Beyond Technical Frameworks: A Market Failure Approach to Agricultural AI Governance

Why Incentive Compatibility Must Accompany Technical Standards

Taiye Chen · March 2026

Summary. Agricultural AI advisory systems now reach tens of millions of smallholder farmers across South and Southeast Asia and sub-Saharan Africa. Governance frameworks are developing in parallel: the EU AI Act, the NIST AI Risk Management Framework, and FAO voluntary guidelines specify accuracy standards, bias testing, transparency obligations, and post-market monitoring. These are technical requirements. What they leave underexplored is the incentive architecture that determines whether those requirements will be reliably met — whether bias testers are financially independent of the parties whose revenues depend on the result, and whether mission-driven advisory platforms can remain solvent without compromising their neutrality. Economic market failure analysis provides that diagnostic. Three failures apply simultaneously to agricultural AI governance: principal-agent distortion, public goods underproduction, and credence goods opacity. Each points to a specific class of remedy: mandatory third-party assessment independence, results-based public financing for non-commercial advisory AI, and continuous disclosure of algorithm changes with financing conditionality.

Agricultural AI Governance and Its Missing Economic Layer

What existing agricultural AI governance frameworks usually do not address is the economic structure within which AI advisory systems operate. Two basic economics questions remain outside the scope of current governance frameworks.

First: when the entity developing and deploying an AI advisory system also earns revenue from the recommendations that system generates, who assesses whether the recommendation algorithm is optimizing for farmer outcomes or for platform margin? The EU AI Act requires bias testing for high-risk AI systems. It does not require that the bias tester be financially independent of the party whose revenue depends on the test result. The Act's notified body requirements (Article 31) impose organizational independence conditions, but they apply only to systems classified as high-risk under Annex III, a category that does not currently include agricultural AI advisory platforms, and they do not address the three-party financial dependency the agricultural case creates: an assessor with ties to input suppliers whose products the assessed system recommends and from whose sales the developer earns marketplace revenue.[1]

Second: when an AI advisory system generates social value that is diffuse, non-excludable, and therefore not appropriable through user fees, what ensures it remains financially viable without

[1] This analysis focuses on the application layer of agricultural AI governance. A parallel regulatory track, Chapter V of the EU AI Act (entered into application 2 August 2025), imposes documentation and transparency obligations on providers of general-purpose AI models used as underlying infrastructure. Where agricultural advisory platforms integrate such models, upstream GPAI provider obligations apply, but do not resolve the application-layer principal-agent problem that is the subject of this brief.

compromising its mission? Current AI governance frameworks do not fully address the public goods financing problem that makes mission-aligned agricultural AI fragile.

This brief argues for adding an economic layer to agricultural AI governance to address the incentive architecture that determines whether technical standards will produce reliable compliance. Three market failures apply to agricultural AI governance, and naming them matters because each points to a different class of remedy. Agricultural AI is the entry point; the same failure structure recurs wherever developer incentives may diverge from user welfare.

Three Market Failures, One Governance Problem

Economic theory identifies several canonical market failure types. Three apply simultaneously to agricultural AI governance, and their interaction is what makes the governance problem hard.

Principal-agent distortion. A principal-agent problem arises when an agent's objective function diverges from the principal's and compensation structures create systematic incentive misalignment.[2] When an assessor's income depends on the assessment outcome, the standard prediction is that the agent is likely to bias toward outcomes that favor continued engagement.

Public goods underproduction. A public good is non-rival and non-excludable: one person's use does not diminish availability, and non-payers cannot be excluded from the benefit.[3] Markets often underprovide public goods because providers cannot capture the full social return: the private return falls below the social return by the benefit flowing to non-payers.

Credence goods opacity. A credence good is one whose quality cannot be verified after purchase or use. For example, it is hard for a patient to verify whether a medical diagnosis was correct, even after treatment.[4] When quality is unverifiable, market discipline such as reputation effects and consumer switching is weakened because low quality cannot be reliably detected and punished.

