

When Noise Becomes Structure: Microbiomes, Projection Failure, and the Recovery of Ecological Organization

Flyxion

Independent Researcher

Abstract

A 2026 commentary in *npj Science of Plants* argues that modern plant breeding programs systematically underestimate a major source of biological variation by failing to account for the root-associated microbiome. The authors propose expanding the classical $G \times E$ framework to a GEM model incorporating genotype, environment, and microbiome as interacting components shaping plant phenotype. This essay takes that proposal as a case study in a more general epistemological and ontological pattern. The epistemological dimension concerns what happens when formal models treat unmeasured causal structure as residual noise, and how that structure is eventually recovered when residuals become too organized to ignore. The ontological dimension concerns a deeper problem that the GEM framework does not fully address: whether the organism was the right unit of analysis to begin with. Drawing on the work of Deleuze and Guattari as interpreted for social theory by Bogard, this essay argues that the plant is not a subject that possesses a microbiome but a temporary stabilization within a larger ecological assemblage in which microbiome, soil, and environment are co-constitutive rather than external. The microbiome did not disappear from breeding trials because it was unimportant. It disappeared because the formal framework treated a derivative unit—the organism—as if it were a primary and self-sufficient substance. The same pattern appears across disciplines whenever formal models impose molar boundaries on systems whose causal organization is fundamentally molecular and processual. A set of mathematical appendices develops the projection, trajectory-basin, and field-theoretic interpretations of these arguments in formal terms.

1 Introduction

In a field trial, a plant grows. It germinates, extends its root system into soil, recruits bacteria and fungi from the surrounding environment, and eventually produces a measurable yield. The breeders recording that yield assign it to the plant's genotype, to the environmental conditions of the site, and to the interaction between the two. This accounting is not arbitrary. It reflects a century of quantitative genetics, a well-developed theoretical apparatus, and genuine predictive success. The $G \times E$ framework has allowed breeders to select for traits, to anticipate which genotypes will perform well across environments, and to improve agricultural productivity substantially.

But the plant growing in that field trial is not simply a genotype responding to an environment. It is an assemblage. The rhizosphere surrounding its roots contains populations of bacteria, archaea, and fungi that alter how it acquires nitrogen and phosphorus, regulate its hormone balance, suppress pathogenic organisms, and modulate its tolerance to drought and salinity (Oyserman et al., 2021). These microbial communities are not incidental to plant function. They are causally active at every level of plant development. When a genotype yields well in one location and poorly in another, some of that variation may reflect differences in local microbial guilds rather than differences in rainfall, temperature, or soil chemistry. The model records the discrepancy as $G \times E$ interaction. The microbiome disappears into the residuals.

A 2026 commentary in *npj Science of Plants* argues that this disappearance is not merely a technical limitation to be corrected by adding sequencing steps to existing protocols (Araujo et al., 2026). It reflects a structural feature of how breeding knowledge has been organized. The argument is worth taking seriously not primarily because it is new—the role of rhizosphere microbiota in plant performance has been studied for decades (Dwivedi et al., 2025; Escudero-Martinez & Bulgarelli, 2023)—but because it identifies a pattern that recurs across scientific disciplines. This essay examines that pattern at two levels. At the epistemological level, the question is what happens when formal models systematically discard organized causal structure by assigning it to residual variance. At the ontological level, the question is deeper: whether the plant was ever the appropriate unit of analysis, or whether the organism is itself an artifact of the projection—a molar simplification of what is, at the level of efficient causation, a molecular and processual assemblage.

2 The Projection and Its Residuals

The classical model of plant phenotype can be stated compactly. Observed performance is partitioned into a genetic component, an environmental component, and a term capturing the degree to which genetic effects depend on environmental context:

$$P = \beta_G G + \beta_E E + \beta_{GE}(G \times E) + \varepsilon$$

The ε term is residual variance: the portion of phenotypic variation that the model does not explain. In practice, ε absorbs everything that the model's variables fail to capture, including measurement error, stochastic developmental processes, and any systematic causal factor that the model has not been constructed to represent. The residual is not empty. It is a holding space for unmeasured structure.

What the GEM proposal identifies is that a substantial portion of ε in standard breeding trials is occupied by microbiome variation (Oyserman et al., 2021; Zancarini et al., 2024). Two field locations may share temperature, precipitation, and soil chemistry in all measured respects while differing substantially in microbial composition. When a genotype performs differently at those two locations, the variance model attributes that difference to G×E interaction. This attribution is not technically incorrect in a narrow sense: the genotype is interacting with its environment, and the microbiome is part of that environment. But the attribution is epistemically misleading in a deeper sense. It treats the microbiome's causal contribution as an unstructured environmental fluctuation rather than as an organized, partially heritable, genotype-responsive feature of the system.

