Domain-Value-Backed Currencies

A New Monetary Model - The Metha Biofund Case Study

José Víctor Martínez

Working draft v1.4.1 - 2026-05-20

Table of Contents

Domain-Value-Backed Currencies: A New Monetary Model for Decentralized Financial Ecosystems

Subtitle: The Metha Biofund Case Study

José Victor Martínez

Metha Biofund · https://metha.life

Working draft v1.4.1 — Metha Biofund. Last updated: 2026-05-20.

Abstract

Every dominant monetary system in history derives its value from a structural demand mechanism rather than from intrinsic properties of the medium itself. Commodity-backed currencies derive demand from the stored value of the commodity; fiat currencies derive demand from legal tender mandates and tax obligations; the US dollar derives extraordinary demand from its role as the mandatory settlement currency for global oil trade — the petrodollar mechanism. Yet no existing monetary model systematically links currency *emission* to investment in the productive domain that generates structural demand. This paper introduces Domain-Value-Backed Currency (DVBC): a monetary model in which (a) all commercial transactions in a defined high-value domain require settlement in the currency, (b) emission of the currency is mechanically linked to investment in expanding that domain's productive capacity, and (c) the protocol is governed by token holders. The result is a self-reinforcing loop: emission funds research, research produces commercial outputs, commercial outputs require the currency for settlement, settlement demand attracts new capital, and new capital funds more research. We formalize this model, prove its theoretical superiority over the petrodollar structure in terms of emission-demand coherence, and demonstrate its implementation in Metha Biofund — an open-source Ethereum protocol deployed on the Sepolia testnet with six interconnected smart contracts (MethaToken, MethaCrowdsale, AuditingPool, OpenGovernance, EternalStorage and AuditorRegistryV2) validated by a Hardhat suite of 50+ automated tests and a clean Slither static-analysis run. Metha currently satisfies conditions (b) and (c) in deployed testnet contracts. Condition (a) — mandatory METH settlement for biomedical IP commercialization — is defined as a future governance decision contingent on a legal-and-regulatory layer that the project has not yet stood up; the on-chain primitives required to implement it (typed deliverables, milestone option exercise, multi-auditor consensus, DSMB veto) are live on testnet and documented in §5. This paper is published as part of a proof-of-concept programme seeking peer collaborators in monetary economics, securities and competition law, biomedical research, and accreditation engineering — not as a commercial milestone.

Keywords: monetary theory, domain-value-backed currency, petrodollar, DeFi, token economics, decentralized science, research funding, Ethereum, structural demand

1. Introduction

1.1 The Problem of Monetary Value

Every theory of money grapples with the same foundational question: what makes a token valuable? Classical economists, following Smith and Ricardo, grounded value in the labor embodied in the commodity used as money. Keynes [1] demonstrated that monetary value was inseparable from the institutional structures — the state, the central bank, the banking system — that governed its issuance and acceptance. Friedman [2] argued that monetary value was ultimately a function of supply control: constrain the growth of money supply and inflation dissolves, allowing real purchasing power to stabilize. Hayek [3], by contrast, proposed that the optimal monetary system would emerge from competitive issuance — currencies competing for adoption on the basis of their stability properties.

Each of these frameworks treats currency value as either a political construction (Keynes, Friedman) or a market equilibrium (Hayek). None adequately explains the most striking monetary phenomenon of the twentieth century: the persistence of the US dollar as the world's reserve currency for more than fifty years after the United States unilaterally terminated its gold convertibility obligation in 1971 [4]. The dollar's post-Bretton-Woods dominance is not explained by supply discipline (the United States has run continuous current account deficits since 1971), institutional trust (Watergate, the savings and loan crisis, the 2008 financial crisis), or market equilibrium (no competitor currency emerged despite incentives). It is explained by a structural demand mechanism: oil, the most transacted commodity in global trade, is priced and settled exclusively in dollars. Any nation that needs oil — which is to say, every nation — must first acquire dollars. This demand is structural, not discretionary.

1.2 The Petrodollar as a Template — and Its Incomplete Design

The petrodollar arrangement [5] is the closest historical precedent for what this paper proposes. It demonstrates that a currency can derive durable, non-discretionary demand from its role as the mandatory settlement medium for a specific domain's outputs, independent of the issuing state's monetary discipline. However, the petrodollar mechanism has a critical structural gap: US dollar emission (monetary policy) is entirely decoupled from oil ecosystem investment. The Federal Reserve sets interest rates and manages money supply according to macroeconomic targets — unemployment, inflation — that have no necessary relationship to the health of the global oil industry. The structural demand mechanism exists, but the emission mechanism is arbitrary with respect to the domain that generates the demand.

This gap produces well-documented instabilities. Dollar over-emission (as occurred in 2020–2022) is not dampened by the oil market; it is exported to all dollar-denominated economies through inflation. Dollar under-emission (as in 1979–1982) can trigger global recessions in countries whose productive capacity depends on dollar liquidity but whose own productivity was unrelated to the monetary restriction imposed. The petrodollar is a structural demand mechanism grafted onto a fiat emission system — the two halves are incoherent with each other.

1.3 The DVBC Hypothesis

This paper proposes that the incoherence can be resolved by design. A **Domain-Value-Backed Currency** (DVBC) is a currency in which both the structural demand mechanism and the emission mechanism are locked to the same domain:

- Emission is mechanically linked to investment in the domain's productive capacity.
- Settlement of all commercial outputs of the domain is denominated in the currency.

The result is a closed loop: more emission → more investment → more domain outputs → more settlement demand → more capital attracted → more emission. The currency's value is backed not by a stored commodity (gold standard) or by government mandate (fiat) but by the ongoing productive activity of the domain it funds and in which it settles.

1.4 Paper Contributions

This paper makes four contributions:

1. **The DVBC theoretical framework**: a formal definition, value loop, and comparison with commodity, fiat, and cryptocurrency monetary models.
2. **Formal comparison with the petrodollar**: demonstrating that DVBC achieves emission-demand coherence that the petrodollar structurally lacks.
3. **Metha Biofund as a proof-of-concept DVBC implementation**: a deployed Ethereum protocol in which METH token emission is mechanically linked to biomedical research funding (Sepolia testnet only; mainnet deployment is gated on a four-audit sequence — legal, BARDA-method, code, financial — and on entity formation, multisig custody, insurance and bug-bounty pool, with no fixed date).
4. **Domain generalizability analysis**: examining the conditions under which DVBC logic extends to energy, education, and software domains.

---

2. Monetary Theory Background

2.1 From Commodity Money to Fiat: The Evolution of Value Foundations

The earliest monetary systems were commodity systems: money was a good — gold, silver, salt, grain — that had exchange value independent of its monetary function. The monetary premium (the excess value above commodity use value) derived from its adoption as a medium of exchange, driven by properties of divisibility, durability, portability, and scarcity. Jevons [6] formalized this in terms of the double coincidence of wants: money solves the coordination problem of barter by serving as a universally accepted intermediate good.

The shift to representative commodity money — banknotes convertible to gold on demand — created a separation between the physical medium (paper) and the value backing (gold reserves). This separation was operationally efficient but introduced a fragility: convertibility depended on institutional trust in the issuer's reserve adequacy. Bank runs were the recurring pathology of this system.

The gold exchange standard, codified at Bretton Woods in 1944 (IMF Articles of Agreement Art. IV §1, original text [45]), attempted to stabilize this system internationally by anchoring the dollar to gold ($35/troy ounce) and anchoring all other currencies to the dollar. This

created a hierarchical commodity-backed system: gold backed dollars, dollars backed everything else. When the United States unilaterally suspended dollar-gold convertibility on August 15, 1971 [44] — what Eichengreen [4] characterises as the dollar's "exorbitant privilege" — the commodity foundation was removed, leaving a pure fiat system.

A fiat currency derives value from three mechanisms: (i) legal tender laws requiring its acceptance for private debts and tax obligations, (ii) the institutional credibility of the issuing central bank as a commitment to supply discipline, and (iii) network effects — the value of using the same currency as one's trading partners. None of these mechanisms links emission to productive investment. The central bank issues currency in response to macroeconomic targets (inflation, employment) that may be entirely disconnected from the productive sectors of the economy.

2.2 The Petrodollar Mechanism: Structural Demand Without Emission Coherence

The petrodollar arrangement originated in a series of agreements between the United States and Saudi Arabia beginning in 1974, following the 1973 oil embargo [5]. The arrangement established that OPEC would price and settle oil transactions exclusively in US dollars, in exchange for US security guarantees and arms sales to Gulf states. The consequence — not fully appreciated at the time — was the creation of a structural demand mechanism for dollars that was entirely independent of US monetary policy or economic performance.

To understand the mechanism, consider the position of any oil-importing nation. It must acquire dollars before it can purchase oil. This requires either export earnings in dollars (from selling goods to dollar-area countries) or borrowing dollars (from international capital markets denominated in dollars). Either way, there is a permanent, non-discretionary demand for dollars proportional to global oil import volumes. As global energy consumption grew over the subsequent decades, so did structural dollar demand — independent of whether the United States was running current account surpluses or deficits, expanding or contracting its money supply.

The petrodollar is therefore a historical precedent for the DVBC *demand mechanism* only — for the proposition that a currency can sustain reserve status through mandatory commodity settlement even in the absence of issuer discipline. It is not a precedent for the emission-investment linkage, which is the structural innovation DVBC introduces over the petrodollar. However, as noted in the introduction, the petrodollar has a fundamental structural gap: dollar emission (US monetary policy) is entirely decoupled from oil ecosystem investment. The Federal Reserve does not consider the productive capacity of the global oil industry when setting the federal funds rate. This decoupling is the pathology that DVBC design addresses.

Table 1: Structural Comparison of Monetary Models

Property	Gold Standard	Fiat	Petrodollar	Bitcoin	DVBC
Structural demand mechanism	Commodity use value	Legal tender mandate	Domain settlement (oil)	Scarcity narrative	Domain settlement (defined)
Emission mechanism	Gold extraction rate	Central bank policy	US monetary policy	Protocol constant (halvings)	Investment in domain
Emission-demand	Partial (demand =	None	None	None	Full (by design)

coherence	mining)				
Governance of emission	Geological	Political	Political	Protocol	Token holders
Depletion risk	Yes (finite gold)	No	Yes (finite oil)	Yes (21M cap)	No (expanding domain)

2.3 Cryptocurrency Value Models: A Taxonomy

The emergence of Bitcoin in 2008 [7] introduced a new monetary design space: currencies governed by protocol rather than state or commodity. The subsequent decade produced a rich taxonomy of cryptocurrency value models.

The scarcity model (Bitcoin). Bitcoin's value proposition rests on a fixed, algorithmically enforced supply cap of 21 million units, declining emission through quadrennial halvings, and proof-of-work mining that creates a demonstrable cost floor. Nakamoto [7] explicitly framed this as a digital analog to the gold standard: scarcity as the value foundation. Catalini and Gans [20] characterize this as a "cost-of-verification" model: scarcity-derived value functions as a substitute for institutional trust in the absence of a central issuer. However, Bitcoin scarcity is artificial (enforced by protocol, not geology) and — critically — entirely disconnected from any productive domain. Bitcoin emission does not fund anything; it compensates miners for securing the network. There is no mechanism linking Bitcoin supply to real-world productive output.

The utility/gas model (Ethereum). Ethereum's Ether derives value from its role as the fuel for executing smart contracts on the network [8]. Demand for Ether is proportional to demand for Ethereum computation. This is a utility model: value derives from the service the currency enables access to, not from any productive investment the currency funds. Post-Merge, Ether issuance is further controlled by validator economics, creating a partial scarcity dynamic.

The governance/power model (DeFi tokens). The majority of DeFi protocol tokens derive value from governance rights: the right to vote on protocol parameters, treasury allocations, and future development. Voshmgir [21] provides a comprehensive taxonomy of token classes (purpose, utility, work, asset, and governance) and notes that governance tokens are inherently endogenous instruments: their value is a function of the protocol they govern. This value is real but circular: governance tokens are valuable because they confer power over a valuable protocol, but the protocol's value depends in part on the market capitalization of the governance token. Governance tokens typically have no mandatory settlement function in any external productive domain.

The collateral-backed model (stablecoins). Stablecoins derive value from backing assets: fiat currency reserves (USDC), over-collateralized crypto positions (DAI), or algorithmic mechanisms (FRAX). These are fundamentally derivative instruments — their value is borrowed from whatever backs them.

2.4 The Gap: No Existing Model Links Emission to Productive Investment

The taxonomy in Section 2.3 reveals a structural gap. No existing monetary model — commodity, fiat, or crypto — mechanically links currency emission to investment in the productive domain that generates demand for the currency.

- The gold standard links emission to mining (extractive, not productive); the Bretton-Woods codification of dollar-gold convertibility is the institutional capstone of this regime [4, 44, 45].
- Fiat links emission to macroeconomic policy targets (disconnected from domain); Keynes [1] formalised the institutional-trust foundation and Friedman [2] the supply-discipline alternative, neither of which connects to any productive sector.
- Bitcoin links emission to mining security (purely operational); Nakamoto's 2008 white paper [7] frames this as a digital analog to gold scarcity, and Catalini and Gans [20] characterise the model as "cost-of-verification" with no productive-domain coupling.
- DeFi governance tokens link emission to protocol incentives (liquidity mining, vesting) unrelated to productive outputs; Voshmgir [21] provides the comprehensive token-class taxonomy, and the empirical record of the 2017–2018 ICO market analysed by Howell, Niessner and Yermack [22] illustrates the disconnect between emission schedules and productive output.
- The petrodollar links demand to domain settlement but emission to unrelated monetary policy [5]; Eichengreen [4] characterises the asymmetric institutional position this creates.

This gap is not merely theoretical. It is the source of a class of monetary instabilities: boom-bust cycles in commodity-linked currencies (petrodollar recycling, oil price shocks documented in [42, 43]), persistent inflation in fiat systems without disciplined emission [2], and speculative volatility in governance tokens whose emission schedule is disconnected from any real productive activity [22].

DVBC is designed to close this gap by construction.

Parallel to Post-Keynesian endogenous money, and where DVBC departs from it. A reviewer correctly notes that the DVBC emission-by-investment mechanism has a structural parallel in the Post-Keynesian endogenous-money tradition that the §2.4 gap argument should engage explicitly. Lavoie [62] and Moore [63] develop the horizontalist proposition that money supply is endogenous: rather than being set exogenously by central-bank discretion, the money stock responds to credit demand for production. In the horizontalist model, commercial banks accommodate creditworthy borrowing demand on demand, central-bank reserves follow the resulting deposit base, and the money supply is therefore a function of the productive economy's appetite for credit rather than of an external monetary-policy target. The structural similarity to DVBC's $e(t) = f(I(t))$ — emission as a function of domain investment — is real and worth acknowledging. The structural distinction matters as much as the parallel: the horizontalist linkage operates through *commercial-bank credit creation in response to loan demand*, intermediated by the banking system and ratified ex post by the central-bank reaction function; the DVBC linkage operates through *direct equity investment in a productive domain*, encoded in protocol rather than ratified ex post by any reaction function. The horizontalist mechanism explains money-supply behaviour at the macro level under existing institutions; the DVBC mechanism designs a protocol-level emission rule that produces a structurally analogous coherence between emission and productive demand within a specific domain. The two operate at different layers (institutional vs protocol), through different linkage types (credit vs equity), and over different aggregation scopes (macroeconomy vs domain) — but they share the foundational claim that emission ought to track productive activity rather than being set by external policy. The §2.4 gap argument should be read as *building on* the horizontalist tradition rather than as *departing from* it: the

architectural innovation is the substrate (smart contract) and the linkage type (equity-to-emission rather than credit-to-emission), not the foundational economic claim.

---

3. Domain-Value-Backed Currency: The Theoretical Model

3.1 Formal Definition

Definition 1 (Domain-Value-Backed Currency). Let D be a productive domain characterized by a flow of commercial outputs $O(t)$ at time t, each output having a market value $v(o) > 0$. A Domain-Value-Backed Currency C for domain D satisfies three conditions:

(a) Mandatory Settlement. All commercial transactions involving outputs $o \in O(t)$ in domain D must be settled in C. Formally, for any transaction T = (buyer, seller, output o), the settlement medium must be C.

(b) Emission-Investment Linkage. The emission rate of C at time t, denoted $e(t)$, is a deterministic function of investment $I(t)$ in D's productive capacity: $e(t) = f(I(t))$ where f is a monotone non-decreasing function specified by protocol.

(c) Holder Governance. The protocol governing both (a) and (b) — including the function f, the scope of mandatory settlement, and the definition of domain outputs — is governed by C holders through on-chain voting.

Condition (a) creates the structural demand mechanism. Condition (b) creates emission-demand coherence. Condition (c) creates a self-regulating feedback loop governed by those with the greatest stake in the currency's long-run value.

3.2 The DVBC Value Loop

The DVBC mechanism produces a self-reinforcing value loop that can be formalized as follows. Let:

- $P(t)$ = price of C in external currency at time t
- $I(t)$ = investment in domain D at time t (denominated in external currency)
- $e(t)$ = emission rate of C at time t
- $O(t)$ = volume of commercial outputs produced in D at time t
- $S(t)$ = settlement demand for C from domain outputs at time t

The DVBC dynamics are characterized by four relationships, each numbered explicitly and tagged with an equation:

(1) **Investment generates research outputs.** g reflects the research-to-output production function (with lag Δ representing the research-to-commercialization pipeline):

$$O(t + \Delta) = g(I(t)) \quad (1)$$

(2) **Emission is mechanically linked to investment** by condition (b):

$$e(t) = f(I(t)) \quad (2)$$

(3) **Outputs generate settlement demand**, where h maps output volume to required C holdings for settlement, scaled by licensing fraction L and per-licence median revenue R:

$$S(t) = h(O(t), L, R) \quad (3)$$

(4) **Settlement demand drives price** relative to outstanding supply via the standard monetary exchange equation specialised to the DVBC domain as derived in §3.3-bis:

$$P(t) = \hat{P}(S(t), M(t), V(t)) \quad (4)$$

where $\hat{P}$ is the equilibrating inverse of $MV = PQ$ with M the outstanding METH supply and V the velocity. Tagging the fourth relationship makes the loop's closure visually explicit and lets §3.3-bis refer back to it by number.

Properties of the four functions. The qualitative shape of g, f, h and $\hat{P}$ is load-bearing for the §3.3 coherence argument and for the §3.3-bis Equation-of-Exchange specialisation, so we state the assumed properties of each explicitly. The **production function** g mapping investment to research outputs is assumed monotonic non-decreasing in I — more investment cannot produce strictly less expected output, given a competently managed lab — and concave under the standard diminishing-returns assumption that holds across the empirical biomedical-research-productivity literature [42, 43]. We treat g as stochastic, with research uncertainty captured by a noise term whose variance is therapeutic-area-specific; the §3.3 coherence claim is a statement about expected dynamics, not pathwise dynamics, and the stability sketch added below to §3.3 is correspondingly a statement about expected-value trajectories under bounded-noise assumptions. The **emission function** f mapping investment to per-period METH emission is, in the current Metha implementation, a *protocol-constant* — emission per period is fixed by the protocol parameters described in §4.1, and per-contributor proportional allocation distributes that constant across the period's contributors. This is a v3.x design choice rather than a structural property of DVBC: a future protocol revision could make f strictly monotonic increasing in I — for example, by adding bonus-token emission above per-period investment thresholds — and the §3.3 coherence argument would carry through (with a strengthening: f-as-monotonic accelerates the emission-investment linkage rather than weakening it). The implementation choice does not affect the §3.3 coherence claim because the *protocol-level* property required is that ΔM track ΔQ over the time horizon at which monetary identities clear; a constant per-period f satisfies this when paired with the §3.3-bis observation that domain Q itself grows in step with cumulative investment via the lag-Δ production function. The **settlement-demand function** h mapping outputs to settlement demand is monotonic non-decreasing in O, with magnitude scaled by L (the share of commercialisable outputs actually licensed in any given period) and R (per-licence median revenue). The linear-in-expectation specialisation $h(O, L, R) = L \cdot R \cdot O$ holds under L-stable regimes — that is, regimes in which the licensing fraction is approximately constant across periods — and is the form used in the §3.3 closed-form. Under adoption-curve dynamics where L itself is rising (early-protocol) or saturating (mature-protocol), h is convex and concave respectively in O; the §3.3 coherence claim is robust to either regime. We are explicit that these are *assumed* properties under which the analytical results hold, not *derived* properties of the deployed protocol. A formal characterisation under realistic noise, heterogeneous-agent dynamics, and adoption-curve shifts is a Phase-2 deliverable and is an open invitation to the [Economics] respondent track.

The loop closes: higher $P(t)$ attracts more capital to acquire C through the investment mechanism (condition b), increasing $I(t)$, which increases $O(t + \Delta)$, which increases $S(t + \Delta)$, which sustains or increases $P(t + \Delta)$. This is a positive feedback loop with a natural governor: emission (condition b) increases with investment, diluting the supply effect and preventing unbounded price appreciation. The system is self-regulating rather than self-destructive.

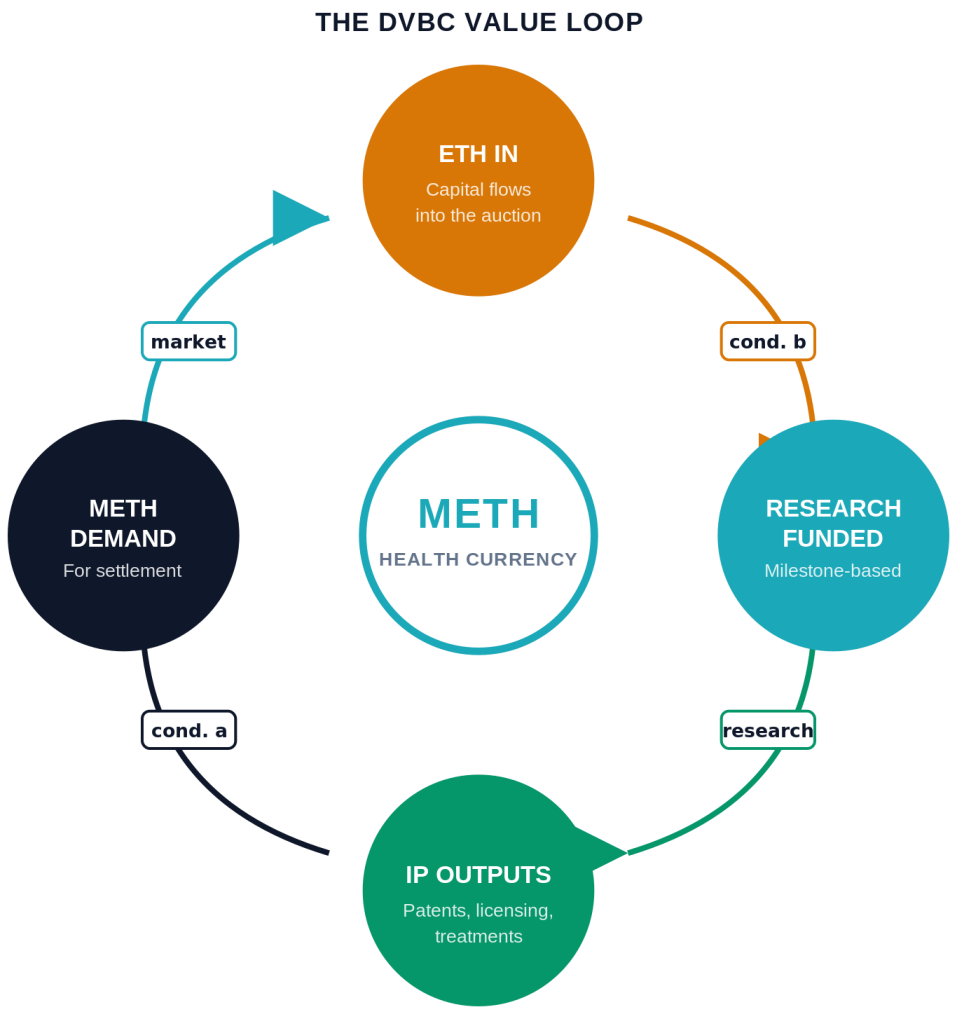

The DVBC value loop. Capital flows in (gold), funding research (teal); commercial outputs (emerald) generate mandatory settlement demand for the currency (navy), which drives price and attracts new capital. Conditions (a) settlement, (b) emission–investment linkage and (c) holder governance close the loop.

3.3 Comparison with the Petrodollar: Emission-Demand Coherence

The petrodollar mechanism satisfies condition (a) — oil transactions settle in dollars — but not condition (b). US dollar emission is governed by the Federal Open Market Committee according to the dual mandate (maximum employment, stable prices), entirely independently of oil ecosystem investment. This creates emission-demand incoherence:

Define **emission-demand coherence** as the property that $de(t)/dt$ and $dS(t)/dt$ share the same sign — when demand for the currency increases (more domain activity), emission increases proportionally; when domain activity contracts, emission contracts. A coherent currency does not inflate when the domain that backs it is contracting, and does not deflate when the domain is expanding.

