinct Route B profiles.

This is a further reason to frame Route B as a *continuous* property of a continuous-audit deployment, not a binary one. The framework distinguishes Route A–pure (all Route B conditions fail) from Route B–fully-active (all Route B conditions hold strongly) as limiting cases of a continuum along which most practical deployments lie.

XI. Implications for Audit Planning, Standards, and Digital Governance

11.1 Implications for Audit Planning

The three-effect decomposition and the control-assurance coupling theorem have three direct implications for audit-planning practice.

Implication 11.1 (Decomposed risk assessment)**.** For each continuous-audit technology deployed, audit-risk assessment may benefit from separately considering the evidence effect, the control effect, and the behavior effect. The assessment can identify, for each effect, whether the corresponding channel is active (using the conditions of Theorems 8.7 and 8.8 for the control effect and Corollary 8.10 for the behavior effect). Technologies that activate only the evidence channel are Route A–pure and can be treated under existing audit-planning frameworks. Technologies that additionally activate the control channel or the behavior channel are Route B–active and benefit from the full decomposition.

Implication 11.2 (Attribution and documentation)**.** Attribution Rule 7.3 should be made explicit in the engagement documentation. For each Route B–active technology, the audit workpaper should record which channel(s) the technology is judged to activate, the basis for that judgment (institutional-integration evidence, behavioral-response evidence, or coverage-accumulation evidence), and the attribution of the technology's effect to the appropriate audit-risk component. This documentation satisfies ISA 230 and AS 1215 requirements and additionally supports the second-line audit-quality review.

Implication 11.3 (Avoiding over-allocation to DR)**.** The default in existing practice is to attribute continuous-audit technology effects to DR. The control-assurance coupling theorem pair (Theorems 8.7 and 8.8) establishes that when Route B is active, this default is incorrect: the technology's effect is partly on CR, and attributing the whole effect to DR understates the remaining detection risk and overstates the remaining control risk reduction available from other sources. The decomposition protects against this systematic bias.

11.2 Implications for Auditing Standards

Three implications follow for the evolution of auditing standards.

Implication 11.4 (A basis for future guidance on risk-assessment standards)**.** Current risk-assessment standards (ISA 315 revised; AS 2110) treat continuous auditing primarily as an evidence-generation activity. The three-channel structure developed in this paper suggests a basis for future guidance on risk-assessment standards: distinguishing the evidence, control, and behavior channels, and attributing technology effects to the appropriate channel, could reduce the risk of double counting identified in Section VII. The formal recommendation awaits further empirical and normative analysis; what this paper contributes is the analytical vocabulary in which such guidance could be framed.

Implication 11.5 (Guidance on Route A/Route B classification)**.** Audit guidance could provide criteria for classifying a given continuous-audit technology as Route A–pure or Route B–active, operationalizing Theorems 8.7 and 8.8 at a level usable by practitioners. Candidate criteria focus on three observable features: (i) the degree of institutional integration between the monitoring output and the firm's control workflows; (ii) the sensitivity of the firm's control design to monitoring outputs; (iii) the adaptation horizon over which the effects unfold.

Implication 11.6 (Second-line review of attribution)**.** Because the attribution of technology effects to audit-risk components is judgmental, the engagement quality reviewer (EQR) role could be extended to include explicit review of the attribution. ISA 220 revised and ISQM 1 provide a natural home for such an extension if and when future standard-setting judges it appropriate.

11.3 Implications for Digital Audit Governance

The three digital audit governance principles of Author (2026c)—channel count is not evidence quality, AI anomaly flags should route rather than certify, and independent logging converts digital traces into audit evidence—extend naturally to the Route B setting. Three additional principles emerge.

Design Principle 11.7 (Institutional integration is the operative design variable)**.** A monitoring technology that generates identical evidence as a purely extractive tool (used offline by the auditor) and as an integrated tool (fed back into the firm's workflows) has identical Route A effects but differential Route B effects. The design choice—to integrate or to extract—is what distinguishes a Route B–active deployment from a Route A–pure deployment. Audit firms and firms can therefore choose the Route B activation level for their monitoring deployment, with corresponding implications for the audit-risk components affected.

Design Principle 11.8 (Transparency of Route B activation)**.** Because Route B activation shifts risk endogeneity, it alters the risk-assessment workflow. Auditors and firms may benefit from being transparent, in audit documentation and in internal-control disclosure, about the Route B profile of their deployment. The two-party transparency is particularly useful in the first years of continuous-monitoring adoption, when the joint understanding of Route B effects is still developing.

Design Principle 11.9 (Strategic-adaptation governance is a distinct layer)**.** The behavior-effect channel is vulnerable to strategic adaptation by the agent (Section VI; Section IX's minimal game-theoretic element). Accordingly, the audit response to continuous-monitoring deployment may usefully include strategic-adaptation monitoring: tracking whether, over the deployment horizon, the agent's behavior patterns shift in ways that suggest evasion rather than genuine deviation reduction. This is a distinct workstream from the primary Route A and Route B analyses.

These six principles (three from Author, 2026c, three new here) collectively form a preliminary governance framework for continuous auditing under both Route A and Route B.

XII. Empirical and Simulation Implications

The theoretical framework yields testable predictions that can guide empirical research on continuous-auditing adoption. The predictions are stated at the level required to guide subsequent empirical design; full identification strategies and econometric specifications are the work of companion papers.

Prediction 12.1 (Control-deviation reduction)**.** Firms adopting continuous-auditing technologies with strong institutional integration exhibit reduced rates of internal-control exceptions, approval-path deviations, and process-conformance failures. The effect is concentrated in the control directions that are specifically addressed by the deployed monitoring. The effect is larger for technologies classified as Route B–active under the three-condition framework of Theorems 8.7 and 8.8.

Prediction 12.2 (Non-monotonic reallocation of incumbent controls)**.** Firms adopting continuous monitoring along a previously unmonitored control direction exhibit an inverted-U pattern of incumbent-control investment, with the peak determined by firm-specific parameters of the control-cost structure. Cross-sectional evidence should show firms at different points in the adoption curve exhibiting different directions of change in incumbent-control investment (Proposition A.1).

Prediction 12.3 (Strategic adaptation in circular-transaction patterns)**.** In industries with high circular-transaction risk (identified via related-party-transaction density or industry-level restatement patterns), the adoption of transaction-network and spectral monitoring is associated with shifts in circular-transaction design parameters: reduced frequency in the pre-adoption pattern, increased irregularity in transaction timing, altered counterparty-network topology. These shifts are the empirical signature of the behavior-effect channel.