The expanded model proposed by Oyserman et al. adds microbial composition as an explicit term:

$$P = \beta_G G + \beta_E E + \beta_M M + \beta_{GE}(G \times E) + \beta_{GM}(G \times M) + \beta_{EM}(E \times M) + \beta_{GEM}(G \times E \times M) + \varepsilon$$

The GEM framework should therefore be understood as both a correction and a limitation. It corrects a significant omission by recovering microbiome structure from the residuals of the G×E model, and in that sense it represents a genuine epistemic advance. Yet it remains committed to an ontology in which genotype, environment, and microbiome are treated as separable interacting entities. The framework expands the number of boxes; it does not question the box-like structure itself. The deeper lesson of the microbiome literature may be that these categories are themselves derivative abstractions from a more continuous ecological process. Genotype, environment, and microbiome may be useful

decompositions of a coupled system in the way that wave, current, and shoreline are useful decompositions of a coastal system: practically adequate, but not ontologically primitive. In this sense, GEM may represent not the endpoint of the transition the microbiome literature is driving but an intermediate stage in a broader movement from organism-centered explanations toward assemblage- and process-centered accounts of biological organization. The rest of this essay traces the logic of that broader movement.

The more immediate question is why the microbiome remained in ε for so long given that its biological significance was not unknown. The answer begins with the structure of the projection itself and deepens into a question about the ontological commitments embedded in that projection. Appendix A develops the projection interpretation in formal terms; the present section confines itself to the conceptual argument.

3 The Structure of Projection Failure

A model is a projection. It maps a high-dimensional phenomenon onto a lower-dimensional representation, necessarily discarding information in the process (Box, 1976). Whether that discarding is epistemically justified depends on whether the discarded information is random with respect to the questions the model is designed to answer.

When information discarded by a projection is genuinely random, the projection is conservative. It may be imprecise, but it does not introduce systematic error. When discarded information has organized structure that correlates with modeled quantities, the projection introduces bias. Estimates of heritability become inflated or deflated. $G \times E$ interaction effects absorb variance that properly belongs elsewhere. Selection decisions are made on the basis of a distorted representation of what the model is supposed to track. Appendix E states this in the language of fiber collapse: two ecologically distinct states that differ only in their microbiome composition become indistinguishable under the projection, yet diverge in their future trajectories.

The microbiome case fits this pattern precisely. Microbial communities are not randomly distributed across field sites. They have biogeographic structure: they vary with soil type, land use history, climate, and the prior history of what has grown in a location (Revillini et al., 2019). A breeding trial network spanning multiple countries will almost certainly sample environments that differ systematically in microbial composition. When genotype performance is recorded across those environments, some of the variation attributed to environment and to $G \times E$ interaction will in fact reflect microbial effects that are correlated with, but distinct from, the measured environmental variables. The projection has discarded organized structure while treating it as noise.

This kind of failure has a recognizable phenomenology. It proceeds in stages. First,

the model achieves sufficient predictive accuracy to become institutionalized. The $G \times E$ framework is not a rough approximation adopted provisionally; it is embedded in breeding protocols, funding structures, statistical software, and the tacit knowledge of practitioners. Second, anomalies accumulate. Genotypes fail to replicate their performance across environments in ways that the model cannot explain. Heritability estimates vary across studies in ways that seem to exceed what measurement error alone would produce. The residuals develop structure, but that structure is not legible within the model's own vocabulary. Third, a new variable is proposed, positioned as an extension of the existing framework rather than as a replacement, because the existing framework has too much institutional mass to be abandoned. Fourth, the field debates whether the new variable represents genuine causal structure or can be reduced to existing variables with sufficient measurement precision.

The microbiome case is currently in the third and fourth stages. But examining only the epistemological dimension—only the question of what was measured and what was not—misses a deeper level of the problem. The residuals did not merely contain a missing variable. They contained evidence that the primary unit of the model—the organism—was itself a simplification of something more complex.

4 The Wrong Unit: Bodies, Assemblages, and Molar Segmentation

The GEM commentary introduces the concept of the extended plant phenotype when it argues that microbiome recruitment may itself be heritable. Root architecture, root hair density, and root exudate chemistry influence which microbes colonize the rhizosphere (Zancarini et al., 2024). These root traits are genetically controlled. Therefore, different plant genotypes will systematically recruit different microbial communities under comparable soil conditions. The microbiome assembled by a plant is not an independent environmental factor but a partially genotype-dependent phenotypic expression.

