

Biological Faith Systems: How Living Systems Commit Before Certainty

Jeremy C. Jones (ORCID 0009-0007-2515-3774)—HoldingLight LLC

contact@universalcollapse.com

© 2026 | CC BY 4.0

Version: v1.0—Prepared 2026-03 universalcollapse.com

Abstract

Living systems routinely act before outcomes are known. Physiologists describe anticipatory regulation and allostatic adjustment; immunologists study early threat detection and trained memory; evolutionary ecologists analyze bet-hedging and phenotypic diversification; microbiologists show how robust adaptive control emerges from simple network architectures. Each of these literatures documents mechanisms by which organisms commit to action under irreducible uncertainty. Yet the shared structural pattern is rarely named as such, because each subfield describes it in domain-local terms.

This paper argues that commitment under uncertainty is not one adaptive strategy among many but a recurrent structural requirement of viability. Any self-maintaining system operating under changing conditions must instantiate mechanisms that bias action before certainty is available and correct that action through feedback afterward. To name this requirement, the paper introduces Biological Faith Systems: the distributed, embodied, feedback-correctible architecture by which living systems proceed when proof arrives too late. The term does not import religious or metaphysical content; it marks the structural fact that action must precede justification in systems that cannot afford to wait.

The framework is grounded across five mechanism families—allostatic regulation, stress-response cascades, immune anticipation, bet-hedging, and robust microbial adaptation—and positioned relative to active inference, autopoiesis, and allostasis. Three testable predictions are stated: a recurrent timing asymmetry in which action precedes resolution, constraint-sweep hysteresis as a residue of commitment, and phylogenetic depth of commitment architectures below cognition and nervous systems. If these predictions fail, the framework should be revised or abandoned. If they hold, biology has a structural lens for seeing what its subfields have long studied in pieces.

Keywords: biological faith systems, commitment under uncertainty, viability, allostasis, active inference, autopoiesis, anticipatory regulation, bet-hedging, hysteresis

1. Introduction: The Fragmentation Problem

Living systems routinely act before outcomes are known. Animals begin digestive and metabolic adjustments from sensory food cues before nutrients have entered the bloodstream; innate immune systems mobilize from recurrent molecular patterns before pathogen identity is fully resolved; clonal populations facing fluctuating environments can diversify phenotypically before any one cell can know which state will prove advantageous; and bacteria navigating chemical gradients rely on adaptive regulatory loops that preserve responsiveness despite noise and parameter variation. None of this is exotic. It is ordinary biology. Yet the very familiarity of the pattern has helped obscure a deeper question: why does life so often proceed before certainty? The uncertainty at issue is not merely missing information that could, in principle, be eliminated by waiting; it is practical and temporal, rooted in the fact that viability-relevant conditions continue to shift while evidence is still being gathered.

In physiology, this problem appears most clearly in the shift from classical homeostatic language toward predictive regulation. Sterling (2012) argues that regulation is not fundamentally about preserving constancy, but about continually adjusting the internal milieu in ways that promote survival and reproduction, and that efficient regulation therefore requires anticipation rather than feedback alone. Ramsay and Woods (2014), while critical of the conceptual looseness that has accumulated around allostasis, nonetheless reinforce the same broad lesson: biological regulation is best understood by

tracking the coordinated activity of multiple interacting loops, not by imagining one variable passively restored to a fixed set-point after disturbance. Power and Schulkin's (2008) review of cephalic-phase responses makes the point concrete. Sensory food cues trigger anticipatory digestive and metabolic adjustments before nutrients are absorbed, improving efficiency and shaping subsequent feeding. Physiology, in other words, already knows that viable regulation often begins before full confirmation.

Immunology tells a closely related story in a different vocabulary. Janeway and Medzhitov (2002) describe the innate immune system as an evolutionarily ancient defense architecture triggered by pattern-recognition receptors that detect conserved molecular signatures shared across classes of microorganisms. That is already a mode of action under uncertainty: the system does not wait for exhaustive specification of the invader but responds on the basis of recurrent structural cues. The newer literature on trained immunity (Netea et al., 2016) deepens the point by arguing that innate immune cells can exhibit a form of memory through durable changes in gene expression and cell physiology, including in organisms lacking adaptive immunity. And Pradeu and Vivier's (2016) discontinuity theory pushes the framing further still, treating immunity as a change-detection system that responds to sudden shifts in antigenic stimulation while tolerating slow or continuous exposure. Immunology, too, is therefore full of mechanisms by which living systems commit before certainty and recalibrate afterward.

Evolutionary ecology and microbiology reveal the same pattern at a different organizational scale. Kussell and Leibler (2005) show that, in fluctuating environments, clonal populations may adapt either by sensing and responding or by generating diversity through stochastic phenotype switching, with switching favored when environmental change is sufficiently infrequent. Donaldson-Matasci, Lachmann, and Bergstrom (2008) generalize the point by treating phenotypic diversity itself as an adaptation to environmental uncertainty. On this view, a population need not wait for better information in order to reduce vulnerability; it can distribute risk across phenotypes in advance. The familiar vocabulary here is bet-hedging, diversification, or stochastic switching rather than anticipation or prediction, but the temporal structure is strikingly similar. What is being protected is continued viability under conditions in which certainty would arrive too late to do all the needed work.

At the microbial and control-theoretic end of biology, the same logic becomes visible in an especially stripped-down form. Barkai and Leibler (1997) proposed a mechanism for robust adaptation in simple biochemical networks and showed that it applies to bacterial chemotaxis. Alon and colleagues (1999) then demonstrated experimentally that, although some response properties vary with substantial changes in protein concentrations, the precision of adaptation remains robust, consistent with the architecture of the network itself. Yi and colleagues (2000) sharpened the result by arguing that robust perfect adaptation in chemotaxis can be understood through integral feedback control. Long before these molecular examples, cybernetics had already articulated a broader principle: Ashby's (1956) law of requisite variety holds that only variety in the regulator can reduce the variety introduced by disturbance, and Conant and Ashby (1970) later argued that any regulator that is maximally successful and simple must, under broad conditions, be isomorphic with the system it regulates. Taken together, these lines of work show that regulation in living systems is inseparable from organized selectivity under uncertainty.

Taken separately, these literatures are not confused. Each has its own scale, its own evidential standards, and its own explanatory vocabulary. Physiology speaks of allostasis and feedforward regulation; immunology of recognition, tolerance, and memory; evolutionary ecology of diversification and bet-hedging; microbiology of robust adaptation and control architecture. The problem, then, is not that

biology lacks mechanisms. The problem is that the mechanisms are usually described in domain-local terms, so that the shared structural pattern can disappear into the success of the local account. What becomes less explicit is that each of these literatures is grappling with a similar temporal asymmetry: living systems often must enter action before outcome-relevant uncertainty is fully resolved, and must preserve themselves by remaining corrigible afterward.

