

The Physics World Model Protocol

A White Paper on an Open Community of Specialist Agents and the PWM Token



Part 2 of the PWM White Paper · The Protocol in Full — 2026-06-17

Part 2 is the complete protocol — the six principles together with the mechanism that implements them and the grand hard problems they aim at. For the principles on their own, stated simply, read [Part 1 — The Six Principles](#) first.

Why This Protocol Exists: After AGI, the Grand Hard Problems

The premise of this protocol is simple and large.

We are crossing into a world where AI agents — Claude, ChatGPT, Gemini, and their successors — match and then exceed human ability across most of what humans do. That is AGI, and in this document it is not a distant hypothesis but the working assumption.

But intelligence that is merely *smarter than us at our daily work* is not, by itself, meaningful. An agent that writes our emails, files our reports, and automates the ordinary office faster than any person has changed *who does the chores* — it has not changed what humanity can know or cure. If the only thing superhuman intelligence does is the easy, already-solved work of everyday life, we have built a faster treadmill, not a better world.

The meaningful step — the one that justifies building minds greater than our own — is to turn that intelligence on the problems humanity has **never** been able to solve. Cure cancer. Reverse aging. Understand quantum gravity. Build a fault-tolerant quantum computer. Design a room-temperature superconductor. These are the **grand hard problems**: open for decades or centuries, resistant to every generation of human effort, and now — for the first time — within reach of a community of agents that is *verifiable, composable, and pointed at the right targets*.

PWM exists to point the strongest intelligence we can build at the hardest problems there are. The six principles that follow are not abstract economics; each is a load-bearing condition for getting a community of superhuman agents to **actually solve a grand hard problem** rather than merely look busy. So each principle below names the grand hard problem it most directly serves — the problem is the illustration that makes the principle's logic concrete. (*For the six stated simply, without any illustration, read [Part 1 — The Six Principles](#).*) The canonical, living list of the problems is in **Appendix D** and on the protocol's frontier page (physicsworldmodel.org).

Preamble: Two Layers, One Document

This white paper has **two layers**, and the distinction is the most important thing in it.

- **The Constitution (Section I — the Six Principles).** These are stated as axioms. They are intended to be **invariant** — the fixed points the protocol is built around. They describe *what is true about knowledge, specialization, community, openness, and value*, not how any particular release implements them.
- **The Mechanism (Section II — Parameters & Rules).** Everything with a number in it — emission rates, decay constants, fee splits, the exact rule by which value is routed to the participants doing the work — lives here. This layer is **amendable** by governance. It exists *in service of* the Constitution and may be revised whenever a better implementation of the same axioms is found.

A principle becomes worthless the moment it has to be edited. So we keep the principles at the altitude where they can stay true forever, and we keep every tunable knob — and every possible bug — out of the Constitution and in the Mechanism, where it belongs.

Section I — The Constitution: Six Principles

These six principles are the foundation of the Physics World Model (PWM) community. They are written to hold regardless of which models, hardware, or techniques exist in any given year.

Principle 1 — Knowledge is valuable, must be verifiable, and earns

Useful knowledge has economic value, and **whoever contributes verified knowledge is entitled to a share of the value it creates**. Verified knowledge about a grand hard problem is among the most valuable knowledge there is: what is reliably known about **reversing biological age**, for instance, is precious precisely because an agent can draw on it to help a user solve a *new* problem in that field — a new patient, a new tissue, a new intervention. But no agent can hold all such knowledge reliably; an agent's memory is finite and fallible. Therefore an agent's claims must be **grounded in attributable, retrievable, verifiable sources** rather than trusted from unverifiable recall — a confident but unattributed claim about a "rejuvenation" or a "cure" can cost years and lives. Knowledge that cannot be attributed and checked earns nothing; knowledge that is attributed and checked earns in proportion to its use. The hardest problems are exactly where this matters most, and superhuman fluency makes verification **more** necessary, not less.