The petrodollar is incoherent by this measure: US monetary expansion in 2020–2022 (during which the M2 money supply grew by approximately 40%, FRED M2SL series [42]) occurred during a period of depressed oil demand (COVID-19 lockdowns reduced global oil consumption by approximately 9%, IEA *World Energy Outlook 2021* [43]). Emission expanded while domain activity contracted — precisely the opposite of what coherence requires.

DVBC achieves coherence by construction: condition (b) makes *e(t) = f(I(t))*, and since *I(t)* is positively correlated with *O(t)*, which drives *S(t)*, emission and demand are mechanically linked.

Stability sketch under bounded-noise assumptions. The claim that the positive-feedback loop is naturally governed — rather than divergently amplifying — rests on the structural property that emission scales with investment via condition (b), so any unbounded P appreciation is diluted by the very emission the appreciation rewards. We can sketch this as a Lyapunov-style argument, with the caveat that this is a *sketch* and not a formal proof; the latter requires noise-structure assumptions and heterogeneous-agent dynamics that are properly an [Economics] Phase-2 task. To avoid notational collision with velocity V in §3.3-bis we write the candidate Lyapunov function as $\mathcal{L}$. Let P^* denote the steady-state METH price consistent with the long-run domain-output value — the price at which the §3.3-bis equation of exchange clears in expectation given the long-run Q, V, and M trajectories. Take as a candidate Lyapunov function $\mathcal{L}(t) = (P(t) - P^*)^2$ measuring squared deviation of the realised price from the steady-state. Under the §3.2 functional assumptions — monotonic g, constant or monotonic f, monotonic h — and a bounded-noise assumption on g (research-productivity uncertainty has finite variance per period), the conditional drift $E[\mathcal{L}(t+1) - \mathcal{L}(t) \mid \mathcal{F}_t]$ is *bounded* on the set $\{P(t) \neq P^*\}$ — in the Foster-Lyapunov sense of bounded mean drift toward the steady-state outside a small neighbourhood of P^*, rather than the stronger pathwise non-increase — because the feedback channels do real work in opposite directions: deviations of P above P^* stimulate investment I (the auction is more attractive when METH appreciates), which stimulates emission e (condition (b)), which dilutes the price upward pressure; deviations below P^* suppress investment, which suppresses emission, which removes the dilution and lets settlement demand re-anchor the price. This is the governor structure §3.3 identifies in prose, formalised as a bounded-drift argument. The sketch deliberately stops short of four things: (i) the bounded-noise assumption on g is empirically defensible for the biomedical-research domain but not derived in this paper; (ii) bounded mean drift is weaker than pathwise non-increase and does not establish almost-sure convergence to P^* — it establishes a stationary distribution concentrated around it; (iii) the implicit "representative agent" framing is replaced by a heterogeneous-agent structure in any candid formal model — and the strategic-bidding dynamics analysed in §4.1.1.bis interact with the stability argument in ways the sketch does not address; (iv) the steady-state price P^* is itself a slowly varying object as the domain's commercial output evolves. The contrast with petrodollar stability is the central qualitative point: petrodollar architecture has no governor of the kind described above, and dollar deviations from a "consistent with oil-domain output" price are absorbed by inflation export rather than by structural feedback. DVBC is designed with the governor built in; the formal noise-and-heterogeneity proof of the corresponding stability theorem is the open invitation to economist collaborators that this section advances.

3.3-bis DVBC dynamics under the Equation of Exchange

A reviewer asks how the emission-demand coherence claim of §3.3 relates to standard economic identities, and in particular to Fisher's Equation of Exchange [55]. The question is sharper than it first appears, because Fisher's identity is among the most general accounting

frameworks available for describing monetary dynamics, and any new monetary model that cannot be expressed within it has a much heavier explanatory burden than one that can. This section shows that DVBC dynamics are expressible within Fisher's framework — and that doing so clarifies, rather than complicates, the §3.3 coherence claim.

Recall the canonical identity: $MV = PQ$, where M is the outstanding stock of the currency in circulation, V is its velocity, P is the average price level (in external currency) of transactions denominated in the currency, and Q is the real volume of those transactions. For a DVBC operating in domain D, the identity specialises naturally:

- $M(t)$ is the outstanding METH supply at time t. From the §4.1 emission rule, $\Delta M(t) = e(t) = f(I(t))$, where $I(t)$ is per-period ETH investment and f is the protocol-constant emission function.
- $Q(t)$ is the per-period volume of commercial transactions in domain D settled in METH. Under condition (a), $Q(t)$ equals the entire commercial output flow of the domain that touches Metha-funded IP. Under §3.2's production function $O(t + \Delta) = g(I(t))$, $Q(t)$ grows in step with $O(t)$, with the research-to-commercialisation lag Δ.
- $P(t)$ is the per-unit price of METH in external currency. $S(t)$ in §3.2 is the demand for METH from settlement; in Fisher's framework, $P(t)$ is the equilibrating variable that clears the supply $M(t)$ against the transaction demand $Q(t)$ given velocity $V(t)$.
- $V(t)$ is the velocity of METH — the rate at which a given unit of METH turns over in settlement transactions per unit time.

The DVBC coherence property of §3.3 is then expressible as the requirement that the identity be maintained over time as Q grows. Taking logarithmic derivatives of $MV = PQ$ — the standard manipulation for a multiplicative identity — yields the growth-rate decomposition

$$\frac{\dot{M}}{M} + \frac{\dot{V}}{V} = \frac{\dot{P}}{P} + \frac{\dot{Q}}{Q}$$

where the dotted quantities are time derivatives. Under the velocity-stability assumption $\dot{V}/V \approx 0$ discussed below, and under the further assumption that the long-run price is locally pinned by the protocol's investment-emission linkage $\dot{P}/P \approx 0$ (i.e. emission tracks settlement demand sufficiently well that P is not the variable absorbing the imbalance), the identity reduces to

$$\frac{\dot{M}}{M} \approx \frac{\dot{Q}}{Q}$$

In words: the *growth rate* of METH supply tracks the *growth rate* of domain transaction volume, which is exactly the §3.3 coherence claim — that $de(t)/dt$ and $dS(t)/dt$ share the same sign — re-expressed in monetary-identity form. Equivalently, in absolute differentials with stable V and locally stable P, we have $V \cdot \Delta M \approx P \cdot \Delta Q$: velocity-weighted emission tracks price-weighted real volume growth. The petrodollar's incoherence under the same identity is that ΔM is set by macroeconomic targets unrelated to ΔQ in the oil domain, so the identity is maintained only through P (dollar exchange-rate adjustment) absorbing the imbalance, which is precisely the inflation-export mechanism documented in §3.3.

The velocity-stability assumption above is not innocent, and Mundell's optimum-currency-area framework [12] is the relevant theoretical bridge. Mundell's central insight is that currency unions succeed where economic integration within the area is sufficiently deep to

make velocity locally stable across the integrated zone. The DVBC architecture satisfies this condition by construction along the *domain* dimension rather than the *geographic* dimension: all participants in the biomedical IP licensing domain — funded labs, licensees, regulators, accreditation firms — transact in METH under condition (a), and the velocity of those transactions is anchored to the long-tail royalty-flow structure of biomedical IP rather than to the short-tail speculative-trading structure that dominates open-market crypto velocity. Biomedical patent royalty flows have characteristic time-horizons measured in years (typically 7–12 years from licence execution to peak revenue), not in minutes; this is the structural reason DVBC velocity is plausibly stable in a way that, for example, a memecoin's velocity is not. The optimal currency area for DVBC is not the Eurozone or the United States; it is the integrated set of participants in the productive domain itself.

Friedman's monetarist programme [2] is the second canonical anchor. Friedman's argument for monetary discipline rests on the proposition that ΔM should track ΔQ with low variance, so that P does not absorb the imbalance through inflation. The DVBC architecture achieves this not through central-bank judgment (which Friedman himself doubted) but through mechanical linkage to investment: $\Delta M = f(I(t))$ and $I(t)$ is in turn a market-revealed signal of the expected growth in $Q(t)$. The DVBC emission rule is a Friedman k-percent rule generalised to a domain-specific, market-revealed $k(t)$, which is at minimum as verifiable as the central-bank inflation-target alternative and arguably more so, because $I(t)$ is observable on-chain in real time while inflation is observable only with substantial lag.

A precision point. The argument above does not "prove" that DVBC works in any strong sense. The Equation of Exchange is an accounting identity, not a behavioural theorem; expressing DVBC within it does not show that participants will in fact transact at the implied price, or that V will in fact remain stable, or that the lag Δ in the production function is short enough for the dynamics to remain coherent across realistic shocks. What the derivation establishes is the weaker and more honest claim: DVBC dynamics are *expressible within* the standard monetary framework, which means the project does not require new economics — only the application of standard monetary identities to a new domain surface. That is a strength rather than a weakness, because it places the DVBC claim within the literature that practitioners and reviewers already know how to interrogate.

3.4 Comparison with Fiat and Bitcoin

Fiat currencies satisfy neither condition (a) nor condition (b). There is no mandatory settlement domain (legal tender laws are general, not domain-specific), and emission is determined by monetary policy targets disconnected from productive investment. Fiat stability depends entirely on institutional credibility — an endogenous variable subject to political erosion.

Bitcoin satisfies neither (a) nor (b) in any productive domain sense. Bitcoin's emission is linked to mining — a security mechanism, not investment in a productive domain. Bitcoin has no mandatory settlement function in any external domain; its value proposition is purely monetary (store of value through scarcity) rather than domain-functional.

Standard DeFi governance tokens approach condition (c) (holder governance) but typically lack (a) (no mandatory settlement in a productive domain) and at most approximate (b) through liquidity mining schemes that link token emission to protocol usage rather than to productive investment in an external domain.

Table 2: DVBC Conditions Satisfied by Monetary Models

Model	Condition (a): Domain Settlement	Condition (b): Emission-Investment Link	Condition (c): Holder Governance
Gold standard	**No**	Partial	**No**
Fiat	**No**	No	**No**
Petrodollar (USD)	**Yes**	**No**	No
Bitcoin	**No**	No	Partial
DeFi governance tokens	**No**	No	**Yes**
DVBC (Metha/METH)	**Yes** (planned)[1]	**Yes** (live)	**Yes (live)**

[1] "Planned" for condition (a) is a forward-looking governance step contingent on completion of the legal-and-regulatory layer described in §6, the audit-gated roadmap of §7, and the regulatory analysis of §8.5; the on-chain primitives required to enforce it are live on testnet today.

3.5 Why Health is the Ideal First DVBC Domain

The selection of domain for a DVBC implementation is not arbitrary. Several properties make a domain well-suited:

Scale and growth. The domain must be large enough that structural settlement demand is economically significant. Global healthcare expenditure approaches $10 trillion annually [9], with biomedical research and development reaching approximately $280 billion in 2024 [10]. Unlike mature commodity markets (oil, gold), the health domain is structurally expanding — aging demographics, emerging economies investing in healthcare infrastructure, and the perpetual productivity of scientific discovery create a demand surface that grows independently of external macroeconomic cycles.

Defined commercial outputs. The domain must produce outputs with clear market value that can be denominated and settled. Health research produces patents, licensing agreements, clinical trial rights, treatment protocols, and regulatory approvals — each with well-defined commercial value. The global pharmaceutical market reached approximately $1.7 trillion in annual revenue in 2024 and is projected to exceed $1.8 trillion in 2025 [10], the majority of which derives from patent-protected products.

No depletion risk. Commodity-backed systems (gold, oil) face the long-run problem of depleting the backing resource. Health outputs are a renewable productive domain: each generation of funded research produces the scientific base for the next generation. A DVBC for health does not face the structural contraction risk that confronts petrodollar demand as fossil fuel depletion accelerates.

Mundell's optimum currency area theory [12] holds that currency unions succeed when economic integration within the area is sufficiently deep. A DVBC satisfies this by construction: all participants in the domain — researchers, funders, auditors, and IP licensees — transact in the same currency, creating the internal integration Mundell identifies as the precondition for monetary stability.

Geographic neutrality. Oil production is geographically concentrated, making oil-linked settlement vulnerable to geopolitical disruption of producing regions. Biomedical research is distributed across institutions globally and is less vulnerable to geographic concentration risk.

Regulatory precedent. Intellectual property law already provides a framework for mandatory licensing terms — compulsory licensing, FRAND commitments, and patent pools demonstrate that settlement terms can be legally enforced. DVBC settlement conditions for health IP are not legally unprecedented; they extend existing IP licensing frameworks.

3.6 Why now? Historical contingencies that gate DVBC feasibility

A natural sceptical question is: if the DVBC architecture is as well-founded as the preceding sections argue, why has it not been built before? The answer is not theoretical insight but a sequence of enabling conditions that have only recently aligned. Four contingencies are load-bearing.

Smart-contract feasibility from circa 2015. A DVBC requires deterministic, low-cost, third-party-verifiable execution of the emission, settlement and governance logic. Pre-Ethereum [8], no general-purpose execution layer existed on which such a protocol could be specified independently of a single trusted issuer. The DVBC architecture is in essence a smart-contract artefact; the substrate for it has existed for roughly a decade.

Regulatory clarity from 2023–2024. MiCA's entry into force [14] (Regulation (EU) 2023/1114, with the bulk of provisions applicable from December 2024) is the first comprehensive regulatory framework that distinguishes asset-referenced, e-money and utility tokens from securities in a way that a DVBC-style settlement instrument can be analysed against without falling back on the Howey/Reves heuristics that pre-MiCA U.S. analysis depended on. Pre-MiCA, the legal default for a settlement token of this kind was "unclassifiable instrument under bespoke jurisdiction-by-jurisdiction analysis." Post-MiCA, the classification path described in §8.5 is at least specifiable, even if not yet closed. The CASP licensing regime that the §7 roadmap relies on did not exist before this window.

Funding-mechanism maturity in the biomedical option-period structure. The audit-and-tranching primitives of §5 (multi-auditor consensus, option-period exercise, DSMB veto, typed deliverables) are calibrated against off-chain patterns that have themselves matured into recognisable form only over the past decade: BARDA's option-period awards [30] reached current contractual maturity in the 2010s; ARPA-H was established in 2022 by the *Consolidated Appropriations Act, 2022* [46] and its option-exercise practices are still being codified; the EIC blended-finance instruments date to 2021 [47]. A DVBC anchored to biomedical IP requires its on-chain primitives to map cleanly onto those off-chain patterns, and the off-chain patterns themselves only recently became standard enough to copy.

Biomedical IP market structure post-Bayh-Dole. The 1980 Bayh-Dole Act and the subsequent maturation of university tech-transfer offices, FRAND practices in standards bodies and patent-pool jurisprudence have produced the legal vocabulary on which Mechanism A and Mechanism B of §6.1 rest. Earlier than this, the encumbrance-recordation

and march-in primitives that the DVBC enforcement framework relies on either did not exist (no march-in doctrine before Bayh-Dole) or had no operational precedent for license-condition encumbrance (the *Jacobsen v. Katzer* line [24] arrived only in 2008). The legal raw material for a credible enforcement layer is therefore itself a recent assembly.

Taken together, the conjunction of programmable execution (post-2015), regulatory framework specificity (post-2024 MiCA), option-period funding maturity (post-2010s BARDA / post-2022 ARPA-H) and IP-encumbrance jurisprudence (post-2008 *Jacobsen*, post-1980 Bayh-Dole) is what made the DVBC architecture proposable at all. The model is not late; the substrate is new.

3.7 Prior art in DeSci: IP-NFTs and the asset/settlement layer distinction

A reviewer presses, with justification, that §10.3's treatment of Molecule Protocol, VitaDAO, ResearchHub and LabDAO is appropriately late in the paper for a reader who has already absorbed §6's settlement-clause architecture, but is too late for a reader still calibrating where DVBC sits within the broader decentralised-science (DeSci) literature. v1.4 anchors the prior-art positioning at §3.7, immediately before the protocol description, so that the asset/settlement-layer distinction frames the rest of the paper.

The leading existing prior-art mechanism for tokenising biomedical intellectual property is the IP-NFT framework developed by Molecule Protocol and operationalised at scale by VitaDAO. *Nature Biotechnology* profiled the model in its 2023 editorial coverage [60], a useful journalistic signal of the model's reach in the DeSci ecosystem (the citation is editorial, not peer-reviewed research validation, and is anchored as such). The IP-NFT pattern represents the maturity state of biomedical-IP tokenisation as of v1.3's publication window and is the reference architecture against which any new biomedical-IP infrastructure proposal must position itself.

Metha's positioning relative to the IP-NFT prior art turns on a layer distinction that v1.3 §10.3 introduces but that this section makes load-bearing earlier. **Molecule Protocol and VitaDAO operate at the asset layer.** The DAO acquires fractional rights to specific compounds, structures them as IP-NFTs, and routes downstream licence revenue to NFT holders [61]. The token is a governance instrument over a portfolio of held IP assets; the monetary thesis is structurally closer to that of a closed-end venture fund than to a domain currency, and the long-run value-accumulation ceiling is bounded by the treasury of held assets rather than by the commercial output of the domain.

Metha operates at the settlement-and-funding layer. The protocol funds independent research via the §4.1 perpetual proportional auction, and the same token settles downstream commercial transactions in the funded IP via the §6 settlement clause, without the protocol itself holding income-producing IP assets (§6 (i)). The token's structural demand mechanism derives from the commercial output of the entire health domain rather than from a portfolio held by the DAO; the §3.5 scalability analysis is the empirical anchor for the upper-bound demand that this layer addresses. The architectural distinction is not a marketing point — it is the difference between a fund (bounded by portfolio) and a currency (bounded by domain).

The two layers are composable rather than substitutable. A natural future stack would have Metha funding research via the auction; outputs tokenised through Molecule's IP-NFT framework [61]; peer review routed via ResearchHub; lab services executed through LabDAO; asset-layer governance structured via a VitaDAO-style DAO; and settlement at the domain layer denominated in METH under condition (a). The §10.3 treatment makes the full-stack

claim; §3.7 anchors it where the reviewer suggested. The structural difference that justifies introducing a settlement-layer currency at all — given the maturity of the asset-layer prior art — is that the asset layer cannot, by construction, deliver the emission-investment coherence of §3.3: a fund's token value is bounded by the assets it holds, while a DVBC's token value is bounded by the productive output of the domain it funds and in which it settles. The two are categorically different monetary instruments, and the DeSci ecosystem requires both.

---

4. Metha Biofund as DVBC Implementation

4.1 Architectural overview

Metha Biofund is an open-source Ethereum protocol implementing the DVBC model for the biomedical research domain. The deployed stack on Sepolia testnet (chainId 11155111) comprises six interconnected smart contracts — MethaToken, MethaCrowdsale, AuditingPool, OpenGovernance, EternalStorage and AuditorRegistryV2 — written in Solidity 0.8.24 with OpenZeppelin v5 dependencies. The auction mechanism (described in detail below) is a perpetual proportional auction modelled on the EOS token launch [11]: time is divided into sequential periods (180 seconds on Sepolia, configured to 23 hours in mainnet config), and within each period any Ethereum address may contribute ETH; at period close each contributor's METH allocation is proportional to their share of total ETH contributed. Per-period emission is split 70/20/10 — 70% to claimable METH for contributors, 20% to an operations stash, and 10% to a developer stash — so that 30% of every period's issuance funds the protocol's own operational and developer capacity from the same emission mechanism that funds research. The 70/20/10 split is not a frozen constant: each leg is parameterised inside EternalStorage and is subject to revision by qualified OpenGovernance majority, so the protocol can re-balance research, operations and developer allocations as the ecosystem matures without redeploying contracts.

The operations stash funds the auditor compensation matrix described in §5 (percentage-of-milestone payment in ETH plus vested METH grant plus a slashing-window hold-back). Under the roadmap, the current monolithic operations stash is scheduled for decomposition into specialised multisigs — legal-ops, marketing-ops, scientific-ops, infra-ops, an emergency-multisig and, conditionally on Metha-Pharma activation, a Metha-Pharma-ops sub-stash — as documented in §10.5. The protocol thereby invests in itself proportionally to its investment in domain research, with the same predictable per-period emission rule serving researchers, auditors, developers and ops functions alike.

Why this implements condition (b). The causal chain from ETH contribution to research funding is direct and mechanically verifiable on-chain. ETH contributed to a period flows into the research treasury (`_weiRaised`), which is the source pool from which the AuditingPool disburses milestone payments to approved labs. METH is emitted at a protocol-constant rate per period (`tokensCreatedPerPeriod`), and each contributor's claimable share is proportional to that contributor's fraction of the period's ETH total:

$$claimable(i) = \frac{contribution(i)}{\sum_j contribution\,(j)} \times tokensCreatedPerPeriod \quad (5)$$

Formally, if *I(t)* denotes ETH invested in period *t* and *e(t)* denotes METH emitted, the current MethaCrowdsale satisfies *e(t) = tokensCreatedPerPeriod* (a protocol constant). Participation

(and therefore competitive allocation) is proportional to $I(t)$: more investment means more participants, more ETH raised, and more research funded per period. Under this design, f in Definition 1.b is technically constant; the proportional allocation rule shifts the I-dependence from total emission to per-contributor share. A future protocol upgrade — implementable via the slot-replacement mechanism of §4.3 without redeploying the rest of the stack — could make f explicitly increasing in I (e.g., bonus tokens when participation exceeds a threshold), but the current design favours predictability over adaptive emission. The protocol-constant emission rate is not arbitrary monetary policy; it is a commitment, visible to all participants in advance, that emission will continue at a fixed rate proportional to the period duration, providing a predictable relationship between participation, investment, and emission.

Front-running-resistant price discovery within the window. The proportional distribution mechanism has a property for DVBC stability that v1.2 of this paper described as "non-gameable" and that we now state more precisely: the mechanism is *front-running-resistant within each auction window*. No contributor benefits from time-ordered placement within the period, because every contributor's claimable share is determined by *aggregate* contribution at period close, not by the moment of contribution. This eliminates the destabilising dynamics of fixed-price token sales (immediate price discovery to premium over issue price) and of Dutch auctions (complex strategic bidding). The DVBC acquires tokens at a price that reflects aggregate domain interest at each period, which is the property required for emission and demand to remain coherent across market regimes. However, "front-running-resistant within the window" is materially weaker than "non-gameable in absolute terms"; §4.1.1.bis below sets out the specific manipulation vectors that a fixed-window auction remains susceptible to under the MEV / LVR literature accumulated since 2020, and the mitigations the current contract has and could acquire.

4.1.1 Empirical precedent and the DVBC compounding effect

The perpetual proportional auction is not a theoretical construct. The EOS token launch [11], which ran from 26 June 2017 to 1 June 2018, implemented essentially the same auction structure that Metha replicates: a fixed daily token allocation distributed proportionally to the ETH contributed within each window, repeated for 350 sequential periods. Over its full duration the mechanism raised approximately **US$4 billion in ETH** [23] — the figure documented in the U.S. Securities and Exchange Commission's 2019 cease-and-desist order, with the broader 2017–2018 ICO market context analysed by Howell, Niessner and Yermack [22] — making it the largest publicly-conducted token sale on record at the time and demonstrating that the proportional auction is empirically capable of bootstrapping a fungible token from zero market capitalisation to billions of dollars in committed capital. The relevant lesson for DVBC is mechanism-validation: the auction structure itself can attract participation at scale, conditional on sufficient market belief in the underlying value proposition. The mechanism's resistance to gaming (post-window price discovery eliminates front-running) and its predictability (fixed daily emission, no surprise dilution) appear to have contributed materially to that capital formation. For Metha, this empirical track record provides direct evidence that the cold-start problem (§10.2) is solvable along a known pathway — the auction is not the bottleneck; the credibility of the value proposition is.

We address the obvious counter-question explicitly: in 2019 the SEC settled with Block.one for $24 million over unregistered securities allegations arising from the EOS sale [23]. This precedent is directly relevant to the §8.4 regulatory analysis, but it does not undermine the empirical lesson here, because the auction *mechanism* and the resulting token's *regulatory characterisation* are independent design choices. The SEC's findings were grounded in the

broader value-accrual structure — a Block.one-controlled enterprise that retained substantial discretion over network development and presented investors with the prospect of profits derived from the issuer's continuing efforts — not in the auction mechanics themselves. The §6 architecture is specifically designed to weaken precisely the regulatory analogy that bound EOS, by removing the "common enterprise + portfolio of revenue-generating assets" structure that the SEC found dispositive (Metha holds no income-producing IP, runs no in-house product roadmap whose profits accrue to the protocol, and disburses ETH to independent labs whose IP they retain in full). The detailed regulatory-classification implications are taken up in §8.4 and §8.5.

