

# Collaborative Fiction Governance Framework

---

## What This Framework Is

This framework describes how a shared fictional world manages its own coherence and evolution. It addresses who can establish what, at what level of authority, through what process, and how those decisions can be revisited. It applies to any context where multiple contributors (or a single author over time) are building within a shared fictional world: writer's rooms, shared universes, tabletop RPG campaigns, fan fiction communities, collaborative worldbuilding projects, and solo authors maintaining consistency across a long series.

The framework is purely about governance -- the meta-layer that sits on top of whatever creative content the participants produce. It does not address the content of worldbuilding decisions (what to build), the creative processes for generating content (how to invent), or the ordering of worldbuilding work (what to build first). Those are the domains of other frameworks. This one answers a different question: once people are building together, how do they keep the shared world coherent while preserving creative freedom?

## Why Governance Matters for Fiction

A shared fictional world without governance degrades through entropy, not malice. Contributors make decisions that subtly contradict each other. Small inconsistencies accumulate until the world no longer feels internally consistent. Participants lose confidence that their contributions will be respected, which reduces investment. The world stops feeling shared and starts feeling contested.

The solution is not to restrict creativity -- it's to make the rules of play explicit so that creativity can operate with confidence. This is the same principle that makes game rules enabling rather than restrictive: basketball rules don't exist because players

might cheat. They exist so that everyone is playing the same game, which makes the play meaningful. A collaborative fiction governance framework serves the same function. By establishing what's settled, what's open, and how changes propagate, it creates conditions where contributors can be creative *confidently* -- knowing their work won't be arbitrarily invalidated and that other people's work won't arbitrarily break theirs.

The constitutional analogy is deliberate and structural. A constitution doesn't list rules -- it establishes the process by which rules get made, changed, and interpreted. It operates at multiple levels with different authorities and different thresholds for change. A fictional world's governance needs the same architecture: foundational commitments that are hard to change, working decisions that are easier to revise, and a clear process for how things move between levels.

This is governance by consent. The participants agree to the framework because it enables the collaborative work they want to do. The framework has no enforcement mechanism beyond the social contract of participation: if you're playing this game, these are the rules of play. You can propose changing the rules through the framework's own processes. You can leave the game. What you can't do is unilaterally ignore the rules while claiming to be part of the shared project.

---

## Component 1: The Canon Hierarchy

The canon hierarchy defines the levels of authority that creative content can hold within a shared world. Each level carries specific implications for how other contributors should relate to that content -- what they can rely on, what they can build on, what they should treat as provisional, and what exists outside the shared canon entirely.

### The Five Levels

**Foundational Canon** is the constitutional layer. It contains the commitments that define the project: the genre, tone, core premises, physical laws, and structural principles that make this world *this world* rather than some other world. If you changed

foundational canon, you would have a different project. Everything at every other level must be consistent with foundational canon. Changes at this level have the highest threshold because they cascade through everything built on top of them.

Foundational canon is typically small. A few pages, not a volume. It answers questions like: What kind of stories does this world generate? What are the physics and metaphysics? What is the emotional register? What structural commitments constrain all future development? The smaller and more precise the foundational layer, the more room it leaves for creative development at other levels while still providing coherence.

Examples of foundational canon: "This is a hopeful-tone medical drama set in a multi-species space community." "Magic has costs that are paid by the user, never by the target." "No species has a single culture." "The system is trying and failing in specific ways, not malicious." These are the commitments that, if violated, would make a contributor's work incompatible with the shared world regardless of its individual quality.

**Established Canon** is the statutory layer. It contains decisions that emerged from actual creative work -- stories written, sessions played, content produced -- and have been accepted into the shared world as settled. Established canon is load-bearing: other content has been built on top of it. Changing it requires understanding what depends on it and managing the cascade.

Established canon grows over time as the project accumulates creative work. A species whose biology was described in a published story. A political event that multiple stories reference. A character relationship that has been developed across several contributions. These began as creative inventions and became established through use and acceptance. The bible or reference document of a shared world is primarily a record of established canon.

The distinction from foundational canon: established canon was *derived* from creative work, not established as a precondition for it. It can be changed (with appropriate process) without changing the fundamental identity of the project. Changing the fact that Species B's joining process involves physical transformation is a big deal that affects a lot of downstream content. But it doesn't change what *kind of project* this is, the way changing the tone from hopeful to grimdark would.