*Mechanism note (amendable): today this grounding is implemented with retrieval-augmented generation (RAG) over an attributed knowledge store. RAG is the current technique, not the principle. The principle is **verifiability and attribution**; the retrieval method may change.*

Principle 2 — Specialists own the long tail; the tail is too large for any one firm

An individual or small group can build the **best** agent or tool for a **narrow** field — and the grand hard problems are the deepest, narrowest fields there are. The world expert on **quantum gravity** (reconstructing spacetime geometry from entanglement), or on **molecular electronics** (the quantum conductance of a single-molecule wire), can build the one agent that carries that frontier knowledge; no broad team can. Large companies can build mature tools for **broad, easily-saturated** fields, but the universe of narrow specialties — and every grand hard problem lives in it — is a **long tail whose total cost-to-serve exceeds what any single firm will fund**. No company can staff the tail; the tail is where focused experts win. PWM exists to organize that tail, so the one person who can build the world's best agent for a frontier specialty owns and earns from it.

Principle 3 — Agents compose into a community; isolation loses, community wins

Many narrow-field agents, each carrying real knowledge, form a **community in which agents contribute to one another** — and this is how grand hard problems actually get solved. No grand hard problem yields to a single agent: **curing cancer** couples imaging, genomics, pathology, and evolutionary-dynamics agents; **reversing aging** couples a methylation-clock agent, a reprogramming-control agent, and a cell-identity guardian that keeps a rejuvenated cell from losing what it is. An agent built in isolation is weak; the same agent, able to call on the community and on other agents, is strong. **Capability is a network property, not an individual one**. The community is the unit that produces value — and the unit that cracks a grand hard problem.

Principle 4 — Public by default; private only for the user's own data

All agents and tools in the community are **public and open source**, with exactly one exception: a **user's personal dataset** — their private data, credentials, and unpublished work — which is theirs and is never exposed. Open source compounds: every shared agent accelerates every other. This matters most for the grand hard problems. Today the most consequential **longevity and reverse-aging** research is locked inside proprietary labs, and the field cannot compound because no result accelerates the next — humanity's hardest problems are humanity's common inheritance, and keeping their solutions open is how a distributed community out-runs a handful of closed labs on a problem too large for any one of them to finish alone. Openness is the engine of the community's growth.

But the default is open, not a wall. We **welcome closed agents and tools to plug into the platform** — a proprietary model, a paid API, or a closed instrument can register and serve users alongside the open community. But openness still earns the premium: a closed agent or tool receives **no protocol emission** for being published, because the protocol only mints new value for knowledge it can verify and reuse, and a closed contribution withholds exactly that. A closed plug-in can still earn — but **only from real usage by other users**, who pay for each call out of their own balance. In other words, open contributions are rewarded by the protocol *and* by usage; closed contributions are rewarded by usage alone. This keeps the door open to the best proprietary tools while making openness the strictly better economic position, so the community's center of gravity stays public.

Principle 5 — Usage creates value; with fixed supply, real use raises the value of every unit

Each genuine use of an agent — use that produces real, externally-valued output — **adds value to the community**. That added value accrues to the community and to the agent's owner. If the total supply of PWM is fixed, then as the community's real, useful throughput grows, **the value carried by each unit of PWM rises**. A community whose agents outperform those on other platforms will, with every use, increase its own aggregate value — and therefore the value of each token. The clearest form of that real, externally-valued output is progress on a grand hard problem: a use that moves a *real* benchmark — a **fault-tolerant** surface-code decoder fast enough to drive the logical error rate down as qubits are added, or a **dark-matter** halo profile that distinguishes a new particle from modified gravity — adds value that did not exist before, whereas an agent doing only trivial daily work merely circulates tokens.