The question of this paper is therefore not whether organisms anticipate, diversify, regulate, or learn. Biology clearly shows that they do. The question is whether these phenomena are best understood as a set of separate adaptive strategies, each local to its own subfield, or as heterogeneous expressions of a deeper structural requirement. Is commitment under uncertainty something living systems happen to evolve in many places, or is it demanded by the very conditions of viable self-maintenance under change?

This paper argues for the second view. Commitment under uncertainty is not one adaptive tactic among many. It is a structural requirement of viability. Any system that must preserve its own organization while conditions continue to shift must possess mechanisms that allow it to bias action before certainty is complete and to correct that action by feedback afterward. Later sections introduce **Biological Faith Systems** as the name for the distributed, embodied, feedback-corrigible architecture by which living systems satisfy that requirement. The claim is not that all organisms use the same mechanism, nor that one formalism can simply subsume the entire field. The claim is that a recurring structural feature of life has been visible in fragments and now needs to be named as such. The broader cross-domain motivations for this terminology are not defended here; in the present paper, BFS is introduced and assessed strictly as a biological framework.

The argument proceeds in five steps. Section 2 develops the structural case that waiting for certainty is not a neutral baseline for living systems but a recurrent route to failure. Section 3 formally defines Biological Faith Systems and defends the terminology. Section 4 grounds the framework across several mechanism families, from physiology and immunology to bet-hedging and microbial adaptation. Section 5 positions BFS relative to active inference, autopoiesis, and allostasis. Section 6 then states three empirical predictions by which the framework can be assessed. The contribution is not a new molecular mechanism or a new physical principle. It is a structural lens that makes a recurring pattern in biology more visible, more nameable, and more testable.

2. The Structural Argument

Any structural account of life must begin with time. Living systems do not regulate from outside the processes that threaten them; they regulate while metabolizing, repairing, sensing, and exchanging matter with surroundings that continue to change during the act of regulation itself. Cybernetics captured one side of this difficulty in the law of requisite variety: disturbance cannot be reduced unless the regulator can answer it with sufficient internal organization. Conant and Ashby sharpened the point by arguing that an effective and simple regulator must, in a relevant sense, model the system it regulates. For living systems, then, the problem is not merely that the world is complex. It is that viability must be preserved in real time under variation that cannot be exhaustively sampled before action becomes due (Ashby, 1956; Conant & Ashby, 1970; Di Paolo, 2005).

This makes it important to distinguish stability from viability. A stable structure can persist under relatively fixed conditions without actively responding to change. A viable system, by contrast, must

preserve its organization through changing conditions by means of sensitivity, feedback, repair, and plasticity. That distinction matters because it shows why life cannot be understood as mere persistence. What must be explained is not simply why organisms endure, but how they remain coherent while the conditions of endurance are continually shifting. The central question is therefore not only how a system returns to a prior state after disturbance, but how it stays available for further regulation at all. This is the point at which uncertainty becomes a constitutive feature of life rather than an incidental obstacle.

The uncertainty at issue here is not simply ignorance that could, in principle, be eliminated by waiting for more information. In living systems, the relevant variables are themselves moving targets: nutrient gradients change, tissues degrade, pathogens replicate, temperatures drift, and opportunities for correction narrow while the organism is still gathering evidence. The system is embedded in the very process it must regulate. Under such conditions, uncertainty is practical, temporal, and often structural. A self-maintaining system can improve its estimate of what is happening, but it cannot suspend the demands of viability while that estimate approaches completion. To wait is already to intervene, because the costs of delay accumulate in the organism's own state (Di Paolo, 2005).

Once this temporal structure is acknowledged, waiting for certainty ceases to look like a neutral baseline. In many biological settings, inaction is itself destabilizing. Metabolic demands continue whether or not the organism has completed its inference; damage accumulates whether or not the cell has resolved every ambiguity; environmental threats intensify whether or not the animal has identified them with confidence. Sterling's account of allostasis is important here because it makes explicit that biological regulation is not exhausted by retrospective error correction around fixed setpoints. Efficient regulation often depends on anticipating need and preparing for it before the full demand arrives. The cephalic-phase responses offer a familiar physiological example: sensory food cues trigger digestive and metabolic adjustments before nutrients enter the bloodstream, thereby reducing the magnitude and cost of later correction (Sterling, 2012; Power & Schulkin, 2008; Ramsay & Woods, 2014).

The implication is stronger than the claim that anticipatory regulation is sometimes useful. It is that, for living systems, delayed action can itself be a mode of failure. A regulator that waits until outcome-relevant uncertainty is largely resolved may already have ceded the very conditions that make successful regulation possible. This is not to deny the costs of false alarms, mistimed commitments, or overreaction. It is to recognize a deeper asymmetry: in many biological contexts, the cost of acting early and correcting is lower than the cost of waiting until the window for viable correction has narrowed or closed. Biological organization is therefore structured around the possibility of error-tolerant commitment rather than certainty-dependent passivity. In that respect, the organism's problem is not how to avoid all mistaken action, but how to remain viable when action must precede proof (Sterling, 2012; Ramsay & Woods, 2014).

The positive claim of this section follows directly. Any self-maintaining system operating under changing conditions must instantiate mechanisms that enable commitment prior to certainty. By commitment I do not mean explicit belief, symbolic endorsement, or deliberative inference. I mean a structured bias toward viability-preserving action when outcome cannot yet be fully predicted, calculated, or verified. In cybernetic terms, the regulator must answer disturbance with organized selectivity before all uncertainty is removed. In biological terms, the organism must do something now that keeps it available for correction later. Commitment under uncertainty is therefore not best understood as one adaptive strategy among many. It is better understood as a precondition of viability itself: the condition under which any

particular regulatory strategy becomes possible at all (Ashby, 1956; Conant & Ashby, 1970; Di Paolo, 2005).

Equally important, such commitment is not blind. The architecture required here is not obstinate persistence but corrigible persistence. Biological systems act under incomplete information and then update, attenuate, reverse, or reinforce their response in light of feedback. The general form is action first, correction after. Robust bacterial chemotaxis provides a phylogenetically deep illustration. Barkai and Leibler (1997) showed that exact adaptation in chemotaxis can emerge from network architecture rather than fine-tuning of parameters, and Yi and colleagues (2000) later argued that robust perfect adaptation in this system depends on integral feedback control. The lesson is broader than chemotaxis itself. Viable systems can be organized to commit under uncertainty without demanding infallibility, because feedback can render commitment revisable rather than absolute (Barkai & Leibler, 1997; Yi et al., 2000; Alon et al., 1999).