This argument is presented as an extension of the classical phenotype concept: the plant has a microbiome the way it has a root architecture. But the logic of the argument actually undermines the classical concept rather than extending it. If the plant's genetic program determines what microbial communities it recruits, and if those microbial communities in turn determine much of the plant's functional performance, then the causal boundary between the plant and its microbiome is not a clean line but a reciprocal entanglement. The plant is partly constituted by the microbiome it recruits. The microbiome is partly constituted by the plant that recruits it. Neither is fully prior to or independent of the

other.

This problem has been recognized in biological theory under the heading of the holobiont. Gilbert, Sapp, and Tauber argued that the individual organism, long treated as the fundamental unit of biological identity and selection, is more accurately understood as a symbiotic assemblage whose functional capacities depend on microbial partners throughout development and across generations (Gilbert et al., 2012). Bosch and McFall-Ngai extended this argument to propose the metaorganism as the relevant unit of evolutionary analysis: a composite entity whose physiology, immunity, and development are constituted through ongoing interactions with its associated microbial communities (Bosch & McFall-Ngai, 2014). The heritability of microbiome composition has since been demonstrated across multiple host systems, indicating that the symbiotic assemblage is not merely a phenotypic response to environment but a partially heritable feature of the host-microbiome system (Morris & Bohannan, 2024).

This is where the work of Deleuze and Guattari, as interpreted for social theory by Bogard, becomes relevant. The problem being encountered in plant biology—the problem of where to draw the boundary around the relevant causal unit—is precisely the problem that Deleuze and Guattari address in their treatment of bodies and assemblages. For Deleuze and Guattari, bodies are not primary substances with fixed identities. They are heterogeneous composites whose components enter into composition with other bodies, both living and nonliving, to form variable assemblages (Bogard, 1998). The subject is not the ground of its own constitution but an effect: a composition of organs, membranes, nerves, and physiochemical processes, but also tools, means of nourishment and transport, and the materials of production and consumption (Bogard, 1998).

Transposed to plant biology, the implication is not merely that the plant has a microbiome that should be measured. The implication is that the plant-as-organism is already a derived and partial stabilization of a more primary ecological assemblage. What breeders have been selecting is a molar unit—the individual plant, bounded and measurable—that is itself produced by and dependent upon molecular processes that extend well beyond the organism's physical boundary.

Bogard follows Deleuze and Guattari in distinguishing between molar and molecular levels of organization (Bogard, 1998). Molar segments are the large-scale, institutionally recognized divisions that scientific practice operates with: organisms, species, genotypes, environments. Molecular processes are the more primary flows, forces, and connections from which molar forms are produced and by which they are continuously sustained. The important point is that molar forms do not reduce to molecular processes, but neither are they independent of them (Anderson, 1972). A molar segment is a crystallization

of molecular dynamics, and it constantly flows and leaks back into those dynamics. Breeding programs work at the molar level. They select individual organisms, measure their outputs, and infer the genetic contributions to those outputs. This is productive work. But it systematically underestimates the degree to which the molar unit it works with is constituted by molecular processes that are not intrinsic to any individual. The rhizosphere microbiome is one of those processes. It is not a property of the plant in the way that root architecture is a property of the plant. It is a relation between the plant and the surrounding ecological field. When that relation is ignored, the breeding program is not merely missing a variable; it is treating a relational, field-dependent entity as if it were a self-sufficient substance.

Appendix D develops a network-theoretic formalization of this segmentarity argument, representing the plant-soil system as a graph whose nodes include plant tissues, microbial taxa, nutrient pools, and chemical signals, and whose edges represent exchange relations among them.

5 Why Omitted Structure Accumulates: Institutions and Over-coding

There is a recognizable reason why models tend to omit causally important structure and then resist correcting that omission even after the structure's importance becomes apparent. The reason is that models succeed before they fail, and that success generates institutional commitments that outlast the conditions that justified them.

The $G \times E$ framework succeeded. It captured enough of the causal structure of plant performance to allow productive breeding. Crop yields improved. Genotypes were selected. The framework generated returns on investment for the institutions that adopted it, which gave those institutions reasons to continue using it and to train practitioners in its methods. By the time the limitations of the framework became apparent—by the time the residuals became so structured that they could not plausibly be attributed to random noise—the framework had accumulated enormous institutional and epistemic inertia.