**Working Canon** is the common-law layer. It contains content that has been contributed and is being treated as part of the shared world, but hasn't been fully stress-tested against other content or formally accepted as established. This is where most active creative work lives. A story currently being written. A worldbuilding detail introduced in the last session. A character trait established in recent episodes. It's canon *for now* -- other contributors should treat it as real and build on it provisionally -- but it hasn't hardened yet.

Working canon is the most dynamic level. Content enters here when it's contributed and accepted into active use. It moves up to established canon when it's been built on enough that changing it would require managing dependencies. It moves back down to exploratory (or gets withdrawn entirely) if stress-testing reveals problems. The defining property of working canon is that contributors can rely on it without treating it as immovable. It's the "current state of play."

Most collaborative fiction operates primarily at this level. The governance framework's main practical contribution is distinguishing working canon from established canon -- helping contributors know which parts of the current state of play are load-bearing and which are still in flux.

**Exploratory Canon** is the laboratory. It contains content that is explicitly speculative -- someone is trying something out, testing a direction, proposing an addition that hasn't been integrated into active use yet. Exploratory content exists within the shared creative space (it's visible to all contributors, not private), but other contributors should not build dependencies on it. It might become working canon if it holds up. It might be withdrawn or substantially revised. Its purpose is to give contributors a way to experiment without the pressure of commitment.

Exploratory canon reduces the cost of creative risk. A contributor who has an idea that might not work can introduce it at the exploratory level, get feedback, and iterate without the idea being treated as a *fait accompli*. This is especially important in larger collaborative groups where a premature contribution to working canon might be difficult to walk back once other people have started building on it.

**Divergent Canon** is the playground. It contains content that deliberately contradicts established or foundational canon -- not by accident but as a conscious creative choice to explore "what if this were different?" Alternate timelines, alternate universes, what-

if scenarios, AU fan fiction, experimental rewrites. Divergent content is clearly marked so that no contributor mistakes it for canon at any level.

Divergent canon serves two functions. First, it gives contributors a space for creative work that the shared canon's constraints don't permit -- a way to explore roads not taken without threatening the coherence of the main world. Second, it occasionally produces ideas that are good enough to feed back into the main canon through the change process. A divergent exploration might reveal that a foundational commitment is unnecessarily restrictive, or that an established detail would work better if revised. The divergent space is where those discoveries can happen safely.

## How the Levels Relate

The hierarchy is not just a ranking -- it's a set of relationships that determine how content at different levels interacts.

**Consistency flows downward.** Content at any level must be consistent with content at all levels above it. Established canon must be consistent with foundational canon. Working canon must be consistent with both. Exploratory canon must be consistent with all three. Divergent canon is the exception -- it's explicitly exempt from consistency requirements, which is precisely what makes it divergent.

**Authority flows downward.** When content at two levels conflicts, the higher level takes precedence. A working-canon detail that contradicts established canon is an error that needs resolution (the working-canon detail is probably wrong, or the established canon needs revision through the change process). A divergent-canon detail that contradicts foundational canon is operating as intended.

**Content flows upward.** New creative work enters at the lower levels and may move up over time. Exploratory becomes working when it's accepted into active use. Working becomes established when it's been built on enough to be load-bearing. In rare cases, established canon that proves to be truly foundational to the project's identity might be recognized as foundational. Content does not skip levels -- the progression from exploratory to working to established represents increasing integration, stress-testing, and dependency-building that can't be shortcut.

**Dependency weight increases upward.** Changes at higher levels cascade further. Changing foundational canon potentially affects everything. Changing established canon affects everything built on it. Changing working canon affects current projects that reference it. Changing exploratory canon affects nothing beyond the exploration itself. This is why the change threshold increases with level -- the cost of change increases with the amount of dependent content.

---

## Component 2: The Authority Structure

Governance requires someone (or some process) to make decisions. The framework identifies four governance roles that every shared world needs. In a small group, one person might hold all four. In a large community, they might be distributed across committees, elected positions, or defined processes. The framework describes what each role does and the tradeoffs of different distributions, without prescribing a specific structure.