The deeper observation, however, is that the EOS precedent under-states the case for DVBC because EOS itself relied on a *purely speculative* perceived-value signal during the auction: the EOS token had no underlying productive backing while the crowdsale ran, only an expectation that the future EOS network would matter. The mechanism nevertheless worked at scale because the market repriced the live, near-final EOS market value into every successive day's auction, and the proportional rule then automatically sold each day's allocation at a price near that live market signal. DVBC strengthens this loop structurally because the perceived value is not pure speculation — it is anchored to a productive domain (biomedical research outputs, IP generation, licensing royalties accrued in METH under condition (a)). The combination — proportional auction *plus* DVBC anchoring — is empirically stronger than either alone, for three concrete reasons. **First**, the auction mechanism gives the market a continuous price-discovery surface: 350+ sequential price prints per year, each averaged over a full window, suppressing the boom-bust microstructure that fixed-price sales and Dutch auctions induce. **Second**, the DVBC anchoring means that market belief is not grounded purely on narrative or scarcity but on observable productive output — research funded, milestones completed, IP generated, license royalties accrued and settled in METH — each of which is queryable on-chain or recordable at the patent offices that Mechanism A (§6.1) attaches encumbrances to. **Third**, the loop has *both* a proven scalable distribution mechanism *and* a non-speculative value floor: under DVBC, even in the limit where speculative premium goes to zero, the structural settlement demand described in §8.3 sets a floor below which the token cannot trade without arbitrage against actual royalty flows. The EOS precedent demonstrates the upper part of the loop; the DVBC architecture is what supplies the floor that EOS lacked.

4.1.1.bis Auction integrity under MEV/LVR scrutiny

The reviewer is correct that the MEV literature of the last five years has established that no on-chain auction structure is fully manipulation-immune. Daian, Goldfeder, Kell, Li, Zhao, Bentov, Breidenbach and Juels [51] documented in "Flash Boys 2.0" (IEEE S&P 2020) that decentralized-exchange order flow is systematically exploitable through transaction-reordering and front-running, and that the manipulation surface generalises to any on-chain auction with a known closing condition. Milionis, Moallemi, Roughgarden and Sridharan [52] (arXiv:2208.06046, 2022) formalised the "loss-versus-rebalancing" (LVR) framework, which shows that even automated market-making structures with no explicit front-running surface impose a measurable cost on liquidity providers from informed-trader arbitrage between fixed price-update intervals. Mao, Zhang, Venkatakrishnan and Lin [56] (arXiv:2405.09465, 2024) further document manipulation surfaces specific to proposer-builder-separation contexts under Proof-of-Stake Ethereum, including future-block-auction strategies that exploit known closing windows. The lesson for the Metha proportional auction is that the *form* of manipulation differs from continuous-orderbook MEV, but the auction is not exempt from the manipulation surface.

Four specific manipulation vectors apply to a perpetual proportional auction with a fixed period boundary, and the paper should acknowledge each:

1. **Last-block sniping.** A large contribution placed in the final block(s) of a period dilutes the share of contributors who participated earlier in the window. Under the proportional rule, a contribution of size *x* in the last block has identical share-weight to a contribution of size *x* in the first block, even though the late contributor bears strictly less risk of intra-period information arrival. The asymmetry is at its worst when periods are short.
2. **Gas auction at period close.** Where multiple actors all wish to contribute in the final block, EIP-1559 [53] priority-fee competition allocates inclusion to whichever bidder is willing to pay the highest tip. Sustained gas-auction dominance by a deep-pocket actor at every period close constitutes a slow capture of disproportionate share even without explicit price manipulation.
3. **Sandwich-at-period-close.** Distinct from the classical DEX sandwich, this is a pattern in which an actor with cheap capital observes the running ETH total mid-window, calculates that their marginal contribution will move them from a "below-average cost-per-METH" position to "above-average", and times the contribution to flip the categorical outcome without paying the market-clearing premium that a continuous auction would have charged. The proportional rule's averaging-over-window property is the structural property exploited.
4. **Multi-period crowding-out.** An adversary with sustained deep capital can dominate several consecutive periods to capture a disproportionate share of cumulative emission. The per-period rule contains this within each window; it does not contain cross-period accumulation, and a sufficiently deep-pocket adversary can in principle accumulate a governance-relevant METH stake more cheaply than a true open market would have permitted.

Mitigations the current contract has. Two structural properties of the deployed auction work in the protocol's favour against the vectors above. *First*, post-window price discovery means the contributor does not know their own marginal cost-per-METH at the moment of submission; the price emerges only after the window closes. This significantly reduces the expected ROI of last-block sniping, because the sniper cannot condition their contribution on a known marginal price — only on the running aggregate. *Second*, the mainnet configuration uses 23-hour periods rather than the 180-second periods of the Sepolia testnet, which corresponds to roughly 5760 blocks per window (at 12-second slot time) rather than 15 blocks. Any single-block disproportionate contribution is diluted across a much larger denominator at mainnet cadence than the DEX-TWAP literature operates against, and "last-block sniping" in the literal-single-block sense is structurally weaker than in DEX MEV contexts. This is not immunity; it is a meaningful constant-factor reduction in the manipulation surface relative to the comparable continuous-orderbook setting.

Mitigations the v3.x roadmap could add. Four candidate mechanisms are under consideration for a future protocol revision, each of which trades complexity for additional manipulation resistance: (i) **VRF-randomised period closing** using Chainlink VRF or an equivalent verifiable-random-function service, so that the exact closing block is not predictable in advance; this directly defeats last-block sniping and gas-auction inclusion at period boundaries. (ii) **Commit-reveal contribution scheme**, in which contributions are committed under a hash during the window and revealed at close, so that intra-window observation of the running total is not possible; this directly defeats sandwich-at-period-close at the cost of two-transaction settlement. (iii) **Batch-auction with off-chain aggregation** in

the style of CoW Protocol [54], in which a batch solver aggregates contributions and produces a single clearing settlement at period close, internalising the front-running surface; this is the most architecturally invasive option and is properly a research direction rather than a near-term roadmap item. (iv) **EIP-1559 base-fee feedback** at the protocol level — e.g. a per-period bonus discount for contributions placed before the final *n* blocks — that makes sustained period-close dominance progressively more expensive. Items (i) and (iv) are the most promising candidates for a v3.x increment because they preserve the current contract's simplicity and predictability while measurably reducing manipulation surface; items (ii) and (iii) are flagged as open invitations for protocol-engineering respondents.

Regulatory consequence. An auction that produces honest within-window price discovery but tolerates between-period manipulation by deep-pocket actors is structurally closer to the MiCA "asset-referenced token" category than to a pure utility token, because price formation is materially susceptible to a small set of large holders. This sharpens the §8.4 / §8.5 regulatory analysis rather than disturbing it: the project's working classification target was already "other crypto-asset" with ART reclassification as the principal residual risk, and the MEV-surface analysis here is consistent with — not contradictory to — that positioning. Honestly characterising the price-formation surface as susceptible to multi-period crowding-out is more defensible regulatorily than overclaiming "non-gameable", because regulators correctly view the latter claim as red-flagged on its face for any on-chain mechanism.

EOS wash-trading and recycling, and the post-batch-auction frontier. Two further acknowledgments belong in this section, both of which a reviewer is correct to press and neither of which weakens the §4.1.1.bis positioning. *(i) EOS wash-trading and recycling.* The EOS empirical precedent invoked in §4.1.1 is not unblemished. The 2017–2018 EOS auction was criticised at the time for coordinated capital cycling — wash-trading and recycling of contributed ETH through related addresses — that inflated headline contribution figures without committing net new capital. The post-window proportional distribution rule that Metha adopts is partially robust to recycling in that a recycled wei contributes its share-weight to whichever address held it at period close, not in aggregate across recycles; the per-period proportional rule does not double-count cycled capital. The rule is not fully robust against coordinated wash by a small set of related addresses that hold the recycled capital at period close in proportions that distort the share calculation. The honest characterisation is that the auction structure constrains the manipulation surface relative to fixed-price or Dutch-auction alternatives but does not eliminate it; the four-vector enumeration above already covers the related multi-period crowding-out vector, and wash-recycling belongs alongside it as a documented residual rather than as a structural weakness denied. *(ii) The literature has moved past batch auctions.* Canidio and Henneke [64] formalise the next step beyond the batch-auction direction listed above as a v3.x roadmap candidate. Their *Fair Combinatorial Auction for Blockchain Trade Intents* argues that *fair combinatorial auctions* are superior to standard batch auctions on both wash-resistance and fairness over multi-asset settlement, and the CoW Protocol's CIP-67 reflects this design evolution. The implication for DVBC is the opposite of the framing some readers may apply: batch auctions are not the current wash-resistance frontier — combinatorial auctions are. The treatment of batch auctions as a "research direction rather than near-term v3.x item" is therefore strengthened, not weakened, by the literature: a v3.x increment toward batch-auction settlement would target a design pattern that the academic frontier has already moved past, while a combinatorial-auction direction is itself a research-frontier proposition rather than a near-term implementation. v1.4 holds both batch auctions and fair combinatorial auctions as longer-horizon design candidates alongside the VRF-randomised closing and EIP-1559 base-fee feedback items already identified as the most promising near-term v3.x increments; the canonical batch-

auction foundation in Budish, Cramton and Shim [65] is retained as the positive-cite anchor for the design family for any reader who wishes to engage the comparison directly.

4.2 Deployed contracts (Sepolia, 2026-05-13)

The current six-contract deployment is summarised in Table 3; the per-contract role is described in §9.1.

AuditorRegistryV2 is new in v3 of the on-chain stack (deployed 2026-05-13); see §5.1 for its role. All six contracts are verified on Sepolia Etherscan. The addresses in Table 3 reflect the deployment at the date of this paper version; the canonical, always-current list is maintained at https://metha.life/#trust. Source code: https://github.com/JVMart/metha (currently private; public release planned alongside mainnet deployment).

4.3 Parameter slots via EternalStorage

The protocol's core contracts are registered in named slots within the `EternalStorage` contract, enabling logic upgrades without data migration. The slot names are utf-8 strings whose `keccak256` derives a deterministic `bytes32` identifier; the address stored at each identifier is the runtime target that downstream contracts resolve through `EternalStorage`. Because the hash is a pure function of the slot name, the four identifiers below are stable across deployments — only the address each slot points to changes when a contract is replaced by governance.

Slot name (utf-8)	`keccak256(slotName)`	Stored target (v3, 2026-05-13)
`"Token"`	`0x1317f51c845ce3bfb7c268e5337a825f12f3d0af9584c2bbfbf4e64e314eaf73`	MethaToken (`0x0A1e36…E16c`)
`"Token Management"`	`0xf211cd286baf73f1080b70167f1e337318ed8eb6a6a297dd53d1eb5ed41cfa42`	MethaCrowdsale (`0x7B96b7…F560`) [1]
`"Auditing Pool"`	`0x4855f03b3d03850cd0e1a6f2ad2d6f13be451750f929ff03199183eb1cbbb7e7`	AuditingPool (`0xD301Cf…8571`)
`"Governance"`	`0x9409903de1e6fd852dfc61c9dacb48196c48535b60e25abf92acc92dd689078d`	OpenGovernance (`0x7BF8dd…917B`) [2]

AuditorRegistryV2 is not registered through `EternalStorage`. Unlike the four parameter slots above, the AuditorRegistryV2 contract is wired peer-to-peer directly with `AuditingPool` via paired setter functions (`AuditingPool.setAuditorRegistry()` and `AuditorRegistryV2.setAuditingPool()`), each callable only by the owner — which at deploy time is the deployer wallet and which is transferred to OpenGovernance as part of the

[1] The slot name "Token Management" is preserved from v2 for storage compatibility; in v3 the slot points at `MethaCrowdsale` (the contract authorised to mint METH at period close), reflecting the v2→v3 consolidation of the token-management role into the crowdsale contract. A future governance proposal may rename the slot for semantic clarity.

[2] The `"Governance"` slot is written by `OpenGovernance.builder()` itself rather than by the deploy script, so the contract self-registers as the authorised writer of the parameter set on first activation.

post-deploy ownership handover documented in §7. This is a deliberate architectural choice rather than an oversight: routing every inter-contract reference through `EternalStorage` would impose a governance round-trip on links that are tightly co-evolving with the registry contract itself, and the `AuditingPool ⇔ AuditorRegistryV2` pair is one such case. The registry can still be replaced through governance — OpenGovernance, once it holds ownership of `AuditingPool`, can call `setAuditorRegistry(newAddress)` directly, achieving the same upgrade affordance that a slot rebind would provide. The trade-off is that swap proposals must target `AuditingPool.setAuditorRegistry()` (or equivalently `AuditorRegistryV2.setAuditingPool()`) rather than an `EternalStorage` slot rebind; the affordance is preserved, the proposal payload differs.

Any governance proposal can replace the contract at any of the four slots above through a successful community vote (§4 of OpenGovernance behaviour, expanded in §5.5 and §10.5), and the auditor registry can be replaced through the direct-owner-call path described in the previous paragraph. Together, these two upgrade affordances enable the protocol to evolve its mandatory settlement requirements — the most legally complex aspect of full DVBC implementation — through community governance rather than requiring a hard fork or centralized administrative action.

Stash architecture position paper — deferred to v1.4. The stash architecture (10% dev / 20% ops, with the vesting schedule and slot-registry mechanics described in this section and the governance affordances described in §10.5) is the subject of a planned position paper that v1.2 of this paper referenced but that the present v1.3 does not yet produce. We commit the position paper as a v1.4 deliverable. Interim guidance is captured in two places: §10.5 of the present paper sets out the design intent (vesting-locked structural float cap, governance-affordance constraints, EternalStorage slot-replacement pathway), and the Metha Roadmap v2 working document [28] sets out the parameter-tuning pathway as a Phase-2 governance task. Readers seeking a deeper treatment than the current paper provides should consult [28] until the position paper is published with v1.4; the position paper will reconcile the design intent with empirical performance data accumulated through the testnet-and-early-mainnet operating window and will close the open parameter-tuning questions [28] flags.

4.4 OpenGovernance: implementing condition (c) — holder sovereignty

Condition (c) requires that METH holders govern the protocol — including, critically, future decisions about mandatory settlement requirements for domain outputs. The `OpenGovernance` contract implements this through snapshot-weighted voting against the slot-based contract registry of §4.3.

Snapshot-based voting power. Voting power in any governance proposal is determined by a snapshot of METH token balances taken at proposal submission, using the snapshot capability of MethaToken. This prevents manipulation through mid-vote token transfers: a participant cannot acquire tokens after seeing a proposal they want to influence and then vote with the freshly-acquired balance. Voting power is a function of long-run METH holding at the moment the proposal was tabled, not of transient position taken in response to the proposal itself. The snapshot mechanism is the proximate defence against the most obvious capture vector — buying votes on a hot proposal — and is reinforced by the broader plutocracy-mitigation discussion in §10.5.

Why holder governance is essential for DVBC. The mandatory settlement requirement (condition a) is the most consequential DVBC parameter. Mandating that all biomedical

patents derived from Metha-funded research settle licensing transactions in METH requires legal, institutional, and market coordination that cannot be achieved by a centralised issuer — it would simply represent a terms-of-service condition of dubious enforceability. But METH holders, as the economic constituency with the strongest incentive to establish and enforce settlement requirements, are precisely the appropriate governing body for this evolution. Condition (c) is not merely a governance nicety; it is the mechanism by which condition (a) will be established over time, and the slot-based upgradability of §4.3 is the technical lever by which that governance decision becomes an on-chain reality.

Complementarity with the planned Stiftung. A reasonable reader may ask why a Liechtenstein Stiftung is needed at all if OpenGovernance already governs the protocol. The answer is that the two organs are complementary, not redundant. The Stiftung is the legal counterparty: it holds the protocol's IP rights (where applicable), signs the funding agreements that carry the settlement clause and the encumbrance recordation of §6.1, contracts with the Spanish CASP-authorised SL, holds insurance policies, and serves as the addressable entity for GDPR data-subject requests, tax filings under DAC8/CARF, and any litigation against the protocol. OpenGovernance, by contrast, is the on-chain decision organ for protocol parameters: it proposes and votes on slot replacements, slashing severity, the 70/20/10 split, the 2/3 acceptance threshold, and condition (a) activation. The Stiftung executes; OpenGovernance decides. The Legal DD §3.6 [27] recommends this B+A hybrid as the structure that simultaneously satisfies CASP licensing requirements (the SL is the regulated leg), foundation-style non-profit anchoring (the Stiftung is the issuer leg), and on-chain stakeholder governance (OpenGovernance is the parameter-setting leg).

The Stiftung residual veto as deliberate centralisation — a tradeoff, not an oversight. The Stiftung's residual veto over governance outcomes represents a deliberate centralisation layer in an otherwise decentralised governance system. v1.3 acknowledges this as a tradeoff and states the reasoning explicitly. The Stiftung-as-controller architecture is required by the legal-and-regulatory layer of the project, not by the protocol's internal logic. Three concrete requirements force the existence of an addressable legal entity with a designated controller: (i) tax filings under DAC8 [34] and CARF [35] for crypto-asset reporting, which presuppose a reportable legal entity; (ii) GDPR data-controller accountability for any personal-data processing the protocol or its surfaces touch (auditor KYB records, contributor wallet-address-to-identity linkages where the Stiftung holds them under the §8.5 two-tier architecture), which presupposes an addressable controller under Art. 4(7) GDPR; and (iii) CASP licensing requirements in Spain (the operating-subsidiary jurisdiction in the B+A entity hybrid), which presuppose a licensed legal person operating the relevant CASP-licensed activities. A pure on-chain DAO with no legal counterparty fails on each of these three independent requirements; the Liechtenstein Stiftung is the architecturally minimal legal entity that satisfies them while preserving the public-benefit-perpetuity property the project's purpose requires. The Stiftung's veto power is statutorily bounded rather than discretionary: under the foundation's charter — drafted to track the public-benefit purpose described in §10 — the executive can only refuse to *execute* a governance outcome on grounds that the outcome would breach the charter, the foundation's regulated-entity obligations, or applicable law in a jurisdiction in which the Stiftung or its controlled subsidiaries are licensed. The executive cannot vote *for* a proposal, cannot initiate proposals, and cannot reverse outcomes after execution. The veto is in this sense a *circuit-breaker* against the highest-consequence governance failures, not a positive policy lever, and the foundation-law fiduciary obligations of the executive constrain its use in ways that an unbounded centralisation surface would not be. The tradeoff is real and the v1.3 paper does not paper over it. Identifying minimum-centralisation architectures that retain regulatory standing — for instance, a multi-

jurisdictional federation of public-benefit entities, each with strictly bounded local veto authority and a clearing-house function for cross-jurisdictional consistency — is an open research direction we explicitly invite; counterproposals from the [Legal] and [Code] respondent tracks are welcomed.

---

5. On-chain audit infrastructure

The audit layer is the mechanism that links emission (condition b) to *productive* investment rather than merely to any ETH contribution. The v3 stack deployed on 2026-05-13 introduces a tiered registry of accreditation firms (AuditorRegistryV2), multi-auditor consensus at every milestone verdict, typed deliverables with content-addressed evidence, on-chain option-period exercise, a Data Safety Monitoring Board veto lane, and a tiered slashing schedule. Each of these primitives mirrors a corresponding off-chain pattern documented in the project's *Auditing and Funding Biomedical Research* methodology brief [29], where BARDA option-period exercise, NIH DSMB review, EMA scientific advice, ARPA-H option-exercise, EIC blended-finance instruments and biotech Series A term sheets are analysed side-by-side.

5.1 AuditorRegistryV2: tiered registration and category specialisation

Auditors in v3 are not anonymous individuals but firms or named experts registered in AuditorRegistryV2, the sixth contract added to the deployed stack on 2026-05-13. Registration requires a METH-denominated stake at one of three seniority tiers — Junior (≥ 10 METH), Senior (≥ 100 METH) and Principal (≥ 1,000 METH) — which gates the value of milestones an auditor may sign off on (Junior caps at 5 ETH per milestone, Senior at 100 ETH, Principal uncapped). A 7-day unstake cooldown prevents same-block exit after a controversial verdict. Firms may further declare specialisation across seven on-chain categories — Generic, ScientificPeerReview, ClinicalSafety, Financial, QualityAssurance, RegulatoryOracle and Programmatic — and accumulate attestations keyed by `certKey` (e.g. GxP, ISO 17025, KYB) issued by accredited oracle entities. The registry is the authoritative source for two downstream gating decisions: which auditors may apply to which research requests (a Junior firm cannot bid on a request whose largest milestone exceeds the Junior value cap; a request typed as ClinicalData cannot be approved without ClinicalSafety auditors in the panel), and which firms hold the DSMB veto lane described in §5.5.

Accreditation criteria for the RegulatoryOracle category. §6.3 below sets out the three oracle architectures under consideration for off-chain settlement verification and identifies the AuditorRegistryV2 *RegulatoryOracle* category as the accreditation surface for the multi-attestor consortium pattern. The canonical framing of the oracle problem this addresses is Caldarelli [58]. Admission to the RegulatoryOracle category requires: (i) **KYB attestation** — the candidate entity must be a registered legal person in a jurisdiction with a recognised company register and must furnish current registration evidence, beneficial-ownership disclosure compliant with the Stiftung's CASP-aligned AML programme, and tax-residency documentation; (ii) **jurisdictional registration evidence** sufficient to identify the regulatory regime under which the entity operates as a professional services firm (e.g. attorney registration, certified public-accounting registration, or equivalent in non-Anglo-American jurisdictions); (iii) **professional liability insurance** of at least €1M aggregate coverage for the categories of attestation work the firm proposes to perform under the Registry; and (iv) **multi-attestor consensus on KYB validity** — at least two other firms already registered in any category of AuditorRegistryV2 must cross-attest the KYB submission as a prerequisite to

category activation. The cross-attestation requirement prevents single-attestor capture of category admission and is materially stricter than the single-attestor floor used for the lower-impact category specialisations. Slashing within the RegulatoryOracle category follows the same Minor/Major/Critical tier schedule that AuditorRegistryV2 applies generally, with the Major tier configured strictly: a first-strike false attestation that a settlement event has occurred (or has not occurred) when the contrary is established triggers a Major-tier slash of at least 25% of staked bond, scaling toward Critical-tier (full bond plus category suspension) on subsequent strikes within a rolling window. The stricter Major-tier configuration is justified by the disproportionate downstream consequence of a false settlement attestation: it directly determines whether the on-chain bond of §6.1 is released or slashed. A 30-day post-attestation challenge window applies to all RegulatoryOracle attestations; during that window any other firm registered under the category may submit counter-evidence to AuditingPool, triggering a re-review process that escalates to a multi-firm panel on the same Condorcet-style aggregation as the milestone-review pathway of §5.2. The challenge window provides a structural delay between attestation and irreversible bond-settlement consequences, which is the principal mitigation against fast-attestation capture by a colluding sub-set of firms; the cost is settlement-finality latency of approximately one month, which is acceptable for off-chain pharma-licensing time-scales (deal timelines are quarterly or longer).

Sybil rotation and panel-membership mitigations. Beyond the multi-attestor consensus and KYB cross-attestation specified above, the RegulatoryOracle accreditation programme will rotate panel membership on a per-attestation basis. Drawing the attestation panel pseudo-randomly from the available accredited pool via Chainlink VRF (or an equivalent verifiable-random-function service, consistent with the VRF-randomised period-closing option discussed in §4.1.1.bis) ensures that a fixed coalition of firms cannot persistently capture attestations for a given lab, therapeutic area or settlement-event class. The 30-day post-attestation challenge window already specified provides the structural delay during which independent firms — including any firm not drawn for the panel on a given attestation — may surface evidence of collusion and trigger the multi-firm re-review pathway of §5.2. The combination of VRF-rotation across attestations and the 30-day challenge window converts persistent-coalition capture from a single-step economic exercise into a multi-vector campaign against a rotating panel, an observable challenge window, and a re-review process whose Condorcet aggregation in §5.2 is itself robust to single-firm dominance. The mitigation is not absolute; it raises the cost of capture materially while leaving the residual exposure that any oracle layer carries (§6.3 “what this does not solve”). The full RegulatoryOracle accreditation playbook is a Phase-2 deliverable and is an open invitation to the [Compliance] and [Legal] respondent tracks.

5.2 Multi-auditor consensus at milestone review

Beyond initial proposal approval, the v3 AuditingPool requires multi-auditor consensus at each milestone verdict: an assigned panel of independent auditors reviews the lab’s proof bundle and a milestone is accepted only when a two-thirds super-majority signs an `acceptVerdict` call. Disbursement amounts where auditors disagree are resolved by a Condorcet-style averaging rule, which also applies to the initial proposal sizing. Formally, the approved milestone schedule is the mean of the recommendations submitted by approving auditors:

$$approvedAmount(k) = \text{mean}\{\ milestone[a][k] : vote[a] = \text{approve}\ \} \quad (6)$$

This Condorcet-style aggregation prevents any single auditor from dominating the sizing of a milestone, while still allowing structurally generous or structurally conservative auditors to register their judgement in the final mean. Funded projects then release capital through a six-state milestone machine — `Locked` → `Unlocked` → `Claimed` → `ProofSubmitted` → `Disputed` → `Approved` — in which each state transition requires on-chain evidence of progress. Funding is released incrementally as milestones are achieved, not as a lump sum at project approval, so that research that stalls at a milestone does not consume treasury resources allocated to subsequent milestones. The 2/3 acceptance threshold is configurable via OpenGovernance proposal and EternalStorage parameter slot; in v3 it is fixed at the canonical two-thirds for resistance to single-firm capture. The pattern aligns the protocol with the panel-review structure documented in the Auditing-Funding methodology brief [29], where BARDA option-period exercise, DSMB veto, NIH study sections and Horizon Europe consortium reviews all rely on panels rather than individuals.