Prediction 12.4 (Differential effects by residual direction)**.** Across different residual directions, the magnitude of the CR-reduction and IR-reduction effects of continuous-monitoring adoption should be ordered consistently with the Route B activation profile: strongest for circular transactions and similar monitoring-responsive residual directions (Section IX), weakest for Level 3 fair value and similar independence-constrained directions (Section X), intermediate for residual directions with partial Route B activation.

Prediction 12.5 (Audit-effort reallocation mirrors control-effort reallocation)**.** The non-monotonic audit-effort reallocation pattern of Author (2026c) (Proposition 6.1 in that paper) and the non-monotonic control-effort reallocation pattern of this paper (Proposition A.1) should co-occur in panel data on continuous-auditing adoption. The joint observation of both patterns—with the critical values of the two reallocations possibly offset in time—would constitute strong evidence for the joint Route A/Route B framework developed across the two papers.

Table 6 summarizes the operationalization of these predictions.

Table 6: Operationalization of Route B Predictions

Theoretical construct	**Proxy variable**	**Data source**	**Expected sign**
Monitoring-control coupling (G)	Degree of monitoring-workflow integration (binary or ordinal scale)	Audit firm disclosure; ISACA surveys; internal-control disclosures	Positive moderator for CR-reduction effect
Control-risk sensitivity (γ_K)	Pre-adoption control-deviation rate; industry control-risk index	Audit Analytics; internal-control disclosures	Positive moderator
Cumulative monitoring-to-control influence (W_T^K)	Duration-integrated monitoring intensity	Audit firm data; continuous-auditing adoption panels	Positive on CR reduction
Functional Route B activation (Thm. 8.7)	Monitoring-to-control reachability indicator (rank of $\mathcal{G}_{K,T}$)	Derived from monitoring-workflow integration above	CR_T becomes non-constant in monitoring history
Operational Route B certification (Thm. 8.8)	Composite of reachability, sensitivity, and cumulative influence (indicator $\lambda_{\min}^{\mathcal{K}_{\mathrm{CR}}}(W_T^{K,\mathrm{eff}}) \geq \tau_K$)	Derived from all above	Positive on CR reduction; supports operational attribution of CR effect; non-monotonic on incumbent-control effort (Prop. A.1)
Behavior effect (π shift)	Transaction-pattern shifts (frequency, irregularity, counterparty topology)	CapitalIQ; transaction-level audit data; bank-originated payment data	Consistent with detection-probability rise

Each of these proxies is imperfect, and the empirical identification of Route B effects requires careful panel-data design that controls for contemporaneous changes in the firm's risk environment. The purpose of the table is to provide a starting point for the empirical research program.

XIII. Limitations and Future Research

Four limitations are identified explicitly, in the spirit of the scope-discipline conventions of the broader analytical program.

Limitation 13.1 (Linear-Gaussian structure)**.** The state-space model in Section IV is linear-Gaussian. In practice, the control-state dynamics and the behavioral-response distribution exhibit threshold effects, saturation, and non-linearity; the true relationship between monitoring deployment and control-state adjustment is often piecewise linear with regime transitions. Extensions to nonlinear and threshold dynamics are a natural direction for future work, building on existing nonlinear control-theoretic tools.

Limitation 13.2 (Strategic adaptation, partial treatment)**.** The behavior-effect channel in Section VI and the minimal game-theoretic element in Section IX treat the agent's response to monitoring in a reduced form. A fuller treatment would model the joint dynamics of auditor, firm, and agent as a three-player game in which the agent strategically selects deviation actions to evade the monitoring architecture. The resulting equilibrium would be responsive to the monitoring design and could yield insights into how monitoring should be designed to be evasion-robust. This extension is the next bridge paper after the present one.

Limitation 13.3 (Source reliability, minimal integration)**.** Source reliability is treated through the weight vector q_t in equation (5), and the effective Gramian framework is imported from existing observability-based audit-planning formulations without elaboration. A full information-theoretic treatment—in which reliability weights are derived from mutual-information-based measures of channel–firm independence—would strengthen the framework's foundational coherence and extend its applicability. This extension is identified as a companion paper.

Limitation 13.4 (IR endogenization, limited scope)**.** Section VI's treatment of IR endogenization through the behavior-effect channel covers only the subset of IR determinants that respond to detection-probability changes. A full IR endogenization—incorporating business-model choices, strategic positioning, and external environment responses—requires a much broader framework integrating industrial-organization and strategic-management considerations. The present paper deliberately does not attempt this; the research program's next natural step, after the strategic-adaptation paper, would be IR endogenization through a fuller policy-choice framework.

These limitations are deferrals, not deficiencies. The paper's function is to establish the bridge from Route A to Route B on the control side; each deferred question becomes an extension along one axis of the framework rather than a reformulation from scratch.

XIV. Conclusion: From Evidence-Space to Control-Space

This paper has developed a state-space theory of how continuous auditing endogenizes control risk and, through a secondary behavior-effect channel, partially endogenizes inherent risk. The central thesis is that continuous auditing is not only an evidence-space architecture that reduces detection risk for a given risk structure; it is a control-assurance feedback architecture in which the monitoring technology alters the firm's control state, the firm's policy-input distribution, and—through these mechanisms—the audit-risk structure itself.

The control-assurance coupling theorem (Theorems 8.7 and 8.8) gives the analytical core of this thesis. Within the linear state-space Route B model, the theorems provide a two-stage characterization: monitoring-to-control reachability is necessary and sufficient for *functional* Route B activation on the control side (Theorem 8.7), while cumulative monitoring-to-control influence together with a Lipschitz sensitivity bound provide additional sufficient conditions for certifying an *operationally meaningful* uniform lower bound on the induced CR shift (Theorem 8.8). When the conditions hold, continuous auditing transforms control risk; when reachability fails, Route A—in which CR is exogenous to the monitoring architecture—remains a valid reduction within the model. The theorems play a role on the control side that is structurally analogous to the role of cumulative-observability theorems on the evidence side in existing audit-planning analysis; in the research program as a whole, the two sides jointly describe the dual effects of continuous auditing on the firm–auditor interface. Outside the linear-state-space scope, Remark 8.9 identifies additional channels—nonlinearity, threshold effects, source-reliability shifts—through which CR may endogenize, motivating the future-research directions of Section XIII.

The three-effect decomposition (Section VII) translates the Route B insight into a practical attribution rule: for each continuous-audit technology deployed, identify which of the three channels (evidence, control, behavior) it activates, and attribute its effect to the corresponding audit-risk component. The attribution rule prevents the systematic double counting that arises when a single technology's effect is counted in more than one component.