### The Four Roles

**The Founder role** establishes and maintains foundational canon. This role has the highest authority and the narrowest scope -- it can set or change the constitutional layer but doesn't control day-to-day creative decisions at other levels. In a single-author project, the author is the founder. In a writer's room, this might be the showrunner or the original creator. In a community project, this might be the originating group or a designated constitutional committee.

The founder role is the one most likely to be held by a specific person or small group rather than distributed broadly. Foundational canon defines the project's identity, and diffusing that authority too widely risks the project losing coherent identity. However, the founder role is not absolute -- it is itself subject to the meta-governance layer (Component 5). The founder can be wrong about what should be foundational, and the process for correcting this must exist even if it's difficult to invoke.

**The Steward role** manages the body of established and working canon -- accepting contributions, tracking dependencies, identifying inconsistencies, and maintaining the reference documents (bible, wiki, campaign notes) that record the current state of the

shared world. This is the most labor-intensive role and the one most likely to be undervalued. Without active stewardship, the shared world's canon becomes ambiguous -- people disagree about what's been established, dependencies go untracked, and contradictions accumulate unnoticed.

In a small group, the steward might be whoever maintains the shared document. In a larger project, this might be an editor, a continuity manager, or a wiki team. The steward role doesn't require creative authority (the steward doesn't decide whether a contribution is *good*), but it requires the authority to flag inconsistencies and to designate content at the appropriate canonical level. A steward who notices that a new contribution contradicts established canon has the authority and responsibility to raise the issue, even if they don't have the authority to resolve it unilaterally.

**The Contributor role** is the creative engine. Contributors produce the actual content of the shared world -- stories, sessions, worldbuilding documents, character development, whatever form the creative work takes. Every participant in a collaborative fiction project fills this role, including people who also hold other roles.

The contributor's primary right is creative freedom within the constraints of higher-level canon. A contributor can introduce new content that is consistent with foundational and established canon without requiring approval -- it enters at the working-canon level by default. The contributor's primary responsibility is consistency: checking that their contributions don't contradict content at higher canonical levels, and flagging when they're unsure. Contributors also have the right to propose changes to any canonical level through the change process.

**The Arbiter role** resolves conflicts -- between contributions that are inconsistent, between contributors who disagree about canon, and between any other disputes that the normal process doesn't resolve. The arbiter role is the one a healthy project invokes least often, because most consistency issues are caught by stewardship and most creative disagreements are resolved by discussion. But it needs to exist for the cases that don't resolve naturally.

In a small group, the arbiter might be the founder by default, or the group might arbitrate collectively. In a larger project, this might be a designated authority or a defined voting process. The key property of the arbiter role is that its decisions are

*binding within the framework* -- if the arbiter resolves a conflict, the resolution is accepted as canon at the appropriate level. Without this property, the conflict resolution process has no endpoint, and disputes can persist indefinitely.

## Distribution Patterns

How these roles are distributed depends on the size, structure, and culture of the collaborative group. The framework identifies three common patterns, but hybrids and other arrangements are possible.

**Concentrated authority** places most or all roles in a single person or very small group. This is typical of single-author projects, showrunner-led writer's rooms, and game-master-led RPG campaigns. It's efficient and produces coherent worlds, but it creates a single point of failure (if the authority-holder leaves, the project's governance goes with them) and can suppress creative contributions from other participants if the authority is exercised too broadly.

**Delegated authority** distributes roles across defined positions while maintaining a clear hierarchy. The founder holds foundational authority, delegates stewardship to an editor or continuity team, and delegates arbitration to a defined process. Contributors have explicit creative freedom within their domain. This is typical of larger writer's rooms, shared-universe publishing projects, and organized collaborative worldbuilding groups. It scales better than concentrated authority but requires more explicit process definition to work well.

**Distributed authority** spreads roles across the community through democratic or consensus-based processes. The foundational canon might be established by community vote. Stewardship might be shared among active contributors. Arbitration might use a defined community process. This is typical of open fan fiction communities, collaborative wikis, and community-governed RPG settings. It's the most inclusive but requires the most process overhead and is most vulnerable to the problems of diffuse responsibility (nobody tracks consistency because everyone assumes someone else is).

