The Agricultural Viability Frontier

When AI Expands What Farming Can Economically Do

Johnny Kao

Independent Researcher, Tokyo, Japan

ORCID: https://orcid.org/0009-0004-1446-6001

Working Paper

Abstract

Artificial intelligence in agriculture is commonly evaluated through changes in yield, input use, labor cost, profitability, sustainability, or adoption. These measures are necessary, but they can miss a distinct economic effect: a technology may change whether a production configuration is economically viable at all. This paper develops the Agricultural Viability Frontier framework by formalizing the concept of an agricultural frontier introduced in Kao (2026). Building on research on digital agriculture, technology adoption, intensive and extensive supply responses, farm viability, and the economics of autonomous machinery, the paper distinguishes optimization from frontier expansion. Optimization improves a crop-location-scale-production configuration that was already viable; frontier expansion changes membership in the viable set.

The framework defines a production configuration, an economically viable set conditioned on technology and the surrounding economic environment, and frontier-expanding technological change. It further proposes a Frontier Expansion Test centered on four questions: whether the activity would occur in substantially the same configuration without the technology, which constraint is binding on viability, whether the technology materially relaxes that constraint, and whether complementary conditions remain sufficient. Recent evidence on autonomous field machinery, labor scarcity, robotic harvesting, and autonomous greenhouse control illustrates why technical success or positive return on investment is not by itself evidence of frontier expansion.

The contribution is integrative rather than a claim that viability boundaries or extensive-margin effects are new: the paper extends a prior geographic use of a farm viability frontier into a technology-conditioned, configuration-level framework for evaluating when AI makes agriculture better and when it makes additional agriculture economically possible.

Keywords: agricultural AI; agricultural viability frontier; digital agriculture; economic viability; technology adoption; autonomous machinery; labor scarcity; frontier expansion

Author note: This working paper develops the Agricultural Viability Frontier framework by formalizing the concept of an agricultural frontier introduced in Kao (2026).

1. Introduction

Artificial intelligence is entering agriculture through machine vision, autonomous equipment, decision support, climate control, robotics, and increasingly integrated digital production systems. The economic case for these technologies is usually expressed in familiar terms: higher yield, lower labor requirements, reduced chemical or water use, greater precision, improved reliability, or higher profit. That emphasis is well grounded. Recent research documents substantial economic and resource-use effects from digital agricultural technologies while emphasizing that returns vary across technologies, production systems, farm sizes, and regions (Finger, 2023; Papadopoulos et al., 2024; Khanna et al., 2024; Lan & Ban, 2025).

Yet an improvement in an operating farm and a change in whether a farm can operate are not the same economic event. Consider a harvesting robot. On a greenhouse that is already profitable and fully staffed, the robot may reduce cost and improve the margin. On another greenhouse, customers and capital may be available but management may refuse to build the next site because the operation cannot reliably hire enough harvest labor. The same robot can then have a different effect: it may remove the constraint that prevented expansion. The first case is an optimization effect. The second is a viability effect.

This distinction is the central idea developed in this paper. Kao (2026) describes the agricultural frontier as the economic boundary between production that makes business sense and production that does not, and proposes a counterfactual question: without the technology, would the farm happen at all? The book also distinguishes cost-down technologies, which make viable farms better, from technologies that remove a constraint that previously made production unviable. The present paper converts that argument into a scholarly framework with explicit definitions, a lightweight set representation, a diagnostic test, and an empirical research agenda.

The argument is not that agricultural economics has ignored boundaries, constraints, transformation, or extensive-margin effects. Babcock (2015) reviews the distinction between intensive responses such as yield changes and extensive responses such as changes in land use. Finger (2023) explicitly extends the digital-agriculture conversation beyond efficiency to substitution and redesign, including the possibility of new production systems. Gray (2000), in a World Bank study of Kazakhstan, used the term farming viability frontier for a primarily geographic boundary beyond which cultivated agriculture could not be sustained without subsidies, and discussed policies that could extend that frontier by reducing costs and improving profitability. These are important precedents. The narrower gap addressed here is analytical: current evaluation of AI-enabled agricultural technology rarely asks, in a common counterfactual framework, whether a technology changes the set of production configurations that are economically viable.

The paper therefore asks: when does AI merely improve the economics of agricultural production that is already viable, and when does it expand what agricultural production is economically viable? The contribution is threefold. First, it defines the Agricultural Viability Frontier at the level of a production configuration rather than only a geographic boundary. Second, it formalizes the Optimization-Frontier Distinction as a change in membership in the economically viable set under different technology states. Third, it proposes a Frontier Expansion Test that places binding constraints and complementary conditions at the center of empirical evaluation. The framework is intentionally light. It is not a calibrated production model and does not assign universal weights or thresholds. Its purpose is to make the counterfactual logic explicit enough to guide research, technology evaluation, and comparison across heterogeneous agricultural settings.