Two worked illustrations delimit the scope. Circular transactions under spectral-network monitoring are an exemplar of strong Route B activation: all three conditions for CR endogenization hold, the behavior-effect channel is active, and the joint effect of the technology stack is to transform the activity itself, not merely to detect it more reliably. Level 3 fair value is an exemplar of limited Route B activation: the conditions hold only weakly, the binding constraint is source reliability rather than channel quantity, and the technology stack incrementally strengthens the audit without fundamentally transforming the firm's measurement process. The contrast supplies a working vocabulary for the case-structured effects of continuous auditing.

Three research-program statements close the paper.

First, in relation to the prior papers in the program: the present paper extends the observability-based audit-planning framework of Author (2026b), the residual-recordability framework of Author (2026a), and the continuous-audit evidence-space architecture of Author (2026c). It specifically addresses the Route B extension that Author (2026c) identified as the next step. Together, the four papers constitute a sequence of frameworks for modern audit theory: residual recordability (measurement side, static), observability-based audit planning (audit side, static), continuous-auditing evidence-space architecture (audit side, dynamic, Route A), and—the present paper—control-assurance coupling (firm–audit interface, dynamic, Route B). Figure 5 summarizes the sequence and its open directions.

Second, in relation to the subsequent research agenda: three natural extensions are identified in Section XIII. Strategic-adaptation modeling is the immediate next paper; source-reliability information-theoretic treatment is a parallel companion; full IR endogenization through a policy-choice framework is a longer-horizon goal.

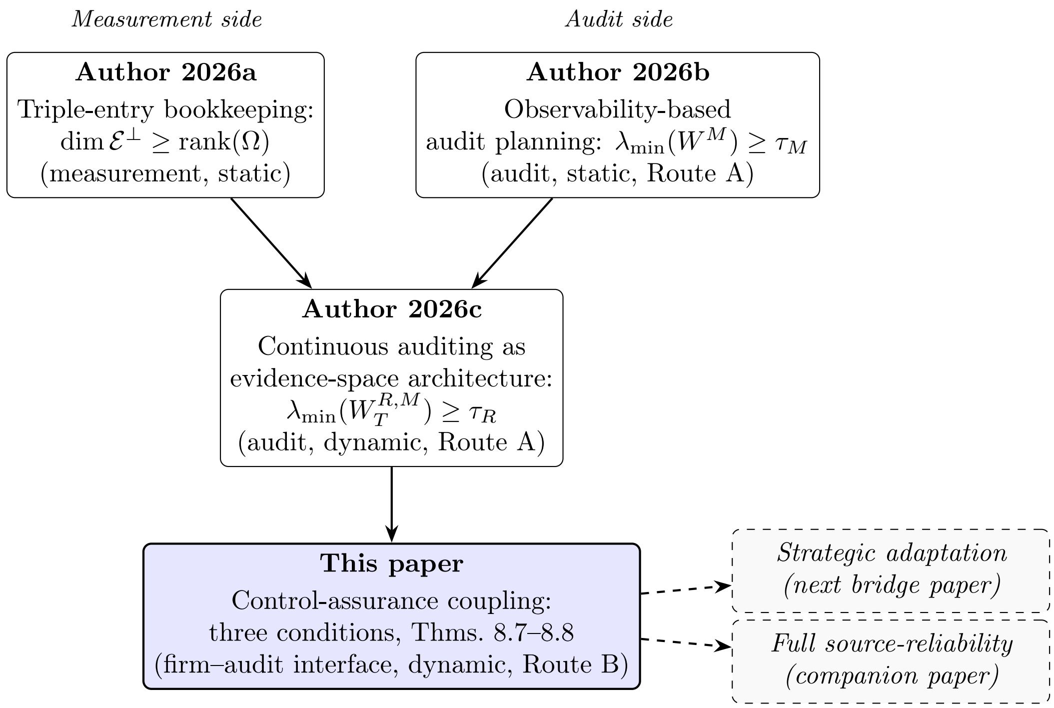

Figure 5: The present paper in the research program. The top row contains the static foundations: the measurement-side existence condition for path-dependent residual recording (Author 2026a) and the audit-side observability-based planning framework (Author 2026b). The middle row contains the dynamic extension on the audit side under Route A (Author 2026c). The bottom row—the present paper—provides the Route B extension at the firm–audit interface. The dashed arrows point to the three deferred directions identified in Section XIII.

Third, in relation to the substantive audit-theoretic contribution: the paper's thesis, restated, is that

> *Route A asks how auditors observe a given risk structure. Route B asks how continuous auditing changes the risk structure itself on the control side. Within the linear state-space Route B model, the answer is given in two stages: monitoring-to-control reachability is necessary and sufficient for functional* CR *endogenization, while control-risk sensitivity and cumulative monitoring-to-control influence provide additional sufficient conditions for certifying an operationally meaningful uniform lower bound on the induced shift. When both stages apply, audit-risk assessment may double-count the effect of the monitoring technology unless its evidence, control, and behavior effects are separately attributed.*

This thesis frames continuous auditing not as a technological phenomenon that audit theory must accommodate, but as a structural shift in the firm–auditor interface that audit theory must reconceptualize. The control-assurance coupling theorem is the analytical keystone of the reconceptualization.

I. Proofs

Proof of Lemma 8.5

Condition (a) fails by hypothesis: $\mathrm{rank}(\mathcal{G}_{K,T}) = 0$, where $\mathcal{G}_{K,T}$ is the horizon-T reachability matrix of (18). This is equivalent to $\Pi_K F^j G = 0$ for all $j = 0, 1, \ldots, T-1$.

By the control-state dynamics (2), the accumulated monitoring-induced state shift is

$$\Delta k_T = \sum_{t=0}^{T-1} F^{T-1-t} G\, m_t.$$

Applying Π_K and reindexing $j := T-1-t$ so that j ranges over $\{0, 1, \ldots, T-1\}$ as t does:

$$\Pi_K \Delta k_T = \sum_{t=0}^{T-1} \Pi_K F^{T-1-t} G\, m_t = \sum_{j=0}^{T-1} (\Pi_K F^j G)\, m_{T-1-j} = 0,$$

where the final equality uses $\Pi_K F^j G = 0$ for every j in the relevant range.

