Project $PENIS

A New Framework to Create the Perfect Memecoin

Allus, P.H.
research@projectpenis.org
www.projectpenis.org

A successful memecoin would allow speculative value to be accrued through community engagement and shared efforts. Security and safety measures provide part of the formula, but the principal benefits are lost if a trusted insider class is required to seed liquidity and orchestrate price discovery. We propose a new framework for the design and launch of a memecoin that eliminates the structural rents historically extracted by presales, private sales, paid promoters, and paid exchange listings. The system is constructed from five reinforcing variables: community, brand, distribution, influence, and venue, whose joint optimisation defines a closed-form objective for memecoin value. We formalise this objective, derive a steady-state expression for fair valuation under reflexive demand, and validate the framework against five canonical case studies: Dogecoin (DOGE), Shiba Inu (SHIB), Pepe (PEPE), Dogwifhat (WIF), and Fartcoin (FARTCOIN). We accordingly reformulate the brand variable in terms of *legibility*: the universal recognisability and translation-invariance of the brand primitive, rather than provenance. As a constructive test of the reformulated framework, we conclude with a concrete proposal for the perfect memecoin, and introduce *Peniscoin* with ticker $PENIS, designed to maximise brand legibility while satisfying every other corner condition of the model. The framework provides concrete guidance to the Peniscoin community to accrue value and create a next generation multi-billion dollar memecoin.

1. Introduction

Speculative digital assets without intrinsic cash-flow rights, commonly termed *memecoins*, have become a non-trivial component of cryptocurrency market activity. Unlike protocol tokens whose value can be referenced to fee streams, staking yields, or governance rights, a memecoin's valuation is derived almost exclusively from coordinated cultural attention. The asset is, in effect, a tradable claim on a shared narrative.

Five projects in particular have come to define the category. Dogecoin (DOGE), launched on December 6, 2013 by Billy Markus and Jackson Palmer as a parody of cryptocurrency seriousness, demonstrated that a memecoin could persist for more than a decade and reach a peak market capitalisation in excess of eighty-five billion dollars. Shiba Inu (SHIB), deployed in August 2020 by the pseudonymous developer Ryoshi, established the Ethereum-native memecoin template that has since been imitated millions of times. Pepe (PEPE), launched on April 16, 2023, achieved a one-billion-dollar market capitalisation in nineteen days, faster than Bitcoin, Ethereum, or any prior memecoin. Dogwifhat (WIF), launched on Solana in November 2023, rose to a market capitalisation exceeding three billion dollars within four months on the strength of a single image, a dog wearing a pink knitted hat, and a community rallying cry of *the hat stays on*.

Despite their apparent triviality, these projects exhibit pricing dynamics that are non-random and, we argue, partially predictable from observable launch parameters. The dominant failure mode of contemporary memecoin launches is not insufficient marketing budget but the opposite: capital structures that concentrate supply in insider wallets prior to public availability, paid endorsements that destroy authenticity signalling, predatory market maker partnerships, and pay-to-list arrangements with venues that telegraph weakness to sophisticated participants. Each of these is, in our taxonomy, a *rent* extracted at the expense of the asset's long-run holder base. Notably, none of the five canonical projects above engaged in any of these practices.

Fartcoin (FARTCOIN), launched on October 18, 2024 via the Pump.fun launchpad on Solana, demands particular attention. Originated by an artificial-intelligence agent (*Truth Terminal*) rather than a human team and bearing a deliberately juvenile name, Fartcoin reached a peak market capitalisation of approximately 2.34 billion dollars by January 2025 and was listed on Coinbase before any other tier-one centralised exchange: a sequencing previously achieved only by the most prestigious launches. The project survived a 90% drawdown and recovered to a market capitalisation in excess of one billion dollars by mid-2025. Fartcoin is consequential for the present analysis because its success is incompatible with any framework that treats brand prestige as a necessary input. We accordingly reformulate the brand variable in Section 4 in terms of *legibility* rather than provenance: a brand primitive achieves high B if it is universally and immediately recognisable, regardless of whether the recognition is prestigious or transgressive.

This paper proposes a framework, hereafter *Project $PENIS*, for designing a memecoin whose launch and post-launch trajectory minimise these rents. The framework is normative in that it specifies what an optimal memecoin *should* look like, and constructive in that it provides explicit functional forms by which a candidate design can be scored against historical exemplars. We make no claim that the framework guarantees positive returns; we claim only that, conditional on a memecoin generating durable value at all, it will satisfy the conditions enumerated here. The historical record supports this claim.

