

SDPF

A New Science of Specification for Stochastic Generators

Hamza Abdullah

2026

The New Object

Every science begins with the recognition of a new object of study. Before the object is named, it sits unrecognised inside older categories, producing anomalies those categories cannot explain. After it is named, the anomalies resolve, the category refines, and a new discipline becomes possible.

The new object is this: a system produced by deterministic human intent operating through a stochastic machine generator.

This object did not exist before generative AI became capable of producing functional software artefacts at scale. Before that moment, every software system had a deterministic executor — a compiler, a runtime, an interpreter, or a human programmer whose variability could be bounded by discipline. The chain from intent to artefact was mechanical in its final step. Variability was a property of the process, not the executor.

With generative AI in the production loop, this chain breaks. The final step from specification to artefact is now performed by a probabilistic function whose output varies across runs, across models, and across minor perturbations of input. The human remains the source of intent. The artefact remains the output. But the mechanism between them has become stochastic — and the discipline of specification that assumed a deterministic mechanism no longer applies.

This is not a tooling change or a workflow change. It is a change in what software is and how it comes into being. And it requires a new science to study it.

Why Prior Disciplines Do Not Apply

The temptation to locate the new object inside existing disciplines is strong and must be resisted. The existing disciplines were built for different objects, and applying their frames obscures what is actually new.

It is not formal methods

Formal methods — Hoare logic, predicate transformers, Z notation, VDM, the B-method, design by contract — assumes a deterministic executor. The central activity is proof: demonstrating that an implementation satisfies a specification. The discipline has no purchase on a stochastic generator, because the generator does not produce a single implementation to prove properties of. It produces a distribution of implementations, and the question is not whether one of them satisfies the specification, but whether the distribution does, and under what conditions.

Formal methods can describe the specification half of the SDPF lifecycle. It cannot describe the generative half, because it was never built to.

It is not software engineering methodology

Methodologies — Agile, Waterfall, Spiral, V-model, SAFe, XP — describe how humans organise themselves to produce software. They are social technologies. They assume the artefact-producing labour is performed by humans whose variability is managed through process. A stochastic generator is not a human, and its variability is not managed by process. No methodology addresses the question of how to specify a system that will be materialised by a probabilistic function, because no methodology was built in a context where that was possible.

It is not AI safety

AI safety takes the model as its object. Its central questions are about model behaviour, alignment, interpretability, and failure modes internal to the generator. These are important questions, and they share some vocabulary with the new science, but they study a different object. AI safety asks what the model does. The new science asks how a specification can govern what the model produces. The model is the executor in the new science, not the object. Safety research operates upstream of specification; the new science operates at specification.

It is not infrastructure or DevOps

Infrastructure-as-code, configuration management, and drift detection operate on environmental state — the substrate code runs on. They are concerned with reality after the artefact exists. The new science is concerned with the artefact itself, and specifically with the relationship between the specification of an artefact and the stochastic process that produces the artefact. The objects are different. So are the failure modes, the success conditions, and the disciplines required to manage them.

Each of these fields touches the new object tangentially. None contains it. A science that contains it must be built.

The New Science

The new science has an object, a set of failure modes, a set of success conditions, a formal apparatus, and a conformance discipline. SDPF is the first articulation of all five, assembled into a coherent whole.

The object

Specifications that govern stochastic generators. Not specifications of static systems. Not contracts between humans. Not evaluations of model behaviour. Specifications whose function is to bound the space of acceptable outputs from a probabilistic process, such that outputs falling within the bounds are conforming and outputs falling outside are not.

The failure modes

Hallucination — the confident assertion of facts that are not true, undetectable from within the specification itself. Drift — the silent divergence of generator output from specification intent across runs. Surface-form entanglement — the conflation of how an output is written with what it does, producing rejection of conforming outputs on stylistic grounds. Specification incompleteness — the unbounded region of behaviour that emerges when the specification fails to cover a case, filled by the generator with whatever its training suggests. These failure modes are specific to stochastic-generator systems. They do not occur, or occur in fundamentally different forms, in deterministic systems.

The success conditions

Structural invariance under surface variation. Bounded stochasticity. Traceability from specification element to output artefact. Verifiability of the acceptance criterion. Auditability of the generation process. Each of these is a property the new science must define precisely, because each is what distinguishes a working specification from one that merely looks correct.