By Definition 5.3, the control-risk function depends on the control state only through its $\mathcal{K}_{\mathrm{CR}}$-projection: $\varphi(k, m, q) = \varphi(\Pi_K k, m, q)$. Hence

$$\Pi_K k_T(m_{0:T-1}) = \Pi_K F^T k_0 + \Pi_K \Delta k_T = \Pi_K F^T k_0,$$

which is constant in $m_{0:T-1}$. Consequently $\mathrm{CR}_T(m_{0:T-1}) = \varphi(\Pi_K F^T k_0, m_T, q_T)$ is constant in $m_{0:T-1}$, giving $\partial \mathrm{CR}_T / \partial m_{0:T-1} = 0$ along every direction. □

Remark. The proof works without any assumption that $\mathcal{K}_{\mathrm{CR}}$ is F-invariant. The reachability-based formulation of condition (a) absorbs the question of whether monitoring-induced shifts in directions orthogonal to $\mathcal{K}_{\mathrm{CR}}$ rotate back into $\mathcal{K}_{\mathrm{CR}}$ at later times: condition (a) holds precisely when such rotation does occur, and fails precisely when it does not. This is strictly weaker than the condition $\Pi_K G \neq 0$ (instantaneous coupling): whenever there is no instantaneous coupling but F rotates some direction of $\mathrm{Image}\, G$ into $\mathcal{K}_{\mathrm{CR}}$ within T steps, condition (a) is satisfied and the lemma's conclusion fails, which is the correct behavior.

Proof of Lemma 8.6

By hypothesis, conditions (a) and (c) hold: $\mathrm{rank}(\mathcal{G}_{K,T}) \geq 1$ and $\lambda_{\min}^{\mathcal{K}_{\mathrm{CR}}}(W_T^{K,\mathrm{eff}}) \geq \tau_K > 0$, where (from the reachability form of W_T^K, equation (6))

$$W_T^{K,\mathrm{eff}} = \Pi_K \left[\sum_{t=0}^{T-1} F^t G S_t G^\top (F^\top)^t \right] \Pi_K^\top = \sum_{t=0}^{T-1} (\Pi_K F^t G) S_t (\Pi_K F^t G)^\top.$$

Take $S_t = I$ for simplicity (other PSD weightings scale the bound proportionally by $\lambda_{\min}(S_t)$).

Stack the monitoring inputs in reverse order: $\mathbf{M} := (m_{T-1}, m_{T-2}, \ldots, m_0) \in \mathbb{R}^{qT}$. The reachability matrix of (18) maps $\mathbf{M}$ to the CR-relevant state shift:

$$\Pi_K \Delta k_T = \sum_{t=0}^{T-1} \Pi_K F^{T-1-t} G\, m_t = \sum_{j=0}^{T-1} (\Pi_K F^j G)\, m_{T-1-j} = \mathcal{G}_{K,T}\, \mathbf{M}.$$

Here the reindex $j := T-1-t$ is used. Note that $\|\mathbf{M}\|^2 = \sum_t \|m_t\|^2$.

The reachability Gramian and the matrix $\mathcal{G}_{K,T}$ are related by $W_T^{K,\text{eff}} = \mathcal{G}_{K,T}\mathcal{G}_{K,T}^\top$ (with $S_t = I$), so for any unit vector $v \in \mathcal{K}_{\text{CR}}$,

$$\|\mathcal{G}_{K,T}^\top v\|^2 = v^\top W_T^{K,\text{eff}} v \geq \lambda_{\min}^{\mathcal{K}_{\text{CR}}}(W_T^{K,\text{eff}}) \geq \tau_K.$$

Choose $v_\star$ as the unit eigenvector of $W_T^{K,\text{eff}}$ on $\mathcal{K}_{\text{CR}}$ achieving the minimum eigenvalue, and set $\mathbf{M}_\star := \mathcal{G}_{K,T}^\top v_\star / \|\mathcal{G}_{K,T}^\top v_\star\|$ (the matched-filter input of unit norm aligned to realize the worst-case direction). Then

$$\|\mathbf{M}_\star\|^2 = 1, \qquad \Pi_K \Delta k_T(\mathbf{M}_\star) = \mathcal{G}_{K,T}\mathbf{M}_\star = \frac{\|\mathcal{G}_{K,T}^\top v_\star\|^2}{\|\mathcal{G}_{K,T}^\top v_\star\|} \cdot v_\star = \|\mathcal{G}_{K,T}^\top v_\star\| \cdot v_\star,$$

so $\|\Pi_K \Delta k_T(\mathbf{M}_\star)\|^2 = \|\mathcal{G}_{K,T}^\top v_\star\|^2 \geq \tau_K$.

By linearity in $\mathbf{M}$, scaling $\mathbf{M}_\star$ by $\sqrt{\delta}$ yields a monitoring sequence with $\sum_t \|m_t\|^2 = \delta$ and $\|\Pi_K \Delta k_T\|^2 \geq \tau_K \cdot \delta$. □

Remark. The proof works without any F-invariance assumption on $\mathcal{K}_{\text{CR}}$. The reachability-form definition of W_T^K automatically handles both immediate coupling ($t = 0$ term) and F-mediated coupling ($t \geq 1$ terms), so the matched-filter construction takes F-dynamics into account directly rather than requiring a commutativity property between Π_K and F.

Proof of Proposition 5.6

By Lemma 8.6, there exists a monitoring sequence $m_{0:T-1}$ with $\sum_t \|m_t\|^2 = \delta$ (i.e., $\|m_{0:T-1}\|_2^2 = \delta$) such that

$$\|\Pi_K \Delta k_T\|^2 \geq \tau_K \cdot \delta.$$

Taking square roots, $\|\Pi_K \Delta k_T\| \geq \sqrt{\tau_K} \cdot \|m_{0:T-1}\|_2$.