2. Existing Economic Lenses on Digital and Autonomous Agriculture

2.1 Efficiency, profitability, adoption, and redesign

A large literature evaluates digital agriculture through production and resource-use performance. Finger (2023) distinguishes efficiency, substitution, and redesign pathways through which digital innovation can contribute to more sustainable and resilient agricultural systems. Papadopoulos et al. (2024), reviewing more than one hundred peer-reviewed studies, report substantial benefits across recording and mapping, guidance, variable-rate technologies, robotics, and farm-management information systems. Khanna et al. (2024) emphasize that AI-based digital technologies can alter farmers' decision space as well as the economics of adoption. Lan and Ban (2025) reach a broadly positive but conditional conclusion in a global meta-analysis: average economic benefits are positive, but their magnitude varies materially by technology, farm size, region, and development context.

A second literature asks why farmers adopt or reject these technologies. Khanna et al. (2024) argue that AI-enabled technologies differ from earlier precision tools because they can alter farmers' decision space through prediction, reinforcement learning, autonomous equipment, and recommendation systems. The economic implications therefore extend beyond the engineering performance of the tool to incentives, risk, information, farm heterogeneity, and the structure of decisions. Kroupová et al. (2025), in a meta-analysis of technology-adoption models, find recurring effects from age, education, farm size, technical literacy, information, and expected profitability. These studies are essential because a technology cannot affect agriculture if it is not adopted.

A third literature is explicitly transformative. Klerkx and Rose (2020) frame Agriculture 4.0 technologies as potentially game-changing innovations that may create diverse transition pathways and new agricultural systems. Finger (2023) organizes digital innovation around efficiency, substitution, and redesign. Redesign can include production systems that would not exist in their previous form. It would therefore be inaccurate to characterize the existing literature as concerned only with marginal efficiency. The conceptual issue is instead that transformation, adoption, and performance are different questions from viability membership. A production system can be radically redesigned while remaining economically inferior to an existing alternative; conversely, a relatively narrow technology can leave the production system largely unchanged while removing the one constraint that prevented expansion or continuation.

Table 1. Adjacent economic lenses and the remaining viability question

Literature lens	**Primary question**	**What it captures**	**Remaining viability question**
Performance / precision agriculture	How does technology change yield, cost, inputs, profit, or environmental outcomes?	Intensive performance and resource-use effects.	Would the configuration have occurred without the technology?
Adoption	Who adopts, when, and under what incentives or constraints?	Diffusion, heterogeneity, expected profitability, information and risk.	What changes after adoption in the set of viable configurations?
Redesign / transition	How can digital technology reshape production systems?	System transformation, substitution, new organizational or technical arrangements.	Does redesign cross an economic viability boundary or only change how production is organized?
Intensive / extensive supply response	Does response occur through yield/intensity or through land/area?	Supply response and land-use margins.	Can viability change through scale, organization, crop, or supervision even without a land-area change?
Autonomous machinery economics	When is autonomy profitable and how does it affect farm size or machinery choice?	Technology-specific feasibility, cost curves, supervision and labor effects.	Is positive profitability an optimization effect or evidence that a previously non-viable configuration became viable?

2.2 From the extensive margin to a viability frontier

The intensive-extensive margin distinction is the closest familiar economic analogy. Babcock (2015) describes agricultural supply response as occurring through higher yield or more intensive use of existing land on one hand and changes in harvested area on the other. Technologies can affect both. The distinction is useful, but it is not identical to the framework proposed here. A viability shift may occur without any movement in the geographic land frontier. A farm may continue rather than exit, a greenhouse company may build a second facility, a small irregular field may become profitable with different machinery, or a remote site may operate under a new supervisory arrangement. These are extensive in a broader organizational sense, but they are not always changes in cultivated land area.

Gray (2000) provides an important historical antecedent. In the context of post-Soviet Kazakhstan, he defined a farming viability frontier as the geographic limits beyond which cultivated agriculture could not be sustained without subsidies. The frontier could differ by farm type, and policy could move it by changing land markets, input and output markets, costs, and profitability. That formulation already contains two ideas retained here: viability is economic rather than purely biophysical, and the boundary is conditional on the organization and environment of production.

The contribution of the present paper is therefore not the phrase "viability frontier" itself. It is the configuration-level, technology-conditioned formalization and the counterfactual test used to distinguish optimization from frontier expansion. Gray's formulation is an explicit antecedent and is treated as such throughout this paper.