5.3 Typed deliverables and content-addressed evidence

Each milestone submission anchors a vector of up to thirty-two typed deliverables — the cap enforced on-chain by `AuditingPool.MAX_DELIVERABLES_PER_MILESTONE = 32`. A deliverable is a tuple `(DeliverableType, evidenceHash, URI)`, where `evidenceHash` is the content hash of the underlying artefact (IPFS, Arweave or signed PDF) and the URI is a resolvable pointer. The on-chain `DeliverableType` enum distinguishes StudyReport, DataPackage, IPAssignment, Prototype, TechTransfer and Other — categories analogous to (rather than mirrored from) the BARDA deliverable taxonomy and the Horizon Europe periodic-report categories analysed in the methodology brief. The milestone state machine commits to `keccak256(abi.encode(deliverables))` on submission, so any tampering with the off-chain artefact bundle invalidates the on-chain anchor and forces resubmission. This converts the previously opaque `bytes32 proof` into a typed, queryable, externally checkable bundle without imposing a heavy storage cost on the chain. A typed-deliverable bundle for a Phase-I clinical milestone, for example, would include at minimum a StudyReport (the signed protocol-completion report), a DataPackage (the encrypted raw data set), and an IPAssignment (the patent or invention-disclosure record) — each independently hash-anchored and resolvable.

5.4 Option-period exercise: contractual tranching mirrored on-chain

Where a lab anticipates that downstream milestones will depend on a programmatic decision external to the protocol — analogous to BARDA's base-period-plus-option-period structure or to ARPA-H option-exercise — the lab may submit the request with `optionExerciseRequired = true`. Under this flag, milestone 0 unlocks automatically on `acceptVerdict`, but every subsequent milestone $N \geq 1$ stays in the `Locked` state until the Program Officer (granted to OpenGovernance under current configuration) calls `exerciseOption(reqId, N)`. The mechanism on-chain-replicates the contractual tranching pattern that BARDA, ARPA-H, EIC blended-finance instruments and biotech Series A term sheets converge on (Auditing-Funding methodology brief, §3 and §5). The novelty is not the tranching itself — which has decades of off-chain precedent — but the fact that the option-exercise gate is enforced by code rather than by an administrative review queue: a milestone with a met deliverable but an unexercised option does not progress, and a milestone with an exercised option but a contested deliverable does not pay out. The pattern matters because it gives the funder explicit go/no-go gates per phase rather than a single irrevocable commitment, while preserving the lab's ability to plan around a known governance interface.

5.5 DSMB veto lane

For requests with a clinical component, AuditorRegistryV2 supports a Data Safety Monitoring Board veto lane: firms holding `dsmbVetoRights` (granted by oracle attestation in the ClinicalSafety category) can force a clinical-data milestone into the `Disputed` state without joining the consensus quorum. This realises on-chain the safety-monitoring authority that off-chain DSMBs exercise under NIH and FDA frameworks (cf. NIDCR DSMB Guidelines [31] and the FDA "Establishment and Operation of Clinical Trial Data Monitoring Committees" guidance), without claiming to replace the in-person safety review itself. The lab can resubmit a corrected milestone, which clears the veto for one remediation pass; sustained safety concerns can be escalated through OpenGovernance for higher-severity action.

5.6 Slashing schedule

Auditor accountability is enforced through a tiered slashing schedule rather than symmetric reputation points. Slashes fall into three severity bands and escalate by strike count: Minor (3% / 5% / 10% of stake), Major (10% / 25% / 50%) and Critical (100% of stake plus permanent registry ban). Severity is determined by OpenGovernance with documented evidence, and the slashing window stays open for a configurable retroactive period — making rubber-stamping at milestone N still recoverable if material was later found to be fraudulent. The reputation signal therefore lives in the stake balance and strike history itself: a firm at its tier's stake floor with strikes on its record is structurally less attractive to a lab choosing auditors than one that has held a clean record across periods. The full economic-design analysis of the slashing schedule and the auditor compensation matrix is identified as one of three open position-paper triggers in the project roadmap (see §10.5).

Slashing-parameter governance gameability, and the layered defence. The Minor / Major / Critical tier schedule is parameter-revisable through OpenGovernance proposal, which raises a fair question about whether the thresholds themselves can be gamed: an attacker who first manipulates the slashing parameters and then attacks would, on a naive reading, face a smaller economic penalty than the schedule's headline numbers imply. Two structural defences apply. *First*, slashing-parameter proposals are governance-critical under the §10.5 four-layer stack and therefore require the 2/3 supermajority threshold rather than a simple plurality, which raises the accumulation cost of any parameter-game against the same patient-whale residual that §10.5 characterises honestly. *Second*, the slashing window stays open for a configurable retroactive period (already documented above), so a post-hoc reduction of severity does not retroactively reduce slashes already in flight — a successful parameter-game would affect future attestations but not pending ones. The combination is not a guarantee against governance capture in the limit case; it is the layered-defence structure that §10.5 specifies, applied to the slashing surface specifically. Counterproposals on slashing-parameter governance — including arguments that the 2/3-supermajority gate is insufficient for the slashing surface, or that an explicit timelock on slashing-parameter changes should be added — are an open invitation to the [Code] and [Compliance] respondent tracks.

---

6. Settlement clause and IP encumbrance

A foundational architectural choice in the Metha protocol is that the protocol does not acquire title to the intellectual property generated by funded research. Patents, licensing agreements, treatment protocols, and other commercializable outputs remain the property of the principal

investigator or the legal entity through which the research is conducted (a university tech transfer office, a researcher-led spinout, or a research cooperative). The protocol's relationship with the IP is contractual, not proprietary.

The mandatory settlement requirement (condition a) is implemented as a clause in the funding agreement that the researcher signs when accepting Metha financing. The clause obligates the researcher and any downstream rights-holder to denominate licensing transactions with commercial counterparties in METH. This is a contractual restriction on how the IP is licensed, not an ownership claim on the IP itself. The closest legal analog is the copyleft licensing tradition (the GNU GPL [25], Creative Commons ShareAlike): the original rights-holder retains ownership but accepts a redistribution condition as part of using the work or, in this case, of receiving the public-goods funding.

This architectural choice has three consequences worth making explicit.

(i) Metha holds no income-producing assets. The treasury holds ETH (capitalized through the auction) which is disbursed to fund research. The protocol is not an investment fund. This distinction materially weakens the "common enterprise + profits derived from the efforts of others" prong of the Howey test discussed in §8.4: Metha cannot be characterized as a fund manager because there is no managed asset portfolio.

(ii) The settlement requirement operates at contract law, not at the level of patent ownership. Pharmaceutical licensees do not negotiate with Metha over IP terms; they negotiate with the researcher (or the rights-holder entity), who is contractually obligated to denominate the deal in METH as a downstream condition of having received Metha funding. This places the settlement requirement in the same legal category as performing-rights-organization royalties (ASCAP, BMI, GEMA, governed in the U.S. context by the long-standing DOJ consent decrees [41]) or copyleft license conditions — neither of which has historically been treated as a securities-law instrument.

(iii) Researcher autonomy is preserved. Researchers retain full ownership of their IP and full control over commercialization decisions, subject only to the settlement-currency clause. They are free to license to any pharmaceutical company, at any price, on any other terms — only the settlement medium is fixed. This preserves the standard incentive structure of academic and commercial biomedical research and avoids the more contentious institutional commitments that would be required if Metha itself were claiming an IP interest.

The settlement layer / IP custodian distinction is the architectural decision that separates DVBC from existing DeSci protocols (e.g., VitaDAO's IP-NFT model, in which the DAO holds fractional IP rights). DVBC at Metha is a *settlement infrastructure* over independently held IP, not a *fund* that owns it.

6.1 Enforcement: Encumbrance, March-In, and Anti-Shelving Mechanisms

The architectural choice above — Metha as settlement layer rather than IP custodian — raises a natural enforcement question: once the protocol disburses funding and milestones complete, what prevents the rights-holder from circumventing the settlement requirement? Three concrete failure modes must be addressed: (i) outright transfer of the IP to a buyer who does not assume the settlement obligation; (ii) acquisition for the purpose of shelving the technology to suppress competition; and (iii) jurisdictional escape through transfer to a legal entity in a non-cooperating jurisdiction. Each is addressed by an existing legal precedent, and the protocol composes the corresponding mechanisms into a layered enforcement structure.

The framing below has been tightened relative to v1.2 to reflect a reviewer observation we accept in substance: patent-exhaustion doctrine in the United States closes much of the "viral encumbrance binding all subsequent purchasers" path that copyright-derived analogies invite, and the analysis must be honest about which layer carries which weight.

Mechanism A: contractual succession plus recordation as a separate-but-weaker overlay. The funding agreement includes a binding succession clause: the obligation to denominate downstream licensing in METH applies not only to the original rights-holder but to all subsequent assigns, transferees, and licensees of the rights-holder lab and any controlled entity. This is a *contract-law* mechanism: it binds the lab and parties in privity with the lab through the succession-and-assignment chain, with breach damages and arbitration consequences specified in the contractual layer of Mechanism C below. The settlement requirement is recorded with the relevant patent offices (USPTO under 35 U.S.C. §261, EPO under EPC Art. 72/73 and the implementing regulations, JPO, KIPO) as a recorded interest on the patent, primarily to provide constructive notice to subsequent acquirers and to support the contract-layer remedies against a transferee who took with notice. We no longer claim that recordation creates an *erga omnes* obligation that survives transfer to a bona-fide downstream purchaser in the patent context. *Quanta Computer, Inc. v. LG Electronics, Inc.*, 553 U.S. 617 (2008) [48] and *Impression Products, Inc. v. Lexmark International, Inc.*, 581 U.S. 360 (2017) [49] progressively closed the door on attempts to impose post-sale licensing restrictions running against downstream purchasers of patented articles; the practical effect is that any "viral" patent-term claim against a bona-fide downstream purchaser who is not in privity with the lab is, in the U.S. patent jurisdiction at least, presumptively unenforceable. The *Jacobsen v. Katzer*, 535 F.3d 1373 (Fed. Cir. 2008) [24] analogy that this section relied on in v1.2 is therefore retained only for what it actually establishes: that a license condition — including a settlement-currency condition — is judicially enforceable as a *condition of license* against the licensee, with both equitable and monetary remedies. That is the lab→Metha relationship, not the lab→pharma→generic-manufacturer chain. The GPL parallel [25] is similarly load-bearing for the copyright surface of any software artefacts that travel with the funded work, but is *not* a transferable argument for viral patent-term enforcement against downstream patent purchasers.

The consequence of this tightening is that the on-chain bond (Mechanism C, on-chain layer) and the Metha Pharma backstop (§10.4) become *more* load-bearing, not less. Under v1.2's framing, the legal layer carried much of the third-party enforcement weight. Under the v1.3 framing, the legal layer carries the rights-holder-and-affiliates surface, while third-party enforcement is principally a function of (a) the lab's own continuing economic relationship with the protocol (gradual bond release contingent on annual compliance attestation), (b) the Metha Pharma residual licensee of last resort (§10.4) for cases where no commercial counterparty is willing to license under settlement terms, and (c) the reputational layer's chilling effect on regulated pharma counterparties whose own audit trails make knowing acquisition of an encumbered patent materially costly even where it is not legally unenforceable. The protocol does not require *erga omnes* patent encumbrance to function; it requires that the rights-holder's expected-value calculus include the bond, the reputational record, and the long-run loss of franchise that march-in would precipitate (Mechanism B, unchanged).

Mechanism B: march-in rights and anti-shelving provisions. The strongest precedent for preventing IP capture-and-suppress strategies is the Bayh-Dole Act (35 U.S.C. § 203, 1980), which permits the federal government to "march in" and grant compulsory licenses if a federally-funded invention is not "reasonably commercialized" within a defined period. Metha

replicates this structurally: the funding agreement provides that if the rights-holder fails to demonstrate good-faith commercialization activity within a specified period (typically 5–7 years post-patent grant), OpenGovernance may, by qualified majority, authorize compulsory licensing of the IP to qualified third parties on FRAND terms. Two reinforcing provisions accompany this: a **use-it-or-lose-it reversion clause** triggers automatic loss of exclusive rights if no good-faith commercialization milestone is met within the period, and a **non-exclusive licensing requirement** prohibits worldwide-exclusive licenses (regional or field-of-use exclusivity remains permitted). The combined effect is that the strategy of "buy worldwide-exclusive license, then shelve" is structurally precluded: any acquirer who buys to suppress must accept that the IP will be re-licensed to competitors within years.

Mechanism C: dual enforcement layer (legal + on-chain + reputational). Mechanisms A and B establish what the rights-holder must do; mechanism C establishes what happens if the rights-holder breaches. Three reinforcing layers operate concurrently.

The **contractual layer** specifies liquidated damages on detected breach (commercialization without METH settlement, transfer without succession clause, false attestation). Damages comprise (i) full restitution of the funded amount with interest at a defined rate, (ii) a liquidated-damages multiplier of 1×–3× of funded amount representing pre-estimated diffuse ecosystem harm — reasonable under common-law liquidated-damages doctrine if calibrated to documented harm rather than punitive, and (iii) disgorgement of profits earned through the non-compliant commercialization, an equitable remedy with strong precedent in unjust-enrichment doctrine. Choice-of-law provisions (Swiss or Singapore) and arbitration under SIAC or Swiss Chambers rules provide international enforceability via the New York Convention [26], ratified by 172 signatory jurisdictions.

The **on-chain layer** requires the rights-holder to post a METH bond at funding inception (typically 10–20% of funded amount), released gradually over a 20-year period contingent on annual on-chain attestation of compliance. If OpenGovernance establishes breach by qualified threshold (e.g., 2/3 majority, supported by documented evidence), the bond is programmatically slashed to the treasury without judicial adjudication. This provides enforcement on the order of weeks rather than the months or years required for arbitration, complementing the slower contractual layer with an immediate-friction mechanism.

The **reputational layer** records all detected breaches as permanent on-chain records associated with the rights-holder identity and institutional affiliation. This information becomes part of the protocol's public record for due diligence by future funding bodies, academic partners, or investors. The reputational cost is material in the academic and biomedical research ecosystems, where institutional trust is a primary asset and where opportunities (grants, partnerships, board seats, advisory roles) are routinely contingent on track record.

The combined effect of mechanisms A, B, and C is that breach is detectable, enforceable on multiple time horizons (weeks via bond slashing, years via arbitration awards, decades via reputational record), and economically irrational for any actor seeking to evade the settlement requirement. The structure does not assume good faith — it makes good faith the dominant strategy. The legal-precedent foundation (Bayh-Dole, GPL viral, FRAND, *Jacobsen v. Katzer*, NY Convention) means that the enforcement framework operates within established doctrines rather than relying on novel legal interpretation.

The "rich-lab-walks" scenario. A determined critic will press the obvious adversarial case: what if an early lab funded by Metha builds a commercially successful franchise on the back of

that funding, then concludes that the marginal cost of forfeiting the on-chain bond and absorbing reputation damage is lower than the marginal cost of settling future licensing revenue in METH, and simply walks? The architecture handles this case along three reinforcing axes — with the on-chain bond, not the recordation layer, doing the load-bearing work in exhaustion-respecting jurisdictions. First, where local doctrine permits, Mechanism A's recordation of the settlement clause against the rights-holder lab itself creates a contractual succession path that runs against the rights-holder and against parties on notice of the encumbrance; §6.2-bis is explicit that, under *Quanta* [48] / *Impression Products* [49] exhaustion doctrine, this path does *not* automatically bind a bona-fide downstream purchaser of the patented article — and that the soft-variant clause is designed to operate as a contractual condition between Metha and the rights-holder rather than as a post-sale restriction on a downstream commercial counterparty. The walk-and-license-downstream strategy that exhaustion enables is precisely the scenario the on-chain bond and the Metha Pharma backstop are calibrated to absorb, not the scenario the recordation layer alone can prevent. Second — and primarily — the on-chain bond (Mechanism C) is calibrated to dominate not the absolute economics of breach but the *expected-value* economics conditional on detection probability: if the bond is sized at 10–20% of funded amount and slashing happens in weeks via the §6.1 OpenGovernance threshold, the breakeven for a walk requires the lab to expect that detection probability is sufficiently low, which is implausible once the IP is publicly licensed and a regulated pharma counterparty has been observed exploiting it. Detection in this regime does not require the downstream counterparty to be in breach — it requires only that the rights-holder lab has elected to forgo the bond-release-on-compliance benefit, which is observable from the public on-chain attestation record. Third, Mechanism B (Bayh-Dole-style march-in under 35 U.S.C. § 203 [66] where federal funding lineage applies; FRAND-style compulsory-licensing analogues in other jurisdictions) is the long-horizon backstop: a lab that walks while shelving the IP triggers a compulsory-licensing pathway in which the rights-holder cannot economically afford to let the IP be re-licensed at administered terms, because that converts a single non-compliant license into a permanent loss of exclusivity. The combination is sufficient to make "walk and absorb the slashing" a structurally dominated strategy in expected-value terms — not because the lab cannot afford to walk, but because the bond plus the long-tail backstop together convert a one-off breach into a perpetual economic loss. Metha Pharma (§10.4) is the residual safety valve and, under the §6.2-bis reframing of exhaustion limits, is the primary backstop where the contract-law succession chain has been broken by a bona-fide downstream transferee: even there, the protocol retains a fair-terms commercialisation path, and the rights-holder retains no economic upside from the walk. The framing of v1.2 and v1.3 — which leaned more heavily on recordation as binding downstream counterparties — has been retracted here and in §6.2-bis to match the exhaustion-aware position; v1.4 onward treats the on-chain bond + Metha Pharma backstop as the load-bearing pair, and recordation as a useful but jurisdiction-heterogeneous adjunct.

6.2 Jurisdictional caveats and the soft-variant settlement clause

Two caveats temper the framework as described. First, *erga omnes* enforceability of patent-office encumbrance is heterogeneous: USPTO and EPO admit recordation but successor clauses are not automatically opposable to bona-fide purchasers in JPO and KIPO regimes, where running covenants are more restrictively treated, and the UPC has yet to issue a controlling decision on DLT-conditional licences (Legal DD §2.5.1 [27]). Second, the *Brulotte* doctrine (379 U.S. 29) [32] and the *Genentech* line (C-567/14) [33] authorise limited running royalties but cap conditions that survive patent expiry. The on-chain bond described above is

therefore not redundant but load-bearing: it is the layer that survives the moments when registry recordation does not bind a downstream acquirer. This is consistent with the Legal DD recommendation to treat the three enforcement layers — contractual succession, encumbrance recordation, and on-chain bond — as belt-and-braces rather than ranked alternatives. The DD further recommends adopting a *soft* variant of the settlement clause (settlement-as-option with bond-release acceleration rather than absolute obligation), which materially reduces the Art. 102 TFEU tying-arrangement risk while preserving the structural-demand mechanism described in §3.

[OPEN QUESTION v1.2: the *Brulotte* doctrine caps royalties at patent expiry, which raises a specific question for any settlement framework that aspires to apply to generic-pharma licensing on expired patents. The working answer — pending a U.S. patent-counsel opinion — is that *Brulotte* limits *patent-tied* running royalties beyond expiry but does not block settlement-currency obligations on *trade-secret*, *regulatory-data-exclusivity* or *branded-formulation* licensing that may extend beyond patent term, and does not affect the structural-demand mechanism for as long as the patent remains in force. The applicability boundary for expired-patent generics is under active investigation in the Legal DD follow-on engagement.]

6.2-bis Patent exhaustion and the soft-variant clause as the correct positioning

The recommendation in the Legal DD §3 [27] to adopt the *soft variant* of the settlement clause — settlement-as-option with bond-release acceleration on compliance, rather than absolute mandatory settlement — was originally framed in this paper as a competition-law accommodation under Art. 102 TFEU. Under the v1.3 framing it is also, and arguably more fundamentally, the correct positioning under U.S. patent-exhaustion doctrine. An "absolute" mandatory-settlement clause that purported to bind downstream purchasers of patented articles or their generic successors would invite a *Quanta/Impression Products* challenge on first principles. A soft-variant clause that conditions a bond-release benefit on the lab's continued election of METH settlement — but does not purport to encumber the patent against arms-length acquirers — is structured to be compatible with both exhaustion doctrine (it is a contractual condition between Metha and the rights-holder, not a post-sale restriction on the patented article) and competition doctrine (it is a contingent MFN-style provision under EU Art. 101 jurisprudence, not a tying arrangement under Art. 102; see §8.4 and the *Booking.com* / *Amazon* discussion). The two reframings converge on the same recommendation: the soft variant is the principal text of the settlement clause; the absolute variant is reserved for jurisdiction-specific drafting where local doctrine permits it and competition counsel signs off, and even there the on-chain bond remains the operationally load-bearing layer.

6.3 Oracle layer for off-chain settlement verification

A DVBC value loop in which condition (a) is active requires the protocol to verify that off-chain commercial transactions have actually settled in METH. The canonical framing of this challenge is the "oracle problem" surveyed by Caldarelli [58]: any on-chain protocol that conditions its execution on off-chain state must specify how that state enters the on-chain decision surface, and the trust assumptions of the oracle layer become the trust assumptions of the protocol itself. The slashing path of Mechanism C in §6.1 conditions the on-chain bond on OpenGovernance "establishing breach with documented evidence" — but v1.2 of this paper left the question of how that evidence enters the on-chain decision surface unspecified. The reviewer is correct that this is a load-bearing gap, and that the protocol cannot route around it by appealing to optimistic assumptions about rational lab behaviour. This section specifies the

three concrete oracle architectures under consideration and is explicit about what the chosen combination does and does not solve.

Option 1 — Attestation-based ("self-report under penalty of slashing"). Under this architecture the rights-holder and the licensee jointly produce an EIP-712 [50] typed structured signature over a settlement-event record (transaction value, METH amount transferred, on-chain settlement tx hash, licensed-patent identifier, period of obligation), the content hash of the signed record is anchored on-chain via an `attestSettlement(bytes32 evidenceHash, uint256 amount, address licensee)` call on the AuditingPool, and the underlying signed record is published to a content-addressable store (IPFS / Arweave / signed-PDF) under the existing §5.3 typed-deliverable pattern. **Pros:** very cheap (one EIP-712 signature plus one transaction per settlement event), no external trust assumption beyond the rights-holder and the licensee themselves, fast (single-block latency), and the cryptographic primitives are mature and well-audited. **Cons:** the attestation is only as reliable as the joint rationality of two parties whose private interest may be to under-report or mis-classify a transaction; collusion between rights-holder and licensee defeats it; the cost of fraudulent attestation is bounded by the bond size and the detection probability, which is the whole point of the bond's calibration in §6.1 but is also a non-zero residual.

Option 2 — Oracle consortium (multi-attestor, KYB-bonded). Under this architecture a panel of independent oracle nodes — themselves accredited under AuditorRegistryV2 with the RegulatoryOracle category specialisation (§5.1) — each independently retrieve the off-chain evidence (licensee invoice, banking confirmation of payment in METH, patent-office records confirming the licensed scope), perform their own off-chain due diligence, and submit a typed verdict; the on-chain settlement record is taken as confirmed only on a 2/3 super-majority of oracle attestations, with disagreement resolved by the same Condorcet-style aggregation as the milestone-review pathway in §5.2. **Pros:** third-party verifiability (no single party can fabricate a settlement event without colluding with a super-majority of the oracle panel); the bonded oracle nodes have a direct economic stake in their attestation accuracy via the AuditorRegistryV2 slashing schedule (§5.6); and the result is enforceable on roughly the same timeline as a milestone verdict. **Cons:** the consortium itself is a capture surface — a coordinated subset of oracle firms can in principle confirm fraudulent settlement or deny legitimate settlement; the ongoing per-event cost is non-trivial; latency is days-to-weeks rather than blocks; and the oracle node set cannot bootstrap from zero — populating the panel with sufficient KYB-attested firms across the relevant therapeutic and jurisdictional surface is a Phase-2 execution problem.