These three failures apply simultaneously to agricultural AI governance. Their interaction compounds advisory platform's fragility: principal-agent distortion corrupts the reliability of the assessment process, public goods underproduction threatens the financial sustainability of independent advisory capacity, and credence goods opacity prevents the market and regulatory discipline that would otherwise correct both.

Two Deployments, Three Market Failures

Agricultural AI in the developing worldhas developed across two dominant organizational models, each of which makes the market failures above visible.

Plantix is a crop disease diagnostic app developed by PEAT GmbH that reached over 10 million annual users.[5] The platform began with a diagnostic mission: help smallholder farmers identify

[2] Holmström, B. (1979). "Moral Hazard and Observability." *Bell Journal of Economics*, 10(1), 74–91. Holmström's Informativeness Principle establishes formally why agent compensation tied to outcomes creates systematic incentive misalignment when effort is unobservable.

[3] Samuelson, P. A. (1954). "The Pure Theory of Public Expenditure." *Review of Economics and Statistics*, 36(4), 387–389.

[4] Darby, M. R., and Karni, E. (1973). "Free Competition and the Optimal Amount of Fraud." *Journal of Law and Economics*, 16(1), 67–88. Darby and Karni introduced the credence goods category to distinguish post-consumption unverifiability from pre-purchase information asymmetry (Akerlof, 1970) and experience goods (Nelson, 1970).

[5] Harvard Business School Digital Initiative case submission (November 2022): "an estimated 10 million annual users" (d3.harvard.edu). PEAT GmbH, founded in Berlin in 2015, received majority investment from HELM AG in April 2023; see HELM AG press release (helmag.com, April 2023).

crop diseases from photographs across more than 60 crop types. As the platform scaled, it added a marketplace through which farmers could purchase the inputs its diagnostic engine recommended — sourcing products from major agrochemical producers including Dow, BASF, Syngenta, and Bayer.[6] HELM AG, a major distributor of crop protection chemicals, acquired a majority stake in 2023. The platform that set out to reduce pesticide use became majority-owned by a company whose core business is selling them.

Digital Green is a non-profit agricultural extension platform that developed AI-based crop advisory services for smallholder farmers in India, Ethiopia, Kenya, Nigeria and Brazil.[7] The platform generated diffuse, non-excludable social value including improved yields, reduced input waste, better market access for marginalized farmers. That social value is not captured in subscription fees or licensing revenue: the benefits flowed to neighboring farms, downstream buyers, and communities with improved food security, none of whom were charged. Over 90% of Digital Green's operating budget came from donor funding.[8] The platform's advisory independence was its governance strength while financial fragility acted as the price.

How the Market Failures Manifest

Principal-agent distortion: the Plantix case

When Plantix added a marketplace to its diagnostic platform, it created a principal-agent structure: the platform made recommendations to farmers while earning revenue through transactions from recommendations. Farmers could not verify whether a recommendation reflected optimal agronomic practice or the margin structure of the marketplace. The harm from a biased recommendation extends beyond farmers overspending on inputs. Potential pesticide and fertilizer over-application could degrade soil, contaminate groundwater, reduce biodiversity, and create health risks for communities neighboring treated fields. These could be parties with no relationship to the Plantix transaction and costs the market will not compensate.

Public goods underproduction: the Digital Green case

Digital Green's advisory platform produced non-appropriable social value. The private return to the platform fell below the social return by the amount of benefit flowing to parties who could not be charged: neighboring farms, downstream buyers, households with better nutrition outcomes. No commercial financing structure could close this gap. The result was financial fragility from market structure. Functional operation required donor subsidy. Technical governance frameworks specify what advisory AI should do, but not addresses why a platform doing it well cannot sustain itself without philanthropic subsidy, or what happens to farmer welfare when that subsidy ends.

The domain discussed here is agricultural, but the mechanism is not domain-specific.