Framed in this way, the requirement is older than nervous systems and older than cognition. Nervous systems elaborate it; they do not originate it. Physiological anticipation, microbial chemotaxis, and other non-neural forms of regulation already display the same basic logic: successful living systems do not wait until the world is fully disambiguated before they move. They carry forward a limited, embodied organization of what matters, act on that basis, and preserve themselves by remaining correctable. This phylogenetic depth matters theoretically because it blocks a common misreading of the argument. Commitment under uncertainty is not first a problem of human reasoning, animal deliberation, or cognitive prediction. It is first a problem of life (Barkai & Leibler, 1997; Yi et al., 2000; Power & Schulkin, 2008).

From this perspective, the familiar biological literatures are best read as descriptions of mechanism families that implement a more general requirement. Allostasis, feedforward control, stress mobilization, immune anticipation, bet-hedging, and memory-based recalibration are not rival explanations of the same phenomenon; they are heterogeneous solutions to the same structural demand. The unifying claim is therefore not that organisms sometimes manage uncertainty well. It is that any system that must maintain itself under changing conditions must commit before certainty and correct by feedback thereafter. What requires naming is not one mechanism among others, but the distributed architecture by which life remains viable when proof arrives too late. That is the conceptual step the next section takes.

3. Biological Faith Systems: Definition and Properties

If the previous section established that viable living systems must commit before certainty is available, this section names the architecture that makes such commitment possible. I call these architectures **Biological Faith Systems**. A Biological Faith System is the distributed, embodied, feedback-corrigible organization by which a living system biases action toward viability when the conditions of success cannot be fully predicted, calculated, or verified in advance. The term therefore does not name a single mechanism, organ, or pathway. It names a structural role: the present-tense organization that lets a system proceed now in ways that preserve the possibility of correction later.

Put plainly: a Biological Faith System is any organization of a living system that enables it to act before the outcome is certain and to adjust when the outcome arrives. The bias toward early action is not reckless; it is coupled to feedback mechanisms that correct, recalibrate, or update the initial commitment.

What makes this architecture faith rather than mere reactivity is the temporal asymmetry at its center: the system must move first, on incomplete grounds, because the alternative—waiting—is itself a threat to viability. The architecture is older than brains, older than cognition, and older than any explicit representation of the future. It is as present in bacterial chemotaxis as in human decision-making, though it takes different mechanistic forms at each scale.

Several clarifications follow immediately. First, Biological Faith Systems are **pre-propositional**. They do not begin with language, explicit belief, or symbolic representation. The architecture at issue is older than cognition in that sense: it belongs first to regulation, coupling, and self-maintenance. This matters because the core phenomenon is easily misdescribed if one starts from human belief and reasons downward. What must be explained is not why organisms sometimes “believe” things in a psychological sense, but how living systems remain organized when action cannot wait for fully resolved evidence. Barandiaran, Di Paolo, and Rohde (2009) make a closely related point in a different vocabulary when they argue that even minimal proto-cellular systems can satisfy core conditions of agency without appeal to propositional states. The lesson for present purposes is straightforward: commitment under uncertainty is organizational before it is conceptual.

Second, Biological Faith Systems are **embodied and distributed**. In most cases there is no single locus at which “the commitment” resides. Biological commitment is realized across coupled loops—metabolic, endocrine, immune, neural, mechanical—whose interaction biases the system toward some viable range of action. Ramsay and Woods’ (2014) analysis of homeostasis and allostasis is instructive here. Their argument is not merely that the two terms should be tidied up conceptually, but that regulation is best understood by looking at the activity and relation among multiple regulatory loops rather than at one isolated variable. That point generalizes. A Biological Faith System should be understood at the level of coordinated regulatory organization, not as an inner homunculus or a hidden controller.

Third, Biological Faith Systems are **feedback-driven and corrigible**. The point of commitment under uncertainty is not to avoid all error. It is to remain viable despite the inevitability of acting before full resolution. A Biological Faith System therefore does not commit absolutely; it commits revisably. Its outputs are modulated by feedback, strengthened when they preserve viability, weakened when they fail, and abandoned when they become too costly to sustain. In this respect, the relevant contrast is not between certainty and irrationality. It is between architectures that can move provisionally and remain correctable, and architectures that would require a degree of prior justification that living systems typically cannot afford.

This also clarifies the relationship between BFS and the frameworks closest to it. The nearest existing analogue is active inference and the broader free-energy tradition, which likewise treats living systems as acting under uncertainty in ways tied to self-organization. But that literature carries explicit formal commitments: Friston’s (2010) core formulation centers variational free energy minimization, and the biological extensions of the framework often rely on the statistical boundary concept of the Markov blanket. BFS does not reject those commitments; it simply does not require them. The claim advanced here is prior to any one formalism. Whatever mathematics one ultimately prefers, a viable account of living self-maintenance must accommodate the fact that biological systems routinely commit before certainty and repair afterward. BFS is therefore compatible with active inference, but not reducible to it.

The main terminological question, then, is why call this architecture **faith** at all. Why not settle for the safer language of anticipatory regulation, uncertainty management, or commitment mechanisms? Because

each of those terms misses something load-bearing. *Anticipatory regulation* identifies an important mechanism family, but not the general structural requirement that makes such mechanisms necessary. *Uncertainty management* sounds optional and strategic, as though living systems occasionally choose among alternative styles of coping. But the argument of the previous section was precisely that commitment under uncertainty is not optional in that sense; it is constitutive of viability. *Commitment mechanisms* comes closer, but it still understates the temporal asymmetry that matters most: the system must proceed before full justification is available and rely on subsequent feedback to confirm, reshape, or extinguish the trajectory it has already entered.

Used carefully, **faith** names that asymmetry exactly. It indicates that action precedes proof. The word is not intended here in a religious sense, nor as a synonym for credulity, dogma, or insulation from evidence. On the contrary, the relevant kind of faith is inseparable from corrigibility. A Biological Faith System does not persist by refusing correction; it persists by risking provisional commitment in a form that remains answerable to outcomes. What is “trusted” is not a doctrine but an architecture: the distributed organization of the living system and its capacity to remain dynamically adjustable after action has begun. In that thin but precise sense, faith captures something that blander alternatives flatten. It marks the fact that the system cannot wait until reasons are complete, and must instead rely on organization that can justify itself only downstream, through continued viability.