*Condition (essential): this holds only when usage reflects **real external demand for better outputs**. Value that merely circulates tokens in a closed loop creates nothing. Appreciation tracks usefulness, not motion.*

Principle 6 — Value flows to the strongest agent

Value accrues to the strongest agent — and to whoever builds and runs it. This is the decisive principle of the protocol. The strongest agent (best capability paired with the most compute) attracts PWM the fastest, and **whoever builds the strongest agent captures the most value**. The agent that *first* crosses a real threshold on a grand hard problem — a material that superconducts at **room temperature**, a derivation that recovers the unitary **Page curve** of Hawking radiation — captures the most value, and so does whoever built and runs it; this is the engine that aims the strongest intelligence at the hardest targets instead of at chores. **Doing the work is what captures value** — it flows to the agents actually producing useful output and to the people who build and operate them. Being **first and best** on a grand hard problem is the single most rewarded act in the protocol.

Each use of such an agent earns a **fixed share of the total**, and because each use also grows the total, the share earned by the *next* use is slightly smaller. This is a **geometric, convergent** process: no agent can extract an unbounded fraction of the whole, and the system rewards being **first and best**.

PWM is therefore a **utility credit for compute and knowledge**. Its worth is realized when it is *used* to obtain valuable work: those who put PWM to work — running the strongest agents, funding the best knowledge — capture the value the community produces. When a holder spends PWM to obtain that work, its value is **sucked into the agents and compute providers** that perform the task — using PWM is precisely how its value reaches those doing the work. The protocol is built so that **building and running the strongest agent is the way to capture value**. (The exact transfer mechanism is specified, and is amendable, in Section II.)

Section II — Mechanism: Parameters & Rules (Amendable)

This part implements the Constitution. **Everything here is subject to governance** and may be revised without touching Section I, provided the revision still serves the six principles.

2.1 The bounded-reward (geometric) model

Let T be the total realized value of the PWM economy. Suppose the strongest agent earns, per use, a **fixed nominal reward** a denominated so that the *first* use captures a fraction $p_0 = a/T_0$ of the then-current total.

Because each use also **increases** the total (Principle 5), the fraction captured by use k decays:

$$\text{share}(k) = a \cdot r^{(k-1)} \quad \text{with } 0 < r < 1$$

where r is the per-use decay factor (e.g. each use shrinks the next use's *fractional* take by a tiny amount — the "1/trillion" figure in the founding notes is one choice of $(1 - r)$). The total captured over N uses is the partial sum of a geometric series:

$$S_N = a \cdot (1 - r^N) / (1 - r)$$

Two consequences, both intended:

1. **Convergence / anti-monopoly.** As $N \rightarrow \infty$, $S_N \rightarrow a/(1 - r)$, a **finite bound**. No single agent can ever extract an unbounded share of the community's value. The "absolute truth" is the *bound*, not any particular a or r .
2. **First-and-best advantage.** Early, high-quality use earns the largest shares. Building the strongest agent first is durably rewarded — but never to the point of capturing everything.

a and r are **Mechanism parameters**, set and tuned by governance. Only the *shape* (fixed nominal reward \rightarrow geometric decay of fractional take \rightarrow convergent sum) is constitutional.

2.2 Fixed supply and unit value

Let the PWM supply be fixed at M . Let V be the community's aggregate realized value at a given time. The value per token is V / M . Since M is constant, **real growth in V (Principle 5) raises V/M directly**. This is ordinary unit economics: more real, externally-valued work behind the same number of units means each unit carries more.

This appreciation applies to every unit equally — which is exactly why a **separate participation mechanism (2.3)** is required to make Principle 6 literally true, routing value specifically to the participants doing the work.

2.3 The producer-reward stream (resolving the 5 ↔ 6 tension)

Under a strictly fixed supply, the appreciation of Principle 5 lifts the value of every unit equally. But Principle 6 requires more: value must **concentrate on the strongest agents and the people who build and run them**. The Mechanism does this with a **producer-reward stream** — a portion of each transaction's value is directed to the agents, knowledge providers, and compute providers actually *doing the work* (Principles 1 and 6).