The present framework generalizes that logic in two directions. First, the unit of analysis is a production configuration rather than only a point in geographic space. Second, the frontier is explicitly conditioned on technology. The resulting concept should not be confused with a technical-efficiency or stochastic production frontier, which asks how close observed output is to a best-practice production possibility. The Agricultural Viability Frontier instead asks whether a configuration clears an economic hurdle under a given technology and environment. It is a boundary of viable choices, not a measure of productive efficiency.

2.3 Labor scarcity and autonomous equipment as frontier-relevant evidence

Labor-saving innovation provides a particularly clear mechanism because agricultural labor can be either a cost or a hard capacity constraint. Gallardo and Sauer (2018) review the induced-innovation tradition and the long history of agricultural technologies responding to labor scarcity and relative factor prices. Recent evidence shows that this mechanism remains active. Win, Rutledge, and Maredia (2025) find that California farmers facing labor shortages respond through higher wages, cultivation changes, labor contractors, and, with a lag, greater adoption of labor-saving technology. Kuroiwa, Veettil, and Gupta (2026) find that rice farmers in rural India adopted direct seeding in response to pandemic-related migrant labor scarcity while maintaining average profits close to those from transplanting, although with different downside risks.

Research on autonomous field machinery moves closer to the viability question. Lowenberg-DeBoer et al. (2021) show that autonomous equipment can alter farm cost curves and allow smaller arable farms to approach minimum unit production costs. Al-Amin et al. (2023) find that autonomous crop machines can farm small and irregularly shaped fields profitably in scenarios where large conventional equipment performs poorly. Al-Amin et al. (2024) extend this logic to strip cropping, a production configuration whose conventional mechanization is constrained by high labor requirements, and model conditions under which autonomous machinery can make the system economically attractive. Shockley et al. (2022) show the reverse mechanism: on-site supervision requirements can erode the economic advantage of autonomous machines and, in some scenarios, make them no longer commercially viable. Most directly, Strine, Fiechter, and Lowenberg-DeBoer (2025) find that conventional mechanization remains more profitable when labor is reliable, whereas large-scale autonomous machines can become economically viable under severe labor shortage; the authors identify farm expansion as a likely early pathway for economic advantage.

Together, these studies show that agricultural technology effects are conditional on the constraint environment. They also reveal the need for a common classification. A paper can demonstrate positive return on investment, lower cost, or technical profitability without showing that production would not otherwise occur. Conversely, the largest economic effect of autonomy may appear precisely where the relevant baseline is not a cheaper human-operated farm, but a farm, expansion, or production configuration that would not be undertaken at all.

3. The Agricultural Viability Frontier Framework

Definition 1 (Agricultural production configuration). An agricultural production configuration p is a materially specified combination of what is produced, where it is produced, at what scale, with what production system, and under what organizational or supervisory arrangement. A compact representation is:

$$p=\langle c, l, s, m, o\rangle$$

Here c denotes crop or production output, l location, s scale, m production system or method, and o the organizational or supervisory arrangement. The representation is deliberately schematic. The relevant dimensions can differ by application, but the unit must be defined tightly enough that a counterfactual comparison is meaningful. “Tomato production” is too broad; a specific crop-location-scale-system arrangement can be evaluated for viability.

For illustration, consider a commercial tomato greenhouse in a remote region. The crop, location, production scale, greenhouse system, and supervisory arrangement together define the configuration. Under a baseline technology state, the operation may be technically feasible but economically non-viable because it requires scarce expert supervision on site. If AI-enabled monitoring and control reduce the amount of local expert presence required, the same configuration may cross the viability hurdle. The relevant question is therefore not whether AI improves greenhouse performance in general, but whether it changes the economic viability of this specified configuration.

Definition 2 (Economically viable set and Agricultural Viability Frontier). Let T denote the available technology state and E the surrounding economic environment, including relevant output and input prices, wages, capital costs, market access, policy, regulation, and other conditions that affect economic return. Let R(p; T, E) denote the expected risk-adjusted net economic return to configuration p after relevant operating and capital costs, and let H(p, E) denote the context-specific viability hurdle. The economically viable set is:

$$V(T,E)=\{p\in P: R(p;T,E)\geq H(p,E)\}$$

The Agricultural Viability Frontier is the boundary separating configurations inside V(T, E) from configurations outside it. The notation is conceptual rather than empirically calibrated. It does not assume that all decision-makers share a single required return, nor that uncertainty can be reduced to one stable number. The purpose is simply to make clear that viability is conditional: the same physical farm can move into or out of the viable set when technology, prices, labor availability, financing, or policy changes.

