What is claimed is:

1. A computer-implemented method for cryptographically enforced governance of an autonomous agent, comprising:

   receiving, at an execution substrate, a proposed action associated with an agent object, the proposed action being associated with one or more policy references including one or more canonical aliases provided by at least one of the agent object, the execution substrate, or a governing context;

   resolving the one or more policy references to obtain candidate external policy objects;

   filtering the candidate external policy objects based on one or more freshness constraints including at least one of a validity window, a revocation state, or an anti-rollback monotonicity constraint;

   verifying authenticity of at least one external policy object remaining after the filtering using cryptographic verification;

   determining, prior to enabling performance of the proposed action, including prior to instantiating, activating, or admitting use of an execution context or capability context, whether the proposed action is authorized under the verified external policy object; and

   permitting the execution substrate to enable performance of the proposed action only if authorized, otherwise deterministically denying the proposed action as a valid non-execution outcome.

2. The method of claim 1, wherein determining whether the proposed action is authorized comprises evaluating scope metadata that defines applicability of permissions or prohibitions to at least one of an action class, an execution substrate class, a trust zone, or a lineage class.

3. The method of claim 1, wherein determining whether the proposed action is authorized comprises evaluating a lineage continuity requirement for the agent object prior to permitting mutation, delegation, or propagation, and denying the proposed action upon failure of the lineage continuity requirement.

4. The method of claim 1, further comprising recording, in an append-only audit record, an entry corresponding to at least one of the resolution, the freshness filtering, the cryptographic verification, the authorization determination, or the denial.
5. The method of claim 1, wherein denying comprises refusing to instantiate the execution context and emitting a denial result configured to be stored in a memory field of the agent object.
6. A non-transitory computer-readable medium storing instructions that, when executed by one or more processors of an execution substrate, cause performance of the method of claim 1.
7. The non-transitory computer-readable medium of claim 6, wherein the instructions further cause denial of the proposed action upon detection of at least one of an unauthorized override attempt, a stale-policy downgrade attempt, or an unresolved lineage fork.
8. A system for cryptographically enforced governance of cognition-native semantic agents, comprising:

   a semantic agent object comprising semantic fields including at least an intent field, a memory field, a lineage field, and a policy field, wherein the policy field comprises one or more policy references including one or more canonical aliases;

   a policy resolution component configured to resolve, at runtime, the one or more policy references stored in the policy field to obtain one or more policy objects external to the semantic agent object;

   a verification component configured to verify authenticity and validity of the one or more policy objects using cryptographic verification; and

   a governance gate operatively coupled to an execution substrate and configured to, prior to enabling performance of a proposed action of the semantic agent object, including prior to instantiating, activating, or admitting use of an execution context or capability context for the proposed action, deterministically permit or deny the proposed action based on at least the verification and an authorization determination under the one or more policy objects, wherein the proposed action comprises at least one of execution, mutation, delegation, or propagation, and wherein denial of the proposed action by the governance gate results in non-execution as a valid system outcome.

9. The system of claim 8, wherein the one or more policy objects are external to the semantic agent object and are immutable absent an override authorized under the one or more policy objects.
10. The system of claim 8, wherein the authorization determination is performed independently of inferred intent scoring, model alignment scoring, predicted outcomes, introspection, or interpretability outputs associated with the semantic agent object.
11. The system of claim 8, wherein the governance gate denies the proposed action without enabling performance of the proposed action, including without instantiating, activating, or admitting use of the execution context, when the cryptographic verification fails.
12. The system of claim 8, wherein the verification component is configured to evaluate freshness constraints of the one or more policy objects, including at least one of a validity window, a revocation state, or an anti-rollback monotonicity constraint, and wherein the governance gate denies the proposed action upon a freshness failure.
13. The system of claim 8, wherein the one or more policy objects include scope metadata defining applicability of permissions or prohibitions to at least one of an action class, an execution substrate class, a trust zone, or a lineage class.
14. The system of claim 8, wherein the governance gate is configured to deny a mutation or propagation action when the lineage field fails to demonstrate continuity under a lineage constraint defined by the one or more policy objects.
15. The system of claim 14, wherein a descendant semantic agent object of the semantic agent object inherits at least one constraint defined by the one or more policy objects and applicable to an ancestor semantic agent object identified in the lineage field.
16. The system of claim 8, further comprising an override mechanism requiring approval by a plurality of authorized participants to supersede a prior policy object with a replacement policy object.
17. The system of claim 16, wherein the replacement policy object includes a cryptographic signature chain linking the replacement policy object to the prior policy object.

18. The system of claim 8, further comprising an append-only audit record configured to record governance-relevant events including at least policy resolution outcomes, verification results, authorization decisions, denials, override events, and non-execution outcomes.

19. The system of claim 8, wherein the policy resolution component is further configured to resolve a plurality of policy references including at least two canonical aliases associated with different governance domains, and to apply a deterministic precedence and combination rule to generate a composite authorization determination, including at least one of conjunctive authorization requiring joint approval, hierarchical precedence based on trust zone or governance tier, or quorum-based satisfaction of required permissions prior to permitting the proposed action.

20. The system of claim 18, wherein the governance gate is further configured to, upon denial of the proposed action, generate a structured denial artifact comprising at least an identifier of each evaluated policy object, a verification result indicator, a freshness evaluation indicator, and a failure basis indicator, and to associate the structured denial artifact with at least one of the semantic agent object, the append-only audit record, or a governing context for machine-verifiable downstream enforcement.