The Deleuzian analysis offers a more precise diagnosis than institutional inertia alone. The $G \times E$ framework did not merely miss the microbiome as a variable. It actively produced the organism as a molar unit of selection, and then reproduced that unit through the institutional practices organized around it. Field trial protocols presuppose the plant as the relevant unit. Statistical frameworks presuppose the plant as the unit of observation. Breeding decisions presuppose the plant as the unit of selection. These presuppositions do not just reflect a pre-existing scientific consensus; they continuously reconstitute the

organism as the natural and obvious unit of analysis, making alternatives difficult to think within the framework's own vocabulary.

This is what Bogard, following Deleuze and Guattari, calls overcoding: the tendency of rigid segmentary systems to saturate their objects with signification, to impose their categorical boundaries as if they were natural boundaries of the phenomena themselves (Bogard, 1998). The organism is not simply a convenient unit of measurement. It is a molar segmentation that breeders have inscribed onto a more fluid ecological process—a cut that interrupts and channels flows of causal influence rather than capturing them in their native organization.

The parallel with other disciplines is not incidental. In economics, the individual household or firm is treated as the relevant unit for many purposes, which causes unpaid labor performed outside market transactions to disappear from the model's accounts. The labor is real and economically consequential, but it does not have a market price, making it invisible to frameworks organized around price signals (Waring, 1988). In cognitive science, the brain or organism is treated as the relevant unit of cognition, causing the causal contributions of body, environment, and material scaffolding to be absorbed into context effects or error terms. In both cases, the unit of analysis was not chosen because it captures a natural boundary in the phenomenon but because it is a tractable surface on which measurements can be made. The choice then generates its own institutional reproduction, making alternative framings progressively more difficult to adopt.

6 The Boundary Problem and the Trajectory Basin

The holobiont concept that appears in the microbiome literature gestures toward a response to this problem. If the plant plus its microbial partners constitutes a functional unit, then selection should target the holobiont rather than the individual organism (Gilbert et al., 2012; Bosch & McFall-Ngai, 2014). This reframing is productive, but it requires confronting a difficulty that the Deleuzian analysis makes explicit.

The organism has a tractable boundary. It can be harvested, weighed, genotyped, and observed. The holobiont does not have a boundary of the same kind. It extends into the soil, the rhizosphere, the seasonal succession of microbial communities, and the history of what has previously grown in a site. It fades into ecological context rather than terminating cleanly. Attempting to bound it operationally—to define the holobiont as the plant plus its microbiome, where the microbiome is the metagenome recovered from root samples at a particular developmental stage—is useful for practical purposes but reintroduces molar segmentation at a higher level of organization. The holobiont is a less distorting molar unit than the organism, but it is still a molar unit.

The Deleuzian response to this situation is not to propose a better boundary but to shift the question (Latour, 2005). Rather than asking where to draw the line around the causal unit, the question becomes: what are the flows, segments, and processes from which the apparent stability of the unit is produced? For Deleuze and Guattari, every structure is simultaneously an event (Bogard, 1998). What appears to be a stable object is a temporary stabilization of processual dynamics, and the dynamics are primary.

Appendix B formalizes this shift using the notion of an admissible trajectory basin. Rather than treating the phenotype as a single outcome, it treats it as a basin of viable trajectories in the full plant-soil-microbiome state space (Kauffman, 1993; Levin, 1998). A genotype is robust not because it produces a high yield under favorable conditions, but because the assemblage it participates in preserves viable developmental trajectories across a wide range of ecological perturbations. The microbiome matters not because it is a variable that affects phenotype, but because it modifies the topology of the trajectory basin: enlarging it by increasing nutrient access and suppressing pathogens, or contracting it by competing for resources and facilitating pathogen colonization. Selection, on this account, acts on the accessibility structure of phenotypic possibility rather than on isolated trait values.

This has a direct implication for breeding. A breeding program that targets the holobiont as its unit of selection is an improvement over one that targets the individual organism. But a breeding program that asks which ecological configurations produce robust and productive plant-microbiome assemblages across soil histories and environmental perturbations is asking a qualitatively different and more adequate question. It treats stability and productivity not as properties of organisms or even holobionts but as features of ecological trajectories: patterns of change that are either accessible or inaccessible given particular configurations of the coupled plant-soil-microbe field.