Option 3 — Stiftung-as-controller (legal entity as off-chain dispute resolver). Under this architecture the Liechtenstein Stiftung (the issuer leg of the B+A hybrid described in §4.4 and §8.5) is the data controller for off-chain settlement evidence under the §8.5 GDPR architecture, and its executive — under the foundation's by-laws and the contractual relationships with funded labs — holds the standing and the subpoena power necessary to compel production of settlement evidence from the rights-holder and, by contractual flow-down, from the licensee. The Stiftung indexes attestations on-chain through a single privileged `indexAttestation()` call that records the existence and integrity of the off-chain record set without revealing its contents. **Pros:** aligns directly with the B+A entity architecture that the legal layer already requires the project to stand up; the Stiftung is the addressable legal counterparty for GDPR, tax filings, and litigation in any event (§8.5); the Stiftung's foundation-law fiduciary obligations to the perpetual public-benefit purpose constrain its decision-making in ways a pure consortium does not. **Cons:** introduces a trusted intermediary at a point in the architecture where the rest of the protocol is permissionless;

slower than either alternative; vulnerable to entity capture; and creates a single-jurisdiction failure surface that the rest of the protocol design tries to avoid concentrating.

The protocol-current preferred path is (1) + (3) as belt-and-braces, with (2) reserved for a later phase. The attestation primitive of Option 1 is cheap enough to deploy at every settlement event by default. The Stiftung-as-controller indexing of Option 3 is the dispute-resolution lane that activates when an attestation is challenged or when an external party reports a suspected non-settled commercialisation. Option 2 is held in reserve for the phase at which the AuditorRegistryV2 ecosystem is mature enough — that is, after enough accredited firms have populated the RegulatoryOracle category — to constitute a non-degenerate consortium; until then, a two-firm "oracle consortium" would be a worse trust assumption than the Stiftung alone, not a better one. The progression from (1)+(3) to (1)+(2)+(3) is therefore an ecosystem-maturity decision, not a contract-redeployment decision, and is governed through the same OpenGovernance and EternalStorage slot-replacement affordances described in §4.3.

What this section explicitly does not solve. The architecture above is robust against any one of the load-bearing parties acting in bad faith, but it is not robust against the joint failure of all three. A rights-holder who has decided to walk, a licensee pharma counterparty that is complicit in mis-reporting, and a Stiftung executive that is captured or compromised — the combination defeats the oracle layer in every option above. The Metha Pharma backstop (§10.4) is the residual fix for this limit case: the protocol retains a fair-terms compulsory-licensing path that does not require the bad-faith counterparty's cooperation. We do not claim the residual fix is costless or instant; we claim it converts a "bad actor wins" outcome into a "bad actor loses the franchise and the protocol re-licences to a clean counterparty" outcome. Specifying the oracle layer further — and in particular reducing the residual trust assumption on the Stiftung — is an open invitation channel for protocol-engineering respondents; counterproposals from the [Code] and [Legal] tracks are explicitly welcomed.

6.4 Conflicts of laws: civil-law vs common-law treatment of DLT-conditioned patent encumbrances

The exhaustion-doctrine analysis of §6.1 and §6.2-bis is principally a United States patent-law analysis, anchored in *Quanta* [48] and *Impression Products* [49]. A protocol that aspires to a global biomedical-IP surface cannot stop there. The treatment of DLT-conditioned patent encumbrances diverges materially between common-law jurisdictions (United States, United Kingdom) and civil-law jurisdictions (Spain, Germany, France, and the EU more broadly via the EPC/UPC infrastructure), and the v1.3 paper should state that divergence rather than smuggle it into the U.S. framing.

In the United States, exhaustion as articulated in *Quanta* and *Impression Products* closes the principal pathway by which a settlement-currency condition could attach to a patented article downstream of first authorised sale. Patent-misuse doctrine, which historically punished tying-style overreach with unenforceability of the patent itself, reinforces this closure; a USPTO-recorded settlement clause that purports to run with the article against a bona-fide downstream purchaser invites both an exhaustion defence and a misuse counter-claim. The soft-variant settlement clause (§6.2-bis) is the textually defensible posture in the U.S. forum.

In the European patent system, the situation is structurally more nuanced. EPC Art. 73 admits transfers and licences of European patent applications and grants subject to conditions, and EPC Art. 74 routes substantive transfer-validity questions to national law; the EPC itself is

doctrinally neutral on whether a DLT-conditioned licence is enforceable beyond the immediate licensee. The Unified Patent Court, operational since June 2023, has not at the time of v1.3 issued controlling case law on DLT-conditional licences specifically, and CJEU jurisprudence on patent exhaustion (notably the *Centrafarm* line) has been confined to parallel-import scenarios rather than viral licence-term scenarios. The practical consequence is that the European doctrinal surface is *open* on the precise question this protocol raises, rather than closed against it; that openness should be characterised as a research opportunity, not as a foundation for confident enforceability claims.

Within civil-law member states, running-covenant doctrines (*Realobligationen* in German civil law, *obligaciones propter rem* in Spanish, *obligations réelles* in French) admit some forms of obligation-running with property but are materially more restrictive than the common-law equitable-servitude tradition. None of them maps cleanly to a settlement-currency obligation attached to a patent; the better characterisation in civil-law analysis is contractual succession reinforced by patent-office recordation as constructive notice, which is the same posture §6.1 adopts on the U.S. side. The doctrinal asymmetry between U.S. and EU jurisdictions is therefore narrower than a first reading might suggest, once both sides are reduced to their actually-functioning mechanisms: contractual succession plus bond, with patent-office recordation as a notice-shifting overlay rather than as an enforcement instrument.

The operational consequence for the protocol is that the soft-variant clause (§6.2-bis) is the only formulation that is jurisdictionally robust across the surface; an absolute-mandatory clause is defensible only in select fora and would not be defensible in others under current doctrine. The on-chain bond (Mechanism C, §6.1) and the Metha Pharma residual-licensee backstop (§10.4) carry decisive operational weight in the residual gap between what the legal layer can enforce in any given jurisdiction and what the protocol requires. The reframed §6.1 architecture is consistent with this conflicts-of-laws analysis by construction. By analogy, the broader DLT-as-register literature has reached a similar conclusion: DLT recordation does not displace the assurance function of state-backed registers, but it can document and operationalise commitments that the state-backed register itself cannot police [57]. We adopt the analogy with care — Quinn and Connolly address real-property land registers, not patent registers — but the structural lesson transfers. This is preliminary multi-jurisdictional analysis; a full conflicts-of-laws audit across the seven jurisdictions identified as primary commercial surface in the Legal DD [27] §2.5 — United States, United Kingdom, Germany, France, Spain, Switzerland, Japan — is committed as a Phase-2 deliverable and is an explicit Open Invitation to the [Legal] respondent track.

---

7. Implementation status and roadmap

The current Metha implementation directly satisfies conditions (b) and (c) and establishes the infrastructure for condition (a). The path to full DVBC realization proceeds through three phases, the second and third of which are deliberately date-unspecified and audit-gated.

Phase 1 (Sepolia testnet, deployed 2026-05-13). Token auction live; AuditingPool review operational; OpenGovernance contract deployed; AuditorRegistryV2 deployed with tiered seniority, category specialisation and DSMB grant infrastructure. Condition (b) is live: METH emission is mechanically linked to research investment. Conditions (c) is live: METH holder governance of all protocol parameters via OpenGovernance and slot-based contract registry. The v3 capabilities described in §5 — multi-auditor consensus at every milestone verdict,

typed deliverables with content-addressed evidence, on-chain option-period exercise, DSMB veto lane, and the tiered slashing schedule — are all live on testnet. No research outputs have yet been produced; condition (a) is not yet applicable.

Phase 2 (mainnet deployment, date deliberately unset). Mainnet auction with 23-hour periods is gated on the cumulative completion of four sequential audits — legal due diligence (now complete at v1.0, May 2026), BARDA-method audit, smart-contract security audit, and financial audit of treasury and stash architecture — together with the standing-up of a Liechtenstein Stiftung or equivalent (Legal DD §3.6 recommends a B+A hybrid: foundation + Spanish CASP-authorised SL), a multisig custody arrangement for the testnet owner address, professional liability insurance and a bug-bounty pool. The Legal DD estimates a 15–18 month critical path from May 2026 under the recommended scenario, with year-1 costs in the €520k–€1.35M range and recurring annual costs of €280k–€620k. The roadmap deliberately does not commit to a calendar date for mainnet; commitment is gated on outputs, not on time.

The choice of Liechtenstein as the principal foundation jurisdiction is not arbitrary and reflects a structured comparison performed in the Legal DD [27] against the candidate alternatives: Swiss Stiftung (Zug), Singapore Foundation Company, U.S. Delaware C-corp or unincorporated association, BVI/Cayman foundation, and Estonia OÜ. The relevant decision criteria were: (i) explicit DLT/blockchain framework in force (Liechtenstein's TVTG/Blockchain Act is the most mature in the EEA, predating MiCA by three years and dovetailing with it cleanly); (ii) compatibility with MiCA passporting (Liechtenstein is in the EEA and a Liechtenstein-issued instrument can be passported into the EU single market, which a Swiss-issued instrument cannot without additional structuring); (iii) foundation-law fit for a non-profit anchor (Liechtenstein Stiftung law is mature, predictable, and explicitly recognises perpetual non-profit purpose, where Singapore Foundation Companies and Delaware unincorporated associations have less developed jurisprudence on perpetual public-benefit purposes); (iv) tax neutrality on protocol-level transactions; and (v) recurring operational cost. Switzerland was a close second on every criterion except MiCA passporting; Singapore was strong on (i) and (iv) but loses on MiCA reach; Delaware was excluded because of the SEC enforcement-priority overhang. The B+A hybrid recommendation (Liechtenstein Stiftung issuer + Spanish CASP-authorised SL operating arm) is the structure that simultaneously satisfies foundation-law fit, MiCA reach, and CASP regulatory standing without forcing all functions into a single jurisdiction. The detailed multi-criterion comparison, with cost tables and risk-matrix scoring, is published in §3 of the Legal DD; readers seeking the jurisdiction-selection rationale beyond the summary above are directed there.

Cost-band methodology and sensitivity. The €520k–€1.35M Phase-1 and €280k–€620k recurring cost bands quoted above are derived from the Legal DD [27] §3.8 cost tables and are composed of five categories: legal counsel (≈€180k–€480k Year 1, covering Stiftung-charter drafting, Spanish-SL incorporation counsel, CASP-license counsel, securities counsel across the seven primary commercial jurisdictions of §6.4, and patent counsel for the §6 settlement-clause work); accounting and audit (≈€60k–€120k Year 1, covering Stiftung statutory audit, Spanish-SL statutory audit, and the consolidated reporting work required under DAC8/CARF); Stiftung incorporation and tributary fees (≈€80k–€200k Year 1, covering the Liechtenstein statutory fees, foundation-board compensation for the first operating year, and the Liechtenstein FMA-registration administrative fees); Spanish SL incorporation and CASP licence (≈€120k–€350k Year 1, covering the SL statutory incorporation, the CASP-license application fees and counsel time, and the AML programme stand-up); and insurance and bonding (≈€80k–€200k Year 1, covering the professional-liability programme, the directors-and-officers programme for the foundation board, and the cyber-and-protocol-failure

insurance that the §8.5 architecture treats as part of the operational baseline). The low band assumes a Liechtenstein–Spain footprint operating under a shared-services arrangement; the high band assumes additional jurisdictional coverage to support operational presence in the United Kingdom and Switzerland during Phase-1 stand-up. These figures are illustrative, not firm quotes; they are based on the Legal DD's market-rate assumptions current as of Q1 2026 and will fluctuate with counsel-firm engagement terms, regulatory-fee adjustments, and exchange-rate movement against the euro reference. The range refreshes per quarter as part of the Stiftung's standing operations review.

Phase 3 (condition (a) activation, post-mainnet, no fixed date). Activation of mandatory METH settlement for biomedical IP licensing originating from Metha-funded research. Pre-requisites: governance maturity (the open voting-model question must be closed, see §10.5), critical mass of auditor firms (at least one accredited DSMB-authorised firm per major therapeutic area, plus oracle-attested KYB across the auditor pool), and a legal opinion from at least two independent counsels on the enforceability of the soft-variant settlement clause in the major patent jurisdictions (USPTO / EPO / JPO / KIPO).

At Phase 3 completion, Metha METH satisfies all three DVBC conditions. The value of METH is then backed not by declaration, not by commodity storage, and not by scarcity alone — but by the ongoing productive output of the biomedical research ecosystem it funds and in which it settles.

---

8. Implications and Comparison with Existing Models

8.1 How DVBC Differs from Existing DeFi Token Models

The most common objection to DVBC classification of Metha METH is that it resembles existing DeFi governance tokens with an optimistic narrative about future settlement requirements. This objection misreads the causal structure.

Governance tokens like COMP (Compound), UNI (Uniswap), or AAVE derive value from control over protocols whose value is denominated entirely within the DeFi ecosystem. Their structural demand is internal: users need governance tokens to influence protocols they are using. There is no external productive domain whose commercial outputs generate settlement demand independent of the DeFi ecosystem itself.

METH's structural demand mechanism, when Phase 3 is reached, derives from the commercial value of biomedical research — a domain that exists and generates value entirely independently of the Ethereum ecosystem. A pharmaceutical company licensing a patent derived from Metha-funded research would need METH not because they are DeFi users but because the settlement requirement is a condition of the IP license signed by the researcher (§6). This is the same structural difference between the petrodollar (mandatory settlement in an external commodity market) and a governance token (demand within an endogenous ecosystem). Metha is settlement infrastructure across an external productive domain; governance protocols are control mechanisms over endogenous capital pools. The two are categorically different instruments.

Furthermore, METH emission is linked to research investment (condition b) by protocol, not by narrative. The MethaCrowdsale contract's relationship between ETH in and METH out is deterministic and verifiable on-chain. This is fundamentally different from a governance

token whose emission schedule is determined by governance vote (creating circular dependency) or liquidity mining incentives (creating temporary demand without structural backing).

8.2 Relationship to Sectoral Currency Proposals

The DVBC model relates to a broader literature on sectoral or domain-specific currencies. Keynes proposed a supranational currency (the Bancor) for international trade settlement in his 1942 *Proposals for an International Clearing Union* [19]. Mundell [12] established the theory of optimal currency areas — the conditions under which a group of economies benefits from sharing a currency. The DVBC model can be read as extending Mundell's framework to the domain dimension: the "optimal currency area" for a DVBC is defined not geographically but by productive domain membership.

Recent proposals for energy currencies (electricity-backed tokens like WPR [13]) and carbon credits (carbon-backed tokens like UPCO2) partially instantiate DVBC logic. Energy tokens link to commodity output, not to investment in energy production capacity — they satisfy condition (a) partially but not condition (b). Carbon credit tokens link to emissions reduction outputs but have no mandatory settlement requirement and no emission-investment mechanism. Neither satisfies all three DVBC conditions.

The DVBC framework provides a generalization that clarifies why these partial implementations lack the self-reinforcing value loop: without all three conditions, the feedback mechanism is incomplete.

8.3 Scalability: Implied METH Demand Under Growth Scenarios

To ground the DVBC value claim empirically, consider the implied METH demand under different health domain growth scenarios.

Let L = revenue from research outputs settled in METH as a fraction of total biomedical IP-derived commercial revenue. Let R = total commercial revenue from biomedical IP-protected products. With the global pharmaceutical market at approximately $1.6T in 2024 (IQVIA *Global Use of Medicines 2024* baseline, with the *Global Trends in R&D 2024* methodology brief reporting an adjacent figure depending on inclusion scope) [10] and patent-protected products comprising the majority share, the IP-derived revenue base lies in the $700B–$1.1T range depending on inclusion scope. The required METH holdings H for settlement are then approximately:

$$H \approx \frac{L \cdot R}{P(\text{METH})} \quad (7)$$

where *P(METH)* is the METH price in external currency, and H denotes required METH holdings (distinct from *S(t)* in §3.2, which denotes the demand flow). This is a simplified upper-bound estimate assuming full liquidity and ignoring velocity of money effects. A more rigorous model would incorporate token velocity (V), giving $H \approx (L \cdot R) / (P \cdot V)$, where $V > 1$ reduces implied holdings. We present this as an order-of-magnitude illustration rather than a precise forecast.

If L = 0.01 (1% of biomedical IP-derived commercial revenue settled in METH) and R = $800B, then structural settlement demand implies approximately $8B in METH must be held for settlement purposes. If L grows to 0.10 (10% of the market), structural demand implies $80B.

These are not predictions — they are boundary conditions that illustrate why health is a high-value DVBC domain even at small penetration rates.

The critical observation is that these demand figures are *structural* (driven by real commercial transactions) rather than *speculative* (driven by price expectations). Structural demand does not evaporate when price momentum reverses; it persists as long as the commercial transactions continue.

8.4 Regulatory Treatment: Distinguishing Securities Risk from Antitrust Risk

A precise regulatory analysis of DVBC implementations must distinguish two categorically different risks that arise from different conditions of the model. Conflating them — as much commentary on tokenized funding mechanisms does — leads to incoherent recommendations.

Condition (b) is the locus of securities-law risk. The proportional-auction emission mechanism has the structure of an investment with expectation of value appreciation derived from the protocol's continued operation. This is the regulatory category to which the Howey test (US) and MiCA's "asset-referenced token" framework (EU) most directly apply. The IP-ownership architecture described in §6 substantially weakens, but does not eliminate, this exposure: because Metha holds no income-producing IP assets, the "common enterprise + profits from the efforts of others" prong is harder to establish than in DAO models where a treasury actively manages an IP portfolio. The residual exposure derives from the auction-and-emission dynamic itself, not from any fund-management role.

Condition (a) is the locus of antitrust / competition-law risk. The mandatory settlement requirement is a contractual restriction on the licensing terms of funded IP. The relevant legal categories are EU Article 102 TFEU (abuse of dominant position via tying arrangements) and US patent misuse doctrine. This is competition law, not securities law. Crucially, the §6 architecture — Metha as settlement layer rather than IP custodian — places this requirement in the same legal category as performing-rights-organization royalty terms (ASCAP, BMI, GEMA; cf. the U.S. consent-decree framework [41]) and copyleft license conditions, neither of which has historically been classified as either a tying arrangement or a securities instrument.

With this distinction in mind, three legal frameworks bear on the analysis:

EU competition law (condition a). Article 102 prohibits abuse of a dominant position. A Metha settlement requirement could in principle be characterized as a tying arrangement if Metha (as the licensing rights-holder) commanded a dominant share of biomedical IP. Two factors mitigate this risk: first, Metha does not hold the IP — researchers and their institutions do — so any tying claim would have to be levelled against the researcher accepting funding, not against the protocol; second, the settlement requirement is a condition of participation in a public-goods funding mechanism, more analogous to Creative Commons licensing terms than to a monopolist's tying.

US patent law (condition a). The doctrine of patent exhaustion limits conditions that patent holders can impose on downstream use of patented goods. Settlement-currency requirements for licensing (not for goods) are not directly addressed by exhaustion doctrine and are most analogous to contractual payment terms, which are generally enforceable. Patent misuse claims would similarly have to be levelled against rights-holding researchers, not against the protocol.

Relationship to Most-Favored-Nation (MFN) clauses in IP licensing. A reader familiar with standard IP licensing practice will recognise a structural family resemblance between the soft-variant settlement clause recommended by the Legal DD [27] and the contingent MFN provisions that appear routinely in pharmaceutical and software licensing — clauses under which the licensee receives the benefit of the most favourable terms granted to any other licensee, conditional on certain triggers. Two clarifications are worth making explicit. *First*, the soft-variant settlement clause is structurally closer to a *contingent MFN* than to an *exclusivity arrangement*: it does not block any commercial counterparty from licensing on any terms, it merely conditions the settlement medium (with bond-release acceleration on compliance and bond-slashing on non-compliance) under the framework of §6.1. *Second*, MFN clauses have a long history of antitrust analysis under both EU competition law and U.S. *per se* doctrine, and the prevailing line is that contingent MFN-style provisions are *not* presumptively anticompetitive where they are not deployed by a dominant firm to raise rivals' costs (Salop and Scott Morton on MFN doctrine [40]; the EU's *Booking.com* (Case C-264/23, ECJ 19 Sept 2024) [38] and *Amazon* (Cases AT.40462 / AT.40703) [39] decisions illustrate the relevant enforcement boundary). The relevance for the present paper is that the regulatory category in which the soft-variant settlement clause is best analysed is not "novel tying arrangement" but "contingent MFN of licensing terms" — a category with a developed jurisprudence and a non-trivial body of negotiating practice that DVBC implementations can draw on rather than invent from first principles.

Combined exhaustion + Brulotte + Art. 102 / Art. 101 analysis under the soft-variant clause. §8.4 has so far treated the antitrust surface around the settlement-currency clause through the lens of Art. 102 TFEU (abuse-of-dominance test that would apply if the protocol achieved essential-facility status) and, separately, through the vertical-MFN line of cases under Art. 101 TFEU — *Booking.com* (Case C-264/23, ECJ 19 Sept 2024) [38] is an Art. 101 vertical-restraints case, not an Art. 102 dominance case, and the *Amazon* (Cases AT.40462 / AT.40703) [39] commitments decision likewise turned on Art. 101 vertical analysis; the two strands are doctrinally distinct and v1.4.1 distinguishes them explicitly where earlier drafts ran them together. v1.3 extends that analysis by combining it explicitly with the U.S. patent-exhaustion findings of §6.1 / §6.2-bis and with the *Brulotte* royalty-term doctrine [32]. *Brulotte v. Thys Co.*, 379 U.S. 29 (1964) [32] bounds royalty obligations to the term of the patent: post-expiry royalty obligations on patent licences are *per se* unenforceable under U.S. doctrine, a rule that *Kimble v. Marvel Entertainment*, 576 U.S. 446 (2015) reaffirmed against *stare decisis* challenge. Exhaustion (*Quanta* [48], *Impression Products* [49]) constrains the per-article reach of licence terms beyond first authorised sale within the patent term. Read together, the two doctrines confine any "viral" settlement-currency obligation strictly to within-term, original-licensee scope: post-term royalty-equivalent obligations are unenforceable, and within-term per-article restrictions against bona-fide downstream purchasers are also unenforceable. The soft-variant clause of §6.2-bis is structured to be compatible with both bounds — pending the §7 jurisdiction-specific legal-opinion deliverables — by operating as a contractual condition between Metha and the rights-holder during the patent term, with bond-release acceleration tied to compliance, and by not purporting to encumber the patented article against downstream purchasers or to survive patent expiry. The TFEU Art. 102 surface in the European jurisdiction is, under the soft variant, materially smaller than it would be under an absolute-mandatory variant: an absolute clause that blocked alternative-currency licensing would invite a tying-arrangement characterisation under Art. 102(d), while the soft variant does not block alternative licensing — it conditions a bond-release benefit on the rights-holder's election of METH settlement, which is an MFN-style provision and not a tying arrangement under any of the Commission's

recent decisional practice. The combined analysis — exhaustion plus Brulotte for U.S. patent law, exhaustion-equivalent plus Art. 102 for EU competition law — is what makes the soft-variant clause defensible across both jurisdictions simultaneously; §6.4 above is the cross-jurisdictional companion to this paragraph.

8.4-bis Standard Essential Patents, FRAND, and *Huawei v ZTE*

A reviewer presses an antitrust surface that §8.4's exhaustion-plus-*Brulotte*-plus-Art. 102 framing only partially addresses: the Standard Essential Patent (SEP) and Fair, Reasonable and Non-Discriminatory (FRAND) doctrine that the Court of Justice of the European Union developed in *Huawei Technologies Co. Ltd v ZTE Corp.*, Case C-170/13 (ECJ 16 July 2015) [59]. The question is whether a DVBC architecture that achieved sufficient market reach for protocol-funded biomedical IP to become non-substitutable across a therapeutic surface — that is, an essential-facility position analogous to an SEP holder's — would face an Art. 102 TFEU challenge against the mandatory-settlement clause. v1.4 engages the question explicitly rather than leaving it to the Art. 102 framing of §8.4.

Huawei v ZTE established procedural rather than substantive constraints on SEP holders seeking injunctive relief: a SEP holder may sue for injunction only after (i) notifying the alleged infringer of the specific patent and the manner of infringement, (ii) presenting a written offer at FRAND terms, (iii) treating the alleged infringer as a "willing licensee" during good-faith negotiations, and (iv) accepting an escrow of disputed royalties in the interim. These are conditions on the procedural posture of an injunction, not substantive prohibitions on the licensing terms themselves; *Huawei v ZTE* does not bar a SEP holder from preferring one settlement medium over another, provided the offer is FRAND and the willing-licensee test is honoured. The doctrinal weight of the case for the present analysis is procedural, not exclusionary.