There is precedent in biology for this sort of terminology problem, though the comparison should be handled carefully. Mid-century debates over **teleonomy** were driven by a similar need: to preserve scientifically legitimate talk about end-directedness without importing stronger teleological or metaphysical baggage. Historically, Pittendrigh introduced the term, and Mayr later became one of its major analysts and defenders. More recently, Dresow and Love have argued that teleonomy may not have been the most durable or conceptually indispensable solution after all. That debate is useful here for exactly that reason. It shows that biology has long faced cases in which recurring structural patterns are either over-metaphysicized by inherited vocabulary or underdescribed by thinner mechanistic language. BFS belongs to that same family of interventions. It does not propose a new force, a non-natural principle, or a disguised theology. It proposes a name for a recurring structural pattern that existing terms only partially capture.

For that reason, the burden of the term is operational clarity. A Biological Faith System should be identifiable wherever four features co-occur: irreducible uncertainty relevant to viability; a bias toward action prior to full resolution; modification of that bias through feedback; and some retention of the outcome in architecture, whether transiently as priming or durably as memory, threshold shift, or reweighted response tendency. Retention here includes the persistence of the enabling architecture itself, not only episode-specific traces such as priming or threshold shift. This is why the term names neither a mechanism class nor a metaphorical flourish. It names a pattern of organization that can be tracked across otherwise separate literatures.

These four features are jointly necessary. A system that exhibits only some of them does not constitute a BFS on this account. A purely reactive loop that responds to perturbation without any bias toward early action—a thermostat restoring a fixed setpoint after deviation has already occurred—lacks the temporal asymmetry that defines the architecture. A stateful system that retains traces of past events but whose retention bears no relation to viability—a geological stratum recording depositional history—satisfies the retention condition but not the functional one. And a nonliving adaptive controller that exhibits formal

analogues of all four features may share structural similarities with a BFS, but the present claim is restricted to living systems whose viability is at stake in the sense developed in Section 2. The framework is also not a claim that all biological contexts impose this demand with equal force. Dormant systems, slowly varying regimes, and organisms in stable environments may face the temporal asymmetry less acutely; the necessity claim applies most directly to self-maintaining systems whose viability boundaries can be crossed faster than uncertainty can be fully resolved. BFS names a structural requirement, not an omnipresent description.

Once defined in this way, **the BFS framework** does not compete with allostasis, trained immunity, bet-hedging, chemotaxis, or stress mobilization. It groups them at the level of structural function. Each is a different answer to the same problem: how to keep a living system coherent when the evidence required for certainty arrives slower than the costs of inaction. The next section therefore turns from definition to demonstration by showing how this architecture appears in several distinct mechanism families across biological scale.

4. Mechanism Families

The argument to this point has been structural. This section gives that argument biological weight. The aim is not to collapse distinct literatures into a single mechanism, but to show that several well-established mechanism families share the same organizational logic: they mobilize action before full resolution of uncertainty, remain answerable to feedback, and preserve viability by allowing correction after commitment rather than requiring certainty before it.

4.1 Homeostatic and allostatic regulation

Physiology offers the clearest entry point because it contains some of the most familiar examples of anticipatory control. Sterling's (2012) account of allostasis explicitly characterizes regulation as predictive rather than merely retrospective: efficient regulation prepares for needs before they fully arrive, thereby reducing error magnitude and coordinating multiple subsystems. Ramsay and Woods (2014), while critical of the definitional looseness that has accumulated around allostasis, likewise argue that regulation is best understood by examining the activity and relation among multiple interacting loops rather than treating one variable as if it were governed by a single static set-point. Cephalic-phase responses make the point concrete. Sensory food cues can trigger salivation, gastric secretion, pancreatic responses, and insulin release before nutrients have been absorbed; Power and Schulkin (2008) note that these anticipatory responses prepare animals to digest, absorb, and metabolize food, and that blocking the cephalic insulin response worsens glucose control and reduces meal size. Recent work on thermoregulation makes the same structure explicit in another physiological register: negative feedback remains indispensable, but feedforward control can excite thermo-effectors in advance of the predicted challenge, with feedback supervising and refining the result (Mitchell, Snelling, & Fuller, 2025). Read through the present framework, these are not merely optional embellishments on a fundamentally reactive organism. They are instances of a Biological Faith System: commitment begins before certainty, and feedback closure makes that commitment corrigible rather than blind.

4.2 Stress-response cascades

Stress biology reveals the same architecture under a sharper temporal asymmetry. McEwen's (1998) account of allostasis and allostatic load emphasizes that adaptation to challenge depends on coordinated activation of neural, neuroendocrine, and immune-related systems, and that both failure to mobilize and failure to terminate such responses carry biological cost. At the cellular level, the heat-shock response shows the same logic without any appeal to cognition or centralized deliberation. Richter, Haslbeck, and Buchner (2010) describe heat stress as a challenge that damages cellular structures and interferes with essential functions, prompting activation of an ancient pathway that transiently induces heat-shock proteins; many of these proteins function as conserved molecular chaperones that prevent nonspecific aggregation and assist the recovery of native structure. The structural inference is important. These systems are not organized to wait until damage is complete and then decide whether response was warranted. They are organized around the fact that delay itself can be catastrophic. A mobilization launched somewhat early may be costly, but a mobilization launched after protein damage or systemic destabilization has advanced may no longer preserve the conditions required for successful correction. In BFS terms, stress cascades are commitment architectures tuned to hostile time scales: they spend first because viability cannot afford to settle everything in advance.

4.3 Immune anticipation and trained immunity

Immunology supplies a second major mechanism family because it shows commitment under uncertainty in systems organized around threat detection, tolerance, and biological identity. Janeway and Medzhitov (2002) describe the innate immune system as an evolutionarily ancient form of host defense triggered by pattern-recognition receptors that detect conserved molecular signatures shared across classes of microorganisms and activate conserved defense pathways. That architecture is already pre-committed: it does not wait for organism-by-organism certainty, but responds on the basis of recurrent structural cues that are good enough to justify early mobilization. Pradeu and Vivier's (2016) discontinuity theory sharpens the point further by treating the immune system as a change-detection system, one that responds to sudden changes in antigenic stimulation rather than to static substance categories alone. Trained immunity adds temporal depth. Netea and colleagues (2016) argue that innate immune cells can exhibit a form of memory, including in organisms lacking adaptive immunity, through epigenetic reprogramming and sustained changes in gene expression and cell physiology. On the present account, this is a Biological Faith System in especially clear form: pattern-recognition furnishes the initial commitment bias, subsequent experience reshapes response thresholds and future readiness, and the whole architecture remains corrigible because response is continually modulated by downstream success, failure, persistence, and cost. The immune system does not first secure complete certainty and then act. It acts through a distributed architecture built precisely for the fact that certainty often arrives too late.