7 Recovery Rather Than Discovery

It is worth being precise about what the GEM commentary actually claims. It does not claim that plant breeders were unaware of soil microbiology. It does not claim that rhizosphere ecology was unknown or that the importance of nitrogen-fixing bacteria was a secret. The biological knowledge existed. What was missing was the integration of that knowledge into the formal quantitative framework used to make breeding decisions.

This distinction matters because it identifies where the problem resides. The problem is not primarily a knowledge deficit; it is an integration deficit, and behind that, an ontological commitment that shaped what could be integrated. The microbiome was causally present in every field trial that recorded $G \times E$ interaction. It contributed to every

yield measurement. It was not absent from the phenomenon; it was absent from the model. And because it was absent from the model, it was absent from the decision process. Breeders could not select for microbiome recruitment capacity because their statistical apparatus had no variable capable of representing that capacity. They could not identify genotypes whose performance was microbiome-contingent rather than robustly self-sufficient because their framework treated the microbiome as part of the environmental background rather than as a co-constitutive element of the system being selected.

The microbiome case is therefore a recovery rather than a discovery. A discovery is the identification of something new. A recovery is the recognition of something that was always there but excluded from formal representation by the architecture of the framework used to study it. The history of science contains many recoveries that were initially framed as discoveries, because the framework that excluded the structure in question had no vocabulary for describing the exclusion itself. Solow's growth residual, which was initially interpreted as a measure of technological progress, was later understood to absorb unmeasured variation in human capital, institutional quality, and knowledge spillovers: organized structure, not noise, excluded by the model's choice of variables (Solow, 1957). The social determinants of health spent decades in the residuals of epidemiological models focused on individual behavioral risk factors before acquiring sufficient causal specificity to inform policy. The pattern is the same in each case: omitted structure accumulates in residuals, residuals develop recognizable organization, and eventually the framework must expand its boundary or be replaced.

This is what Bogard, following Deleuze and Guattari, would describe as the tendency of representational frameworks to mistake their own boundaries for the boundaries of what they represent (Bogard, 1998). When a model organizes itself around molar units, it does not merely simplify the phenomena. It produces those units as the natural objects of inquiry, making it progressively more difficult to ask whether a different organizational level might be more causally adequate. The microbiome's invisibility in breeding programs was not a failure of attention but a structural consequence of the molar segmentation that constituted the breeding framework in the first place.

A complementary lens on this invisibility comes from the decision-science literature on noise. Kahneman, Sibony, and Sunstein distinguish between bias, which refers to systematic deviation in a consistent direction, and noise, which refers to unwanted variability that the model leaves unexplained (Kahneman et al., 2021). A central lesson of their analysis is that apparent noise may sometimes indicate hidden structure rather than irreducible randomness: what appears as scatter is often the signature of unmeasured factors operating in a patterned way. The contemporary microbiome literature suggests

exactly this possibility within plant breeding. Variability previously attributed to environmental noise may, in significant part, arise from structured differences in rhizosphere microbial communities. The challenge is therefore not merely statistical but ontological: determining whether unexplained variance reflects genuine randomness or the omission of causally relevant system components.

From this perspective, the GEM framework represents more than the addition of a new explanatory variable. It represents a reduction in epistemic noise through expansion of the system boundary. What was previously treated as residual variance becomes partially legible once microbiome dynamics are incorporated into the model. The resulting improvement is not simply predictive accuracy but a refinement of what counts as the relevant unit of analysis—a boundary correction rather than a parameter correction.

The history of plant breeding thus illustrates a recurring scientific pattern: models initially succeed by compressing complex systems into tractable variables; unexplained variance accumulates within residual categories; and subsequent advances arise not from more sophisticated statistics alone but from identifying previously hidden sources of organization within those residuals. The microbiome is a contemporary example of a broader transition from treating variability as noise toward treating it as structured ecological information. Whether breeding programs can respond adequately to that transition depends not only on acquiring new measurement tools but on revising the ontological commitments that determined what counted as signal and what counted as background in the first place.

8 Implications and Conclusion

For plant breeding specifically, the GEM framework suggests several reorientations. The most immediate is measurement: systematic characterization of rhizosphere microbiomes in breeding trials, linked to genotype performance data, would allow the confounding of microbial and environmental effects to be partially disentangled (Week et al., 2025; Mueller & Linksvayer, 2022). The more consequential reorientation is conceptual. If the plant's capacity to recruit beneficial microbial partners is a heritable trait, then that capacity is a legitimate target of selection. Breeders could in principle select for genotypes that reliably associate with nitrogen-fixing bacteria in low-input environments, that recruit disease-suppressive microbial consortia under pathogen pressure, or that maintain rhizosphere microbiome composition under drought stress (Zancarini et al., 2024; Dwivedi et al., 2025).