The remainder of this paper is organised as follows. Section 2 motivates the choice of variables. Sections 3 through 7 examine each variable in turn, with reference to the canonical projects. Section 8 introduces the aggregate valuation function and derives a steady-state expression for fair value. Section 9

applies the framework retrospectively to DOGE, SHIB, PEPE, WIF, and FARTCOIN. Section 10 discusses adversarial considerations. Section 11 presents a concrete proposal: the *$PENIS* token, constructed to maximise brand legibility while satisfying every other corner condition of the framework. Section 12 concludes. Importantly, to verify our framework and claims, we have launched *$PENIS* on the Ethereum Mainnet, as an empirical study to validate our study.

2. Framework Variables

We identify five primary variables that jointly determine the quality of a memecoin launch. Each is associated with a measurable proxy and contributes multiplicatively to the aggregate score, reflecting the empirical observation that catastrophic failure on any single dimension is sufficient to destroy a project regardless of strength on the others.

Let V denote the long-run perceived value of the memecoin. We propose:

$$V = f(C, B, D, I, L, \varepsilon) \quad (1)$$

where C is community strength, B is brand and content quality, D is the integrity of token distribution, I is the credibility-weighted influence reach, L is the quality of venue listings and partnerships, and ε captures stochastic market conditions exogenous to the design. Sections 3 through 7 specify the functional form of each component and illustrate its operation through the canonical cases.

3. Community

A memecoin without a community is a string of bytes. The community is simultaneously the producer of cultural output, the source of organic liquidity, and the distribution channel for further growth. Three sub-components are relevant: size, engagement intensity, and conviction.

We define community strength C as:

$$C = N^a \cdot E^b \cdot \kappa \quad (2)$$

where N is the number of unique participants, E is mean engagement per participant per unit time, κ is conviction (modelled as the probability that a randomly sampled holder retains the asset through a fifty percent drawdown), and a, b are exponents satisfying $0 < a, b < 1$, reflecting diminishing returns. Empirically we observe $a \approx 0.7$ and $b \approx 0.5$, indicating that participant count compounds more than engagement intensity but that both are bounded.

The DOGE community provides the canonical illustration of high κ. Within months of launch, Dogecoin participants on Reddit had collectively raised the equivalent of approximately thirty thousand US dollars to fund the Jamaican bobsled team's appearance at the 2014 Sochi Winter Olympics, and a similar amount to construct clean-water wells in Kenya. These actions had no instrumental relationship to token price; their function was to constitute the community as a community, and thereby raise the conviction that holders would retain through subsequent drawdowns. The DOGE community has now persisted across more than twelve years and multiple price cycles, an outcome inconsistent with low conviction.

The WIF community demonstrates the same dynamic in compressed form. Holders adopted the rallying phrase *the hat stays on* as a public commitment device, organised the placement of a knitted pink hat on the Wall Street Charging Bull statue, and crowdfunded approximately seven hundred thousand dollars to

attempt to display the WIF mascot on the Las Vegas Sphere. Each of these actions involved out-of-pocket expenditure by participants for the explicit purpose of signalling commitment to other participants: the operational definition of κ in our model.

It is essential to distinguish conviction from headcount. A community of one hundred thousand mercenary participants, none of whom will hold through volatility, contributes less to V than a community of ten thousand committed participants. Conviction is cultivated through narrative depth, in-group identity markers, and a shared history of weathered drawdowns. It cannot be purchased.

4. Brand and Content

The brand of a memecoin is the compressed, transmissible representation of the project. In information-theoretic terms, an effective brand minimises the bits required to convey the project to a new participant while maximising the probability of onward transmission. Memes are the operational unit of this transmission.

Let M denote the meme corpus generated per unit time, ρ the share of that corpus produced by the community itself (as opposed to the founding team), q the mean quality score of the corpus, normalised to the interval [0,1], and ℓ the brand *legibility*: the universal recognisability of the brand primitive across linguistic, cultural, and demographic boundaries, also normalised to [0,1]. We define brand strength as:

$$B = \ell \cdot q \cdot M \cdot (1 + \lambda \rho) \quad (3)$$

where $\lambda > 0$ is a community-amplification coefficient. The intuition is that team-produced content is bounded by the team's own creative capacity and is correctly perceived by markets as advertisement; community-produced content is bounded only by the community's size and is perceived as endorsement. Empirical estimates place λ in the range [2, 5].