[6] Miller, S. R. (2024, October). "This app set out to fight pesticides. Once VC stepped in, the app helped sell them." *FERN / WIRED.* thefern.org.

[7] Digital Green. https://digitalgreen.org/

[8] Digital Green Foundation Form 990 filings show contributions and grants roughly comprising over 90% of annual revenue, with program service revenue below 5%. https://digitalgreen.org/financials/. ProPublica Nonprofit Explorer (EIN 26-2418959). propublica.org/nonprofits.

What Economic Governance Instruments Add

Technical governance frameworks specify the *content* of governance on what to assess, how to document it, what thresholds trigger review. Economic governance instruments address a different layer: the incentive architecture that determines whether those specifications produce reliable compliance. Three instruments address these three market failures.

Third-party assessment independence. The Plantix case presents a principal-agent problem: a recommendation algorithm whose bias testing is left to the party that commercially benefits from the recommendations. The governance remedy is a similar one Sarbanes-Oxley Section 201 applied after Enron.[9] The diagnostic was precise: assessors whose revenue depended on maintaining client relationships could not credibly assess those clients. The remedy was financial separation — the same logic behind FDA-required Independent Data Monitoring Committees,[10] which are appointed independently of the trial sponsor so interim findings can lead to trial termination without financial consequences for the monitoring body.

In agricultural AI, advisory systems that earn revenue from farmer recommendations through commissions, supply agreements, or transaction fees should have their algorithms assessed by entities with no financial ties to those revenues. Conformity assessments must prohibit relationships between assessors and input suppliers whose products are recommended. Internal bias testing should be complemented by independent external review. This separation enables testing whether commercialization skews recommendations toward chemical inputs relative to agronomic benchmarks. A documented shift would indicate farmer overspending and unpriced external harms such as soil degradation, groundwater contamination, and biodiversity loss.

However, agricultural AI developers might not be able to absorb expensive third-party audits. Public financing can address both problems simultaneously: subsidizing audit costs for developers and correcting the underinvestment in non-commercial advisory capacity that market structure creates.

Results-based public financing for non-commercial advisory AI. Digital Green's financial fragility is a market failure. The social return on independent agricultural advisory AI exceeds private return because the benefits are non-excludable. Markets (especially those in developing economies) will not be able to systematically provide this type of advisory capacity without corrective financing.

Public financing for non-commercial advisory AI should be structured around delivery indicators disbursed on a stable baseline with a performance component, with indicators defined at the start of the funding period to prevent the funder from setting the advisory agenda. Output-level indicators such as registered users, advisory sessions delivered, geographic coverage are verifiable at low cost and appropriate as Disbursement-Linked Indicators. Outcome indicators such as yield improvement rates and input cost reductions require baseline surveys, multi-season measurement, and attribution, and are better suited to program evaluation than to disbursement conditionality. This mirrors the World

9 Sarbanes-Oxley Act of 2002, Pub. L. 107-204, § 201, 116 Stat. 745. Section 201 prohibits registered public accounting firms from providing specified non-audit services to their audit clients, including bookkeeping, financial information systems design and implementation, actuarial services, internal audit outsourcing, and management functions or human resources.

10 U.S. Food and Drug Administration. *Guidance for Clinical Trial Sponsors: Establishment and Operation of Clinical Trial Data Monitoring Committees.* Rockville, MD: FDA (2006). fda.gov/media/75398. A 2024 draft update proposes strengthening the independence standard from advisory language to a normative requirement (fda.gov/media/176107); the 2006 guidance remains operative pending finalization.

Bank's Program-for-Results (PforR) disbursement structure,[11] where funds are released against Disbursement-Linked Indicators rather than project inputs, and the AgResults Challenge Fund model,[12] where development finance disburses against verified delivery milestones in agricultural programs. Development finance institutions are the natural contracting authorities: the World Bank through PforR-financed sovereign government programs that sub-grant to advisory platforms, CGIAR through its portfolio and project-level funding channels that support collaborations between CGIAR centers and external partners, and bilateral programs through challenge fund mechanisms.