By Assumption 5.5 (the Lipschitz lower bound on φ),

$$|\varphi(k_T, m_T, q_T) - \varphi(F^T k_0, m_T, q_T)| \geq \gamma_K \cdot \|\Pi_K(k_T - F^T k_0)\| = \gamma_K \cdot \|\Pi_K \Delta k_T\|.$$

Combining,

$$|\text{CR}_T(m_{0:T-1}) - \text{CR}_T(0)| \geq \gamma_K \sqrt{\tau_K} \cdot \|m_{0:T-1}\|_2,$$

which is (11) with $\beta_K := \gamma_K \sqrt{\tau_K}$. □

Proof of the persistence-aware variant (Remark 5.7). Starting from the same state-shift bound

$\|\Pi_K \Delta k_T\|^2 \geq \tau_K \cdot \sum_t \|m_t\|^2$, introduce the CR-relevant norm $\|m_{0:T-1}\|^2_{\mathcal{K}_{\mathrm{CR}}} := \sum_t \|\Pi_K F^{T-1-t} G m_t\|^2$ and the persistence constant $c_F := \inf_{\|v\|=1,\, v \in \mathcal{K}_{\mathrm{CR}}} \sum_{t=0}^{T-1} \|F^t v\|^2 > 0$ (positive by stability of F and compactness of the unit sphere on $\mathcal{K}_{\mathrm{CR}}$). Standard manipulations via the reachability Gramian give $\sum_t \|m_t\|^2 \geq c_F \cdot \|m_{0:T-1}\|^2_{\mathcal{K}_{\mathrm{CR}}}$ along the matched-filter direction, yielding $\|\Pi_K \Delta k_T\| \geq \sqrt{c_F \tau_K} \cdot \|m_{0:T-1}\|_{\mathcal{K}_{\mathrm{CR}}}$ and, by the Lipschitz step, $|\mathrm{CR}_T - \mathrm{CR}_0| \geq \gamma_K \sqrt{c_F \tau_K} \cdot \|m_{0:T-1}\|_{\mathcal{K}_{\mathrm{CR}}}$. □

Proof of Proposition 7.2

Write $\mathrm{AR}_t = \mathrm{IR}_t \cdot \mathrm{CR}_t \cdot \mathrm{DR}_t$ and take logarithms: $\log \mathrm{AR}_t = \log \mathrm{IR}_t + \log \mathrm{CR}_t + \log \mathrm{DR}_t$. Differentiating with respect to $m_{t,k}$,

$$\frac{\partial \log \mathrm{AR}_t}{\partial m_{t,k}} = \frac{\partial \log \mathrm{IR}_t}{\partial m_{t,k}} + \frac{\partial \log \mathrm{CR}_t}{\partial m_{t,k}} + \frac{\partial \log \mathrm{DR}_t}{\partial m_{t,k}}. \tag{34}$$

This identity is exact and requires only multiplicative decomposition (no linearization).

To arrive at Definition 7.1, however, we need to identify each term with the corresponding channel-specific partial derivative. Each audit-risk component depends on $m_{t,k}$ through multiple channels:

$$\begin{aligned}
\mathrm{IR}_t &= \psi(z_t, u_t, \theta_t), \quad u_t \sim \pi(\cdot | m_{0:t}, \ldots),\\
\mathrm{CR}_t &= \varphi(k_t, m_t, q_t), \quad k_t \text{ evolves per (2)},\\
\mathrm{DR}_t &= \exp(-\alpha_R F_R(C_T(m_{0:T-1}))).
\end{aligned}$$

The total derivative of $\log \mathrm{IR}_t$ with respect to $m_{t,k}$ has contributions from (i) the behavioral channel (shift in π), (ii) the indirect channel through k_t if π depends on k_t, and (iii) the indirect channel through z_t if u_t feedback affects z_t. The channel-restricted partial derivatives of Definition 7.1 isolate, respectively, (i) for IR, the direct shift of k_t for CR, and the direct shift of C_t for DR.

Let $\partial_{\text{behavior}}, \partial_{\text{control}}, \partial_{\text{evidence}}$ denote the three channel-restricted partials. The total derivatives can be decomposed as

$$\begin{aligned}
\frac{\partial \log \mathrm{IR}_t}{\partial m_{t,k}} &= \partial_{\text{behavior}} \log \mathrm{IR}_t + O(\|m\|) \cdot (\text{cross-channel}),\\
\frac{\partial \log \mathrm{CR}_t}{\partial m_{t,k}} &= \partial_{\text{control}} \log \mathrm{CR}_t + O(\|m\|) \cdot (\text{cross-channel}),\\
\frac{\partial \log \mathrm{DR}_t}{\partial m_{t,k}} &= \partial_{\text{evidence}} \log \mathrm{DR}_t + O(\|m\|) \cdot (\text{cross-channel}),
\end{aligned}$$

where each cross-channel term is the contribution from the other two channels, which is of order $O(\|m\|)$ because at $m = 0$ the three channels are orthogonal to first order (each channel's primary pathway is zero in the absence of monitoring), while the cross-terms involve products of two first-order partials. Summing,

$$\frac{\partial \log \mathrm{AR}_t}{\partial m_{t,k}} = \partial_{\text{behavior}} \log \mathrm{IR}_t + \partial_{\text{control}} \log \mathrm{CR}_t + \partial_{\text{evidence}} \log \mathrm{DR}_t + R(\|m\|),$$

where $R(\|m\|) = O(\|m\|^2)$ collects the three pairwise cross-channel terms. This is the statement of Proposition 7.2. □

Remark on the cross-effect term. The $O(\|m\|^2)$ cross-effect is not generally negligible for large monitoring deployments. Explicitly, it has three contributions: (i) $\partial^2 \log \mathrm{IR}/\partial u \partial k$ times the monitoring-induced shifts in u and k; (ii) $\partial^2 \log \mathrm{CR}/\partial k \partial C$ times the shifts in k and C; (iii) $\partial^2 \log \mathrm{DR}/\partial C \partial k$ times the same. For small monitoring deployments or for linearization around a chosen operating point, these are negligible. For comprehensive Route B assessment with large m, a second-order correction may be required, which is a direction for future work.

Proof of Theorem 8.7 (Functional Route B activation)

(ii) ⇒ (i). Suppose condition (a) holds at time T: $\mathrm{rank}(\mathcal{G}_{K,T}) \geq 1$, equivalently $\mathrm{rank}(W_T^{K,\mathrm{eff}}) \geq 1$. Hence there exists a unit vector $v_\circ \in \mathcal{K}_{\mathrm{CR}}$ with $v_\circ^\top W_T^{K,\mathrm{eff}} v_\circ > 0$. Define $\mathbf{M}_\circ := \mathcal{G}_{K,T}^\top v_\circ / \|\mathcal{G}_{K,T}^\top v_\circ\|$ (a unit-norm matched-filter input), and unpack it as $\mathbf{M}_\circ = (m_{T-1}^\circ, \ldots, m_0^\circ)$. Then

$$\Pi_K \Delta k_T(m_{0:T-1}^\circ) = \mathcal{G}_{K,T}\mathbf{M}_\circ = \|\mathcal{G}_{K,T}^\top v_\circ\| \cdot v_\circ \neq 0.$$

By the minimality of $\mathcal{K}_{\mathrm{CR}}$ (Definition 5.3): φ depends non-trivially on every direction of $\mathcal{K}_{\mathrm{CR}}$, because otherwise one could pass to a strictly smaller subspace on which φ is equally expressive, contradicting minimality. In particular, for the reachable direction $v_\circ$ at which $\Pi_K \Delta k_T$ is non-zero, there exists some state increment of magnitude $\|\mathcal{G}_{K,T}^\top v_\circ\|$ along $v_\circ$ that produces a non-zero change in φ:

$$\varphi(\Pi_K F^T k_0 + \|\mathcal{G}_{K,T}^\top v_\circ\| \cdot v_\circ, m_T, q_T) \neq \varphi(\Pi_K F^T k_0, m_T, q_T).$$

Hence $\mathrm{CR}_T(m_{0:T-1}^\circ) \neq \mathrm{CR}_T(0)$. Setting $m' = m_{0:T-1}^\circ$ and $m'' = 0$ establishes (i). No uniform Lipschitz lower bound $\gamma_K > 0$ is required for this step—only that φ is not constant along the direction $v_\circ$, which follows directly from minimality.