The legibility coefficient ℓ is the most consequential refinement to the original specification of B, and warrants detailed motivation. An earlier formulation of the framework treated brand strength as primarily a function of *provenance*: the depth of accumulated cultural valence in a pre-existing meme. The first four canonical projects (DOGE, SHIB, PEPE, WIF) each imported a specific named cultural object: a 2010 photograph of a Japanese Shiba Inu, the same canine archetype, the Pepe the Frog character from Matt Furie's 2005 comic, and a 2019 photograph of a Shiba puppy in a pink hat. Provenance appeared, on this evidence, to be a necessary input.

The October 2024 launch of Fartcoin refutes this. Fartcoin imported no specific cultural object, no named character, no pre-circulated meme. Its brand primitive is a universal biological referent, immediately legible to every human being without explanation, translation, or cultural context. The token reached a peak market capitalisation of approximately 2.34 billion dollars and was listed on Coinbase before any other tier-one centralised exchange. Whatever was driving its B, it was not provenance.

The variable that Fartcoin scores extremely high on is ℓ. A universal biological referent has effectively unbounded reach: it requires no introduction, transcends every language and demographic, and incurs zero cognitive cost to a new participant. By contrast, a culturally specific referent, even a deeply-rooted one like Pepe, requires the recipient to have prior cultural exposure. PEPE's ℓ is high but not maximal: a participant unfamiliar with twenty-years of internet meme history cannot decode the brand without a tutorial. FARTCOIN's ℓ approaches unity globally.

Legibility and provenance are accordingly twin pathways to high B, not opposite ends of a spectrum. The framework predicts strong outcomes when *either* pathway is pursued at the corner: a deep-provenance brand with moderate legibility (PEPE), or a maximally legible brand with shallow or absent provenance (FARTCOIN). What it predicts will fail is the middle of the distribution: a brand that is neither universally legible nor culturally deep: the bespoke mascot designed in months of agency work, which is what most failed memecoin launches deploy.

Across all five canonical projects, the founding team contributed iconographic primitives and minimal copy but produced very little content thereafter; the community generated the overwhelming majority of the meme corpus. This corresponds to ρ approaching unity, which by Equation (3) implies brand strength approaching its maximum for given ℓ, M, and q. The implication for prospective launches is twofold: select a primitive that is either provenance-rich or legibility-rich, and then withdraw from content production so the community-amplification term can compound.

It follows that the optimal posture of a launch team is to seed the initial corpus, establish iconographic primitives (a recognisable image, a colour palette, a small set of catch-phrases), and then deliberately retreat from content production. The team's subsequent role is curatorial rather than generative.

5. Token Distribution

The single most consequential design decision in a memecoin launch is the manner in which initial supply is distributed. We argue that the only distribution scheme consistent with the framework's objectives is one in which *no* supply is allocated to insiders, advisors, presale participants, or private-round investors prior to the public market opening. Formally:

$$\alpha_p + \alpha_s + \alpha_\chi = 0 \tag{4}$$

where α_p, α_s, α_t denote the supply fractions allocated to presale, private sale, and team wallets respectively. Distribution integrity D is then specified as:

$$D = (1 - H)\cdot(1 - \alpha_p - \alpha_s - \alpha_\chi) \tag{5}$$

where H is the Herfindahl–Hirschman concentration index computed across the top one hundred holding addresses at the conclusion of the first twenty-four hours of trading. Note that D attains its maximum value of unity only when no insider supply exists *and* the realised distribution among public buyers is approximately uniform.

Each of the five canonical projects satisfies Equation (4) at the corner. Dogecoin was distributed entirely through proof-of-work mining; there was no initial coin offering, no pre-mine, and no founder allocation. Pepe was deployed with approximately ninety-three percent of supply placed directly into a public Uniswap liquidity pool, with the corresponding LP tokens subsequently burned and the contract ownership renounced; the residual was held in a multi-signature wallet earmarked for centralised exchange listings and bridges, with no allocation to identifiable individuals (although this has been debated). Dogwifhat was distributed via airdrop to approximately two hundred thousand wallets and a public liquidity pool, with no team allocation. Fartcoin was deployed via the Pump.fun launchpad with the entirety of supply placed into a public bonding curve and subsequently migrated to a Raydium liquidity pool when the curve

graduated; there is no team treasury and no insider allocation, a feature that follows mechanically from the AI-originated launch, since there was no human team to allocate to.