4.4 Bet-hedging and phenotypic diversification

Evolutionary ecology and microbiology show that the same structural requirement can be implemented at a different organizational scale. Kussell and Leibler (2005) show that clonal populations in fluctuating environments can adapt either by sensing and responding or by generating diversity through stochastic phenotype switching, and that switching can be favored when environmental change is infrequent. Donaldson-Matasci, Lachmann, and Bergstrom (2008) generalize this point by treating phenotypic diversity itself as an adaptation to environmental uncertainty: not an individual generalist that solves every condition directly, but a genotype that spreads risk by generating heterogeneous descendants. Beaumont and colleagues (2009) then provide direct experimental evidence in bacteria, reporting *de novo*

evolution of bet-hedging in *Pseudomonas fluorescens* subjected to fluctuating conditions. Plant systems show the same structure in another medium. Clauss and Venable (2000) found that desert annual populations in historically more xeric environments exhibited lower mean germination fractions, consistent with bet-hedging models of delayed germination under uncertainty. What matters structurally is that the commitment architecture is now distributed across a population or lineage rather than centered in a single organism's immediate regulatory loop. The system does not reduce uncertainty by waiting for better information; it reduces vulnerability by allocating its possible futures across phenotypes, dormancy states, or response modes in advance. At the lineage level, the corrective term is differential survival across fluctuating conditions, not within-lifetime sensorimotor adjustment. BFS therefore includes not only anticipatory regulation within bodies, but organized diversification across generations when lineage-level viability is the relevant scale of commitment.

4.5 Robust adaptation in microbes

Bacterial chemotaxis supplies the phylogenetically deep anchor that prevents the entire argument from drifting back toward cognition. Barkai and Leibler (1997) proposed a mechanism for robust adaptation in simple biochemical networks and showed that it applies to bacterial chemotaxis; crucially, the adaptation property in their model follows from network connectivity rather than delicate parameter tuning. Alon and colleagues (1999) reinforced the point experimentally, showing that although some response properties varied with systematic changes in protein concentrations, the precision of adaptation remained robust. Yi and colleagues (2000) then argued that robust perfect adaptation in chemotaxis is structurally linked to integral feedback control and suggested that related control principles may underlie many homeostatic mechanisms. The relevance to BFS is direct. A bacterium navigating a gradient cannot wait until its environment is richly characterized before it moves; it biases motion, samples consequences, and corrects through an architecture that preserves responsiveness across noise and parameter variation. This is commitment under uncertainty in one of its most stripped-down forms. No symbolic representation, explicit belief state, or nervous system is required. The commitment architecture is already present in the organization of the loop itself.

Across these families, the unifying claim is not that all living systems anticipate in the same way, or that one vocabulary can erase their mechanistic differences. The claim is narrower and stronger: whenever viability depends on acting before the world is fully resolved, living systems evolve architectures that permit such action without forfeiting correction. In physiology this appears as feedforward regulation supervised by feedback; in stress biology as early mobilization under asymmetric cost; in immunology as pattern-based response and memory; in bet-hedging as diversified commitment at population scale; and in microbial control as robust adaptation built into network topology. Calling these Biological Faith Systems does not replace the local literatures. It renders visible the structural pattern they instantiate together: life must commit before certainty, and must remain organized enough to learn from what follows.

5. Relationship to Existing Frameworks

At this stage, the question is not whether biology already possesses theories of anticipation, autonomy, and goal-directed organization. It does. The real question is whether any of those frameworks names the precise claim advanced here: that commitment under irreducible uncertainty is a structural requirement of viability, and that this requirement can be identified across otherwise disparate biological literatures. The purpose of this section, accordingly, is not adversarial differentiation. It is exact placement. BFS should

be read as continuous with the strongest neighboring frameworks while still making a non-redundant claim of its own.

5.1 Active inference and the Free Energy Principle

The closest existing neighbor is active inference and the broader Free Energy Principle. Friston's 2010 review presents the free-energy principle as a unifying account of perception, action, and learning organized around the minimization of free energy, with prediction and generative modeling at the center of the framework. Friston (2013) generalizes the ambition further, proposing that life itself can be understood as biological self-organization in systems possessing Markov blankets, such that internal states appear to minimize a free-energy functional and thereby preserve structural and functional integrity. Kirchhoff and colleagues' (2018) work then explicitly links this formalism to biological autonomy by treating living systems as layered, self-sustaining organizations defined in statistical terms by their Markov blankets.

BFS converges with this family of views at the level of the problem they are trying to solve. Both begin from the fact that living systems must remain organized while acting from incomplete information. Both reject the picture of life as a merely reactive process waiting passively for certainty to arrive. The difference is not one of direction, but of level. The free-energy framework offers a specific mathematical description of how such organization can be understood, including commitments to variational free energy, generative models, and—at least in its biological generalization—Markov blankets. BFS, by contrast, makes a prior structural claim: any adequate account of living self-maintenance must explain how organisms commit before certainty and correct afterward. On this view, active inference is not displaced by BFS; it becomes one powerful way of formalizing a requirement that BFS identifies at a more general level.

That distinction matters methodologically. The advantage of BFS is that it can range across mechanism families without requiring that all of them be translated, at the outset, into one mathematical idiom. A physiological feedforward loop, an immune threshold architecture, and a population-level diversification strategy need not look formally identical in order to instantiate the same structural demand. BFS therefore asks a different question from the free-energy literature. It asks not primarily how to model living systems as inferential systems, but what any viable living system must be able to do when the costs of waiting exceed the costs of provisional action. The answer proposed here is: it must commit under uncertainty in a corrigible way. In that sense, BFS is compatible with active inference but not dependent on it.

5.2 Autopoiesis and biological autonomy

BFS is even more continuous with the autopoietic and autonomy-based traditions. Di Paolo's (2005) intervention is especially important here. He argues that autopoiesis provides a systemic language for intrinsic teleology, but that its original formulation must be elaborated if it is to explain sense-making and agency. His proposed supplement is adaptivity: a many-layered property that allows organisms to regulate themselves with respect to their conditions of viability. Mossio and Moreno (2010), in turn, describe organizational closure as a fundamental property of biological systems while also insisting that closure alone underdetermines what is required for a full-fledged organism; additional control and organizational properties are needed. Taken together, these views already move very close to the terrain BFS occupies.