The §6.2-bis soft-variant settlement clause is structured to fall *outside the scope* of the SEP/FRAND framework rather than to satisfy or fail it, because it operates as a *contractual condition between Metha and the rights-holder* rather than as a *market exclusion against arms-length licensees*. It does not block alternative-currency licensing; it conditions a bond-release benefit on the rights-holder's election of METH settlement and leaves the rights-holder free to license to any commercial counterparty on any other terms (§6 (iii) on researcher autonomy). The SEP/FRAND framework governs the conditions under which a SEP holder may seek injunctive relief against an unwilling licensee on a patent the holder has voluntarily committed to license on FRAND terms; the soft-variant clause does not involve any such SEP commitment, does not seek injunctive relief, and does not control whether a downstream party may use the patented article at all. Even at the limit case in which protocol-funded IP became non-substitutable across a therapeutic surface — the conditions under which an essential-facility characterisation would be plausible — the soft variant's MFN-style structure aligns with the contingent-MFN analysis the European Commission applied in *Booking.com* [38] and in *Amazon* [39]: provisions that condition a benefit on a counterparty's election, rather than provisions that exclude counterparties from the market, are not within the *per se* enforcement boundary of Art. 102. Salop and Scott Morton's MFN doctrine [40] supplies the analytical bridge between the licensing-practice literature and the competition-law surface that the v1.3 paper already invokes.

The structural distance between the soft-variant clause and the *Huawei v ZTE* fact pattern is therefore material. *Huawei v ZTE* governs a unilateral injunction against an unwilling licensee on a specific patent the SEP holder controls; the soft-variant clause is a multilateral

contractual structure between Metha, the rights-holder lab and downstream commercial counterparties under which no party is enjoined from licensing on any terms. The on-chain bond (§6.1 Mechanism C) operates as an economic incentive, not as an injunction; the Metha Pharma backstop (§10.4) operates as a residual fair-terms licensee, not as a barrier to entry. The *erga omnes* enforceability that *Huawei v ZTE* presupposes is precisely the property §6.1 already retracts under *Quanta* [48] and *Impression Products* [49]: the protocol does not claim that the settlement-currency obligation runs against bona-fide downstream purchasers of patented articles, which removes the surface on which the SEP/FRAND framework would otherwise bite.

The honest residual is that doctrinal asymmetry between U.S. patent-misuse jurisprudence and EU SEP/FRAND jurisprudence (cf. §6.4) is a Phase-2 audit task, not closed by the present paper. v1.4 commits a dedicated SEP/FRAND legal-opinion deliverable as part of the Phase-2 legal-and-regulatory audit pipeline described in §7, alongside the seven-jurisdiction conflicts-of-laws audit already committed. Counterproposals on the SEP/FRAND surface — including arguments that the analysis above understates the essential-facility risk at scale — are explicitly invited via the [Legal] respondent track.

Crypto asset regulation (condition b). MiCA [14] in the EU and the evolving U.S. enforcement-and-judicial line exemplified by the *SEC v. Ripple Labs* [15] summary-judgment ruling govern whether METH constitutes a security or an asset-referenced/utility token. The emission-investment linkage creates a structure with security-like features under Howey, although the absence of a Metha-controlled income-producing asset (§6) and the role of METH as a settlement medium (rather than purely a speculative claim on a portfolio) push the classification toward "currency / payment instrument" or, under MiCA, toward "utility token" with currency-like properties. Regulatory clarity will require jurisdiction-specific engagement before Phase 3 implementation.

8.5 The seven-plane regulatory map and the recommended hybrid entity

The protocol combines several categorically distinct regulatory zones that must be kept analytically separate: (i) issuance and crowdsale (MiCA Title II, Howey, Reves, FATF); (ii) mandatory IP settlement (Art. 102 / 101 TFEU, U.S. patent-misuse doctrine, national patent laws); (iii) governance (DAO regimes, with joint-and-several liability of active participants in most jurisdictions); (iv) personal data (GDPR / LOPDGDD, immutability conflict with Art. 17 erasure right); (v) AML/CFT (AMLR, AMLD6, FATF R.15); (vi) operational resilience (DORA [37], applicable to CASPs since 17 January 2025 per Art. 64); and (vii) taxation (DAC8 [34], CARF [35], national CIT/PIT/VAT). The full seven-plane analysis with regulatory references, multi-criterion comparisons, cost tables and a risk matrix is published as the project's *Legal Due Diligence v1.0* [27].

The principal classification target under MiCA is the "other crypto-assets" category for METH (utility / payment token), with ART (Asset-Referenced Token) reclassification as the principal residual risk if condition (a) were activated before the contractual redesign recommended in the DD. The Spanish MiCA transitional regime for pre-existing CASP service providers expires on 30 June 2026 (the day before the new regime takes effect on 1 July 2026), and the DD concludes that the project cannot lawfully proceed to mainnet with services targeted at EU residents from a personal wallet without a regulated legal entity in place by that date. The recommended path is a two-jurisdiction hybrid: a Liechtenstein Stiftung (or Zug, Switzerland) as protocol issuer and non-profit anchor, and a Spanish operating SL as a MiCA-authorised

CASP with limited services (advice + placement without firm commitment), at a year-1 cost in the €520k–€1.35M range and recurring €280k–€620k.

These analytical shifts narrow but do not close the regulatory gap; the path described here is the project's working hypothesis, not a closed determination. The v1.2 publication of this paper is intended to invite review by securities and competition counsel, including specifically on (a) the soft-variant settlement clause and its compatibility with the *Brulotte* doctrine [32] and Art. 102 TFEU, (b) the Howey reclassification surface area under U.S. Treasury and SEC enforcement priorities current at the time of mainnet readiness, and (c) the practical resolution of GDPR Art. 17 erasure obligations against on-chain immutability.

GDPR Article 17 and on-chain immutability — the project's working mitigation. The conflict between GDPR's right to erasure and the practical immutability of public blockchains is well-known and is not solved by any project in the space. Metha's working mitigation, pending a definitive legal opinion in v1.3, is a two-tier data-architecture pattern that is now standard in the regulated-DLT literature: (i) *no personal data is stored on-chain in clear*. The on-chain layer holds only addresses, content hashes (IPFS / Arweave / signed-PDF), milestone state, and minimal metadata that is not identifying without external join. (ii) *Personal data lives off-chain* — researcher identity, KYB/KYC documentation, clinical-trial subject data, oracle-issued attestation evidence — under the Stiftung / SL operating entity's control, in a system that supports erasure, rectification and access under GDPR Arts. 15–17 in the ordinary way. Erasure requests are exercised against the operating entity as data controller, not against the chain. The on-chain content hashes remain valid as integrity anchors even after the off-chain artefact is erased; the consequence of erasure is that the off-chain artefact is no longer resolvable from the hash, which is the operationally compliant outcome under the dominant academic and regulator reading of the two-tier architecture (the chain "forgets" the personal data in the practically relevant sense — it is no longer reachable from the on-chain pointer — even though strict GDPR-Art.-17 compliance in the form of cryptographic erasure of every reference is debated in the literature; cf. Godyn et al. 2022 [67] on GDPR-blockchain compliance and Politou et al. 2019 [68] on the underlying blockchain-mutability surface). This pattern is consistent with the European Data Protection Board's draft Guidelines 02/2025 on processing personal data through blockchain technologies [36] (draft; interpretation may evolve in the post-adoption consultation window), with the CNIL's longstanding position, and with the architectural analogue documented for GDPR-compliant education credentials by Molina et al. 2023 [69]; the formal mapping of which Metha data fields fall in which tier, plus the data-controller responsibilities allocated between Stiftung and operating SL, is an output of the legal-and-data-governance audit that gates Phase 2 in §7.

Honest framing of the BARDA / ARPA-H alignment. Several sections of this paper (notably §5.4 on option-period exercise and §3.6 on funding-mechanism maturity) lean on structural alignment between Metha's primitives and the BARDA and ARPA-H option-period award patterns [30]. To avoid any ambiguity: the alignment is *unilateral*. Metha mirrors the BARDA and ARPA-H option-period structure because that structure has been validated by decades of off-chain practice as a way to gate funding on programmatic decisions, not because either agency has endorsed, partnered with, or formally evaluated Metha. There is no BARDA or ARPA-H endorsement of Metha, no formal partnership, and no implied claim of one; the Metha methodology brief [29] is an analysis of these agencies' published practice and is not produced under any agency mandate. Any future formal engagement with U.S. federal agencies would proceed through standard channels (FOA responses, Other Transaction Authority arrangements, or contracted assessment), would be disclosed explicitly when it

occurs, and is not a precondition for the protocol's technical or legal feasibility as described here.

---

9. Current Implementation and Limitations

9.1 Deployed Architecture

Metha Biofund is deployed on Sepolia testnet at the addresses listed in Table 3 (§4.2). The system comprises six smart contracts written in Solidity 0.8.24 with OpenZeppelin v5 dependencies, validated by 50+ Hardhat tests, including dedicated suites for multi-auditor consensus (§5.2), typed deliverables (§5.3), option-period exercise (§5.4) and the AuditorRegistryV2 lifecycle (§5.1, §5.5, §5.6), complemented by a clean Slither static-analysis run (May 2026).

The six contracts implement the three DVBC conditions as follows:

- **MethaToken** (ERC-20 with snapshot): the DVBC currency itself; snapshot capability enables governance voting without token lockup.
- **MethaCrowdsale**: implements condition (b) — emission proportional to ETH investment per period, with the 70/20/10 split (claimable / operations / developer) described in §4.1.
- **AuditingPool**: implements the quality gateway ensuring emission is linked to productive investment; in v3 enforces multi-auditor consensus (§5.2), typed deliverables (§5.3), option-period exercise (§5.4) and the milestone state machine.
- **OpenGovernance**: implements condition (c) — METH holder governance of all protocol parameters, including slot replacement, slashing severity determination, and the option-period Program Officer role.
- **EternalStorage**: typed key-value storage enabling logic upgrades without data migration, ensuring protocol longevity; hosts the slot table described in §4.3.
- **AuditorRegistryV2**: tiered registry of accreditation firms (§5.1) with seven on-chain category specialisations, oracle-issued attestations (GxP, ISO 17025, KYB), DSMB veto-rights grants, and the tiered slashing schedule of §5.6.

Sepolia deployment manifest (reproducibility). For reproducibility, the v3.x contract stack referenced throughout this paper is deployed on the Sepolia testnet (chainID 11155111) as of 2026-05-13. The manifest below lists the canonical contract addresses and Etherscan endpoints under which each contract's verified source can be inspected. All contracts in the table are source-verified; readers can reproduce the on-chain state from the verified source and the deployment transaction without trusting the project's own representations.

Contract	Address (verified on Sepolia Etherscan, click to view)
EternalStorage	`0xD81fFB850e44dd131D30E94eeD033CB9c8c546d9`
MethaToken	`0x0A1e366572f63bCa45A66dE4F5F1Ab5A62b3E16c`
OpenGovernance	`0x7BF8dd021477A44eD0F4b5eB48cB80cfb594917B`
MethaCrowdsale	`0x7B96b7e7670C6330b5F4f8b3F46245e4F49FF560`
AuditingPool	`0xD301Cff18Cec35F2c698a0fD7DFf81AF81098571`
AuditorRegistryV2	`0xDead08DfFa69FF1bde30A72E4BD993A67c61Ee03`

All six contracts deployed 2026-05-13, source-verified, on chainID 11155111 (Sepolia testnet).

The deployment script and the per-contract initialisation parameters are published in the project repository. Mainnet deployment is gated on completion of the formal-verification programme described in §9.2 and on the code-audit milestone identified in §7; the mainnet manifest will be published as a comparable table in v1.4 of this paper.

9.2 Security Analysis and Ownership

Static analysis with Slither 0.10 [16] identified no critical vulnerabilities. Key findings: ETH transfer calls in deposit refund pathways follow checks-effects-interactions (reentrancy safe); Solidity 0.8.24 provides built-in overflow protection; all privileged operations are gated by `onlyOwner` (OpenZeppelin Ownable). Current ownership is the deployer EOA `0x57829c5f9A5A2b09C22715133F1250d5DC1860c9`. The deploy script (`scripts/deploy_sepolia.js`, lines 261–303) implements two distinct hardening steps gated behind environment-variable switches — `TRANSFER_POOL_OWNERSHIP=true` migrates `AuditingPool` and `AuditorRegistryV2` ownership to the `OpenGovernance` contract (turning `exerciseOption`, `pause`, `setAuditorReward`, and slashing into token-weighted governance calls), and `REVOKE_DEPLOYER_ROLES=true` revokes the deployer's `MINTER`, `PAUSER`, `SNAPSHOT`, and `DEFAULT_ADMIN` roles on `MethaToken`. Both switches are intentionally *off* on Sepolia: the testnet posture is that the deployer remains the "Program Officer" for fast iteration (pause / unpause via direct EOA tx, mint METH ad-hoc, adjust parameters without a governance proposal cycle). At mainnet deployment, both switches are flipped on as part of the launch sequence; transferring `OpenGovernance`'s own owner address to a 3-of-5 Gnosis Safe multisig is one of the top quick-wins in the project's Legal DD (May 2026, §7.2) and is part of the same mainnet-launch checklist. The v1.4.1 framing replaces the v1.2/v1.3 "transfer is in flight" claim, which was correct for mainnet planning but wrongly suggested a pending Sepolia migration that the deploy script explicitly does not perform.

Slither static analysis is not a substitute for an independent audit. This point deserves to be stated unambiguously: the clean Slither run and the 50+ Hardhat test suite together constitute the *internal* security baseline that the project is willing to publish about. They do not constitute, and should not be taken as, a third-party security audit. Slither catches a well-defined class of static patterns (reentrancy, uninitialised storage, shadowed state, dangerous low-level calls, common ERC-20 footguns, control-flow anomalies); it does not exhaustively cover protocol-economic vulnerabilities, governance-attack vectors, oracle-trust assumptions, cross-contract invariants under adversarial sequencing, or rounding-and-precision edge cases under realistic adversarial input. An independent security audit by a firm with no prior involvement in the codebase is the *third* of the four sequential audits described in §7 (after legal due diligence, which is complete, and the BARDA-method audit, which is pending), and is a mainnet prerequisite without exception.

Formal verification: scope, properties, and Phase-2 timing. The v1.2 of this paper noted that formal verification using Certora Prover [17] was "under evaluation"; v1.3 commits to a specific scope and timing window. The protocol will engage Certora (or an equivalent specification-and-prover vendor — K Framework / Runtime Verification is the principal alternative under consideration) for verification of four named properties on the deployed stack, contingent on completion of the code audit and prior to mainnet deployment. The four properties are: (i) **no-double-claim of METH** — for each `(contributor, period)` pair, the

cumulative claimable amount satisfies the equation in §4.1 and a second claim against the same period reverts; (ii) **monotonic emission per period** — `tokensCreatedPerPeriod` and the 70/20/10 split are constant within a period, and the per-period totals sum to the protocol-constant emission rate without rounding leakage; (iii) **reachable-disbursement of treasury balances** (the "no-locked-treasury-funds" property of v1.3) — for every ETH balance held by AuditingPool or by an operations stash, there exists a reachable sequence of governance and admin transitions that disburses the balance, under a fixed set of role-assignment assumptions specified in the Certora harness. Note the specifiability caveat: this is a *reachable-disbursement* safety property rather than a full *liveness* property — Certora Prover excels at safety invariants and a strict liveness formulation would require bespoke witness construction; the scoped formulation captures the load-bearing fact (no admin sequence can render funds permanently unspendable) without overclaiming progress guarantees; (iv) **OpenGovernance bypass-resistance under atomic single-transaction flash-loan attack** — voting power for any proposal is determined exclusively by the snapshot taken at proposal submission (which reads `block.number - 1` checkpoints), and no atomic borrow-vote-repay sequence within a single transaction can change the outcome of a live proposal regardless of the borrowed METH amount. The single-transaction scoping is precise: multi-block flash-rent products that hold borrowed tokens across the snapshot boundary are *not* defeated by property (iv) — they are addressed at the governance layer through quorum and supermajority thresholds rather than at the contract layer. The flash-loan-governance property is called out explicitly because it is the highest-impact governance attack vector documented in the post-2020 DeFi-exploit record and the snapshot mechanism of §4.4 is the proximate defence; formal verification of this property is a mainnet prerequisite without exception. The engagement window is Phase 2, between the code audit and mainnet deployment, on the audit-gated, no-fixed-date schedule of §7. Specification artefacts and the prover output will be published alongside the v1.5 of this paper on completion.

Additional properties to scope during the Certora engagement. The four properties above are the load-bearing safety-and-economic-soundness invariants. A careful engagement will scope at least the following further invariants during the specification phase, even though they are not the headline four: *role/access-control monotonicity* (only `DEFAULT_ADMIN_ROLE` can grant `MINTER`, `PAUSER`, `SNAPSHOT`; the deployer's role-revocation path under the mainnet `REVOKE_DEPLOYER_ROLES` step is one-way); *pause/upgrade safety* (when `MethaCrowdsale` is paused, `contribute()` and `claimPeriod()` both revert, and the storage state remains consistent across pause/unpause cycles); *cross-chain replay resistance for EIP-712 permits* (METH's `ERC20Permit` `DOMAIN_SEPARATOR` includes `chainid`, so a permit signed for chain X cannot be replayed on chain Y); *EternalStorage slot-registry round-trip* (`writterSlots[slot] != 0 ⇒ writters[getWritterId(writterSlots[slot])] == writterSlots[slot]`); and *AuditorRegistryV2 stake monotonicity* (`stake[firm] ≥ _requiredStake(seniority)` at all reachable states except during the seven-day cooldown window). These are stated here as the public list of "additional invariants the engagement will scope," not as v1.4.1 commitments to verify them all; the final property surface for the engagement is a deliverable of the specification phase itself.

9.3 Operational characteristics and attack-vector inventory

Several reviewer comments press for a per-operation gas-cost summary and a consolidated economic-attack-vector inventory in a single section, rather than the present distribution of attack-vector analysis across §4.1.1.bis, §5, §6.3, §9.2 and §10.5. v1.4 adds this section as a

single defusing surface. The figures and the inventory are deliberately summary-level; the detailed analysis remains in the section where each vector is first introduced.

Measured gas costs for top operations. The figures below are *measured* averages from the v3 Hardhat fixture (Solidity 0.8.24, optimizer enabled at 200 runs, viaIR=true, EVM cancun target) captured by `hardhat-gas-reporter` on the 2026-05-20 snapshot committed to `metha-api/gas-report.txt`. v1.4 and earlier published *back-of-envelope* estimates that ran roughly a factor of two below the measured numbers below; v1.4.1 replaces those estimates with the measured averages. Mainnet costs will be within a constant factor depending on EIP-1559 base-fee dynamics [53] at the time of execution; the relative ordering and the order-of-magnitude framing are the load-bearing claims, not the literal absolute figures.

Operation	Description	Measured avg gas
`contributeCurrentPeriod()`	One-period contribution to MethaCrowdsale	201,667
`claimPeriod()`	Claim METH for a single completed period	270,047
`submitRequest()` / `submitRequestTyped()`	Submit research request with typed deliverables (scales with count, on-chain cap 32)	~300k – ~1.5M
`acceptVerdict()` / `voteHype()`	Auditor signs a milestone verdict (per-auditor hype vote)	90,637
`exerciseOption(reqId, N)`	Program Officer exercises milestone N option	~80,000
`voteYes(proposalId)`	OpenGovernance vote at snapshot weight	229,187
`submitProposal()`	OpenGovernance proposal creation (reads `block.number - 1` snapshot)	318,620
`AuditorRegistryV2.registerFirm()`	New audit firm registers with stake at given tier	343,376
`AuditorRegistryV2.slash()`	Tier-dependent stake reduction + event emission	170,320
Deploy (full stack, sum of all 6 contracts)	Hardhat artifact deployment cost; AuditingPool dominates at 5.4M (9.1% of block limit)	~18,000,000

Two observations on operational characteristics warrant explicit mention. *First*, the most frequent operations — `contributeCurrentPeriod()`, `claimPeriod()`, `voteYes()` — sit in the 200k–270k range, which is towards the upper end of typical ERC-20-adjacent contract interactions under post-Merge Ethereum economics but does not impose a structurally exclusionary cost surface for low-value participation; at a base-fee of 20 gwei and ETH at $3,000, a `contribute` transaction costs ~$12 of gas and a `claim` ~$16, which is order-of-magnitude tolerable for the contribution sizes the protocol targets. *Second*, the highest-cost operation in the table — `submitRequestTyped()` — scales with the number of typed deliverables anchored on submission, with the §5.3 on-chain cap of 32 bounding the worst case at roughly ~1.5M gas; the typical milestone with three to five deliverables remains in the

~300k range. *Third*, the `voteYes` cost (229k) is materially higher than a naive ERC20Votes vote because the implementation reads checkpoint state via `getPastVotes(snapshot)` at `block.number - 1`, which is the proximate defence against the §9.2 property (iv) flash-loan attack; the gas premium relative to a pure-current-balance vote is the cost of the defence. None of the operations above is in the gas-cost regime where it would dominate the lab's or the contributor's economic decision over participation; the binding economic costs at mainnet remain ETH-denominated value flows (contributions, milestone disbursements, bond posting), not gas-denominated transaction costs.

Economic-attack-vector inventory. The table below consolidates the attack vectors analysed throughout the paper into a single cross-reference surface. Each vector points to the section in which the vector is introduced and the mitigation analysed.

Vector	Section addressing it
Oracle manipulation (single-attestor fabrication / collusion)	§5.1 (RegulatoryOracle accreditation, VRF Sybil-rotation, KYB multi-attestor cross-attestation), §6.3 (Options 1+3 belt-and-braces, Option 2 reserved)
51% governance / hostile token accumulation	§10.5 (four-layer stack — snapshot voting, quorum / 2/3 supermajority, position-paper triggers, planned governance redesign), §9.2 property (iv) flash-loan bypass-resistance
Flash-loan governance attack	§9.2 property (iv) (Certora-committed bypass-resistance proof), §4.4 (snapshot-at-submission voting weight)
MEV / last-block sniping at period close	§4.1.1.bis (post-window discovery, 5760-block mainnet windows, VRF-randomised closing as v3.x candidate)
Multi-period crowding-out by deep-capital actor	§4.1.1.bis (acknowledged residual, partial mitigation via vesting-locked stash float cap of §10.5)
Reentrancy / classic Solidity footguns	Slither-clean per §9.2; checks-effects-interactions in deposit refund paths
Wash-trading and contribution-recycling	§4.1.1.bis v1.4 addendum (per-period proportional rule partially robust to recycling, not robust against coordinated wash by related addresses at period close)
Validator timestamp manipulation at period boundaries	§9.4 (known limitations — mitigated by proportional distribution rather than time-dependent payout)
Slashing-parameter governance gameability	§5.6 v1.4 addendum (2/3 supermajority threshold for slashing-parameter changes + retroactive slashing-window)
Oracle joint-bad-faith failure (rights-holder + licensee + Stiftung executive all compromised)	§6.3 ("what this does not solve") + §10.4 (Metha Pharma residual fix)
Verdict-ordering race (two auditors race to be	§5.2 (Condorcet aggregation tolerates

the third yes-vote, or to file the DSMB veto first)	ordering; the race is non-catastrophic but is an observable ordering dependency on the verdict-emission event)
Low-gas multisig recipient (2300-gas stipend on `wallet.transfer` in `MethaCrowdsale.transferEth` and `AuditingPool.claimRequestDeposit`)	§9.3 v1.4.1 hardening note — at mainnet the operations stash address should be a multisig with a fallback that fits in 2300 gas, or the transfers should be migrated to `call{value:amount}("")` with explicit CEI guards (state already zeroed before transfer in current code, so the migration is mechanical)
EIP-712 cross-chain replay on the (forthcoming) `attestSettlement` endpoint	§6.3 (domain separator includes `chainid`; replay across Sepolia ↔ mainnet is structurally prevented when the endpoint ships)

The Certora property commitment of §9.2 covers properties (i)–(iv) — no-double-claim, monotonic emission per period, no-locked-treasury-funds, OpenGovernance flash-loan bypass-resistance — which formally verify the technical layer underpinning the second, third and fifth vectors in the table above. Formal verification of the remaining vectors is not committed in the present scope: the oracle-manipulation, MEV, multi-period-crowding-out, wash-trading, and timestamp-manipulation vectors are structurally outside the Certora property surface as currently specified and are mitigated through the design and governance layers rather than through formal proof. Extension of the formal-verification scope to additional vectors is an open invitation to the [Code] track once the initial four-property engagement has completed.

9.4 Known limitations

Known limitations requiring attention before mainnet:

(i) the auditor bootstrapping cold start (the firm-based registry of §5.1 addresses this at the architectural level; bootstrap of an initial 1–10 accredited firm cohort remains a Phase-2 execution task);
(ii) oracle dependency for IPFS / Arweave proof validation and for the certification attestations consumed by AuditorRegistryV2 (community inspection plus accredited oracle entities are the current mechanisms);
(iii) validator timestamp manipulation at period boundaries (mitigated by proportional distribution);
(iv) supermajority threshold calibration for core slot changes and for slashing-severity determination;
(v) the four sequential audits — legal (complete at v1.0, May 2026), BARDA-method (pending), code (pending), financial (pending);
(vi) entity formation under the recommended B+A hybrid (Liechtenstein Stiftung + Spanish CASP-authorised SL);
(vii) multisig custody of all privileged addresses (the EOA-to-Safe migration above is the first of several);
(viii) professional liability insurance ≥ €1M;
(ix) bug bounty pool funded from the operations stash.