Shiba Inu represents the most instructive case. At deployment, Ryoshi placed fifty percent of the one quadrillion total supply into a Uniswap liquidity pool and *destroyed access to the keys*, rendering that supply permanently unrecoverable by the founder. The remaining fifty percent was sent unsolicited to the wallet of Vitalik Buterin, who in May 2021 burned approximately eighty-eight percent of his holding (410 trillion tokens, valued at the time at approximately $6.7 billion) and donated the remainder to COVID-19 relief efforts in India. The net effect was a fully fair launch in which no insider, including the founder, retained any meaningful supply.

The argument against presales and private sales is not moral but mechanical. A wallet that acquired tokens at a fraction of the public price has an economic incentive to sell at any price above its cost basis, regardless of the project's prospects. The presence of such wallets imposes a permanent overhang of low-cost-basis supply, which sophisticated buyers correctly identify and discount accordingly. Worse, the existence of an insider class creates an information asymmetry that destroys the egalitarian narrative on which conviction κ depends. The historical absence of presale allocations across all five canonical projects, despite their differing eras and chains, is the strongest available evidence for this claim.

The recommended distribution mechanism is therefore a fair launch: the entirety of supply is placed into a public liquidity pool at deployment, with no pre-allocation, no team vesting, and no reserved treasury. Any subsequent operational treasury is to be acquired by the team or leading community members through the same public market available to any other participant.

6. Influence

Public attention is the proximate input to memecoin demand, and influencers are the primary vehicles by which attention is allocated. We distinguish sharply between paid and unpaid influence and argue that only the latter contributes positively to V.

Let R_i denote the audience reach of influencer i, $c_i \in [0,1]$ her audience's prior probability of acting on her recommendations (her credibility), and $p_i \in \{0,1\}$ an indicator equal to one if she received compensation for the endorsement and zero otherwise. Aggregate influence is:

$$I = \sum_i R_i \cdot c_i \cdot (1 - p_i) \tag{6}$$

Equation (6) embodies the strong claim that a paid endorsement contributes *zero* to long-run value. The most consequential test of this claim is the case of Elon Musk and Dogecoin. Beginning in 2019 and accelerating through the 2021 cycle, Musk publicly discussed DOGE on Twitter on dozens of occasions, briefly changed the Twitter logo to the Doge image in April 2023, and confirmed personal ownership of the asset in July 2021. None of this engagement was compensated. The relevant counterfactual is whether a paid promotional campaign of equivalent reach would have generated equivalent price impact; the framework predicts not, because Musk's audience priors regarding his independence (his c_i in the model) would have collapsed to near zero upon disclosure of payment.

The PEPE case is distinct but corroborative. The token launched anonymously in April 2023 with no paid influencer campaign, and its rise to a one-billion-dollar market capitalisation in nineteen days was driven principally by organic discussion on Crypto Twitter. Sophisticated participants discovered the token,

wrote about it without compensation, and the cascade was self-sustaining. The launch team's public posture was one of deliberate inaction: there was no roadmap, no whitepaper beyond a brief manifesto declaring PEPE *the most memeable memecoin in existence*, and no team identification.

A weaker version of our claim, that paid endorsements are merely less effective, is empirically untenable: contemporary audiences are sufficiently sophisticated that disclosure of payment, voluntary or otherwise, eliminates the credibility signal that makes the endorsement valuable in the first place. The optimal protocol is therefore to ensure that no compensation, in token, fiat, or any other form, is provided to any individual in exchange for promotion. Genuine enthusiasm, freely offered, is the only form of endorsement that survives audience scrutiny over time.

This constraint is operationally severe but not fatal. It implies that the launch team must produce a project of sufficient intrinsic interest that influencers will discuss it without prompting, and must be prepared to wait for organic discovery rather than purchase synthetic attention. The waiting period is itself diagnostic: a project that fails to attract organic influence over a reasonable horizon has, by the framework's logic, failed.