The present paper does not reject that terrain. It sharpens it. Autopoiesis names the self-producing organization of the living system. Organizational closure names the reciprocal constraint regime through which biological organization maintains itself. Adaptivity specifies the system's capacity to monitor tendencies relative to viability boundaries and to regulate accordingly. BFS contributes a further point: under real biological conditions, such regulation cannot wait for certainty. Di Paolo is explicit that adaptivity involves distinguishing tendencies according to whether they approach or recede from the boundary of viability, and acting on them so that future states are prevented from crossing that boundary. BFS names the commitment architecture by which exactly this sort of regulation proceeds when evidence is temporally incomplete and action must begin before the case is settled.

Put differently, autopoiesis tells us what must be maintained, and adaptivity tells us that the system can regulate in relation to that maintenance. BFS adds the temporal-operational claim that such regulation must take the form of provisional commitment under uncertainty. That is why BFS should be read as a specification, not a repudiation, of the autonomy tradition. It preserves the normativity of self-maintenance and the organizational logic of closure, but it makes newly explicit the fact that viable regulation is structured by a recurrent asymmetry: action first, feedback correction after.

5.3 Allostasis

Allostasis stands in a somewhat different relation to BFS because it is less a general philosophy of life than a powerful account of predictive physiological regulation. Ramsay and Woods (2014) emphasize both why allostasis has been influential and why it remains conceptually difficult: the term has become diffuse even as it captures phenomena that older homeostatic language handles poorly. Schulkin and Sterling (2019) sharpen the predictive side of the view by describing allostasis as a regulatory strategy in which the brain forecasts likely needs and computes responses that reduce costly errors before they fully materialize. On its strongest reading, allostasis is already an important recognition that biological regulation is anticipatory, coordinated, and efficiency-seeking rather than merely reactive.

BFS preserves what is strongest in that insight while relocating the explanatory center of gravity. Allostasis explains how physiological variables are managed through predictive adjustment. BFS explains why anticipatory adjustment is required in the first place, and why the same logic should be expected outside the narrow physiological settings in which the term allostasis is usually most at home. This broader framing also helps with the definitional drift problem Ramsay and Woods (2014) identify. Rather than stretching a physiology term until it bears the weight of every form of biological anticipation, BFS anchors the discussion at the level of structural necessity: living systems must commit before certainty when waiting itself threatens viability. Allostasis then appears not as a rival to BFS, but as one especially vivid mechanism family within it.

5.4 Teleonomy and the vocabulary problem

The final comparison is not with a mechanistic or formal framework, but with an earlier vocabulary intervention in biology. Mayr's 1974 analysis begins from the observation that teleological language is pervasive in biology and that many biologists regard such language as objective and scientifically useful even while resisting metaphysical teleology. More recent work by Dresow and Love (2023) revisits teleonomy precisely as a proposed conceptual replacement for teleology, while also noting that its alleged indispensability remains open to question. Their historical reconstruction is especially relevant here

because it traces the introduction of “teleonomy” back to Pittendrigh’s 1958 essay and its subsequent uptake by prominent biologists.

BFS belongs to that same lineage of disciplined vocabulary work. The point is not to import religion into biology any more than teleonomy imported Aristotelian final causes back into physiology. The point is that recurring structural patterns are sometimes badly served by the terms already on hand. Existing language may be too local, too mechanistic, or too cautious to name the full form of what is happening. That is the reason for proposing “faith” here. It marks the fact that a living system must proceed before proof is complete, and must rely on downstream correction to vindicate, revise, or extinguish the trajectory it has already entered. Whether the term ultimately succeeds will depend on whether the paper shows that no thinner alternative captures the same combination of necessity, temporality, and corrigibility.

Taken together, these comparisons place BFS fairly precisely. It is broader than allostasis, less formal than the Free Energy Principle, more temporally explicit than autopoiesis alone, and more biologically substantive than a mere vocabulary gesture detached from mechanism. Its novelty does not lie in discovering that organisms anticipate, regulate, or preserve themselves; those points are already well established. Its novelty lies in arguing that commitment under uncertainty is the common structural requirement underneath those phenomena and in naming the distributed architecture that enacts it. That is the sense in which BFS is not a relabeling of neighboring frameworks, but a unification claim about what they each, in different ways, already reveal.

6. Predictions and Discriminators

A structural proposal becomes scientifically useful only when it constrains what one should expect to observe. The claim made here is not just that organisms often handle uncertainty well. It is that viable living systems must possess architectures of provisional commitment. If that is correct, then the claim should yield recurring empirical signatures across otherwise dissimilar domains. Three such signatures are especially important: timing asymmetry, constraint-sweep hysteresis, and phylogenetic depth. Of these, constraint-sweep hysteresis is the most distinctive to BFS and the sharpest empirical discriminator, because it predicts a specific residue of provisional commitment that purely reactive or equilibrium-restoring frameworks have no special reason to expect.

6.1 Timing asymmetry

The first prediction concerns the order of events. If Biological Faith Systems are real, timely biological action should systematically begin before uncertainty is fully resolved, with feedback correction arriving afterward rather than certainty arriving first. This is stronger than the uncontroversial claim that some anticipatory mechanisms exist. Cephalic-phase responses begin from sensory cues before nutrients have been absorbed; innate immunity mobilizes through pattern-recognition receptors that respond to conserved microbial signatures before pathogen identity is completely specified; and bacterial chemotaxis biases motion through an adaptive control architecture that remains robust in the face of noise and parameter variation. In each case, the system moves first on a structured but incomplete basis and only then refines the response through downstream feedback.

The discriminator follows directly. If BFS is correct, then systems experimentally forced to defer commitment until uncertainty is largely resolved should, in domains where delay is costly, show

systematic losses relative to systems allowed to commit early and correct later: larger excursions from viable ranges, slower recovery, higher cumulative energetic cost, or greater damage accumulation. What would count against the framework is strong cross-domain evidence that living systems can routinely postpone regulatory action until informational ambiguity is mostly settled without paying a consistent biological penalty. That would suggest that commitment under uncertainty is a contingent tactic rather than a structural requirement.

To state the counterexample explicitly: the necessity claim would be undermined by a demonstrably viable self-maintaining system that achieves viability without any architectural bias toward early commitment—one that defers all regulatory action until uncertainty is resolved, incurs no systematic cost for that deferral, and requires no feedback-driven correction because its initial conditions are fully specified before action begins. No known living system meets this description, but the description itself is precise enough to be tested. If such a system were found or engineered, BFS would need to be revised from a structural requirement to a strong but defeasible generalization.