Definition 3 (Frontier-expanding technological change). Consider two technology states, T0 and T1, holding the surrounding environment E fixed for the counterfactual comparison. A technology has a frontier-expanding effect for configuration p when p is not economically viable under the baseline technology state but becomes viable under the new state:

$$p\notin V(T_0,E)\wedge p\in V(T_1,E)$$

Equivalently, the frontier-expanding set is $V(T_1, E) \setminus V(T_0, E)$. A technology can improve profits without expanding this set. That distinction is essential because an observed gain in productivity, cost, or return on investment is not sufficient evidence that the frontier moved.

Principle 1 (Optimization-Frontier Distinction). *Optimization* improves a production configuration that was already economically viable. *Frontier Expansion* changes which production configurations are economically viable.

Table 2. Optimization versus frontier expansion

Dimension	Optimization effect	Frontier-expansion effect
Baseline counterfactual	Substantially the same configuration remains viable without the technology.	The configuration is not viable, does not expand, or would exit without the technology.
Primary economic change	Yield, cost, margin, reliability, resource use, or quality improves.	The viable set changes; entry, expansion, continuation, location, scale, or organization becomes economic.
Role of targeted constraint	Constraint may be important but is not binding on viability.	Technology materially relaxes a constraint that is binding on viability.
Empirical signature	Continuous improvement in an existing operation.	A counterfactual change in the discrete decision to produce, expand, continue, locate, or organize production.
Evidence required	Performance comparison can be sufficient.	Requires a credible counterfactual about viability without the technology.

3.1 Binding constraints and complementary conditions

Definition 4 (Binding viability constraint). A constraint is binding on agricultural viability when relaxing it can change the economic decision from non-production to production, non-expansion to expansion, exit to continuation, or one materially different production configuration to another, while other necessary conditions remain sufficiently intact.

This definition explains why identical technology can have different economic effects across farms. Labor may be expensive but available on one operation; harvesting automation then reduces cost. On another operation, labor may be so unreliable that management refuses to add acreage or build another greenhouse. There, the same automation can remove a ceiling. Similar logic can apply to expert knowledge, management capacity, distance from specialist support, reliability, a specific physical bottleneck, financing, or regulation.

Constraint removal is not sufficient by itself. A project that solves labor scarcity can remain non-viable because energy, water, capital, logistics, demand, or crop price is unfavorable. Conversely, a technology may appear to have a weak average effect because many observed farms do not face the targeted constraint, while having a large viability effect in a subset where that constraint is binding. The frontier framework therefore makes heterogeneity central rather than treating it as noise around an average treatment effect.

Technology can also move the frontier inward when it introduces new requirements or when surrounding rules reintroduce the scarce input that autonomy was meant to economize. Shockley et al. (2022), for example, show that human-supervision rules can substantially reduce the commercial viability of autonomous field equipment. The relevant object is therefore not AI capability in isolation but the full technology-in-environment configuration.

4. The Frontier Expansion Test

A viability framework is useful only if it can discipline empirical claims. The Frontier Expansion Test is a four-question diagnostic for evaluating whether a reported agricultural technology effect should be interpreted as optimization or frontier expansion. It is not a statistical estimator and does not replace causal identification. It specifies the counterfactual that empirical work must establish.

1. Baseline viability. Would substantially the same production configuration occur, expand, or continue without the technology? If yes, the default interpretation is optimization unless another configuration-level boundary changes.

2. Binding constraint. What specific factor prevents the baseline decision? A high cost is not automatically binding. The constraint must plausibly determine the produce / do-not-produce, expand / do-not-expand, or continue / exit decision.

3. Causal mechanism. Does the technology materially relax that binding constraint? A technically impressive feature that does not alter the binding factor is unlikely to move the viability frontier.

4. Complement sufficiency. After the targeted constraint is relaxed, are the remaining economic and operational conditions sufficient for the configuration to clear the viability hurdle? Frontier expansion requires the whole business configuration, not only one subsystem, to become economic.

The phrase "substantially the same configuration" matters. A technology may induce a change in scale, location, organization, crop, or production system. If those changes are material to the business model, they should be treated as distinct configurations rather than hidden inside a before-after average. This also prevents overclaiming. A farm may adopt a robot and earn a higher margin, but if it would have planted the same crop at the same scale without the robot, there is no demonstrated frontier effect.