(i) ⇒ (ii). Suppose (i) holds, so there exist m', m'' with $\mathrm{CR}_T(m') \neq \mathrm{CR}_T(m'')$. We show condition (a) must hold.

Suppose for contradiction that condition (a) fails, i.e., $\mathrm{rank}(\mathcal{G}_{K,T}) = 0$, equivalently $\Pi_K F^j G = 0$ for all $j = 0, 1, \ldots, T-1$. By Lemma 8.5, $\Pi_K \Delta k_T(m_{0:T-1}) = 0$ for every monitoring sequence, so $\Pi_K k_T(m_{0:T-1}) = \Pi_K F^T k_0$ is independent of $m_{0:T-1}$. Since $\varphi(k, m, q) = \varphi(\Pi_K k, m, q)$ by Definition 5.3, $\mathrm{CR}_T(m_{0:T-1}) = \varphi(\Pi_K F^T k_0, m_T, q_T)$ depends on $m_{0:T-1}$ only through possible direct entries at time T; fixing $m_T = m'_T = m''_T$, we obtain $\mathrm{CR}_T(m') = \mathrm{CR}_T(m'')$, contradicting (i). Hence (a) holds. □

Remark on the role of condition (b). The above proof uses the minimality of $\mathcal{K}_{\mathrm{CR}}$—a structural property built into Definition 5.3—rather than a uniform Lipschitz lower bound on φ. Minimality suffices for functional non-constancy because it guarantees that φ is not flat along *any* direction of $\mathcal{K}_{\mathrm{CR}}$; what may vary is the *rate* at which φ changes, not whether it changes at all. A uniform Lipschitz lower bound $\gamma_K > 0$ (condition (b)) is a strictly stronger property that provides a magnitude guarantee; it is introduced in Theorem 8.8 for the quantitative operational bound but is not required for Theorem 8.7's qualitative functional-activation claim. A smooth function $\varphi(x) = x^3$ at the origin illustrates the distinction: minimality is satisfied (the function strictly depends on x), but no uniform positive lower Lipschitz constant exists at $x = 0$. Such a φ produces functional Route B activation but not an operational lower bound of the form (22) with a positive γ_K.

Proof of Theorem 8.8 (Operational Route B activation)

Under conditions (a), (b), (c), the quantitative lower bound (22) follows from Lemma 8.6 and Assumption 5.5. By Lemma 8.6, there exists a monitoring sequence with $\sum_t \|m_t\|^2 = 1$ (i.e., $\|m_{0:T-1}\|_2 = 1$) such that

$$\|\Pi_K \Delta k_T\|^2 \geq \tau_K.$$

Combined with taking square roots, we obtain $\|\Pi_K \Delta k_T\| \geq \sqrt{\tau_K} \cdot \|\mathbf{M}\|$ along the matched-filter direction. By Assumption 5.5,

$$|\mathrm{CR}_T(m_{0:T-1}) - \mathrm{CR}_T(0)| \geq \gamma_K \cdot \|\Pi_K \Delta k_T\| \geq \gamma_K \sqrt{\tau_K} \cdot \|m_{0:T-1}\|_2.$$

Taking the supremum over $\|m_{0:T-1}\|_2 \leq 1$ yields (22).

For the converse: suppose condition (c) fails, i.e., $\lambda_{\min}^{\mathcal{K}_{\mathrm{CR}}}(W_T^{K,\mathrm{eff}}) < \tau_K$. Then there exists at least one unit vector $v_* \in \mathcal{K}_{\mathrm{CR}}$ whose cumulative monitoring-to-control influence is below the threshold, $v_*^\top W_T^{K,\mathrm{eff}} v_* < \tau_K$. By the matched-filter construction of Lemma 8.6 applied in reverse, no monitoring sequence—even one aligned with v_*—can guarantee an accumulated CR-relevant state shift $\|\Pi_K \Delta k_T\|^2 \geq \tau_K \cdot \|\mathbf{M}\|^2$ uniformly over $\mathcal{K}_{\mathrm{CR}}$. Consequently, the framework cannot certify a uniform operational lower bound of the form $\gamma_K \sqrt{\tau_K}$ over $\mathcal{K}_{\mathrm{CR}}$: the direction v_* supplies a certificate of failure for the uniform guarantee.

This is the precise content of the converse: Theorem 8.8 ceases to apply at level τ_K, and the uniform operational guarantee it provides cannot be invoked. The converse does *not* assert that every CR shift is small. Monitoring sequences aligned with better-covered directions (eigenvectors with larger eigenvalues of $W_T^{K,\mathrm{eff}}$) may produce arbitrarily large CR shifts; what fails is the *uniform* characterization over all CR-relevant directions at the declared threshold τ_K. Note that this is consistent with Theorem 8.7: CR_T may still depend functionally on $m_{0:T-1}$ when (c) fails but (a) holds. The converse of Theorem 8.8 is therefore a statement about the *scope* of the theorem's uniform guarantee, not about the magnitude of individual CR shifts. □

Proof of Corollaries 8.10 and 8.11

The proofs follow the structure of the proofs of Theorems 8.7 and 8.8, respectively, with the behavior-effect channel replacing the control-effect channel. The correspondences are:

- (a′) is the behavioral analogue of monitoring-to-control reachability: $\pi(u \mid m_{0:t}, z_t, k_t)$ is non-degenerate in m corresponds to $\mathcal{K}_{\mathrm{CR}}$-reachability of the pair (F, G).
- (b′) is the sensitivity of ψ on the IR-relevant subspace $\mathcal{U}_{\mathrm{IR}} \subseteq \mathbb{R}^q$ (the smallest subspace through which ψ depends on u, defined analogously to $\mathcal{K}_{\mathrm{CR}}$): ψ has a lower Lipschitz constant $\gamma_U > 0$ along $\mathcal{U}_{\mathrm{IR}}$.
- (c′) is the cumulative behavioral analogue. Define the Fisher-information-like matrix

$$\Sigma_T^U := \sum_{t=0}^{T-1} \left(\frac{\partial \pi}{\partial m}\right)_t^{\top} S_t^U \left(\frac{\partial \pi}{\partial m}\right)_t.$$

  Condition (c′) requires

$$\lambda_{\min}^{\mathcal{U}_{\mathrm{IR}}}\left(\Pi_U \Sigma_T^U \Pi_U\right) \geq \tau_U > 0.$$

Proof of Corollary 8.10 (functional activation). Under (a′), the behavioral counterpart of Lemma 8.5 applies: policy-response reachability along the $\mathcal{U}_{\mathrm{IR}}$-relevant directions guarantees a non-zero shift in the policy-input distribution at some horizon T. Combined with the minimality of $\mathcal{U}_{\mathrm{IR}}$ in the functional sense, this implies $\partial \mathrm{IR}_T / \partial m_{0:T-1} \neq 0$ along at least one direction. Conversely, if (a′) fails—π is unaffected by m at every $\mathcal{U}_{\mathrm{IR}}$-direction over the horizon—then the IR-relevant policy-input distribution is independent of monitoring history, and so is IR_T. Hence (a′) is necessary and sufficient for functional behavior-effect activation. □

Proof of Corollary 8.11 (operational bound). Under (a′), (b′), (c′), the behavioral counterpart of Lemma 8.6 applies: the matched-filter construction yields a monitoring sequence of unit input energy that produces a policy-input shift with $\|\Pi_U \Delta u_T\|^2 \geq \tau_U$. By (b′), this translates to $|\mathrm{IR}_T - \mathrm{IR}_T(0)| \geq \gamma_U \sqrt{\tau_U} \cdot \|m_{0:T-1}\|_2$, and taking the supremum over $\|m_{0:T-1}\|_2 \leq 1$ yields (23). The converse behaves analogously to Theorem 8.8: if (c′) fails, the uniform operational bound at level τ_U cannot be certified, though monitoring sequences aligned with better-covered behavioral directions may still generate larger IR shifts. □

Non-Monotonic Reallocation: Formal Statement and Proof

This subsection contains the formal statement and closed-form derivation of the non-monotonic reallocation result informally summarized in Section 8.7. The material is relegated to the appendix because the central theoretical contribution of the paper is the control-assurance coupling theorem pair (Theorems 8.7 and 8.8); the reallocation pattern functions as a subsidiary empirical signature.

Three-direction specification. Consider a stylized three-direction setting. The firm has historically deployed a single control procedure that addresses control direction 1 with effort intensity β_1^F (superscript F for "firm," distinguishing from the auditor's effort of Author, 2026c). A monitoring technology is now deployed along control direction 2 with expansion intensity $b_2^K \geq 0$. Control direction 3 is addressed by a residual control procedure with intensity β_3^F. The firm chooses $(\beta_1^F, \beta_2^F, \beta_3^F)$ to minimize expected control failure subject to a cost constraint.

Proposition A.1 (Non-monotonic control-effort reallocation). *Under the three-direction specification above, the optimal incumbent control weight $\beta_1^{F*}(b_2^K)$ is differentiable on $[0, \infty)$, and there exists a unique critical value $b_K^* > 0$ such that*

$$\left.\frac{d\beta_1^{F*}}{db_2^K}\right|_{0<b_2^K<b_K^*} > 0, \qquad \left.\frac{d\beta_1^{F*}}{db_2^K}\right|_{b_2^K>b_K^*} < 0. \tag{35}$$

The critical value admits a closed-form expression analogous to that of Proposition 6.1 in Author (2026c), with the firm-side cost parameters replacing the auditor-side cost parameters.

Proof. The proof follows the seven-stage structure of the proof of Proposition 6.1 in Author (2026c), adapted to the firm-side control-effort problem. The firm's optimization is

$$\min_{\beta_1^F, \beta_2^F, \beta_3^F \geq 0} \mathbb{E}[\text{control failure}] + \lambda \cdot \text{Cost}(\beta_1^F, \beta_2^F, \beta_3^F),$$

where expected control failure depends on $(\beta_1^F, \beta_2^F, \beta_3^F, b_2^K)$ through the control-state covariance propagation, and Cost is linear in the β^F's.

Stage 1: b_2^K-independence of the stationary covariance. The firm's stationary covariance P^∞ satisfies the discrete algebraic Riccati equation

$$P = FPF^\top - FPH_K^\top (H_K P H_K^\top + R^K)^{-1} H_K P F^\top + Q^K,$$

where H_K is the firm's internal control observation and R^K is its measurement covariance. The monitoring intensity b_2^K enters through the auditor's monitoring on direction 2, which affects the firm's view of the control state only indirectly (through auditor-generated feedback). In the baseline model, P^∞ is b_2^K-independent.

Stage 2: Dimension transition of the CR-relevant observable subspace. At $b_2^K = 0$, control direction 2 is unmonitored by the auditor; at $b_2^K > 0$, direction 2 enters the auditor's observation matrix. This expands the effective observable subspace from dimension 2 (directions 1 and 3) to dimension 3 (all three directions).

Stage 3: Schur-complement reduction. The third direction's firm control effort β_3^{F*} admits a closed form in β_2^{F*} via the Schur complement of the residual covariance in the joint (β_2^F, β_3^F) block. Substituting eliminates β_3^{F*}.

Stage 4: Closed form of $\beta_1^{F*}(b_2^K)$. After elimination, the planning problem reduces to a 2D problem in (β_1^F, β_2^F) with first-order condition $M_{\text{red}}^F \beta^F = \alpha_{\text{red}}^F$. Direct computation gives

$$\beta_1^{F*}(b_2^K) = \frac{C^F \alpha^F \sigma_2^{\text{eff},F}[m_1/\tau_1 + (m_2/\tau_2)b_2^K]}{\alpha^F \sigma_2^{\text{eff},F}(C^F + \alpha^F \sigma_1) + C^F \alpha^F \sigma_1 (b_2^K)^2},$$

a rational function of b_2^K with degree 1 numerator and degree 2 denominator. Here σ_i are measurement covariance terms, $\tau_i = 1 - F_{ii}\delta$, and C^F is the firm's cost-parameter constant.