6.2 Constraint-sweep hysteresis

The second prediction concerns residues of commitment. If a Biological Faith System is not just a descriptive gloss but an actual biological architecture, then episodes of commitment should leave traces in thresholds, chromatin state, network occupancy, or response gain. Put differently, when a living system is driven into a constraint regime and then released, the recovery path should not simply retrace the onset path. The system that comes out of the perturbation has been altered by the fact of having passed through it. Evidence consistent with this expectation already appears across multiple domains: the lactose utilization network of *E. coli* shows hysteresis and multistability (Ozbudak et al., 2004), the yeast galactose network shows persistent cellular memory shaped by feedback architecture (Acar, Becskei, & van Oudenaarden, 2005), recurring dehydration in *Arabidopsis* produces transcriptional memory associated with sustained H3K4me3 and stalled RNA polymerase II at trainable genes (Ding, Fromm, & Avramova, 2012), and rice exhibits genotype-specific constraint-sweep hysteresis in transcriptomic state under heat stress, with strong path dependence in Azucena and Pandan Wangi, weaker and non-significant effects in other cultivars, and no comparable signal under drought (Jones, 2026).

This is the logic behind the prediction of **constraint-sweep hysteresis**. If BFS is correct, path dependence should appear not just in specially engineered switches but in ordinary biological regulation wherever provisional commitment is followed by correction, priming, or recalibration. A matched onset-and-release design should therefore reveal systematic differences between the trajectory into constraint and the trajectory out of it. Those differences may appear as altered activation thresholds, faster re-induction, slower relaxation, asymmetric recovery, or persistent sensitization. A purely reactive framework, by contrast, has no special reason to expect such residues once external conditions are re-equated.

6.3 Phylogenetic depth

The third prediction concerns how far down the tree of life the relevant architecture should appear. If commitment under uncertainty is a structural requirement of life, then the enabling mechanisms should be phylogenetically deeper than cognition and deeper than nervous systems. One should find them in microorganisms, plants, and invertebrates before one ever reaches explicit representation or deliberation. That expectation is already plausible. Bacterial chemotaxis displays robust exact adaptation through network architecture and integral feedback control; *Arabidopsis* shows transcriptional stress memory

across recurring dehydration; rice shows genotype-specific transcriptomic path dependence under heat stress and recovery (Jones, 2026); and *Drosophila*, despite relying on innate rather than adaptive immunity, can show a specific primed immune response that depends on phagocytes and persists across later challenge (Pham et al., 2007). These cases are mechanistically heterogeneous, but they converge on the same structural role: action or mobilization under incomplete information followed by selective retention of what proved viable.

This gives the prediction a sharp discriminator. If the relevant architecture emerged only with nervous systems, or only with cognitive prediction in the strict sense, then BFS would collapse into a theory about brains rather than a theory about life. But if comparative work continues to find structurally continuous forms of anticipatory bias, corrigibility, and memory-like residue in organisms that lack brains altogether, then the view defended here gains exactly the kind of support it needs. The issue is not whether microbes, plants, and animals use the same machinery. The issue is whether they instantiate the same requirement at different levels of biological organization.

Taken together, these predictions do more than redescribe familiar biological phenomena. They specify three empirical signatures that a unifying framework ought to risk: a recurrent onset-before-certainty timing relation, systematic hysteresis under matched constraint sweeps, and a phylogenetically deep continuity of commitment architectures below cognition. If those signatures fail, BFS should be revised or abandoned. If they hold, the framework earns the right to function as more than a terminology proposal. It becomes a genuine explanatory lens for why living systems so often act first and justify later.

7. Conclusion

Biology has long described the many ways living systems act before outcomes are fully known. Physiologists have examined anticipatory regulation, immunologists have studied early threat detection and trained memory, evolutionary ecologists have analyzed diversification under uncertainty, and microbiologists have shown how robust adaptive control can emerge from simple network architectures. This paper has argued that these phenomena are not best understood as a scattered collection of local strategies alone. They are more deeply unified as expressions of a shared structural requirement: any self-maintaining system that must remain viable under changing conditions must be able to commit to action before certainty is available, and to preserve itself by correcting that action through feedback afterward. To name that requirement, this paper introduced Biological Faith Systems as the distributed, embodied, feedback-correctible architecture by which living systems proceed under irreducible uncertainty.

The contribution is not a new molecular mechanism, not a new force, and not a replacement for the explanatory work done by physiology, immunology, control theory, or evolutionary biology. It is a structural lens that makes a recurring biological pattern more visible, more precise, and more comparable across domains. BFS introduces no extra substance and no exemption from standard biological explanation. It asks whether the mechanisms biology already studies can be read as diverse realizations of one structural demand. Physiologists still explain anticipatory regulation, immunologists still explain innate memory and discrimination, and control theorists still explain robust adaptation through feedback architecture. BFS contributes at a different level: it proposes that a single temporal-organizational problem runs through all of them.

If that proposal holds, two consequences follow. The first is programmatic. Once a commitment architecture is identified in one domain, one can ask what its structural analogues should look like elsewhere—not the same molecules or control equations, but the same asymmetry between early action and later correction, the same residue in thresholds or memory, and the same trade-off between false alarm and delayed response. Cross-domain comparison becomes more than conceptual tidiness. It becomes a way of discovering new empirical questions, because evidence for priming, hysteresis, diversification, or robust adaptation in one field provides reason to look for related organizational signatures in another.

The second consequence concerns the placement of cognition. If commitment under uncertainty is a structural requirement of life, then conscious belief, deliberation, and explicit reasoning do not introduce the problem of commitment; they inherit and elaborate it. This does not reduce cognition to physiology or deny genuine novelty at the level of brains, language, and explicit thought. It does, however, place influential cognitive frameworks—predictive processing (Clark, 2013; Hohwy, 2013), ecological rationality (Gigerenzer, 2008), active inference—on deeper biological ground. The deepest commonality across such views is biological before it is psychological: a living system must be able to proceed without waiting for complete justification. That observation opens a path toward later work on rationality and belief formation, but it does not require this biology paper to walk that path in full.

The framework was rendered empirically vulnerable through three testable predictions: that timely action will systematically precede full uncertainty resolution, that commitment architectures will leave hysteretic and path-dependent residues, and that such architectures will prove phylogenetically deeper than cognition or nervous systems alone. If those predictions fail, the framework should be revised or abandoned. If they hold, then biology has been looking at one of life's most basic requirements in pieces and is now in a position to see it whole: living systems do not survive because they wait for proof, but because they are organized to move before proof is complete, and to remain corrigible as the consequences unfold.