Observed operation is not identical to economic viability. Owners can choose not to pursue profitable opportunities for reasons outside the modeled return, and unprofitable production can persist temporarily because of subsidies, sunk costs, expectations, family objectives, or financing arrangements. The test should therefore be applied to a clearly specified economic hurdle and supported, where possible, by investment plans, closure or expansion decisions, cost structures, financing conditions, and credible counterfactual evidence rather than by observed adoption alone.

5. Applying the Framework to Recent Evidence

Recent studies illustrate how the same literature can be read differently once the evidentiary question shifts from "does the technology work or pay?" to "did the technology change viability?" Table 3 classifies several examples. The classifications are not claims about the technologies in general; they describe what the cited study design establishes.

Table 3. Frontier relevance of selected recent evidence

Study	Constraint / technology	What the study establishes	Frontier interpretation
Strine et al. (2025)	Severe labor shortage; large-scale autonomous row-crop machinery.	Conventional machinery remains more profitable with reliable labor; autonomy can become viable under severe labor shortage; expansion is identified as an early pathway.	Strong frontier-relevant evidence because viability changes with a binding labor condition; expansion counterfactual is explicit but remains modeled.
Al-Amin et al. (2023)	Small / irregular fields; autonomous crop machines.	Autonomous machines can profitably farm small irregular fields where larger conventional equipment is uneconomic or performs poorly.	Strong candidate frontier effect at the field-configuration level: machinery changes which field configurations can be farmed profitably.
Charlton et al. (2025)	Seasonal harvest labor; robotic apple harvesting.	Estimates robot performance and price conditions under which robotic and manual	Primarily optimization / substitution evidence unless a credible counterfactual shows

Study	Constraint / technology	What the study establishes	Frontier interpretation
		harvesting yield equal profit.	the orchard would otherwise reduce output, fail to harvest, or exit.
Maree et al. (2025)	Skilled greenhouse management; autonomous climate / irrigation algorithms.	All challenge teams achieved profitable strategies in autonomous dwarf-tomato cultivation.	Demonstrates technical and economic capability, not frontier expansion by itself, because human-managed counterfactual non-viability is not established.
Shockley et al. (2022)	Regulatory requirement for on-site supervision of autonomous machines.	Supervision rules reduce profitability and can make autonomous alternatives no longer economically viable.	Shows that the frontier is environment-dependent and can contract when regulation restores the constrained input.
Al-Amin et al. (2024)	High labor requirement in strip cropping; autonomous machinery.	Models scenarios in which autonomy can make strip cropping economically attractive by reducing labor requirements.	Frontier-relevant because autonomy may enable a distinct production configuration otherwise constrained by labor.
Spykman et al. (2026)	Small-scale strip cropping; current commercial crop robot.	Field-lab evidence finds current robot labor requirements and economies of field size still limit the expected small-field advantage.	Counterevidence: frontier expansion should not be inferred from autonomy alone; current deployment may fail the test.

The Strine et al. (2025) result is particularly revealing because the technology itself does not change across the core comparison. With reliable labor, conventional mechanization is economically superior. Under severe labor shortage, autonomous machinery can become viable. The economic role of autonomy therefore depends on whether labor is a priced input or a binding capacity constraint. That is the Optimization-Frontier Distinction in operational form.

Al-Amin et al. (2023) offers another important pattern. The study asks whether autonomous machines alter the economics of small and irregular fields. The reported result is not merely lower cost on a representative field; autonomy changes the profitability of field configurations that conventional large machinery struggles to serve efficiently. In the framework developed here, this is closer to direct evidence that the viable set can change.

The strip-cropping analysis in Al-Amin et al. (2024) sharpens the mechanism. Conventional strip cropping carries higher labor requirements, while the modeled autonomous configuration can change its economics by relaxing that labor burden. This is frontier-relevant because the technology is tied to the feasibility of a distinct production configuration, not merely to a marginal cost reduction within an unchanged system. At the same time, Spykman, Lowenberg-DeBoer, and Gandorfer (2026) provide a useful counterpoint from field-lab evidence: a current commercial crop robot did not yet overcome economies of field size or reduce human labor enough to make small-scale diversified strip cropping economically superior. The contrast illustrates why frontier expansion must be demonstrated rather than inferred from technological capability.

By contrast, the robotic apple-harvest model in Charlton et al. (2025) is economically informative without necessarily demonstrating frontier expansion. A break-even condition between robotic and manual harvesting shows when substitution can preserve profit. Yu et al. (2024) similarly model the timing and profitability of robotic weed management under herbicide resistance, showing how biological and economic conditions shape optimal adoption. These are valuable threshold and adoption results, but a frontier interpretation requires an additional counterfactual: the relevant production, harvest, acreage, continuation, or production-system decision would otherwise not occur. Economic viability of a technology is not automatically evidence that the technology expanded the viable set of agriculture.