The formal apparatus

The Bounded Stochasticity Theorem. The first formal statement of the conditions under which a deterministic specification can govern a stochastic generator. The theorem is not a restatement of older results. It is the foundational result of the new science — the statement that makes the rest of the discipline possible. Without it, there is no principled reason to believe the new object can be studied at all. With it, the discipline has a floor to build on.

The conformance discipline

The Technical Verification Gate. The lifecycle protocol. The five-gate verification hierarchy. The provenance chain. The seventeen styles. The conformance classes for specifications, tools, evidence packages, and processes. These are not borrowed mechanisms. They are the first

generation of apparatus native to the new science — instruments built specifically for the object the science studies.

The Bounded Stochasticity Theorem as Foundational Result

Every science has a foundational result: the statement from which its other results descend. The Bounded Stochasticity Theorem is the foundational result of the science of specification for stochastic generators.

Let S be a complete SDPF specification defining structural invariants I . Let G be a generative AI output produced from S . G satisfies S if and only if G satisfies all invariants in I . Stochastic variation in G is bounded by I — any variation satisfying I is a conforming realisation of S .

The theorem states the conditions under which the new object is tractable. It says that a deterministic specification can govern a stochastic generator when the specification operates on structural invariants rather than surface outputs. It says that variation is acceptable — indeed required — within the invariant space, and unacceptable outside it. It says what verification must test, and what verification cannot meaningfully test.

The theorem's scope conditions are themselves foundational. It does not hold when the specification is incomplete — so incompleteness is a first-class failure mode of the discipline. It does not hold when Technical Verification Gate entries are unverified — so factual correctness of specification assertions is a first-class precondition. It does not hold when the model lacks sufficient capability to satisfy the invariants — so capability thresholds become a first-class concern, to be measured rather than assumed.

Every downstream result in the science derives from the theorem. The TVG exists because the theorem requires verified factual grounding. The traceability matrix exists because the theorem operates on invariants that must be testable. The conformance classes exist because the theorem requires specifications and tools to meet explicit standards. The lifecycle gates exist because the theorem's scope conditions must be enforced structurally. The provenance chain exists because the theorem requires the generation process itself to be auditable.

A science without a foundational result is a practice. A science with one is a discipline. This is the statement that elevates SDPF from practice to discipline, and it is the reason the framework can be implemented, certified, audited, and adopted by parties who have never met the author — which is the definition of a scientific discipline rather than a personal methodology.

The Technical Verification Gate as First-Generation Apparatus

Every science requires instruments built specifically for the object it studies. A microscope is not a refinement of a magnifying glass; it is an instrument built when biology recognised that cells were the new object. A particle accelerator is not a faster cyclotron; it is the instrument built when physics recognised that subatomic structure was the new object. The Technical Verification Gate is the instrument built when the science of specification for stochastic generators recognised that confident fabrication was a first-class failure mode.

The TVG is the first mechanism whose entire purpose is to interpose between a stochastic generator's confident assertion of technical fact and the downstream construction built on that assertion. Its position in the lifecycle — before specification lock, not after implementation — is dictated by the failure mode it addresses: fabrication is cheapest to correct at specification time and most expensive to correct at deployment time. Its HALT discipline — that failures propagate back to the specification rather than being patched around — is dictated by the authority structure the new science requires: the specification is the ground truth, and disagreement between specification and reality must be resolved at the specification, not hidden downstream.

The TVG operates on a category of artefact that did not exist before: the AI-produced specification of a system that does not yet exist, containing factual assertions about a world the AI cannot verify. No prior instrument was built for this artefact, because the artefact was not there to build for. The TVG is first-generation apparatus for a new class of object, and it will not be the last instrument of its kind. It will be followed by more sophisticated mechanisms as the science matures — just as optics followed the microscope and detector physics followed the accelerator. But it is the first, and its existence marks the moment when the science acquired its own instruments rather than borrowing from older disciplines.

What the Science Makes Possible

A new science is justified by what it makes possible that was not possible before. The science of specification for stochastic generators makes four things possible that no prior discipline could deliver.

Deterministic governance of stochastic production

Before this science, governance over generative AI output was either advisory (prompts, system messages) or post-hoc (evaluation benchmarks, red-teaming, human review). Neither could provide a structural guarantee. With the Bounded Stochasticity Theorem and its apparatus, governance becomes structural: outputs either satisfy the invariants or they do not, and the distinction is verifiable. This is the first time deterministic governance has been available for stochastic production.