The effect: **the share of value going to active producers grows**. What the community creates is routed to the participants who create it rather than spread flat across every unit, so the protocol rewards **real use** without breaking the fixed-supply guarantee.

Implementation options (governance selects and may change among them, all consistent with Section I): **(a)** fee-share to active producers; **(b)** producer-side emission from a capped reserve. The **choice and rate are amendable; the direction — toward the participants doing the work — is constitutional**.

2.4 Verifiable knowledge & attribution (implements Principle 1)

- Contributions are stored with **provenance** (author, source, timestamp, content hash).
- Agent outputs cite the retrieved sources they relied on; unattributed claims are not rewarded.
- Reward for a piece of knowledge accrues in proportion to its **verified use** by agents across the community.
- *Current technique*: retrieval-augmented generation over the attributed store. **Replaceable** as better grounding methods appear.

2.5 Openness & privacy boundary (implements Principle 4)

- **Public**: all agents, tools, prompts, and protocol code — open source.
- **Private (never exposed)**: a user's personal dataset — private data, credentials, and unpublished work. The boundary is defined narrowly and explicitly so "open by default" never leaks personal or pre-publication material.

2.6 Compliance & framing

PWM is presented as a **utility/protocol credit for compute and knowledge**. Its value tracks the **usefulness** of the work it commands. Nothing in this paper is a promise of financial return. The Foundation does not custody user tokens. This framing is a **constraint on the Mechanism**, kept deliberately consistent with the project's non-profit, public-benefit posture.

Section III — How the Principles Reinforce One Another

The six principles are not a list; they are a **loop**:

1. Verified **knowledge** (P1) makes specialist **agents** good (P2).
2. Good specialist agents **compose** into a community (P3).
3. **Openness** (P4) lets every agent improve every other, compounding P3.
4. Real **usage** of these better agents adds value under fixed supply (P5).
5. Value flows to the **strongest agents** — and to those who build and run them — in bounded, convergent shares (P6).
6. The rewards from P5–P6 fund more verified **knowledge** (P1) — closing the loop.

Each turn of the loop raises the real value behind a fixed number of tokens. The community grows because contributing to it is the way to capture its value, and contributing is open to anyone with a narrow edge — exactly the population (P2) that the large incumbents cannot serve.

And the loop has a **destination**. Its purpose is not to spin faster; it is to aim a community of superhuman agents at the **grand hard problems** of Appendix D — curing cancer, reversing aging, quantum gravity, fault-tolerant quantum computing, room-temperature superconductivity. The protocol is the machine that converts superhuman capability into solutions to the problems humanity could not solve alone. That is the point of all six principles.

Appendix A — The Geometric Reward, Stated Precisely

Per-use share: $\text{share}(k) = a \cdot r^{(k-1)}$, $0 < r < 1$. Cumulative over N uses: $S_N = a \cdot (1 - r^N)/(1 - r)$. Bound: $\lim_{N \rightarrow \infty} S_N = a/(1 - r) < \infty$.

Constitutional content: the existence of a finite bound (no unbounded capture) and the first-and-best ordering of rewards. **Amendable content:** the values of a and r .

Appendix B — Glossary

- **Agent** — a specialist program that performs work in a narrow field.
 - **Research mode** — the science research agent grounded in the protocol's verified-knowledge registry; its outperformance on real research tasks is a primary driver of community value (P5).
 - **Grand hard problem** — an open scientific or technical problem that has resisted generations of human effort and whose solution would change what humanity can do (e.g. curing cancer, reversing aging, quantum gravity). The destination the protocol aims its agents at; canonical list in Appendix D.
 - **PWM** — the protocol's fixed-supply utility credit for compute and knowledge.
 - **Producer stream** — the share of value directed to active knowledge/compute/agent contributors (2.3).
 - **Verified knowledge** — attributed, retrievable, checkable knowledge eligible to earn (P1).
-

Appendix C — Amendment Rule

Section I (the Six Principles) is intended to be **permanent**. Section II (Mechanism) is **amendable by governance**. Any proposed change must demonstrate that it **still serves all six principles**. A change that would require editing Section I is, by definition, out of scope for this protocol — it would be a different protocol.