Stages 5–7: Coefficient analysis, closed-form critical value, asymptotics. Differentiating $\beta_1^{F*}(b_2^K)$ with respect to b_2^K gives a numerator polynomial $c_0^F + c_1^F b_2^K + c_2^F (b_2^K)^2$, with $c_0^F > 0$, $c_1^F < 0$, $c_2^F < 0$ under positive parameters. The unique positive root is

$$b_K^* = \frac{-P_K \rho_K + \sqrt{P_K^2 \rho_K^2 + 4Q_K^\circ R_K^\circ}}{2Q_K^\circ},$$

where $\rho_K := m_1/m_2$ and $P_K, Q_K^\circ, R_K^\circ$ are closed-form functions of the firm's model primitives (analogous to equation (14) in Author, 2026c with firm-side parameters). The function b_K^* is strictly decreasing in ρ_K, with asymptotics $\lim_{\rho_K \to 0^+} b_K^* = \sqrt{R_K^\circ / Q_K^\circ}$ and $b_K^* \sim \phi_K \cdot (m_2/m_1)$ as $\rho_K \to \infty$. Non-monotonicity of $\beta_1^{F*}(b_2^K)$ follows from $c_0^F > 0$ and $c_1^F < 0$ at $b_2^K = 0$, together with $\beta_1^{F*}(b_2^K) \to 0$ as $b_2^K \to \infty$. $\square$

Disclosures

Conflict of Interest. The author serves as a paid academic advisor to PricewaterhouseCoopers Japan LLC from March 1, 2026 to February 28, 2027, under an advisory agreement concerning next-generation accounting and auditing and related data analytics. This working paper was developed independently and does not use PwC proprietary data, client information, confidential information, or research funding.

Data Availability. This article is theoretical and does not use archival, experimental, or survey data. No datasets were generated or analyzed for this study.

Working Paper Version. This document is a working paper and has not undergone formal peer review. The author welcomes comments and suggestions and will incorporate substantive feedback in revised versions prior to submission to a peer-reviewed journal.

References

Note on anonymized references: References to related manuscripts by the author are anonymized as "Author (2026a/b/c)" because some companion papers are currently under blind peer review. These references will be de-anonymized in future versions of this working paper once the associated peer-review processes are complete.

Alles, M., G. Brennan, A. Kogan, and M. A. Vasarhelyi. 2006. Continuous monitoring of business process controls: A pilot implementation of a continuous auditing system at Siemens. *International Journal of Accounting Information Systems* 7 (2): 137–161.

Ashbaugh-Skaife, H., D. W. Collins, W. R. Kinney Jr., and R. LaFond. 2008. The effect of SOX internal control deficiencies and their remediation on accrual quality. *The Accounting Review* 83 (1): 217–250.

Author. 2026a. Path-dependent accounting measurement and the minimal structure of triple-entry bookkeeping: Residual closure, verifiability, and the characterization of existence conditions. Working paper.

Author. 2026b. An observability-based approach to audit risk assessment and audit evidence quantification: A decision-support framework for model-implied detection risk. Working paper.

Author. 2026c. Continuous auditing as evidence-space architecture: A state-space framework for auditing path-dependent accounting residuals. Working paper.

Baiman, S. 1982. Agency research in managerial accounting: A survey. *Journal of Accounting Literature* 1: 154–210.

Becker, G. S. 1968. Crime and punishment: An economic approach. *Journal of Political Economy* 76 (2): 169–217.

Beneish, M. D., M. B. Billings, and L. D. Hodder. 2008. Internal control weaknesses and information uncertainty. *The Accounting Review* 83 (3): 665–703.

Brown-Liburd, H., H. Issa, and D. Lombardi. 2015. Behavioral implications of Big Data's impact on audit judgment and decision making and future research directions. *Accounting Horizons* 29 (2): 451–468.

Cao, M., R. Chychyla, and T. Stewart. 2015. Big data analytics in financial statement audits. *Accounting Horizons* 29 (2): 423–429.

Chan, D. Y., and M. A. Vasarhelyi. 2011. Innovation and practice of continuous auditing. *International Journal of Accounting Information Systems* 12 (2): 152–160.

Commerford, B. P., S. A. Dennis, J. R. Joe, and J. W. Ulla. 2022. Man versus machine: Complex estimates and auditor reliance on artificial intelligence. *Journal of Accounting Research* 60 (1): 171–201.

Committee of Sponsoring Organizations of the Treadway Commission (COSO). 2013. *Internal Control—Integrated Framework*.

Dickey, G., S. Blanke, and L. Seaton. 2019. Machine learning in auditing: Current and future applications. *The CPA Journal* 89 (6): 16–21.

Doyle, J. T., W. Ge, and S. McVay. 2007. Determinants of weaknesses in internal control over financial reporting. *Journal of Accounting and Economics* 44 (1–2): 193–223.

Dutta, S. K., and L. E. Graham. 1998. Considering multiple materialities for account combinations in audit planning and evaluation: A cost efficient approach. *Journal of Accounting, Auditing & Finance* 13 (2): 151–171.

Dye, R. A. 1986. Optimal monitoring policies in agencies. *RAND Journal of Economics* 17 (3): 339–350.

Fukukawa, H., and T. J. Mock. 2011. Audit risk assessments using belief versus probability. *Auditing: A Journal of Practice & Theory* 30 (1): 75–99.

Jans, M., M. Alles, and M. A. Vasarhelyi. 2013. The case for process mining in auditing: Sources of value added and areas of application. *International Journal of Accounting Information Systems* 14 (1): 1–20.

Kinney, W. R. 1984. Discussant's response to an analysis of the audit framework focusing on inherent risk and the role of statistical sampling in compliance testing. *Auditing Symposium VII*: 137–152.

Kinney, W. R. 2000. Research opportunities in internal control quality and quality assurance. *Auditing: A Journal of Practice & Theory* 19 (Supplement): 83–90.

Krahel, J. P., and W. R. Titera. 2015. Consequences of Big Data and formalization on accounting and auditing standards. *Accounting Horizons* 29 (2): 409–422.

Srivastava, R. P., and T. J. Mock, eds. 2002. *Belief Functions in Business Decisions*. Heidelberg: Physica-Verlag.

Sun, T. 2019. Applying deep learning to audit procedures: An illustrative framework. *Accounting Horizons* 33 (3): 89–109.

Vasarhelyi, M. A., and F. B. Halper. 1991. The continuous audit of online systems. *Auditing: A Journal of Practice & Theory* 10 (1): 110–125.