The autonomous greenhouse evidence provides a related caution. Maree et al. (2025) show that intelligent algorithms can manage a dwarf-tomato crop profitably and at high production levels. That is significant evidence about capability and economic performance. It does not by itself show that the greenhouse configuration was outside the viable set under human management. Frontier claims require a different counterfactual than capability claims.

This distinction changes the burden of proof in a useful way. It does not downgrade efficiency gains, which can be economically large. It prevents two effects from being conflated. A technology may deserve adoption because it saves money without moving the frontier; another may show modest average savings but have disproportionate value where it relaxes a binding constraint. Both are economically important, but they answer different questions.

6. Implications for Agricultural AI Evaluation

First, technology evaluation should identify the denominator behind percentage improvements. A 20 percent labor reduction is economically different when labor is a small, reliably available input than when the operation is already unable to staff an additional production unit. Performance metrics should therefore be paired with a constraint diagnosis. This applies to labor, but also to expert attention, management bandwidth, water, distance, reliability, or other factors that can become binding.

Second, frontier effects are configuration-specific. Claims such as “autonomy expands agriculture” are too broad. A technology can expand viability for small irregular fields while having little effect on large rectangular ones; it can enable expansion in labor-scarce regions while remaining an expensive substitute where labor is abundant. The appropriate empirical object is the interaction between technology and the local constraint environment, not the average technological capability in isolation.

Third, policy can move the frontier independently of engineering progress. Gray (2000) emphasized the role of institutions and markets in shifting the geographic viability of farms. Shockley et al. (2022) shows an analogous mechanism for autonomous machinery: supervision rules can remove part of the labor-saving value. A policy evaluation that measures only technology subsidies or adoption rates can therefore miss whether regulation, financing, infrastructure, liability rules, or market access are the actual binding conditions.

Fourth, the framework separates social value from vendor value. A technology may expand the set of viable agricultural production without creating an attractive technology business if value cannot be captured by the vendor, deployment is too costly, or competition transfers gains elsewhere. Those questions are important but analytically downstream. The present paper stops at agricultural viability; deployment and value capture require separate treatment.

Finally, the framework offers a different way to prioritize technical development. The highest benchmark improvement is not necessarily the most economically consequential improvement. If an additional percentage point of accuracy does not change a production decision, it is an optimization improvement. If a narrower engineering change removes a failure mode that prevented unattended operation, reliable harvest, or remote supervision, its effect on viability can be larger. This does not imply that engineers should optimize only for frontier movement, but it makes explicit the economic mechanism through which technical characteristics matter.

7. Research Agenda

The framework is conceptual, but it generates empirical questions that can be tested. The first requirement is to observe the decisions that standard adoption datasets often omit. Researchers need information not only on farms that adopted a technology, but also on expansions that were postponed, crops that were not planted, fields that were not cultivated, operations that exited, and investments that were rejected because a constraint was binding. Without these counterfactual margins, a viability effect can be misclassified as an ordinary productivity effect or missed entirely.

A minimal empirical design could compare otherwise similar production configurations exposed to an exogenous constraint shock and differential access to a relevant technology. Labor-supply disruptions are a natural candidate. Difference-in-differences or event-study designs could test whether access to automation changes acreage expansion, facility construction, continuation, or crop choice after a labor shock, rather than only changing labor expense among continuing adopters. The labor-scarcity studies of Win et al. (2025) and Kuroiwa et al. (2026) show the value of shocks for identifying adaptation behavior; the next step is to connect those adaptations to discrete viability outcomes.

A second pathway is structural or optimization modeling. The autonomous-machinery literature already models cost curves, machinery choice, labor availability, field size, supervision, and profitability (Lowenberg-DeBoer et al., 2021; Al-Amin et al., 2023; Strine et al., 2025). These models could explicitly map $V(T, E)$ across technology states and report not only the change in profit for configurations that remain viable, but also the configurations entering or leaving the viable set. Sensitivity analysis could then reveal which constraint is actually binding and how robust the frontier movement is to wages, prices, regulation, financing, or technology cost.

A third pathway is prospective investment research. Surveys of growers or agribusiness managers can separate "this technology would improve an operation we will run anyway" from "this technology would change whether we build, expand, continue, or locate the operation." Stated intentions are weaker evidence than realized behavior, but they can identify where to look for frontier effects before sufficient post-adoption data exist.