The deeper reorientation—one that the Deleuzian analysis points toward but the GEM framework does not fully articulate—is to treat the breeding target not as an organism with a microbiome but as an ecological assemblage whose trajectories of development and

production are more or less robust across soil histories, microbial successions, and environmental perturbations. Appendix C develops this reorientation using a field-theoretic notation in which genotype, environment, and microbiome are treated not as independent variables entering an additive equation but as constraints on the evolution of coupled scalar, vector, and accessibility fields whose joint dynamics determine the range of viable plant phenotypes.

The broader implication is one that recurs across the disciplines touched in this essay. Systems whose causal organization is molecular and processual—whose relevant dynamics operate across the boundaries of the molar units that formal models use to represent them—will generate residuals that carry structure, and that structure will eventually force a confrontation with the adequacy of the model’s representational commitments. The appropriate response is not simply to add more variables to the existing framework. It is to ask whether the framework’s choice of unit, and the ontological commitments embedded in that choice, are adequate to the phenomenon the framework is supposed to represent. In many of the most important cases, the answer will be that the unit was chosen for tractability rather than for causal adequacy, and that the accumulated residuals are the system’s way of announcing as much.

The microbiome did not hide from plant breeding. The breeding framework hid the microbiome from itself, and from the organisms it was trying to improve, by insisting that the organism was the right thing to see.

Appendix A: GEM as a Projection of a Higher-Dimensional System

Let P denote an observed plant phenotype—yield, drought tolerance, biomass accumulation, or disease resistance. In the conventional breeding model, phenotype is a function of genotype and environment:

$$P = \beta_G G + \beta_E E + \beta_{GE}(G \times E) + \varepsilon$$

The GEM extension adds the microbiome state M :

$$P = \beta_G G + \beta_E E + \beta_M M + \beta_{GE}(G \times E) + \beta_{GM}(G \times M) + \beta_{EM}(E \times M) + \beta_{GEM}(G \times E \times M) + \varepsilon$$

Both models remain additive in form. From a constraint-field perspective, the deeper issue is that the measured phenotype is a projection of a higher-dimensional ecological state. Let \mathcal{X} be the full plant-soil-microbiome-environment system. Then the observed phenotype is a functional:

$$\pi : \mathcal{X} \rightarrow P$$

The traditional G×E model implements a lossy compression:

$$\pi_{GE} : \mathcal{X} \rightarrow (G, E, P)$$

The GEM model implements a less lossy compression:

$$\pi_{GEM} : \mathcal{X} \rightarrow (G, E, M, P)$$

The microbiome appears as “missing variance” because the original projection discarded structure that remained causally active in the territory. Moving from π_{GE} to π_{GEM} is an improvement in representational adequacy, but it does not change the projective character of the modeling strategy. Some organized structure remains in ε at every level of the GEM hierarchy; the question is only how much.

Appendix B: Phenotype as an Admissible Trajectory Basin

Let \mathcal{X} be a high-dimensional state space whose coordinates include plant genotype, soil chemistry, water availability, microbial composition, root architecture, nutrient flows, pathogen pressure, and developmental history. A plant phenotype is not a single outcome but a basin of admissible trajectories:

$$\mathcal{A}(x_0) = \{ x(t) \in \mathcal{X} : x(t) \text{ remains viable under given constraints} \}$$

where x_0 is the initial state of the plant-soil-microbiome system. Selection acts not directly on isolated traits but on the accessibility structure of this basin. A genotype is robust when it preserves viable trajectories across changing environments:

$$\text{Robustness}(G) \approx \mu(\mathcal{A}_G \text{ across } E \text{ and } M)$$

where μ denotes a measure over the relevant range of ecological perturbations.

The microbiome matters because it changes the shape of \mathcal{A} . Beneficial microbial communities enlarge the set of viable trajectories by increasing nutrient access, suppressing

pathogens, improving water uptake, and stabilizing stress responses. Detrimental communities contract it. In this sense, the microbiome does not merely affect the phenotypic outcome at a given time point. It modifies the topology of possible phenotypes over developmental time.

Selection programs that optimize instantaneous yield are selecting on a single point within \mathcal{A} . Selection programs that optimize robustness are selecting on the measure of \mathcal{A} itself—a qualitatively different and more ecologically adequate objective.