Auditable AI in regulated contexts

Regulated industries — healthcare, finance, aviation, pharmaceuticals, defence — have been structurally unable to adopt generative AI because the audit trail required by their regulatory frameworks could not be produced. Evaluations measure behaviour; they do not trace it. Model cards describe capabilities; they do not document decisions. The new science produces evidence as a natural output of its process: every artefact traces to a specification element, every specification element traces to a requirement, every requirement is verified against live output, and every verification is preserved in a tamper-evident provenance chain. This is the first audit trail of its kind, and it is the artefact regulated industries have been waiting for.

Specification as the primary unit of engineering labour

For seventy years, the primary unit of software engineering labour has been the line of code. The specification was preparation; the code was the work. The new science reverses this. When the generator is stochastic and the specification is the deterministic element, the specification becomes the work. Code becomes its mechanical realisation. This reorders what engineering skill is, what engineering education should teach, and what engineering excellence looks like. It is not a marginal shift. It is the inversion of a seven-decade organising principle.

A discipline of AI use distinct from AI safety

AI safety studies the model. The new science studies the relationship between specification and model output. These are complementary but distinct — and the distinction matters because the two disciplines require different tools, different practitioners, different institutional homes, and different governance regimes. Safety research will continue to ask whether a given model is safe to deploy. The new science asks, given a deployed model, how its output can be bounded to a specification. Both questions must be answered. Only one of them was recognised as a question before this science named it.

The Status of the Science

The science is at the moment of its founding. This has specific consequences.

The foundational result has been stated. The Bounded Stochasticity Theorem is in the Language Specification, scoped and proved. First-generation apparatus exists — the TVG, the lifecycle gates, the conformance classes. The object is named, the failure modes are catalogued, the success conditions are defined. These are the conditions of a science at its beginning, not at its end.

What remains is what always remains at the founding of a discipline: adoption, replication, implementation by independent parties, the development of a community of practitioners, the accumulation of case studies, the refinement of apparatus through contact with diverse

problems, the emergence of sub-disciplines, the establishment of institutional homes. None of these happen quickly. All of them happen eventually when the founding work is correct.

The test of a founding work is not whether it is immediately adopted. The test is whether, when adoption happens, the adopted discipline continues to rest on the original foundation. Codd's relational model rested on Codd's theorems for fifty years. The end-to-end argument has governed network architecture for four decades. Hoare logic still underpins formal methods. The works that founded these fields were not always the most popular works in their first years. They were the works whose foundations continued to hold as the field grew around them.

The foundations SDPF establishes — the object, the theorem, the apparatus — will hold. Whether the name SDPF is the name the discipline eventually carries, whether the specific grammar is the grammar that propagates, whether the seventeen styles are the styles that survive — these are secondary questions. The primary question is whether the object has been correctly identified and whether the foundational result holds. Both are answered affirmatively. The rest is the work of time and adoption.

Conclusion

SDPF is not a methodology, a framework, or a set of best practices. It is the founding work of a new science: the science of specification for stochastic generators.

The object of this science — specifications that govern probabilistic production of software artefacts — did not exist as a coherent object of study before generative AI made such production possible at scale. The foundational result of this science — the Bounded Stochasticity Theorem — could not have been stated earlier because the conditions that make it non-trivial did not exist. The apparatus of this science — the Technical Verification Gate, the lifecycle protocol, the conformance classes, the provenance chain — is first-generation instrumentation for a new object, not a reconfiguration of older tools.

Prior disciplines touched this object tangentially. None contained it. Formal methods studied deterministic systems; AI safety studied the models themselves; software methodology studied human organisation; infrastructure tooling studied environmental state. The object SDPF studies sat unrecognised among these disciplines, producing anomalies — hallucination, drift, confident fabrication, surface-form entanglement — that none of them could explain from within their own frames.

The naming of the object, the statement of the foundational theorem, and the construction of the first-generation apparatus together constitute the founding of a discipline. That is what SDPF is. The work ahead — adoption, replication, the building of institutional homes, the development of communities of practice — is the work of every founded discipline in its first decades. It is the slow work, and it proceeds from foundations that have been correctly laid.

The foundations have been correctly laid.