Three empirical propositions follow from the framework. First, the same technology should generate larger configuration-level effects where the targeted constraint is binding than where it is merely costly. Second, frontier effects should appear in discrete outcomes—entry, expansion, continuation, location, or organization—rather than only in continuous performance metrics. Third, requirements that reintroduce the constrained factor, such as mandatory one-to-one human supervision for autonomous equipment, should attenuate or reverse the frontier effect. These propositions can be tested without converting the framework into a universal scoring model.

8. Limitations

The Agricultural Viability Frontier is a conceptual framework, not an empirically calibrated production model. The viability hurdle is context-specific and may differ across farms because of financing costs, risk tolerance, opportunity costs, subsidies, family objectives, and expectations. The notation is intended to clarify counterfactual structure rather than imply that viability can always be measured with a single stable threshold.

Second, agricultural production can face multiple interacting constraints. Removing one constraint can expose another, and the identity of the binding constraint can change over time. AI may reduce labor requirements while increasing capital, service, connectivity, cybersecurity, or supervision requirements. The framework handles this through the complement-sufficiency step but does not model nonlinear interactions among constraints.

Third, observed operation is an imperfect proxy for economic viability, and causal identification is difficult. Adoption is endogenous, firms self-select into technologies, and unobserved managerial quality may influence both adoption and expansion. Strong frontier claims therefore require more than before-after profitability comparisons.

Fourth, the framework is partial-equilibrium. Widespread adoption can change wages, land values, crop prices, input markets, and competition. A technology that expands viability for individual producers can reduce output prices after diffusion and move the equilibrium frontier again. General-equilibrium and distributional effects are outside the scope of this paper.

Fifth, economic viability is not equivalent to social desirability. A technology may expand the privately viable set of agricultural production while generating environmental costs that are not fully reflected in farm-level returns, including greater pressure on water, land, energy, or ecosystems. Conversely, technologies with positive externalities may appear privately unviable despite broader social benefits. The Agricultural Viability Frontier therefore describes an economic decision boundary under the prices, policies, and constraints represented in E; it should not be interpreted as a welfare frontier unless relevant external costs and benefits are explicitly internalized.

Sixth, although the paper is motivated by AI-enabled agriculture, the logic also applies to non-AI technologies. Irrigation, refrigeration, mechanization, and infrastructure can all change agricultural viability boundaries. AI is distinctive in the kinds of constraints it can relax—especially judgment, supervision, prediction, and autonomy—but the framework should be treated as a technology-evaluation lens rather than as a claim that frontier expansion is unique to AI.

9. Conclusion

Agricultural technology is usually judged by how much it improves production. That remains necessary, but it is not always sufficient. A technology can have a second kind of economic effect: it can change whether a production configuration is viable in the first place.

The Agricultural Viability Frontier framework makes that distinction explicit. A production configuration lies inside or outside a viable set conditioned on technology and the economic environment. Optimization improves a configuration already inside the set. Frontier expansion moves a configuration across the viability boundary by materially relaxing a binding constraint while complementary conditions remain sufficient. The Frontier Expansion Test turns that logic into a practical counterfactual: would substantially the same production, expansion, continuation, location, scale, or organizational configuration occur without the technology?

This framing does not replace research on productivity, adoption, redesign, or the extensive margin. It connects them around a different evaluative question. For AI in agriculture, that question matters because the economic value of intelligence and autonomy may appear not only in lower costs, but in production that becomes manageable, staffable, supervisable, or reliable enough to exist. The relevant measure of technological change is therefore sometimes not how much better farming becomes, but how much more farming becomes economically possible.

AI Use Disclosure

OpenAI's ChatGPT was used only for literature discovery and bibliographic verification. The author independently reviewed all cited sources and remains solely responsible for the final manuscript, its arguments, source selection, and citations.

References

Al-Amin, A. K. M. A., Lowenberg-DeBoer, J., Franklin, K., & Behrendt, K. (2023). Economics of field size and shape for autonomous crop machines. Precision Agriculture, 24, 1738–1765. https://doi.org/10.1007/s11119-023-10016-w

Al-Amin, A. K. M. A., Lowenberg-DeBoer, J., Erickson, B. J., Evans, J. T., Langemeier, M. R., Franklin, K., & Behrendt, K. (2024). Economics of strip cropping with autonomous machines. Agronomy Journal, 116, 572–589. https://doi.org/10.1002/agj2.21536

Babcock, B. A. (2015). Extensive and intensive agricultural supply response. Annual Review of Resource Economics, 7, 333–348. https://doi.org/10.1146/annurev-resource-100913-012424