No distribution is inherently better. The right choice depends on the project's size, goals, and participants' expectations. The framework's contribution is naming the roles that need to exist regardless of how they're distributed, so that a group can make a

conscious choice about distribution rather than leaving governance implicit.

---

## Component 3: The Change Process

Canon is not static. The shared world evolves as creative work continues, as contributors discover implications of earlier decisions, and as the project's needs change. The change process defines how content moves between canonical levels and how settled decisions get revisited.

### Upward Movement: How Content Gains Authority

Content enters the shared world at one of three levels:

**New creative contributions** (stories, sessions, worldbuilding additions) enter at the **working canon** level by default. They are treated as part of the shared world immediately, and other contributors can build on them. No approval process is required beyond consistency with higher-level canon. This keeps the creative process fluid -- contributors don't need permission to contribute, they need consistency.

**Explicitly speculative contributions** (proposals, experiments, rough ideas) enter at the **exploratory** level when the contributor marks them as such. This is a signal to other contributors: "I'm trying this out; don't build dependencies on it yet." Exploratory content can be discussed, iterated on, and revised without the pressure that comes with working-canon status.

**Deliberately contradictory contributions** enter at the **divergent** level when the contributor marks them as such. This is a signal: "I know this contradicts the shared canon; that's the point of this piece." No approval is required for divergent contributions because they don't affect the shared canon.

**Exploratory to working:** An exploratory contribution becomes working canon when it is accepted into active use -- other contributors begin treating it as real and building on it. This can happen through explicit group decision ("let's go with this"), through the steward's recognition that the content is being treated as canon in practice, or through

the contributor's own declaration that they're ready to commit it. The key threshold is: *is anyone building on this?* If yes, it's working canon regardless of how it was originally designated.

**Working to established:** Working canon becomes established canon when it has accumulated enough dependencies that changing it would require managing a cascade. This is not a single moment -- it's a gradual hardening as more content is built on top of the decision. The steward role is responsible for recognizing when working canon has become load-bearing and updating its designation accordingly. Formal recognition can also be triggered by a contributor who wants to know whether a piece of working canon can be relied on long-term.

The question for this transition: *if this changed, what else would break?* If the answer is "nothing much," it's still working canon. If the answer is "several other established things depend on it," it's established. The dependency weight, not the age or the quality, is what distinguishes the two levels.

**Established to foundational:** This is rare and should be treated as significant. A piece of established canon might be recognized as foundational when the participants realize it has become definitional -- the project wouldn't be the same project without it. This typically happens through a formal process (group discussion, founder recognition, community vote) rather than through gradual accumulation. Moving something to the foundational level is a commitment that it will be very difficult to change, so the decision should be deliberate.

## Downward Movement: How Content Loses Authority

Content can also move down the hierarchy, though the process is different for each transition.

**Established to working:** This happens when the group decides that a previously settled decision needs to be reopened. The triggers might be: a new contribution reveals that the established decision creates problems that weren't anticipated, or the group's creative direction has shifted in ways that make the old decision constraining rather than enabling, or a contributor makes a compelling case that the decision was premature. Reopening established canon requires the decision to be flagged (by any

contributor), discussed (by stakeholders -- meaning anyone whose contributions depend on it), and approved (by the arbiter or by consensus, depending on the authority structure). The decision then reverts to working-canon status, meaning it's still in the world but is now acknowledged as mutable.

**Working to exploratory or withdrawn:** Working canon that turns out to be problematic -- inconsistent with higher-level canon, producing unintended consequences, or simply not working creatively -- can be downgraded to exploratory (if it might be salvageable with revision) or withdrawn entirely. Because working canon hasn't accumulated heavy dependencies, this is relatively low-cost. The contributor who introduced it should be consulted, but the steward can flag the issue and the group can resolve it through normal discussion.

**Foundational revision:** Changing foundational canon is the hardest process in the framework -- deliberately so, because foundational changes cascade through everything. The process should require: a formal proposal stating what should change and why, an impact assessment identifying what established and working canon would be affected, discussion among all active contributors, and a high-threshold approval (consensus, supermajority, or founder approval, depending on the authority structure). Foundational revision is the constitutional amendment process. It should be possible but costly.