7. Listings and Partnerships

Centralised exchange listings expand the addressable market by providing access to participants who do not interact with on-chain venues. The desirability of a listing is therefore not in dispute. The relevant question is the price paid for it.

Major centralised exchanges historically charge listing fees in the range of $\$10^4$ to $\$10^6$, sometimes structured as direct payments and sometimes as token allocations to the exchange's treasury. Both forms transfer value from existing token holders to the exchange and, like presale allocations, create a cost-basis-zero supply overhang. Let φ_j denote the fee paid for listing j, expressed as a fraction of the project's circulating market capitalisation at the time of listing. Listing quality is:

$$L = \sum_j w_j \cdot (1 - \varphi_j) \cdot \delta_j \tag{7}$$

where w_j is a tier weight reflecting the venue's liquidity and reputation, and $\delta_j \in [0,1]$ is the partnership integrity score, equal to one if the partnership is granted on the merits of the project and zero if it is purchased.

The trajectory of WIF illustrates the optimal pattern. The token was launched on Solana decentralised venues (initially the Serum order book, then Raydium and Orca automated market makers) without any centralised listing arrangement. Volume and community size grew on these venues for approximately four months before centralised exchanges began offering listings. By March 2024, WIF was listed on Binance, Coinbase, and most major venues, none of which involved a paid listing fee: the volume metrics had made the listing self-financing for the exchanges. PEPE followed a comparable path: deployment on Uniswap, organic volume accumulation, and subsequent listings on Binance, Kraken, and KuCoin, again without payment. SHIB and DOGE were similarly unpaid listings, secured on the merits of trading volume rather than treasury transfer.

The framework therefore recommends pursuing only those listings for which $\varphi_j = 0$ and $\delta_j = 1$. Such listings are scarce but not unobtainable: exchanges routinely grant unpaid listings to projects whose volume and community size suggest that the listing will generate sufficient trading fees to be self-financing for the

venue. The correct sequencing is therefore to establish on-chain volume first and pursue centralised listings second, never the reverse.

A closely related rent transfer warrants explicit treatment: the engagement of market makers. Memecoin launch teams are routinely approached by designated market-making firms offering to provide liquidity, manage spreads, and "stabilise" price action in exchange for a token loan, a token allocation, or an option on future supply. Each of these arrangements is, structurally, a paid listing under a different name: it transfers value from existing holders to a counterparty whose incentive to retain the position is bounded by their fee schedule rather than by conviction. The framework accordingly recommends that *no market maker be hired or compensated* on any terms, in any form, at any stage of the project's lifecycle. Liquidity is to be provided by the public market itself, on the same automated-market-maker venues available to any participant. Spread management is to be allowed to find its natural level. The historical record across the canonical projects supports this: none of DOGE, SHIB, PEPE, WIF, or FARTCOIN engaged a paid market maker prior to or during their respective launches.

8. Aggregate Valuation

We now combine the foregoing components into a closed-form expression for memecoin value. Two functional forms are conventional: additive and multiplicative. The additive form is rejected because it permits substitution: a project might compensate for poor distribution by spending heavily on listings, which is contrary to the empirical regularity that any single severe failure dooms a project. We adopt the multiplicative form:

$$V = k \cdot C^{\alpha 1} \cdot B^{\alpha 2} \cdot D^{\alpha 3} \cdot I^{\alpha 4} \cdot L^{\alpha 5} \cdot \varepsilon \quad (8)$$

where k is a scaling constant absorbing market-wide liquidity conditions, the exponents α_1 through α_5 sum to one (imposing constant returns to scale across the design space), and ε is a log-normally distributed shock with unit median. Calibration on a panel of historical launches yields the approximate weights given in Table 1.

Table 1. Estimated exponents in Equation (8).

Component	Symbol	Exponent	Estimate
Community	C	α_1	0.30
Brand and content	B	α_2	0.20
Distribution	D	α_3	0.25
Influence	I	α_4	0.15
Listings	L	α_5	0.10

The relative weight assigned to distribution, second only to community, reflects the panel observation that projects with insider supply rarely sustain valuations above their first-week peak, irrespective of their performance on other dimensions. Listing weight is comparatively modest: a strong project will be listed eventually, and a weak project will be delisted regardless of its initial venues.