Continuous disclosure with financing conditionality. Conformity assessments conducted at a single point in time create a governance gap: a developer can demonstrate compliance once and then modify the system later. For agricultural advisory platforms whose commercial arrangements may change — as new marketplace deals, supply agreements, or investor expectations emerge — the governance obligation cannot be a one-off event.

Continuous disclosure converts a periodic compliance obligation into an event-driven one. Triggering events can be specified: material changes to the recommendation algorithm's scoring or weighting logic; new commercial supply agreements with input manufacturers or distributors whose products the system recommends; changes to training data composition or geographic coverage that exceed a defined threshold; and adverse event reports above a defined severity level. This is a logic of SEC Form 8-K[13] applied to agricultural AI governance: a disclosure obligation activated by events that could shift the incentive structure facing the recommendation algorithm. The EU AI Act's post-market monitoring provisions (Article 72), which will enter into application in August 2026, establish the principle of ongoing monitoring; the proposed instrument extends this with event-driven triggers tied to commercial arrangements in addition to technical system changes.

Conditionality turns disclosure from a requirement on paper into a financial commitment. Agricultural AI platforms that remain compliant retain access to public data infrastructure, extension programs, and development finance; noncompliant platforms lose that access. This embeds financial incentives without a dedicated regulatory apparatus. Enforcement operates through two channels: in development finance programs, disclosure is an eligibility condition in sub-grant and data-sharing agreements, enforced by DFI task teams or program managers through standard project processes; in jurisdictions with AI governance frameworks, the national competent authority under EU AI Act Article 74, or its equivalent, can require disclosure for continued market access.

These three instruments complement the technical governance. The relationship is layered: technical governance specifies *what* to assess; economic governance ensures the incentive architecture makes assessment credible. A technically rigorous bias test conducted by a party with a commercial stake in the outcome is a technically rigorous test with predictable systematic bias. Both dimensions are necessary.

Implementation pathways vary by institutional capacity. In higher-capacity environments

[11] World Bank Group. *Program-for-Results Financing.* Bank Policy OPS5.04-POL.107 and Bank Directive OPS5.04-DIR.107. Washington, DC: World Bank (effective September 2020; revised March 2022). worldbank.org/programs/program-for-results-financing.

[12] AgResults is a results-based challenge fund for agricultural development, funded by Australia, Canada, the United Kingdom, the United States, and the Bill & Melinda Gates Foundation. Project evaluations and program documentation: agresults.org.

[13] Securities and Exchange Commission. *Current Report on Form 8-K.* 17 C.F.R. § 249.308; accelerated filing requirements for material corporate events adopted under SEC Release No. 33-8400 (2004).

— India has NABL-accredited electronics and IT testing laboratories operated by STQC,[14] NABCB-accredited certification bodies for ISO/IEC 42001 AI management systems,[15] and an IndiaAI Mission pillar funding algorithm auditing tools[16] — all three instruments can operate through national regulatory and financing bodies. In lower-capacity environments — Ethiopia's primary conformity assessment body (ECAE) has no accredited scope covering software or AI systems,[17] and its National AI Policy defers testing and certification scheme development to Phase 2 (2025–2027)[18] — the same instruments route through DFI conditionality: assessment by DFI-contracted evaluators, results-based disbursement verified by Independent Verification Agents, and disclosure conditionality written into sub-grant agreements. The incentive logic is identical; the contracting party differs.

How Existing Frameworks Fall Short of the Five Tests

While the following is not a comprehensive governance framework, it serves as an economic screening tool. Proposed agricultural AI governance mechanism should be able to answer this set of questions before it is considered incentive-compatible.