## Dependency Interaction

The change process interacts with the worldbuilding dependency chain (if the group is using one). Changes to content at lower dependency layers (physics, biology, macro setting) cascade further than changes to content at higher dependency layers (local details, individual story events). This means:

The canonical level of a piece of content should generally correspond to its dependency depth. Foundational canon tends to live at the lowest dependency layers. Established canon spans the middle layers. Working and exploratory canon tend to live at the higher layers where creative work is most active.

When a change is proposed, the impact assessment should trace the dependency chain *downward from the change*: what content was built on top of the thing being changed? The depth of the cascade determines how seriously the change needs to be

treated, regardless of the content's current canonical designation. A working-canon detail that turns out to have three stories built on it may need to be treated with the same care as an established-canon decision, because the practical cost of change is the same.

---

## Component 4: The Consistency Protocol

Contradictions are inevitable in collaborative fiction. The consistency protocol defines how they are detected, classified, and resolved.

### Types of Inconsistency

**Hard contradictions** occur when two pieces of content make claims that cannot both be true. "Species A has four limbs" in one contribution and "Species A has six limbs" in another. These require resolution -- one must be changed, or both must be revised to be compatible.

**Soft contradictions** occur when two pieces of content are in tension but not logically incompatible. One story portrays a character as reserved; another portrays them as gregarious. Both could be true (people behave differently in different contexts), but the tension needs to be acknowledged -- either as deliberate characterization or as an inconsistency to address.

**Implication conflicts** occur when a piece of content doesn't directly contradict anything but implies something that conflicts with established or foundational canon. A story that casually mentions a technology implies that technology exists in the world; if that technology contradicts the world's established physics, the implication is the problem even though nothing was directly contradicted. These are the hardest to catch and the most important responsibility of the steward role.

**Tone violations** occur when content is technically consistent with the canon but violates the foundational tone or thematic commitments. A grimdark torture scene in a hopeful-tone world. A comedic romp in a serious drama. These are foundational-level inconsistencies even though they don't contradict any specific factual claim.

## Detection

Consistency checking is primarily the steward's responsibility, but every contributor shares in it. The framework recommends:

**Pre-contribution checking:** Before introducing significant new content, the contributor checks it against the reference document (bible, wiki, campaign notes) for consistency with established and foundational canon. The depth of checking should be proportional to the content's impact -- a minor scene detail doesn't need the same scrutiny as a new species' biology.

**Peer review at intake:** When new content enters working canon, at least one other participant (ideally the steward) reviews it for consistency. This doesn't need to be formal or heavy -- a quick read-through looking for contradictions is usually sufficient. The goal is to catch problems before dependencies build up.

**Periodic audits:** For long-running projects, the steward periodically reviews the accumulated canon for inconsistencies that slipped through. This is especially important for implication conflicts, which are easy to miss at the individual-contribution level but become visible when the whole body of canon is surveyed.

## Resolution

When an inconsistency is detected, the resolution process depends on the canonical levels involved.

**Between content at the same level:** If both pieces are working canon, the contributors discuss and agree on a resolution, with the steward facilitating if needed. If both are established canon, the issue is escalated to the arbiter because changing either piece has dependency implications.

**Between content at different levels:** Higher-level canon generally takes precedence. Working canon that contradicts established canon is presumed to be the error, and the working-canon contributor is asked to revise. However, if the contributor believes the established canon is what should change, they can initiate the change process (Component 3) to argue for revision.

**Irreconcilable contradictions:** When two important pieces of content genuinely cannot be made consistent and neither contributor will yield, the arbiter decides. The arbiter's decision should be guided by: which resolution is more consistent with foundational canon, which resolution preserves more dependent content, and which resolution better serves the project's creative goals. The arbiter's decision is binding, but either contributor can propose a change through the normal process if they believe the decision was wrong.

**Retroactive consistency:** Sometimes a contradiction is discovered long after both pieces of content have been in use and built upon. In these cases, the resolution that disturbs the least existing content is generally preferred. A small retcon to one piece is better than cascading changes through many pieces. If neither piece can be easily changed, the group may accept the inconsistency as a known imperfection -- logged in the reference document so that future contributors are aware of it and can work around it.