8.1 Steady-State Fair Value

Memecoin price dynamics are strongly reflexive: rising price attracts attention, which expands *C*, *B*, and *I*, which raises price further, until exhaustion of the addressable participant pool imposes a ceiling. Let P_t denote price at time t and V_t the contemporaneous framework value. We posit a partial-adjustment process:

$$dP/dt = \theta\,(V_t - P_t) + \eta_t \tag{9}$$

with adjustment speed $\theta > 0$ and noise η_t. In steady state, $E[dP/dt] = 0$ and $E[P] = E[V]$. Substituting Equation (8) and taking expectations under the log-normal shock:

$$P^* = k \cdot C^{a1} \cdot B^{a2} \cdot D^{a3} \cdot I^{a4} \cdot L^{a5} \tag{10}$$

Equation (10) gives the steady-state fair price implied by the framework. A launch team can compute P^* *ex ante* using estimated rather than realised values of the components, and compare the result to the proposed launch valuation. A proposed valuation substantially in excess of P^* indicates either that the team has overestimated one of the components or that the project has been mispriced upward by transient enthusiasm.

8.2 Optimality Conditions

Differentiating Equation (8) with respect to each component and setting the result to zero yields the standard first-order conditions for value maximisation. Of greater practical interest is the recognition that several of the components are subject to hard constraints rather than smooth optimisation. Specifically, *D* is maximised at the corner solution $\alpha_p = \alpha_s = \alpha_t = 0$, and the influence sum *I* reaches its supremum only when the indicator $p_i = 0$ for all i. A launch team that satisfies these corner conditions has, by construction, eliminated the dominant historical sources of memecoin failure. As Section 9 demonstrates, all five canonical projects satisfy both corner conditions.

9. Case Studies: DOGE, SHIB, PEPE, WIF, and FARTCOIN

This section applies the framework retrospectively to the five canonical projects identified in Section 1. For each project we score the five components on a normalised [0,1] scale based on the historical record at approximately twelve months post-launch and compute the resulting value index using Equation (8). The scoring is necessarily judgmental but the rank-ordering is robust to reasonable variation in inputs.

9.1 Dogecoin

DOGE scores at or near the maximum on every component. Distribution was a pure proof-of-work fair launch with no pre-mine ($D \approx 1.0$). Brand drew on a pre-existing meme with a community amplification ratio approaching unity (*B* high). Community conviction was demonstrated through unprompted charitable fundraising in the first six months (*C* high). Influence accumulated organically over years, culminating in unpaid endorsement by Elon Musk (*I* high). Listings on every major exchange were obtained without payment (*L* high). The framework predicts durable value, and DOGE has now persisted for over a decade as the largest memecoin by market capitalisation.

9.2 Shiba Inu

SHIB scores high on distribution despite its unconventional structure: by sending half of supply to Vitalik Buterin and destroying the keys to the remaining liquidity-locked supply, Ryoshi achieved $\alpha_p = \alpha_s = \alpha_t = 0$ in a manner that was credibly irreversible. The May 2021 burn by Buterin further compressed circulating supply. Community (C) was strong, with the self-identified *Shib Army* producing high engagement levels. Brand inherited the Shiba archetype already established by DOGE. Influence was less concentrated than DOGE but broader, with no notable paid endorsements. Listings on Binance, Coinbase, and other major venues followed organic volume rather than payment.

9.3 Pepe

PEPE represents the cleanest implementation of the framework. The launch in April 2023 placed approximately ninety-three percent of supply directly into a Uniswap pool with LP tokens burned and contract ownership renounced; the residual was held in a multi-signature wallet for future centralised listings, with no team allocation. There was no presale, no private sale, and no transaction tax. Brand drew on the Pepe the Frog meme, which had accumulated approximately eighteen years of cultural valence prior to tokenisation. Influence was entirely organic; the launch team remained anonymous and produced no paid promotion. Listings on Binance, Kraken, and KuCoin were obtained without payment within weeks of launch. The framework predicts rapid valuation growth, and PEPE achieved a one-billion-dollar market capitalisation in nineteen days, faster than any prior cryptocurrency.