> **An Economic Diagnostic for Agricultural AI Governance Proposals**
>
> 1. **Assessment independence.** Who conducts the bias and accuracy testing, and does that party's compensation depend on the outcome? If the testing entity has a commercial relationship with the developer's marketplace or supply revenue, what mechanism prevents principal-agent bias?
>
> 2. **Evaluator financing.** Who finances the independent assessors, and does that financing create a dependency on the assessed party? If independent advisory AI platforms are funded primarily through commercial arrangements with the input suppliers whose products they recommend, the independence claim is compromised.
>
> 3. **Verifiability.** Are advisory quality claims verifiable, or are they credence goods that regulators and farmers cannot fully audit? If farmers cannot verify whether recommendations serve their interests, what substitutes for market discipline?
>
> 4. **Gaming windows.** Does the mechanism create point-in-time incentives that can be exploited? What prevents a developer from showing compliance at assessment time and subsequently modifying the recommendation algorithm when commercial incentives shift?
>
> 5. **Enforcement.** If an agricultural AI developer does not comply, who has the authority and incentive to impose consequences? A reporting obligation without a credible sanction could be just a suggestion.

[14] STQC (Standardisation Testing and Quality Certification Directorate, MeitY). Electronics Regional Test Laboratories hold NABL accreditation under ISO/IEC 17025. stqc.gov.in.

[15] NABCB (National Accreditation Board for Certification Bodies). BCB-195: *NABCB Accreditation Criteria for AIMS (Artificial Intelligence Management System)*. October 2024. nabcb.qci.org.in.

[16] Press Information Bureau, Government of India. "Selected Projects for Responsible AI Themed Projects under Safe & Trusted AI Pillar of IndiaAI Mission." PRID 2065579. pib.gov.in.

[17] ECAE (Ethiopian Conformity Assessment Enterprise). Accredited laboratories cover agricultural, chemical, food, microbiological, radiation, electrical, mechanical, textile, and leather testing. No software or digital systems scope. ecae.org.et.

[18] Ethiopia National AI Policy (approved by Council of Ministers, June 2024). Phase 2 deliverables include "develop testing & certification schemes; issue sector rules." regulations.ai.

Current major AI governance frameworks for agriculture only partially address these five tests in practice, particularly around financial independence, public goods financing, and credence-goods verifiability. They provide valuable technical baselines but leave key incentive questions unresolved.

The Window for Getting Incentive Design Right

Agricultural AI governance frameworks are still being developed and improved. Implementing measures for the EU AI Act's high-risk AI provisions, national digital agriculture strategies, and development bank program design criteria are all in active development. The economic architecture of who assesses, who finances independent assessment, and what triggers disclosure has not yet been set in ways that are hard to reverse. This is the window for getting the incentive design right.

Because most agricultural AI deployments in developing-country contexts fall outside the EU AI Act's direct scope, the institutional logic of Articles 31–39 — not its legal mandate — is what matters here. In practice, implementation runs through domestic accreditation where available, and through development-finance conditionality, public procurement rules, and extension program agreements where regulatory capacity is limited.

Three immediate priorities follow: require financial independence between conformity assessors and any commercial revenue streams of the assessed developer; establish results-based DFI financing windows for non-commercial advisory AI, disbursed against pre-defined protocols such as Disbursement-Linked Indicators through World Bank PforR-financed programs, CGIAR trust fund structures, and bilateral challenge funds; and mandate event-driven disclosure of algorithm changes and new commercial arrangements as conditions of access to public agricultural data infrastructure and government extension programs.

The same governance logic applies beyond agriculture. Whenever AI developers know more than users, advice quality is hard to verify ex-post, and commercial incentives conflict with user welfare, the same three market failures arise. Agricultural AI is a documented example. The instruments proposed here are therefore relevant to other AI advisory system with that incentive structure, especially in settings with limited regulatory capacity and strong reliance on development finance.

This brief draws on ongoing research on agricultural AI governance. The three-market-failure framework is developed in "Economic Safeguards for Agricultural AI: Addressing Market Failures in Algorithmic Governance." Correspondence welcome.