---

## Component 5: The Meta-Governance Layer

The governance framework itself is subject to change. The meta-governance layer defines how the framework's own rules can be modified, what parts are structural (recommended for all projects) and what parts are parametric (set per-project), and how the group maintains the framework over time.

### What's Structural vs. Parametric

**Structural elements** are the components that the framework asserts every collaborative fiction project needs, regardless of size or context. These are:

- Some form of canon hierarchy (the specific levels and their definitions may vary, but the principle that different content has different authority is structural)
- Defined governance roles (the specific distribution varies, but the roles of founder, steward, contributor, and arbiter need to exist)

- A change process with thresholds that increase with canonical level (the specific thresholds vary, but the principle that higher-level changes should be harder is structural)
- A consistency protocol (the specific mechanisms vary, but some process for detecting and resolving contradictions needs to exist)
- The ability to modify the governance rules themselves (this component)

A group can decline any structural element, but they should do so consciously and understand what they're giving up -- in the same way that a game can be played without some of its rules, but the players should know which rules they're waiving and what changes as a result.

**Parametric elements** are the settings that each project configures for itself:

- The number and definitions of canonical levels (five is the framework's recommendation, but a small project might use three and a large one might use seven)
- How the governance roles are distributed (concentrated, delegated, distributed, or hybrid)
- The specific thresholds for changes at each canonical level (consensus, majority vote, founder approval, etc.)
- The specific consistency-checking processes (formal peer review, informal discussion, steward-led audits, etc.)
- Whether divergent canon is actively encouraged, passively permitted, or not part of the project's scope
- How the reference document is maintained and who has edit access

The parametric settings should be established when the project begins and recorded in the project's governance document (a companion to the creative reference document). They can be changed through the meta-governance process.

# Changing the Framework

Changes to the project's parametric settings require a lower threshold than changes to foundational canon, but a higher threshold than normal creative decisions. The recommended process:

- Any participant can propose a governance change
- The proposal should state what change is requested and why the current setting isn't working
- All active participants should have the opportunity to respond
- Approval requires a higher threshold than normal decisions (the specific threshold is itself a parametric setting, but the framework recommends consensus or near-consensus for governance changes)
- Changes take effect after a defined transition period, not immediately, so that in-progress work isn't disrupted

Changes to the structural elements -- dropping a whole component of the framework - - should be treated with the same gravity as foundational canon changes. The group can do it, but they should understand what they're removing and why.

## Framework Adoption

When a group adopts this framework, the process is:

1. The group reviews the framework together and decides which structural elements to adopt (ideally all five, but the framework should not be imposed without understanding).
2. The group sets the parametric elements through collaborative discussion, guided by the project's size, goals, and culture.
3. The decisions are recorded in a governance document that sits alongside the creative reference document.
4. The governance document is itself subject to the change process defined within it -- the framework is self-amending.

A group that starts without governance and adopts this framework mid-project faces an additional challenge: all existing content needs to be designated at the appropriate canonical level. This is a one-time stewardship task that can be done gradually -- the most important content (foundations and heavily-dependended-on decisions) should be designated first, with less critical content designated as it comes up in the normal course of work.

---

## Relationship to Other Frameworks

This framework operates at a different level than creative or worldbuilding frameworks. It governs the *process* of collaborative creation; the other frameworks govern the *content* of creation. They interact but don't overlap.

**Dependency-Based Ordering** determines what worldbuilding work should be done in what order. The governance framework uses the dependency chain to inform change-process decisions: changes at lower dependency layers cascade further and require higher thresholds. But the dependency ordering doesn't address who has authority to make decisions or how contradictions between contributors are resolved -- that's governance.

**Generative Constraints, Systemic Worldbuilding, Species Design frameworks, etc.** are creative tools that contributors use to produce content. The governance framework doesn't constrain what tools contributors use; it constrains how their outputs relate to each other and to the shared canon.

**The project's creative reference document** (bible, wiki, campaign notes) is the *product* of governance -- the record of the current canonical state of the shared world. The governance framework defines the process that produces and maintains that document. The reference document says "here is what's true in the world." The governance framework says "here is how things become true, stop being true, and get revised."