Charlton, D., Devadoss, S., Gallardo, R. K., Luckstead, J., & Vougioukas, S. (2025). Economic viability of robotic fruit harvesters to reduce large seasonal labor demands: Analysis of Gala and Honeycrisp apples. Journal of the Agricultural and Applied Economics Association, 4(1), 70–87. https://doi.org/10.1002/jaa2.70000

Finger, R. (2023). Digital innovations for sustainable and resilient agricultural systems. European Review of Agricultural Economics, 50(4), 1277–1309. https://doi.org/10.1093/erae/jbad021

Gallardo, R. K., & Sauer, J. (2018). Adoption of labor-saving technologies in agriculture. Annual Review of Resource Economics, 10, 185–206. https://doi.org/10.1146/annurev-resource-100517-023018

Gray, J. (2000). Kazakhstan: A review of farm restructuring (World Bank Technical Paper No. 458). World Bank. ISBN 0-8213-4674-1.

Kao, J. (2026). Farming the impossible: How AI is expanding what agriculture can do. ISBN 978-626-01-7158-2.

Khanna, M., Atallah, S. S., Heckelei, T., Wu, L., & Storm, H. (2024). Economics of the adoption of artificial intelligence–based digital technologies in agriculture. Annual Review of Resource Economics, 16, 41–61. https://doi.org/10.1146/annurev-resource-101623-092515

Klerkx, L., & Rose, D. (2020). Dealing with the game-changing technologies of Agriculture 4.0: How do we manage diversity and responsibility in food system transition pathways? Global Food Security, 24, 100347. https://doi.org/10.1016/j.gfs.2019.100347

Kroupová, Z. Ž., Aulová, R., Rumánková, L., Bajan, B., Čechura, L., Šimek, P., & Jarolímek, J. (2025). Drivers and barriers to precision agriculture technology and digitalisation adoption: Meta-analysis of decision choice models. Precision Agriculture, 26, Article 17. https://doi.org/10.1007/s11119-024-10213-1

Kuroiwa, K., Veettil, P. C., & Gupta, I. (2026). Labor scarcity and technology adoption in agriculture: Evidence from rural India during the COVID-19 pandemic. Economic Development and Cultural Change. Advance online publication. https://doi.org/10.1086/741288

Lan, J., & Ban, Q. (2025). The farm-level economic and environmental benefits of precision agriculture technology adoption: A meta-analysis of global evidence. Sustainability, 17(24), 11223. https://doi.org/10.3390/su172411223

Lowenberg-DeBoer, J., Franklin, K., Behrendt, K., & Godwin, R. (2021). Economics of autonomous equipment for arable farms. Precision Agriculture, 22, 1992–2006. https://doi.org/10.1007/s11119-021-09822-x

Maree, S. C., Zhang, P., van Marrewijk, B. M., de Zwart, F., Bijlaard, M., & Hemming, S. (2025). Autonomous greenhouse cultivation of dwarf tomato: Performance evaluation of intelligent algorithms for multiple-sensor feedback. Sensors, 25(14), 4321. https://doi.org/10.3390/s25144321

Papadopoulos, G., Arduini, S., Uyar, H., Psiroukis, V., Kasimati, A., & Fountas, S. (2024). Economic and environmental benefits of digital agricultural technologies in crop production: A review. Smart Agricultural Technology, 8, 100441. https://doi.org/10.1016/j.atech.2024.100441

Shockley, J., Dillon, C., Lowenberg-DeBoer, J., & Mark, T. (2022). How will regulation influence commercial viability of autonomous equipment in US production agriculture? Applied Economic Perspectives and Policy, 44(2), 865–878. https://doi.org/10.1002/aepp.13178

Spykman, O., Lowenberg-DeBoer, J., & Gandorfer, M. (2026). Crop robots as potential enablers of economical and biodiversity-smart small-scale farming. Precision Agriculture, 27, Article 80. https://doi.org/10.1007/s11119-026-10367-0

Strine, J., Fiechter, C., & Lowenberg-DeBoer, J. (2025). The economics of US row crop production with large-scale autonomous machines. Smart Agricultural Technology, 12, 101599. https://doi.org/10.1016/j.atech.2025.101599

Win, M. T., Rutledge, Z., & Maredia, M. K. (2025). Labor shortages and farmer adaptation strategies. Applied Economic Perspectives and Policy, 47(3), 896–913. https://doi.org/10.1002/aepp.13527

Yu, C., Khanna, M., Atallah, S. S., Kar, S., Bagavathiannan, M., & Chowdhary, G. (2024). Herbicide-resistant weed management with robots: A weed ecological–economic model. Agricultural Economics, 55(6), 943–962. https://doi.org/10.1111/agec.12856