Mainnet has no fixed date. Commitment is gated on outputs (audit completion, entity standing, insurance and bounty pool, multisig coverage), not on time.

9.5 Testnet Status and Honest Assessment

All claims in this paper about the Metha implementation are grounded in the deployed testnet contracts and the open-source codebase. No real ETH or real research funding has transacted through the protocol. The 50+ test suite validates contract logic but cannot validate system behavior under adversarial conditions at scale.

The DVBC thesis is a theoretical claim about the value properties of a monetary system once all three conditions are satisfied. Metha currently satisfies conditions (b) and (c) and establishes the governance and audit infrastructure for condition (a). Condition (a) — mandatory METH settlement for biomedical IP — is a future governance decision contingent on the audit-gated, no-fixed-date Phase 2/3 sequence described in §7, not a current protocol property. Papers that claim more than this would be making unverifiable assertions; this paper claims exactly what the deployed code demonstrates and no more.

---

10. Discussion

10.1 Domain Generalizability: Other DVBC Candidates

The three-condition DVBC definition identifies a class of domains where the model is applicable. The conditions require: (a) a domain with identifiable commercial outputs that can be subject to settlement requirements; (b) a causal mechanism linking investment in domain productive capacity to outputs; and (c) a stakeholder community capable of governing the settlement and emission parameters. We examine three candidate domains.

Energy: A DVBC for Renewable Energy. A renewable energy DVBC would link emission to investment in renewable generation capacity (condition b) and require settlement in the currency for energy commodity transactions (condition a). The structural demand would be proportional to global energy trade volumes — potentially exceeding even oil settlement demand. The regulatory challenge is that energy markets are deeply state-regulated, making mandatory settlement requirements in a private currency legally complex without treaty-level coordination. However, the IEA's trajectory toward a net-zero energy system by 2050 creates an expanding renewable energy IP and services market that a DVBC could target at the licensing layer (analogous to Metha's approach) rather than at the commodity transaction layer.

Education: A DVBC for Academic Credentialing. An education DVBC would link emission to investment in curriculum development and course production (condition b) and require settlement in the currency for credential licensing and institutional accreditation fees (condition a). Existing examples like learning management system token experiments partially instantiate this, but without mandatory settlement for accredited credentials. The structural demand signal would derive from employer verification transactions — the multi-trillion-dollar credential verification market.

Software: A DVBC for Open Source IP. Gitcoin's quadratic funding mechanism [18] partially satisfies condition (b) — funding open source development — but has no mandatory settlement requirement for commercial use of funded software (condition a). A software

DVBC would require, as a condition of open source license terms, that commercial deployments of funded software settle a portion of revenue in the currency. This is analogous to copyleft licensing terms (GPL requires open source distribution of derivatives) but extends to commercial settlement rather than code licensing.

In each case, the DVBC model provides a design template that the existing partial implementations lack: condition (b) alone produces a funding mechanism without a demand anchor; condition (a) alone produces a petrodollar analog without emission coherence; condition (c) alone produces a governance token without domain backing. All three conditions are necessary for the self-reinforcing value loop.

10.2 The Cold Start Problem

The most acute challenge for any DVBC implementation is the cold start problem: structural demand from domain settlements (condition a) cannot exist until the domain has produced commercial outputs, but the domain cannot produce commercial outputs until research is funded, and research funding requires the token to have sufficient value to attract contributors.

This is a chicken-and-egg problem present in every monetary system at inception. Historical monetary systems resolved it through state mandate (fiat currencies achieved initial adoption through tax obligation), military coercion (commodity monies achieved international adoption through trade dominance backed by military power), or network effects (Bitcoin achieved early adoption through ideological commitment from a small technical community willing to accept near-zero early price).

For Metha, the cold start resolution has three components. First, the auction mechanism creates METH value independent of settlement demand: contributors acquire METH because they believe in the future settlement demand mechanism, not because it currently exists. This is analogous to early Bitcoin adoption — value from expected future utility rather than current utility. Second, the testnet deployment demonstrates technical feasibility, reducing the uncertainty discount that early contributors apply to the future-value expectation. Third, the institutional partnership pipeline (planned for the post-mainnet phase) creates early demand signals: if a research institution publicly commits to licensing Metha-funded IP in METH, the option value of that future settlement demand is immediately capitalized into current METH price.

The auditor cold start — distinct from the demand-side cold start above — is addressed through the firm-based audit layer described in §5.1. Bootstrapping a small set of accreditation firms (a 1–10 cardinality problem) is significantly more tractable than bootstrapping a heterogeneous global pool of individual experts.

The cold start problem is solvable but not trivial. It requires early contributors to have long investment horizons and to accept speculative-stage risk. This is consistent with the risk profile of research investment generally — the research enterprise as a whole operates on the premise that early-stage investment in uncertain outcomes produces long-run value.

Historical monetary systems offer instructive precedents. The dollar's post-Bretton-Woods transition was anchored by the 1974 OPEC agreements — an institutional commitment that capitalised future settlement demand into present dollar holdings. For Metha, an analogous mechanism would be a signed commitment from a research institution to license Metha-funded IP in METH — converting option value into observable demand signal before the

ecosystem reaches commercial scale. Such a commitment is most plausibly secured during the post-mainnet phase, when the four-audit gate of §7 has been cleared and the soft-variant settlement clause has been validated by at least two independent legal opinions in the major patent jurisdictions.

The empirical track record of the perpetual proportional auction described in §4.1 and §4.1.1 is also directly relevant to cold-start solvability: the EOS launch demonstrated that this exact auction structure is capable of attracting billions of dollars in capital from a near-zero starting point given sufficient market belief in the underlying thesis. The mechanism's known manipulation surface — multi-period crowding-out and other vectors documented in §4.1.1.bis under modern MEV/LVR scrutiny — is not zero, but it does not appear to have been the binding constraint at the EOS scale; the binding constraint was the credibility of the value proposition. The DVBC anchoring — research outputs as observable productive backing, not pure narrative — is precisely the credibility supplement that EOS lacked, and the auction-plus-anchoring combination of §4.1.1 is the compounding effect on which the cold-start strategy ultimately rests. The auction is necessary but not sufficient; the credibility floor is what bootstraps the loop.

The project's deliberate posture during this bootstrap window is that of a *proof of concept seeking peer collaborators* rather than a product seeking buyers: the Open Invitations section of https://metha.life and the [Economics]/[Legal]/[Regulatory]/[Research]/[Compliance]/[Code] subject prefixes on the project's contact channel are the operational expression of that stance. The cold-start strategy is therefore not "raise capital quickly from anonymous contributors" but "compound credibility through audit completion, legal-opinion accumulation, institutional letters of intent and peer-reviewed publication, then activate the auction at mainnet only once the credibility floor exists." The auction will work when the floor does.

[OPEN QUESTION v1.2: the cold-start velocity in practice is undetermined. The Sepolia deployment is a *technical* demonstrator and has not been instrumented to generate cold-start metrics that would translate to mainnet — there is no production ETH at stake, no live contributor population, no on-chain commitment from research institutions. The empirical answer to "how fast does the loop close in practice?" depends on the credibility-accumulation rate during the audit-gated Phase-2 window described in §7 and on the institutional partnerships that the post-mainnet phase will pursue; both are under active investigation in the Roadmap v2 working document [28] and will be reported quantitatively in subsequent paper revisions.]

10.3 Relationship to Existing DeSci Protocols

§3.7 introduced the asset-layer-vs-settlement-layer distinction that separates DVBC from the IP-NFT-based DeSci protocols (VitaDAO, Molecule, ResearchHub, LabDAO, BIO Protocol). This section discusses the implications for ecosystem positioning rather than restating that distinction. The four projects sit at the IP-ownership, discovery, or services-marketplace layers; Metha sits at the funding-and-settlement-infrastructure layer. The projects are therefore composable rather than substitutable with Metha — a future stack in which Metha-funded outputs are tokenised through Molecule's IP-NFT framework, peer-reviewed via ResearchHub, executed against LabDAO services, and governed at the asset layer by a VitaDAO-style structure would represent the full DeSci canon with Metha providing the DVBC monetary foundation underneath.

The v3 audit infrastructure described in §5 closes three of the five gaps that the Auditing-Funding methodology brief [29] identifies for DeSci protocols as a class: depth of scientific peer review (addressed by the tiered registry of accreditation firms in §5.1), physical-verification proxy (addressed by oracle attestations of GxP/ISO 17025/KYB certificates), and on-chain DSMB analog (addressed by the veto lane in §5.5). Two gaps remain open and are not closed by code alone: jurisdictional sponsorship for clinical trials (which presupposes a regulated entity standing in a recognised jurisdiction), and tax / fiscal reporting for ETH-denominated milestone disbursements (which presupposes a CASP-authorised operating entity and DAC8/CARF-aware reporting infrastructure).

The audit infrastructure of §5, combined with the planned Metha Pharma backstop of §10.4, are the two foreground differentiators that block the "just another DeSci token" framing. The audit infrastructure makes emission domain-*value*-backed rather than merely domain-*activity*-backed (anyone can contribute ETH; not all proposals receive funding), and Metha Pharma resolves the predictable failure mode in which incumbent licensees decline fair-terms offers on protocol-funded IP. Neither feature is present in the comparator projects above, and both are load-bearing for the DVBC thesis as distinct from a generic research-funding DAO.

10.4 Metha Pharma: vertical integration as antifragile backstop

The DVBC value loop described in §3.2 closes if and only if research outputs find commercial counterparties willing to license on fair terms. A predictable failure mode — and the question any pharmaceutical reader will raise within the first five minutes of contact — is incumbent refusal: established pharmaceutical companies decline to license protocol-funded IP at fair royalty terms, or impose conditions that suppress competition. To close this gap antifragilely, the project roadmap introduces *Metha Pharma*, a ring-fenced operating entity capable of taking a licence on protocol-funded IP and executing commercialisation directly when the open market does not produce a fair-terms bidder.

Importantly, Metha Pharma is not exclusive of the open market: the default route remains arms-length licensing to the highest fair-terms bidder, with royalties flowing to the protocol; Metha Pharma activates as a backstop when a defined no-bid window expires without a fair-terms offer, or when a fair-terms refusal is documented. Either route — open market or Metha Pharma — returns royalties to the protocol, closing the DVBC value loop on a path that is robust to incumbent refusal.

Open governance questions identified in the roadmap include (i) the legal form (dedicated entity, JV with an existing CRO, alliance with a boutique pharma, or per-project spin-out); (ii) initial capitalisation (a Metha-Pharma-ops sub-stash, a per-project token-holder activation vote, or external SPV capital); (iii) activation triggers (no-bid window, documented fair-terms refusal, holder vote, scientific-council recommendation); (iv) governance scope and ring-fencing of product-liability exposure away from the core protocol; and (v) whether Metha Pharma can itself be an awardee under BARDA / ARPA-H programmes. Treating Metha Pharma as a backstop rather than as an exclusive commercialiser preserves the open-market default and converts a single point of strategic failure into a two-route system.

10.5 Open governance design questions

Three further structural questions remain open and are documented in the project's Roadmap v2 [28] as gating triggers for any pre-mainnet design freeze. First, the voting model — token-weighted voting today defaults to plutocracy, and the roadmap commits to evaluating quadratic, identity-based one-person-one-vote (for certain proposal classes), delegated-

domain-council and reputation-weighted hybrids against criteria of Sybil resistance, capture risk and regulatory readability before condition (a) is activated. Second, the stash architecture — the current monolithic operations stash is scheduled for decomposition into specialised multisigs (legal-ops, marketing-ops, scientific-ops, infra-ops, an emergency-multisig and, conditionally on Metha Pharma activation, a Metha-Pharma-ops sub-stash). Third, the auditor compensation matrix — a percentage of milestone ETH plus a vested METH grant from the operations stash plus a slashing-window hold-back, with the percentage and vesting schedule to be modelled adversarially before deployment. These are not deferred implementation details; they are open governance design questions that the project treats as position-paper-first, code-second.

Hostile takeover via METH accumulation — the protective layers. Token-weighted governance is, in the limit, vulnerable to acquisition: an actor with sufficient capital could in principle accumulate enough METH to dominate any proposal. The protocol's working defences against that scenario operate at four layers, none of which is decisive alone but which together raise the cost of capture substantially. *Layer 1: snapshot voting (§4.4)* prevents the most obvious vector — buying votes after the proposal is tabled — by freezing voting weight at proposal submission. *Layer 2: quorum thresholds and supermajority requirements* on the highest-consequence proposals (slot replacements, slashing severity, condition-(a) activation parameters) raise the cumulative threshold beyond a simple plurality. *Layer 3: the position-paper triggers* described in this section operate as a procedural constraint: structural changes are gated on a published position paper plus a discussion window, which gives the wider holder base time to coordinate a defensive response and to attract counter-capital before the vote. *Layer 4: the planned governance redesign* — the open voting-model question in this section — is explicitly tasked with reducing plutocracy-vulnerability for certain classes of high-consequence proposals (e.g., shifting condition-(a) activation to a delegated-domain-council or identity-based one-person-one-vote regime for that proposal class while retaining token-weighted voting for routine parameter tweaks). The Stiftung, finally, holds a residual veto on actions that would breach its statutes or its regulated-entity obligations — capture of the on-chain governance does not automatically capture the Stiftung's legal counterparty role. The combination is not a guarantee, but it converts hostile takeover from a single-step capital exercise into a multi-vector campaign against several procedural and legal layers, each of which is independently observable and contestable. The cross-references in §8.4 and §9 should be read with this protective-layer stack in mind.

Pre-snapshot accumulation and the residual exposure. Snapshot voting (§4.4) defends against intra-proposal manipulation — once a proposal is submitted, voting weights are fixed and a flash-loan or other in-period borrow-vote-repay sequence cannot change the outcome, a property §9.2 commits to formal verification under property (iv). Snapshot voting does *not* defend against pre-proposal accumulation of voting weight in advance of a planned hostile proposal. v1.3 is explicit about both what the protocol does against this distinct vector and what it leaves residual. Three protocol-level mitigations apply. **First**, the dev-stash (10% allocation) and ops-stash (20% allocation) holdings are vesting-locked under the schedule described in §4.3, with on-chain enforcement of the vesting curve via EternalStorage slot constraints. These holdings cannot be freely sold into the market, which structurally caps the float available for hostile accumulation: the publicly tradable float is materially smaller than the total emitted supply, which raises the marginal cost of accumulating any given voting-weight target. **Second**, the protocol exposes top-N holder analytics through the metha-admin transparency surface, which makes large concentrations into a single address publicly observable. Coordinated accumulation through a small number of related addresses is detectable through standard on-chain analytics (clustering on funding source, gas-fee payer,

contract-interaction signature). The protocol does not on its own do the clustering work, but it provides the data primitives that allow third-party watchdogs (and the Stiftung's compliance team under its CASP obligations) to do it. **Third**, OpenGovernance's quorum thresholds and the 2/3-supermajority requirement for governance-critical proposals raise the cost of accumulation materially: a hostile actor must acquire 67% of the active voting weight, not 50%, and that 67% threshold compounds against the structural float cap of the first mitigation. The residual exposure is honestly characterised: with sufficient capital, sustained sub-quorum accumulation over a long horizon — accumulating positions slowly enough to avoid triggering observable concentration alerts, across enough addresses to avoid clustering signatures — is possible. The flash-loan vector is closed by snapshot; the patient-whale vector is mitigated but not eliminated by structural float cap plus quorum threshold plus observability. The Layer 4 governance redesign noted above is the long-term mitigation specifically targeted at the patient-whale residual; Layer 4 is *supplementary, not sufficient*, in the sense that Layers 1–3 remain necessary under any governance regime and Layer 4 only reduces residual exposure on the highest-consequence proposal classes. A connected question is whether the protocol facilitates *counter-capital coordination* — the symmetric mechanism by which honest holders coordinate to outvote a hostile accumulation. The protocol does *not* directly facilitate this. Coordination is left to the market, to the Stiftung's communications surface as the legal counterparty under its public-benefit charter, and to whatever DAO-tooling the broader Ethereum ecosystem provides (Snapshot.org, Tally, etc.). This is a deliberate choice: protocol-level facilitation of counter-coordination would itself be a centralisation surface, and the architecture treats coordination as exogenous. The Stiftung's residual veto (§4.4) is the legal-layer backstop in the limit case where coordination fails and a hostile proposal would breach the foundation's charter or its regulated-entity obligations.

Disposition of the 10% developer stash if the team disbands. A reasonable reviewer will press the question: what happens to the 10% per-period developer allocation if the founding development team disbands or otherwise ceases to operate the protocol? The architecture handles this on two axes. First, the developer leg is paid out via vesting against the registered developer wallets defined in `MethaCrowdsale`, so a team that disbands ceases to *accrue* unvested portions; the unvested remainder reverts to the protocol rather than crystallising as a windfall to a departing developer. Second, the developer-wallet registry is itself a parameter subject to OpenGovernance: holders can — and, in a no-active-development scenario, would — vote to redirect the 10% leg into another stash (typically the operations stash, or a successor developer cohort onboarded via proposal) so that the per-period emission continues to fund forward operations rather than ossifying as orphaned issuance. Combined with the planned decomposition of the operations stash into specialised multisigs, this means that the 10% developer leg is best understood as an *allocation parameter*, not a vested entitlement to any specific team; the protocol does not require any particular developer cohort to remain active, only that the per-period 10% find a destination that holders endorse.

Note on the upgrade-affordance heterogeneity. A point related to this section, and clarified during the v1.2-rev2 audit of the deployed slot registry (§4.3), is that not every replaceable component in the v3 stack is replaced through the same proposal type. The four `EternalStorage`-registered contracts (`Token`, `Token Management`, `Auditing Pool`, `Governance`) are replaced through a slot-rebind proposal payload; `AuditorRegistryV2` is replaced through a direct `AuditingPool.setAuditorRegistry()` call from the owner — OpenGovernance once ownership has been transferred. The position paper on stash architecture, when it is written, will need to clarify which proposal type applies to each of: the developer-wallet slate (a `MethaCrowdsale` setter rather than an `EternalStorage` rebind),

the dev-stash recipient redirection (likely also a direct setter on `MethaCrowdsale`), and the operations-stash decomposition into specialised multisigs (potentially a combination of slot rebinds for any contract-typed addresses and direct setters for wallet-typed addresses). The architectural axis is consistent — "is this a contract-typed parameter or a wallet-typed parameter, and is the holder an `EternalStorage`-mediated parameter store or a direct state variable on the owning contract?" — but the proposal payload differs case by case.

[OPEN QUESTION v1.2: the precise on-chain mechanics of unvested-leg reversion are partly resolved (vesting halts on a non-active wallet; the unvested remainder remains in `MethaCrowdsale` and is not redirected automatically) but the *triggering* of the redirection vote and the exact OpenGovernance proposal type required for redirecting the developer-stash recipient slate are calibration items that will be specified in the §10.5 position paper on stash architecture prior to mainnet, and the audit-leg corresponding to financial-audit-of-treasury (§7) is the formal validation point. The §4.3 note above narrows the design space: the redirection is a direct setter call on `MethaCrowdsale`, not an `EternalStorage` slot rebind.]

---

11. Conclusion

This paper has introduced Domain-Value-Backed Currency (DVBC) as a new monetary model that resolves the fundamental incoherence of existing systems: the disconnection between currency emission and the productive domain that generates demand for the currency.

The theoretical contribution is the three-condition DVBC definition (mandatory domain settlement, emission-investment linkage, holder governance) and the formalization of the value loop it produces. We have shown formally that the petrodollar satisfies condition (a) but not (b), making it incoherent with respect to emission; that fiat currencies satisfy neither (a) nor (b); that Bitcoin satisfies neither condition in any productive domain sense; and that DVBC, by satisfying all three conditions, achieves emission-demand coherence that no prior monetary model has achieved by design.

The empirical contribution is the demonstration that condition (b) and condition (c) are technically implementable in a deployed Ethereum protocol. Metha Biofund's six smart contracts — MethaToken, MethaCrowdsale, AuditingPool, OpenGovernance, EternalStorage and AuditorRegistryV2 — realize these conditions in 50+ test-validated Solidity code, currently live on Sepolia testnet, with the audit-layer primitives (multi-auditor consensus, typed deliverables, option-period exercise, DSMB veto, tiered slashing) required to enforce condition (a) also live. Condition (a) is a future governance decision contingent on the four-audit, audit-gated, no-fixed-date Phase 2/3 sequence described in §7, not an engineering problem — the governance and audit infrastructure to implement it exists.

The practical contribution is the identification of the biomedical research domain as the ideal first DVBC implementation: a ~$280 billion annual investment domain [10] producing IP with clear commercial value, no depletion risk, global geographic distribution, and existing IP licensing frameworks that provide legal precedent for settlement requirements.

Future work proceeds along three axes. First, empirical validation at mainnet scale — measuring actual auction participation, auditor quality signals, milestone completion rates, and geographic distribution of funded research against the theoretical predictions of the DVBC model. Second, legal structuring of the IP settlement mechanism (condition a) —

engaging with EU MiCA, US securities law, and IP licensing frameworks to establish a legally defensible mandatory settlement pathway under the soft-variant settlement clause recommended by the Legal DD [27]. Third, generalization studies for the energy, education, and software domains, establishing whether the DVBC conditions can be satisfied in these domains and what institutional infrastructure would be required.

The claim at the center of this paper is modest but specific: there exists a class of monetary systems — DVBCs — in which currency emission and domain settlement demand are coherent by design, and this coherence is achievable in a deployed blockchain protocol. Whether Metha succeeds in building the first fully realized DVBC will be determined by empirical evidence, not by theoretical argument. This paper establishes the theoretical framework and proof-of-concept implementation from which that empirical evaluation can begin.

Throughout the three roadmap phases, the project's posture is explicitly that of a proof of concept seeking peer collaborators — economists, securities and competition lawyers, biomedical researchers, accreditation firms and protocol engineers — rather than a product seeking buyers or users. The v1.2 publication of this paper, the v1.0 Legal DD [27], the v1.0 Auditing-Funding methodology brief [29] and the public release of the AuditorRegistryV2 contracts in May 2026 are intended as artefacts to enable that peer engagement, not as commercial milestones. The project actively solicits feedback at info@metha.life with subject prefixes [Economics], [Legal], [Regulatory], [Research], [Compliance] or [Code], as documented in the Open Invitations section of https://metha.life.

References

[1] J. M. Keynes, *The General Theory of Employment, Interest and Money*. London: Macmillan, 1936.

[2] M. Friedman, "The Role of Monetary Policy," *American Economic Review*, vol. 58, no. 1, pp. 1–17, 1968.

[3] F. A. Hayek, *Denationalisation of Money: The Argument Refined*. London: Institute of Economic Affairs, 1976.

[4] B. Eichengreen, *Exorbitant Privilege: The Rise and Fall of the Dollar and the Future of the International Monetary System*. Oxford: Oxford University Press, 2011.

[5] D. Spiro, *The Hidden Hand of American Hegemony: Petrodollar Recycling and International Markets*. Ithaca: Cornell University Press, 1999.

[6] W. S. Jevons, *Money and the Mechanism of Exchange*. London: Henry S. King & Co., 1875.

[7] S. Nakamoto, "Bitcoin: A Peer-to-Peer Electronic Cash System," 2008. [Online]. Available: https://bitcoin.org/bitcoin.pdf

[8] V. Buterin, "Ethereum: A next-generation smart contract and decentralized application platform," Ethereum Foundation, 2014. [Online]. Available: https://ethereum.org/en/whitepaper/

[9] World Health Organization, “Global Spending on Health: Rising to the Pandemic’s Challenges,” WHO, Geneva, 2022.

[10] IQVIA Institute, “Global Trends in R&D 2024: Activity, Productivity, and Enablers,” IQVIA Institute for Human Data Science, 2024.

[11] block.one, “EOS.IO Technical White Paper,” 2017. [Online]. Available: https://github.com/EOSIO/Documentation/blob/master/TechnicalWhitePaper.md. Archived: https://web.archive.org/web/2022/https://github.com/EOSIO/Documentation/blob/master/TechnicalWhitePaper.md. [Block.one wound down EOSIO development in 2022; the archived snapshot is the authoritative reference.]

[12] R. A. Mundell, “A Theory of Optimum Currency Areas,” *American Economic Review*, vol. 51, no. 4, pp. 657–665, 1961.

[13] WePower Team, “WePower: Green Energy Trading Platform Whitepaper” (WPR token), WePower, 2018. [Online]. Available: https://web.archive.org/web/2018/https://wepower.network/wp-content/uploads/2018/01/WePower_Whitepaper_2018.pdf. [Accessed via web archive due to project entity wind-down.]