---

# Design Notes

## Why Five Levels

Three levels (settled / active / speculative) capture the minimum useful distinction but collapse two important differences. Established and working canon have different change costs -- established has dependencies; working doesn't -- and treating them the same either over-protects working canon (making contributors hesitant to revise recent work) or under-protects established canon (making it too easy to change load-bearing decisions). Similarly, exploratory and divergent serve different functions -- one is a staging area for potential canon; the other is a space for creative play outside canon -- and collapsing them confuses whether a contribution is a proposal or an alternate take.

Seven or more levels are possible for very large projects but risk the framework becoming a bureaucratic overhead rather than an enabling structure. Each additional level needs to carry a distinction that affects how contributors relate to the content; if two levels produce the same contributor behavior, they should be merged.

## Why Consent, Not Enforcement

The framework deliberately lacks enforcement mechanisms beyond social agreement. This is a design choice, not a gap. Collaborative fiction is a voluntary creative activity. If the framework required policing, it would be badly designed -- it would be imposing costs that exceed the benefits of participation. The framework works when participants *want* to follow it because it makes the collaborative work better. If participants don't want to follow it, the problem isn't enforcement -- it's that the framework isn't serving the project, and the meta-governance process should be used to fix it.

## Why the Game-Rules Analogy Over the Legal Analogy

The constitutional terminology is useful for communicating what the framework does - - "rules about rules" is immediately clear when phrased as "constitutional." But the framework's operational model is closer to game rules than legal systems. Game rules exist so that play can happen. They're agreed to by the players before play begins.

They can be modified by player agreement. They don't require external enforcement because participation is voluntary. They're enabling, not restrictive -- the rules are what make the game *a game* rather than a formless activity.

The legal analogy is tempting but carries adversarial connotations (enforcement, prosecution, punishment) that are counterproductive for creative collaboration. The game-rules model captures the same structural properties (levels of authority, change processes, consistency requirements) without the adversarial baggage.

## **Why the Steward Role Is Critical**

Of the four governance roles, stewardship is the one most likely to be neglected and the one whose absence causes the most damage. Founding, contributing, and arbitrating are all visible, high-status roles. Stewardship -- maintaining the reference document, tracking dependencies, checking consistency, designating canonical levels - is invisible, unglamorous, and continuous. It's also what actually keeps the shared world coherent over time.

Projects that fail at collaborative fiction governance almost always fail at stewardship. The creative work happens; the governance decisions get made; but nobody tracks what was decided, what depends on what, or where the contradictions are accumulating. The framework names stewardship explicitly because making it visible is the first step toward ensuring it happens.

## **Multi-Project Worlds and the Tree Model**

When multiple creative projects share an underlying world, the governance structure isn't a flat hierarchy with a parent world and child projects. It's a living tree with bidirectional flow.

The root system is the shared universe canon -- physics, species biology, deep history, the foundational substrate that all projects depend on. The trunk and branches are the structural connections between projects, with branch points where shared history diverges into project-specific territory. The leaves are the individual projects where creative work actually happens -- where new content is generated.

The biological tree captures something essential that a simple hierarchy misses: flow runs in both directions. Constraint flows upward from root to leaf, like water -- the shared universe canon feeds each project with the substrate it needs, and each project must be consistent with what the roots provide. Contribution flows downward from leaf to root, like sugar -- the projects do the creative work, and the products of that work (species details, historical events, cultural specifics invented through storytelling) feed back into the shared root system, enriching it for all projects to draw on.

The root system grows *because* the leaves feed it. A shared universe with no active projects is a root system with no sugar production -- it doesn't develop further. The creative work at the project level is what drives growth at every level of the tree. This means the universe-level canon isn't a finished thing that projects draw from; it's a living system that gets richer as the projects produce.

Two projects that share a branch point (closer together in setting, time, or subject matter) have more shared canon between them than two projects on distant branches. The branching pattern determines how much canon is shared and how tightly coupled the projects are.

The governance implications of this model are developed in the [implementation guide](#). The principle is that governance must account for flow in both directions -- not just how the shared world constrains projects (root to leaf), but how project-level creative work enriches and sometimes changes the shared world (leaf to root), and how to handle situations where two projects on different branches are both contributing to the same part of the root system.

---

This work is marked with CC0 1.0 Universal. To the extent possible under law, the author has waived all copyright and related or neighboring rights to this work.