Appendix C: RSVP Field Interpretation of Plant-Microbiome Systems

The Relativistic Scalar-Vector Plenum (RSVP) framework represents physical and ecological systems using three coupled fields defined over a spatial domain Ω and time t :

$$\begin{aligned}\Phi(\mathbf{x}, t) & \text{ scalar resource field} \\ \mathbf{v}(\mathbf{x}, t) & \text{ vector transport and exchange field} \\ S(\mathbf{x}, t) & \text{ accessibility or entropy field}\end{aligned}$$

In the plant-soil-microbiome context, Φ represents material distributions: nutrients, water, biomass, carbon compounds, and chemical signals. The vector field \mathbf{v} represents directed flows: root uptake, microbial exchange, fungal transport, exudation gradients, and pathogen movement. The field S represents the local range of available developmental or ecological trajectories—the accessibility of future states given the current configuration.

The observed phenotype is then a functional of the coupled field history:

$$P = \Pi[\Phi(\mathbf{x}, t), \mathbf{v}(\mathbf{x}, t), S(\mathbf{x}, t)]$$

Rather than treating genotype, environment, and microbiome as independent additive variables, this formulation treats them as constraints on field evolution:

$$\begin{aligned}G & \rightarrow \text{boundary conditions on } \Phi \text{ and } \mathbf{v} \\ E & \rightarrow \text{external forcing conditions} \\ M & \rightarrow \text{local transport and transformation operators on all three fields}\end{aligned}$$

The microbiome is especially important because it affects all three fields simultaneously. It changes the distribution of resources (Φ), redirects material flows (\mathbf{v}), and alters the accessibility of future plant developmental states (S). An additive model of the form $P = G + E + M + \dots$ cannot capture this simultaneity, because it treats the three components as contributing independently to a scalar outcome rather than as jointly shaping the field geometry within which the phenotype develops.

The RSVP interpretation also clarifies why trajectory robustness (Appendix B) is a more natural optimization target than instantaneous yield: robustness corresponds to the stability of the accessible region under perturbations of the field boundary conditions, which is a property of the field geometry rather than of any single field value.

Appendix D: Segmentarity and Microbiome Recruitment

Following a Deleuzian-Guattarian interpretation, the plant is not a closed organism but an assemblage formed through segmented flows. Let the plant-soil system be represented as a network:

$$\mathcal{N} = (V, E)$$

where V includes plant tissues, root zones, microbial taxa, nutrient pools, water channels, and chemical signals, while E represents exchange relations among them. Segmentarity is a partition of this network into functional zones:

$$\sigma : \mathcal{N} \rightarrow \{S_1, S_2, \dots, S_n\}$$

Each segment S_i corresponds to a functional region: root hair zone, rhizosphere community, fungal interface, nitrogen-fixing guild, pathogen-suppressive consortium, or water-access pathway.

A phenotype emerges when these segments are coordinated into a stable functional assemblage:

$$P \approx C(S_1, S_2, \dots, S_n)$$

where C is a coherence operator measuring whether the segments jointly support a viable developmental trajectory. Microbiome recruitment is one of the mechanisms by which the assemblage segments, stabilizes, and maintains itself over developmental time.

This formulation makes explicit what the molar organism concept obscures: the plant is not the agent that recruits a microbiome. It is a partially stabilized sub-network within

a larger assemblage whose coherence depends on the maintenance of exchange relations across the full network, including the microbial nodes that the organism-centric model treats as environmental background.

Appendix E: Projection Failure and Hidden Structure

Let \mathcal{X} be the full ecological system and let \mathcal{Y} be the reduced model used in breeding:

$$\rho : \mathcal{X} \rightarrow \mathcal{Y}$$

In the traditional breeding model, $\mathcal{Y} = (G, E, P)$. If the microbiome is omitted, two distinct ecological states may become indistinguishable under the projection:

$$x_1 \neq x_2 \quad \text{but} \quad \rho(x_1) = \rho(x_2)$$

This fiber collapse means that two systems with different microbial structures appear identical in the reduced model. Projection failure occurs when the omitted difference changes future outcomes:

$$\rho(x_1) = \rho(x_2) \quad \text{but} \quad P_{\text{future}}(x_1) \neq P_{\text{future}}(x_2)$$

The microbiome reveals a failure of representational adequacy. The model did not merely lack data; it collapsed distinct causal structures into the same category, making the subsequent divergence appear as unpredictable environmental noise rather than as the consequence of a systematic and measurable difference in ecological state.