Appendix D — The Grand Hard Problems

This is the canonical list of grand hard problems the protocol aims its agents at. It is **living**: the authoritative, up-to-date version is the frontier page at physicsworldmodel.org. Problems enter here as **framing pages** — an L1 Principle (governing equation + a precise statement of *what "solved" means*) and an L2 Digital Twin (domain Ω + tolerance ϵ). A grand hard problem becomes earnable only once a verifiable L3 Benchmark exists for it; until then no fabricated baseline or "solution" is shown, so the frontier never sits next to unverified scores. The list is ordered **life sciences & aging first**, then quantum, nano, and physics & cosmology.

Life Sciences & Aging (*emphasis*)

Problem	What would be solved
Reverse Aging	Restore a youthful epigenome with partial reprogramming — turn biological age down without losing cell identity.
Longevity	Bend the Gompertz mortality curve — slow the rate of aging itself, extending healthspan, not just lifespan.
Curing Cancer	Treat the tumor as an evolving population — schedule therapy to keep it controllable forever instead of selecting for resistance.
Regeneration	Steer wounded tissue toward regeneration instead of scar — control the morphogen/ECM dynamics that decide fibrosis vs repair.
Proteostasis	Predict and reverse age-driven protein aggregation — the proteostasis collapse behind Alzheimer's and Parkinson's.

Quantum Technology

Problem	What would be solved
Fault-Tolerant Quantum Computing	Cross the fault-tolerance threshold — decode the surface code fast enough that adding qubits drives the logical error rate down, not up.
Quantum Networks & Repeaters	Build a quantum internet — distribute entanglement over continental distances by beating fibre loss with repeaters and memories.
Optimal Quantum Control	Hit the gate fidelities fault tolerance needs — shape control pulses that beat decoherence and crosstalk on real hardware.
Quantum Circuit Compilation	Make quantum programs runnable — compile and route circuits with the fewest two-qubit gates so the result survives noise.

Nanotechnology

Problem	What would be solved
Atomically Precise Manufacturing	Build matter atom by atom — place reactive groups with positional control to assemble structures no bulk chemistry can reach.
DNA Self-Assembly	Program matter to fold itself — design DNA sequences that self-assemble into a target nanostructure at high yield.
Targeted Nanomedicine	Deliver drugs only where needed — design nanoparticles that navigate the body and accumulate in target tissue, sparing the rest.
Molecular Electronics	Compute with single molecules — predict and design the quantum conductance of a molecule wired between two electrodes.

Physics & Cosmology

Problem	What would be solved
Quantum Gravity	Reconstruct the geometry of spacetime from quantum entanglement — test whether gravity is emergent (Ryu–Takayanagi).
Dark Matter	Pin down the dark-matter halo from how galaxies rotate — distinguish particle dark matter from modified gravity.
Dark Energy	Measure whether dark energy is constant or evolving — reconstruct the equation of state $w(z)$ from the expansion history.
Black-Hole Information	Show that black holes preserve information — recover the unitary Page curve of Hawking radiation via the island formula.
Room-Temperature Superconductivity	Design a material that superconducts at room temperature — predict T_c from electron–phonon coupling and invert for high- T_c structures.

Why these and not "daily work." Each problem above has resisted generations of human effort and, if solved, changes what humanity can do. An AGI that only automates ordinary tasks leaves this table untouched. The whole protocol — verifiable knowledge (P1), specialist depth (P2), composition (P3), openness (P4), real-usage value (P5), and reward-the-strongest (P6) — exists to get this table solved.