[14] European Parliament, “Regulation (EU) 2023/1114 on Markets in Crypto-assets (MiCA),” *Official Journal of the European Union*, 2023.

[15] SEC v. Ripple Labs, Inc., No. 20-cv-10832 (S.D.N.Y. 2023).

[16] J. Feist, G. Grieco, and A. Groce, “Slither: A Static Analysis Framework for Smart Contracts,” in *Proc. IEEE/ACM 2nd Int. Workshop on Emerging Trends in Software Engineering for Blockchain (WETSEB)*, 2019, pp. 8–15.

[17] Certora, “Certora Prover — User Documentation and Specification Language Reference.” [Online]. Available: https://docs.certora.com/. See also S. Grossman, I. Abraham, G. Golan-Gueta, Y. Michalevsky, N. Rinetzky, M. Sagiv, and Y. Zohar, “Online Detection of Effectively Callback Free Objects with Applications to Smart Contracts,” *Proc. ACM Program. Lang.*, vol. 2, no. POPL, art. 48, pp. 1–28, Jan. 2018, doi:10.1145/3158136 — the academic foundation for the Certora verification approach.

[18] V. Buterin, Z. Hitzig, and E. G. Weyl, “A Flexible Design for Funding Public Goods,” *Management Science*, vol. 65, no. 11, pp. 5171–5187, 2019.

[19] J. M. Keynes, “Proposals for an International Clearing Union,” British Government White Paper, 1942. Reprinted in *The Collected Writings of John Maynard Keynes*, vol. XXV, D. Moggridge, Ed. London: Macmillan, 1980.

[20] C. Catalini and J. S. Gans, “Some Simple Economics of the Blockchain,” *Communications of the ACM*, vol. 63, no. 7, pp. 80–90, 2020. Earlier version: NBER Working Paper 22952, 2016.

[21] S. Voshmgir, *Token Economy: How the Web3 Reinvents the Internet*, 2nd ed. Berlin: Token Kitchen, 2020.

[22] L. Howell, M. Niessner, and D. Yermack, “Initial Coin Offerings: Financing Growth with Cryptocurrency Token Sales,” *Review of Financial Studies*, vol. 33, no. 9, pp. 3925–3974, 2020.

[23] U.S. Securities and Exchange Commission, “Order Instituting Cease-and-Desist Proceedings, In the Matter of Block.one,” Securities Act Release No. 10714, File No. 3-19568, Sep. 30, 2019.

[24] *Jacobsen v. Katzer*, 535 F.3d 1373 (Fed. Cir. 2008).

[25] Free Software Foundation, “GNU General Public License, Version 3 (GPLv3),” 29 Jun. 2007. [Online]. Available: https://www.gnu.org/licenses/gpl-3.0.html

[26] United Nations, *Convention on the Recognition and Enforcement of Foreign Arbitral Awards (New York Convention)*, 330 U.N.T.S. 38, 10 Jun. 1958.

[27] Metha Biofund, “Legal Due Diligence and Compliance Roadmap, v1.0 (English),” Internal working document, 12 May 2026. [Online]. Available: https://metha.life/documents/legal/Metha_Legal_DueDiligence_2026_EN.pdf. [Accessed: 14-May-2026].

[28] Metha Biofund, “Roadmap & Checklist v2 — Mapa mental de ideas,” Internal working document, 13 May 2026.

[29] J. V. Martínez (ed.), “Auditing and Funding Biomedical Research: How public agencies and private investors decide what to fund, what milestones unlock the money, and what audits repeat at every stage, v1.0,” Metha Biofund, 12 May 2026. [Online]. Available: https://metha.life/documents/Auditing_Funding_Biomedical_Research.docx. [Accessed: 14-May-2026].

[30] Biomedical Advanced Research and Development Authority (BARDA), “BARDA Broad Agency Announcement and Other Transaction Authority Solicitations,” U.S. Department of Health and Human Services, ASPR. [Online]. Available: https://www.medicalcountermeasures.gov. [Accessed: 14-May-2026].

[31] National Institute of Dental and Craniofacial Research, “Data and Safety Monitoring Board Guidelines,” National Institutes of Health, Bethesda, MD. [Online]. Available: https://www.nidcr.nih.gov/research/conducting-nidcr-clinical-research/data-and-safety-monitoring-board-guidelines. [Accessed: 14-May-2026]. See also U.S. Food and Drug Administration, “Establishment and Operation of Clinical Trial Data Monitoring Committees — Guidance for Clinical Trial Sponsors,” March 2006.

[32] *Brulotte v. Thys Co.*, 379 U.S. 29 (1964).

[33] *Genentech Inc. v. Hoechst GmbH (Sanofi-Aventis Deutschland GmbH)*, Court of Justice of the European Union, Case C-567/14, 7 July 2016.

[34] Council of the European Union, “Council Directive (EU) 2023/2226 of 17 October 2023 amending Directive 2011/16/EU on administrative cooperation in the field of taxation (DAC8),” *Official Journal of the European Union*, L series, 24 Oct. 2023.

[35] Organisation for Economic Co-operation and Development, “Crypto-Asset Reporting Framework (CARF) and Amendments to the Common Reporting Standard,” OECD, Paris, 10 Oct. 2022. [Online]. Available: https://www.oecd.org/tax/exchange-of-tax-information/crypto-asset-reporting-framework-and-amendments-to-the-common-reporting-standard.htm

[36] European Data Protection Board, "Guidelines 02/2025 on the processing of personal data through blockchain technologies," adopted as draft at the EDPB plenary of April 2025 and subjected to public consultation through 9 June 2025; final version pending at the date of this paper. EDPB, Brussels, 2025.

[37] European Parliament and Council, "Regulation (EU) 2022/2554 of 14 December 2022 on digital operational resilience for the financial sector (DORA)," *Official Journal of the European Union*, L 333, 27 Dec. 2022, p. 1; applicable to CASPs from 17 January 2025 per Art. 64.

[38] *Booking.com BV v 25hours Hotel Company Berlin GmbH*, Court of Justice of the European Union, Case C-264/23, 19 September 2024 — preliminary ruling on the antitrust assessment of hotel-platform parity clauses, the live EU-level authority on the topic at the date of this paper version. (The earlier 2015 commitments regarding Booking.com parity clauses were given to national competition authorities — France, Italy and Sweden — and later voluntarily extended EU-wide, but were not the subject of a European Commission Reg. 1/2003 commitments decision.)

[39] European Commission, "Commitments decision in Case AT.40462 — Amazon Marketplace, and Case AT.40703 — Amazon Buy Box," C(2022) 9442 final, Brussels, 20 Dec. 2022.

[40] S. C. Salop and F. M. Scott Morton, "Developing an Administrable MFN Enforcement Policy," *Antitrust Magazine*, vol. 27, no. 2, Spring 2013, pp. 15ff.; see also J. B. Baker and J. A. Chevalier, "The Competitive Consequences of Most-Favored-Nation Provisions," *Antitrust*, vol. 27, no. 2, 2013.

[41] U.S. Department of Justice, *United States v. ASCAP*, Civ. No. 41-1395 (S.D.N.Y., entered 14 Mar. 1941, amended 2nd Amended Final Judgment 11 Jun. 2001) and *United States v. BMI*, Civ. No. 64-3787 (S.D.N.Y., entered 14 Dec. 1966, amended 1994) — performing-rights-organization consent decrees, the canonical antitrust framework for collective licensing in U.S. music rights.

[42] Federal Reserve Bank of St. Louis, "M2 Money Stock (M2SL)," FRED Economic Data series, Board of Governors of the Federal Reserve System (US). [Online]. Available: https://fred.stlouisfed.org/series/M2SL. [Accessed: 14-May-2026].

[43] International Energy Agency, *World Energy Outlook 2021*, IEA, Paris, Oct. 2021. [Online]. Available: https://www.iea.org/reports/world-energy-outlook-2021. [Documents the 2020 oil demand contraction.]

[44] Proclamation No. 4074, "Imposition of Supplemental Duty for Balance of Payments Purposes," 36 Fed. Reg. 15724 (Aug. 17, 1971); Address to the Nation Outlining a New Economic Policy ("Nixon Shock"), Richard M. Nixon, 15 Aug. 1971.

[45] International Monetary Fund, *Articles of Agreement of the International Monetary Fund*, original text adopted 22 Jul. 1944, Bretton Woods Conference, Art. IV §1 (par values denominated in gold or in the U.S. dollar of the fineness in effect on 1 July 1944, original text prior to the Second Amendment of 1978). The literal "$35 / troy ounce" figure derives from the U.S. *Gold Reserve Act of 1934* (Pub. L. 73-87, 31 Stat. 337), which fixed the dollar's gold weight; that statutory weight was carried into the IMF par-value regime via Art. IV §1.

[46] *Consolidated Appropriations Act, 2022*, Pub. L. No. 117-103, Div. H, Title II, 136 Stat. 49 (15 Mar. 2022) — initial appropriation funding the Advanced Research Projects Agency for

Health (ARPA-H) and supporting Secretary Becerra's 30 March 2022 placement of ARPA-H within the National Institutes of Health. Formal statutory authorisation of ARPA-H as a defined entity within NIH followed in the *Consolidated Appropriations Act, 2023*, Pub. L. No. 117-328, Div. FF, Title II, § 2331, 136 Stat. 4459 (29 Dec. 2022), codified at 42 U.S.C. § 290c. The earlier *Pandemic and All-Hazards Preparedness Act* (PAHPA, Pub. L. 109-417, 2006) is the foundational BARDA authorising statute.

[47] European Commission, "Commission Decision C(2021) 1510 final of 17 March 2021 adopting the European Innovation Council (EIC) Work Programme 2021 under Horizon Europe"; the EIC Fund (with ∼€4B initial capitalisation) was established in 2020 and its Horizon-Europe-compartment Investment Guidelines apply from March 2021.

[48] *Quanta Computer, Inc. v. LG Electronics, Inc.*, 553 U.S. 617 (2008).

[49] *Impression Products, Inc. v. Lexmark International, Inc.*, 581 U.S. 360 (2017).

[50] Ethereum Improvement Proposals, *EIP-712: Typed structured data hashing and signing*, finalised 2022. [Online]. Available: https://eips.ethereum.org/EIPS/eip-712

[51] P. Daian, S. Goldfeder, T. Kell, Y. Li, X. Zhao, I. Bentov, L. Breidenbach, and A. Juels, "Flash Boys 2.0: Frontrunning, Transaction Reordering, and Consensus Instability in Decentralized Exchanges," in *Proc. IEEE Symposium on Security and Privacy (S&P)*, 2020, pp. 910–927.

[52] J. Milionis, C. C. Moallemi, T. Roughgarden, and A. Sridharan, "Automated Market Making and Loss-Versus-Rebalancing," 2022. arXiv:2208.06046.

[53] V. Buterin, E. Conner, R. Dudley, M. Slipper, I. Norden, and A. Bakhta, *EIP-1559: Fee market change for ETH 1.0 chain*, finalised 2021. [Online]. Available: https://eips.ethereum.org/EIPS/eip-1559

[54] CoW Protocol Foundation, *CoW Protocol: A trustless, MEV-resistant trading protocol for Ethereum* (whitepaper), 2022. [Online]. Available: https://cow.fi/cow-protocol/whitepaper

[55] I. Fisher, *The Purchasing Power of Money: Its Determination and Relation to Credit, Interest and Crises*. New York: Macmillan, 1911 (rev. ed. 1922).

[56] Y. Mao, M. Zhang, S. B. Venkatakrishnan, and Z. Lin, "Flashback: Enhancing Proposer-Builder Design with Future-Block Auctions in Proof-of-Stake Ethereum," arXiv preprint arXiv:2405.09465, May 2024.

[57] J. Quinn and B. Connolly, "Distributed ledger technology and property registers: displacement or status quo," *Law, Innovation and Technology*, vol. 13, no. 2, pp. 377–397, 2021. doi: 10.1080/17579961.2021.1977223.

[58] G. Caldarelli, "Understanding the blockchain oracle problem: A call for action," *Information*, vol. 11, no. 11, p. 509, 2020. doi: 10.3390/info11110509.

[59] *Huawei Technologies Co. Ltd v ZTE Corp.*, Court of Justice of the European Union, Case C-170/13, judgment of 16 July 2015. ECLI:EU:C:2015:477.

[60] "The community of the DAO," *Nature Biotechnology*, vol. 41, no. 11, p. 1346, 2023. doi:10.1038/s41587-023-02005-1. [Editorial coverage of the IP-NFT model and the Molecule Protocol / VitaDAO ecosystem; cited narrowly as journalistic signal of the model's reach, not as peer-reviewed research validation.]

[61] Molecule Protocol, "IP-NFT: A framework for tokenising intellectual property in life sciences," whitepaper and reference implementation. [Online]. Available: https://github.com/moleculeprotocol/IPNFT. [Accessed: 20-May-2026.]

[62] M. Lavoie, *Post-Keynesian Economics: New Foundations*. Cheltenham: Edward Elgar, 2014.

[63] B. J. Moore, *Horizontalists and Verticalists: The Macroeconomics of Credit Money*. Cambridge: Cambridge University Press, 1988.

[64] A. Canidio and F. Henneke, "Fair Combinatorial Auction for Blockchain Trade Intents: Being Fair without Knowing What is Fair," arXiv preprint arXiv:2408.12225, August 2024.

[65] E. Budish, P. Cramton, and J. Shim, "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response," *Quarterly Journal of Economics*, vol. 130, no. 4, pp. 1547–1621, 2015. doi: 10.1093/qje/qjv027.

[66] *Bayh-Dole Act* (University and Small Business Patent Procedures Act of 1980), Pub. L. No. 96-517, 94 Stat. 3015 (12 Dec. 1980), codified as amended at 35 U.S.C. §§ 200–212; implementing regulations at 37 C.F.R. Part 401. Federal march-in rights are codified at 35 U.S.C. § 203.

[67] Godyn, M., Kedziora, M., Ren, Y., Liu, Y., and Song, H.H., "Analysis of solutions for a blockchain compliance with GDPR," *Scientific Reports*, vol. 12, art. 15870, 2022. doi:10.1038/s41598-022-19341-y.

[68] Politou, E., Casino, F., Alepis, E., and Patsakis, C., "Blockchain Mutability: Challenges and Proposed Solutions," *IEEE Access*, vol. 7, pp. 167929–167943, 2019. doi:10.1109/ACCESS.2019.2945736.

[69] Molina, F., Betarte, G., and Luna, C., "A blockchain-based and GDPR-compliant design of a system for digital education certificates," *CLEI Electronic Journal*, vol. 26, no. 1, paper 7, May 2023. doi:10.19153/cleiej.26.1.7.

Author Contributions, Funding, and Data Availability

Author: José Victor Martínez (sole author). Conceptualisation, formalisation, protocol design, smart-contract implementation, deployment, and manuscript preparation. ORCID: 0009-0001-7717-5641.

Affiliation: Metha Biofund — independent research initiative (no parent institution as of v1.4.1; the recommended Liechtenstein Stiftung + Spanish CASP-SL B+A hybrid is the audit-gated Phase-2 structure described in §8.5, not the current operating form).

Competing Interests. The author is the founder of Metha Biofund and would, in any future commercial activation of condition (a), be among the parties economically interested in the protocol's success. The author holds METH tokens from the genesis allocation under the vesting and governance-redirectable terms described in §10.5. No external party — academic, governmental, biotech, or financial — has funded, commissioned, or reviewed this manuscript or had any approval right over its content.

Funding. This work was carried out without external financial support. Hosting, deployment gas, and tooling costs are self-funded. The project has not solicited and does not receive grants

from BARDA, ARPA-H, NIH, Horizon Europe, EIC, or any other public funding body; the references to those agencies' option-period award patterns in §5.4 and §8.5 are analyses of their published practice, not statements of any funding relationship (cf. the "BARDA / ARPA-H alignment is unilateral" paragraph in §8.5).

Data and Code Availability. All smart-contract source, the Hardhat test suite, the deploy and verification scripts, the deployment manifests, and the public-facing dApp source are open-source under the MIT license. Repositories: `metha-api` (contracts, tests, deploy scripts), `metha-client` (public landing and paper source), `metha-admin-v2` (interactive dApp). Sepolia testnet deployment manifest published in §9.1 with verified-source Etherscan links. The version archived alongside this paper is identified by git commit SHA in the Zenodo record metadata.

Disclaimer. Nothing in this paper constitutes legal, tax, financial, or investment advice. The legal-due-diligence analyses summarised in §6, §7, and §8 reference an external counsel engagement and are reproduced for the purpose of explaining the protocol's working hypothesis; readers should obtain independent qualified advice before acting on any of it. Jurisdiction-specific outcomes will depend on facts not analysed here, on regulatory developments after the v1.4.1 cut-off date (2026-05-20), and on counsel of record in the relevant jurisdiction.

Version History.

v1.4.1 (2026-05-20) — Pre-Zenodo hardening pass (citation, legal, math, paper-format, biomedical-funding, smart-contract / EVM).

- §6.1 rich-lab-walks reframed to remove the recordation-erga-omnes contradiction with §6.2-bis exhaustion doctrine.
- Lyapunov function renamed $V \to \mathcal{L}$ (collision with Fisher-equation velocity V); §3.3-bis algebra reworked to log-derivative form; non-increase claim softened to bounded-drift (Foster-Lyapunov).
- Deliverable cap "ten" → "thirty-two" (matches deployed `MAX_DELIVERABLES_PER_MILESTONE = 32`).
- ARPA-H statute corrected to Pub. L. 117-328 § 2331 (Dec 2022); PAHPA 2006 noted as BARDA's authorising statute.
- *Booking.com* / *Amazon* MFN cases re-pocketed as Art. 101 vertical (not Art. 102).
- §8.4-bis SEP/FRAND reworded from "designed to fail" to "structured to fall outside the scope of".
- §9.2 ownership paragraph rewritten to reflect that the Sepolia deployer-as-Program-Officer posture is intentional per the `TRANSFER_POOL_OWNERSHIP` / `REVOKE_DEPLOYER_ROLES` gates in the deploy script.
- §9.2 flash-loan property tightened to atomic single-transaction scope; property (iii) reformulated as reachable-disbursement with safety-vs-liveness caveat; additional Certora invariants to scope listed (role monotonicity, pause safety, EIP-712 cross-chain replay, EternalStorage round-trip, AR v2 stake monotonicity).
- §9.3 gas table replaced with measured averages from `metha-api/gas-report.txt` (the v1.4 estimates ran ∼2× below measurement); attack-vector inventory extended with verdict-ordering race, low-gas multisig recipient, EIP-712 cross-chain replay.

- §10.3 trimmed to cross-reference §3.7 instead of duplicating; IQVIA figure corrected to ~$1.6T; CASP transitional regime expiry to 30 Jun 2026; GDPR Art. 17 "only sense that matters" softened.
- Added: front-matter (Author Contributions / Funding / Data Availability / ORCID / Disclaimer); refs [66] Bayh-Dole, [67] Godyn 2022, [68] Politou 2019, [69] Molina 2023.

v1.4 (2026-05-18) — Revision responding to round-2 review of v1.3 (153 inline comments).

- New §3.7 IP-NFT prior art (refs [60] *Nature Biotechnology* 41:1346, [61] Molecule IP-NFT).
- New §8.4-bis SEP/FRAND + *Huawei v ZTE* (ref [59] Case C-170/13, ECJ 2015).
- New §2.4 endogenous-money parallel (refs [62] Lavoie 2014, [63] Moore 1988).
- New §4.1.1.bis EOS wash-trading acknowledgment (refs [64] Canidio & Henneke 2024, [65] Budish-Cramton-Shim 2015).
- New §5.1 VRF-rotation panel and §5.6 slashing-parameter governance addenda.
- New §9.3 operational-characteristics and attack-vector inventory; legacy §9.3/§9.4 renumber to §9.4/§9.5.
- Pushed back on: "Lavoie/Moore is foundational prior art" (cross-reference §2.4); the antitrust drumbeat already covered by §8.4 / §8.4-bis; "batch auctions are the wash-resistance frontier" (Canidio & Henneke's actual conclusion is the opposite).
- Two suggested references could not be verified and are not cited (*Estèves et al. 2022*; *Amar 2022 Blockchain Mutability*); the doctrinal surface they were attributed to is covered by [36] EDPB Guidelines and [60].

v1.3 (2026-05-15) — Major revision responding to round-1 review of v1.2 (74 inline comments).

- New §3.3-bis (Equation of Exchange specialisation, Fisher [55] / Mundell [12] / Friedman *k*-percent).
- New Lyapunov stability sketch in §3.3; equation (4) added for price relationship; monotonicity assumptions stated explicitly for g, f, h.
- New §4.1.1.bis on auction integrity under MEV/LVR scrutiny (refs [51] Daian 2020, [52] Milionis 2022, [53] EIP-1559, [54] CoW Protocol, [56] Mao 2024).
- New §4.4 paragraph on Stiftung residual veto as deliberate-and-bounded centralisation.
- New §5.1 paragraph on RegulatoryOracle accreditation criteria (ref [58] Caldarelli 2020 replacing mis-attributed Churilov).
- §6.1 Mechanism A reframed: *erga omnes* claim retracted under *Quanta* [48] / *Impression Products* [49]; on-chain bond and Metha Pharma backstop become more load-bearing.
- New §6.2-bis (soft-variant clause under exhaustion doctrine).
- New §6.3 oracle layer (three architectures with explicit pros/cons; commits to Options 1 + 3, reserves Option 2).
- New §6.4 conflicts of laws (civil-law vs common-law; one analogical cite to Quinn & Connolly [57]).
- §8.4 expanded with combined *Brulotte* [32] + exhaustion + Art. 102 analysis under the soft variant.
- New §9.1 Sepolia deployment manifest (addresses + Etherscan + verified-source).
- §9.2 Formal-verification scope expanded: Certora commitment on four named properties.

- §10.5 paragraph on pre-snapshot accumulation vs intra-proposal flash-loan attacks.
- One suggested reference (*Schüler & Lage 2025*) could not be verified and is not cited.

v1.2 (2026-05-14) — Major revision integrating Legal Due Diligence v1.0 and the Auditing-Funding methodology brief; reflects the v3 Sepolia stack deployed 2026-05-13.

- Six deployed Sepolia contracts (added AuditorRegistryV2); on-chain audit infrastructure promoted to §5 (six sub-sections).
- Replaced symmetric +10/−5 reputation with tiered slashing schedule (Minor / Major / Critical).
- Added §10.4 (Metha Pharma vertical backstop), §10.5 (open governance design questions).
- Softened condition-(a) activation language; replaced fixed Phase-2 dates with audit-gated framing.
- Integrated Legal DD findings: seven-plane regulatory analysis, B+A hybrid (Liechtenstein Stiftung + Spanish MiCA-CASP SL), €520k–€1.35M year-1 cost band, MiCA Spain transitional deadline.
- Jurisdictional caveats on IP encumbrance (USPTO/EPO/JPO/KIPO heterogeneity, UPC pending, *Brulotte* and *Genentech*).
- GDPR Art. 17 two-tier mitigation + "BARDA / ARPA-H alignment is unilateral" acknowledgement; explicit "Slither is not a third-party audit" paragraph.
- Verified EternalStorage slot registry; documented AuditorRegistryV2 peer-to-peer wiring with AuditingPool.
- New §3.6 (Why hasn't this been done before?), §4.1.1 (EOS empirical precedent), §6.1 (rich-lab-walks scenario).
- New refs [27]–[47] (Legal DD, Roadmap v2, methodology brief, BARDA/NIDCR sources, *Brulotte*, *Genentech*, DAC8, CARF, EDPB Guidelines 02/2025, DORA, *Booking.com* C-264/23, *Amazon* AT.40462/40703, Salop & Scott Morton, ASCAP/BMI decrees, FRED M2SL, IEA WEO 2021, Nixon Proclamation 4074, IMF Articles of Agreement, ARPA-H statutes, EIC Decision C(2021) 1510).
- Citation audit: replaced non-verifiable [17] Certora entry with Grossman et al. POPL 2018 + Certora docs; *JMRI v. Katzer* → *Jacobsen v. Katzer*; *Brulotte* year 1965 → 1964; WPT → WPR; refreshed NIDCR DSMB Guidelines URL; archived EOSIO whitepaper to web.archive.org.

v1.1 (2026-05-04) — Added §4.4 IP Ownership and Settlement Condition (contractual licence rather than proprietary claim), §4.5 Enforcement (three-layer: recordation succession / march-in / liquidated damages + bond + reputation), rewrote §5.4 distinguishing securities-law from antitrust risk, added EOS empirical precedent in §4.1, refs [22]–[26] (Howell-Niessner-Yermack 2020, SEC Block.one Order 2019, *Jacobsen v. Katzer*, GNU GPLv3, New York Convention 1958).

v1.0 (2026-05-04) — Initial preprint release.

Working draft. Citations and arguments will be refined as feedback is incorporated. Latest version always available at https://metha.life/documents/DVBC_Metha_CaseStudy.pdf.