The mathematical form of this failure is related to, though distinct from, the problem of unobserved confounders in causal inference. In the standard confounder setting, an omitted variable biases coefficient estimates. In the fiber-collapse setting, two systems that differ only in the omitted dimension are literally indistinguishable to the model, so that any subsequent divergence between them is necessarily attributed to stochastic error. The microbiome is not merely an omitted confounder. It is an omitted dimension whose omission collapses the causal fiber of the system—eliminating, by construction, the model’s ability to represent a difference that turns out to be causally decisive.

References

- Araujo, A. S. F., Costa, R. M., & Lopes, A. C. A. (2026). Plant breeding trials should include the belowground microbiome. *npj Science of Plants*, 2, 15.
- Bogard, W. (1998). Sense and segmentarity: Some markers of a Deleuzian-Guattarian sociology. *Sociological Theory*, 16(1), 52–74.
- Bosch, T. C. G., & McFall-Ngai, M. J. (2014). Metaorganisms as the new frontier. *Zoology*, 117(4), 185–186.
- Deleuze, G. (1990). *The Logic of Sense*. Columbia University Press.
- Deleuze, G., & Guattari, F. (1977). *Anti-Oedipus: Capitalism and Schizophrenia*. University of Minnesota Press.
- Deleuze, G., & Guattari, F. (1987). *A Thousand Plateaus: Capitalism and Schizophrenia*. University of Minnesota Press.
- Dwivedi, S. L., et al. (2025). Exploitation of rhizosphere microbiome biodiversity in plant breeding. *Trends in Plant Science*, 30, 1033–1045.
- Escudero-Martinez, C., & Bulgarelli, D. (2023). Engineering the crop microbiota through host genetics. *Annual Review of Phytopathology*, 61, 257–277.
- Gilbert, S. F., Sapp, J., & Tauber, A. I. (2012). A symbiotic view of life: We have never been individuals. *Quarterly Review of Biology*, 87(4), 325–341.
- Kahneman, D., Sibony, O., & Sunstein, C. R. (2021). *Noise: A Flaw in Human Judgment*. Little, Brown Spark.
- Anderson, P. W. (1972). More is different. *Science*, 177(4047), 393–396.
- Box, G. E. P. (1976). Science and statistics. *Journal of the American Statistical Association*, 71(356), 791–799.
- Holland, J. H. (1992). Complex adaptive systems. *Daedalus*, 121(1), 17–30.
- Kauffman, S. A. (1993). *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press.
- Latour, B. (2005). *Reassembling the Social: An Introduction to Actor-Network-Theory*. Oxford University Press.

- Levin, S. A. (1998). Ecosystems and the biosphere as complex adaptive systems. *Ecosystems*, 1(5), 431–436.
- Morris, A. H., & Bohannan, B. J. M. (2024). Estimates of microbiome heritability across hosts. *Nature Microbiology*, 9, 3110–3119.
- Mueller, U. G., & Linksvayer, T. A. (2022). Microbiome breeding: Conceptual and practical issues. *Trends in Microbiology*, 30, 997–1011.
- Oyserman, B. O., Cordovez, V., Flores, S. S., Leite, M. F. A., Nijveen, H., Medema, M. H., & Raaijmakers, J. M. (2021). Extracting the GEMs: Genotype, environment, and microbiome interactions shaping host phenotypes. *Frontiers in Microbiology*, 11. <https://doi.org/10.3389/fmicb.2020.574053>
- Revillini, D., Wilson, G. W. T., Miller, R. M., Lancione, R., & Johnson, N. C. (2019). Plant diversity and fertilizer management shape the belowground microbiome of native grass bioenergy feedstocks. *Frontiers in Plant Science*, 10. <https://doi.org/10.3389/fpls.2019.01018>
- Solow, R. M. (1957). Technical change and the aggregate production function. *Review of Economics and Statistics*, 39(3), 312–320.
- Waring, M. (1988). *If Women Counted: A New Feminist Economics*. Harper & Row.
- Week, B., et al. (2025). Quantitative genetics of microbiome-mediated traits. *Evolution*, 79, 2487–2502.
- Zancarini, A., Le Signor, C., Terrat, S., Aubert, J., Salon, C., Munier-Jolain, N., & Mougel, C. (2024). *Medicago truncatula* genotype drives the plant nutritional strategy and its associated rhizosphere bacterial communities. *New Phytologist*, 245(2), 767–784. <https://doi.org/10.1111/nph.20272>