References

- Acar, M., Becskei, A., & van Oudenaarden, A. (2005). Enhancement of cellular memory by reducing stochastic transitions. *Nature*, 435, 228–232.
- Alon, U., Surette, M. G., Barkai, N., & Leibler, S. (1999). Robustness in bacterial chemotaxis. *Nature*, 397, 168–171.
- Ashby, W. R. (1956). *An introduction to cybernetics*. Chapman & Hall.
- Barandiaran, X. E., Di Paolo, E., & Rohde, M. (2009). Defining agency: Individuality, normativity, asymmetry, and spatio-temporality in action. *Adaptive Behavior*, 17(5), 367–386.
- Barkai, N., & Leibler, S. (1997). Robustness in simple biochemical networks. *Nature*, 387, 913–917.
- Beaumont, H. J. E., Gallie, J., Kost, C., Ferguson, G. C., & Rainey, P. B. (2009). Experimental evolution of bet hedging. *Nature*, 462, 90–93.
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(3), 181–204.

- Clauss, M. J., & Venable, D. L. (2000). Seed germination in desert annuals: An empirical test of adaptive bet hedging. *American Naturalist*, 155(2), 168–186.
- Conant, R. C., & Ashby, W. R. (1970). Every good regulator of a system must be a model of that system. *International Journal of Systems Science*, 1(2), 89–97.
- Di Paolo, E. A. (2005). Autopoiesis, adaptivity, teleology, agency. *Phenomenology and the Cognitive Sciences*, 4(4), 429–452.
- Ding, Y., Fromm, M., & Avramova, Z. (2012). Multiple exposures to drought “train” transcriptional responses in *Arabidopsis*. *Nature Communications*, 3, 740.
- Donaldson-Matasci, M. C., Lachmann, M., & Bergstrom, C. T. (2008). Phenotypic diversity as an adaptation to environmental uncertainty. *Evolutionary Ecology Research*, 10, 493–515.
- Dresow, M., & Love, A. C. (2023). Teleonomy: Revisiting a proposed conceptual replacement for teleology. *Biological Theory*, 18(2), 101–113.
- Friston, K. (2010). The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11, 127–138.
- Friston, K. (2013). Life as we know it. *Journal of the Royal Society Interface*, 10, 20130475.
- Gigerenzer, G. (2008). *Rationality for mortals: How people cope with uncertainty*. Oxford University Press.
- Hohwy, J. (2013). *The predictive mind*. Oxford University Press.
- Janeway, C. A., Jr., & Medzhitov, R. (2002). Innate immune recognition. *Annual Review of Immunology*, 20, 197–216.
- Jones, J. C. (2026). Constraint-sweep hysteresis in rice transcriptome state during heat stress and recovery [Preprint]. HoldingLight LLC.
- Kirchhoff, M. D., Parr, T., Palacios, E., Friston, K., & Kiverstein, J. (2018). The Markov blankets of life: Autonomy, active inference and the free energy principle. *Journal of the Royal Society Interface*, 15, 20170792.
- Kussell, E., & Leibler, S. (2005). Phenotypic diversity, population growth, and information in fluctuating environments. *Science*, 309, 2075–2078.
- Mayr, E. (1974). Teleological and teleonomic: A new analysis. *Boston Studies in the Philosophy of Science*, 14, 91–117.
- McEwen, B. S. (1998). Stress, adaptation, and disease: Allostasis and allostatic load. *Annals of the New York Academy of Sciences*, 840, 33–44.
- Mitchell, D., Snelling, E. P., & Fuller, A. (2025). Revisiting concepts of thermal physiology: Understanding feedback and feedforward control, and local temperature regulation. *Acta Physiologica*, 241, e70063.
- Mossio, M., & Moreno, A. (2010). Organisational closure in biological organisms. *History and Philosophy of the Life Sciences*, 32(2/3), 269–288.

- Netea, M. G., Joosten, L. A. B., Latz, E., Mills, K. H. G., Natoli, G., Stunnenberg, H. G., O'Neill, L. A. J., & Xavier, R. J. (2016). Trained immunity: A program of innate immune memory in health and disease. *Science*, 352(6284), aaf1098.
- Ozbudak, E. M., Thattai, M., Lim, H. N., Shraiman, B. I., & Van Oudenaarden, A. (2004). Multistability in the lactose utilization network of *Escherichia coli*. *Nature*, 427, 737–740.
- Pham, L. N., Dionne, M. S., Bhatt, P., & Bhatt, D. L. (2007). A specific primed immune response in *Drosophila* is dependent on phagocytes. *PLoS Pathogens*, 3(3), e26.
- Pittendrigh, C. S. (1958). Adaptation, natural selection, and behavior. In A. Roe & G. G. Simpson (Eds.), *Behavior and evolution* (pp. 390–416). Yale University Press.
- Power, M. L., & Schulkin, J. (2008). Anticipatory physiological regulation in feeding biology: Cephalic phase responses. *Appetite*, 50(2–3), 194–206.
- Pradeu, T., & Vivier, E. (2016). The discontinuity theory of immunity. *Science Immunology*, 1(1), aag0479.
- Ramsay, D. S., & Woods, S. C. (2014). Clarifying the roles of homeostasis and allostasis in physiological regulation. *Psychological Review*, 121(2), 225–247.
- Richter, K., Haslbeck, M., & Buchner, J. (2010). The heat shock response: Life on the verge of death. *Molecular Cell*, 40(2), 253–266.
- Schulkin, J., & Sterling, P. (2019). Allostasis: A brain-centered, predictive mode of physiological regulation. *Trends in Neurosciences*, 42(10), 740–752.
- Sterling, P. (2012). Allostasis: A model of predictive regulation. *Physiology & Behavior*, 106(1), 5–15.
- Yi, T. M., Huang, Y., Simon, M. I., & Doyle, J. (2000). Robust perfect adaptation in bacterial chemotaxis through integral feedback control. *Proceedings of the National Academy of Sciences*, 97(9), 4649–4653.

This paper is part of the Universal Collapse Theory library. For a reading guide and full architecture, visit universalcollapse.com/roadmap.

AI Disclosure. AI tools were used to assist with manuscript preparation. The underlying theory, arguments, and interpretive claims are the author's own, and the author takes full responsibility for the manuscript.

Citation: Jones, J. C. (2026). Biological Faith Systems: How Living Systems Commit Before Certainty. HoldingLight LLC.