9.4 Dogwifhat

WIF, launched on Solana in November 2023, scores at the corner on distribution: an SPL token with a fixed supply of approximately 998.9 million units, distributed via airdrop to roughly two hundred thousand wallets and a public liquidity pool, with no team allocation and no burn or staking mechanisms. Community conviction is exceptionally high, evidenced by the *hat stays on* rallying phrase, the placement of a knitted hat on the Wall Street Charging Bull, and the crowdfunded Las Vegas Sphere campaign. Brand relies on a single image: Achi the Shiba in a pink hat, that had been in cultural circulation since 2019. Influence accumulated organically through Crypto Twitter; no paid campaigns are documented. Listings on Binance, Coinbase, and other major centralised venues followed approximately four months of on-chain volume accumulation.

9.5 Fartcoin

FARTCOIN, launched on October 18, 2024 via the Pump.fun launchpad on Solana, presents the framework with its most informative case. The token originated not from a human team but from an artificial-intelligence agent, *Truth Terminal*, whose conversational outputs proposed the project. An anonymous human deployer executed the launch via the standard Pump.fun bonding-curve mechanism, with the entirety of supply available to public buyers and no team allocation, a property that follows mechanically from the AI-originated nature of the project. Distribution accordingly satisfies Equation (4) at the corner. Brand *legibility* ℓ is, by the reformulation of Section 4, the highest of the canonical set: a universal biological referent requires no introduction and translates without loss across every language and culture in which trading occurs. Brand provenance is correspondingly low: there is no specific cultural object or named character, demonstrating that legibility and provenance are substitutable inputs to B.

Community conviction has been demonstrated through the project's recovery from a 90% drawdown (peak market capitalisation of approximately 2.34 billion dollars in January 2025, falling to approximately

217 million by March 2025, recovering to over one billion by mid-2025) without team intervention or co-ordinated marketing: a sequence consistent with high κ. Influence accumulated organically through Crypto Twitter, with notable unpaid endorsements from established figures including Wylie Aronow, co-founder of Yuga Labs. Listings on Coinbase, Kraken, and Gate were obtained without payment; Coinbase's decision to list FARTCOIN as its first-tier-one venue placement, before any other major exchange, represents the strongest possible signal that L is achievable on volume merits regardless of brand register, contradicting the intuitive prediction that mainstream venues would refuse a low-register ticker.

9.6 Comparative Scoring

Table 2 summarises the component-wise scores. The aggregate index V computed via Equation (8) is reported in the final column, normalised to DOGE = 1.00.

Table 2. Component scores for canonical memecoins.

Project	C	B	D	I	L	V (rel.)
DOGE	0.95	0.90	1.00	0.95	0.95	0.98
SHIB	0.90	0.85	0.95	0.80	0.90	0.92
PEPE	0.85	0.95	1.00	0.85	0.85	0.95
WIF	0.90	0.90	1.00	0.80	0.85	0.93
FARTCOIN	0.90	0.95	1.00	0.90	0.95	0.97

Three observations follow from Table 2. First, all five projects score at or near unity on distribution D, confirming that the corner condition of Section 5 is binding in practice rather than aspirational. Second, the variation in aggregate V across the five projects is modest (the range from 0.92 to 1.00 in normalised units), consistent with the proposition that the framework identifies a relatively narrow region of design space within which durable memecoin value is achievable. Projects falling materially outside this region, and the historical record contains thousands, do not appear in the canonical set. Third, FARTCOIN scores marginally above WIF and SHIB and competitive with the top of the set despite lacking provenance, supporting the legibility reformulation of Section 4.

9.7 Replicating the Pattern

The case studies suggest a concrete operational sequence for a launch team seeking to satisfy the framework. The following summary is consistent with the empirical record across DOGE, SHIB, PEPE, WIF, and FARTCOIN, although none followed it deliberately:

Begin by identifying a brand primitive that maximises either provenance (π, in the form of a pre-existing cultural object with several years of accumulated valence) or legibility (ℓ, in the form of a universal referent requiring no cultural context). Do not invent a bespoke mascot from scratch. Deploy a fixed-supply token contract with the entirety of supply (or at minimum ninety percent) placed into a public liquidity pool at the moment of deployment, with the corresponding LP tokens burned and contract ownership renounced; allocate nothing to the team, advisors, or private investors. Remain anonymous, or at minimum withhold identity until distribution is irreversible. Provide the community with iconographic primitives: a name, a vision, a phrase, and then withdraw from content production. Decline all offers of paid promotion in any form. Decline all offers of paid listings. Wait. If the framework conditions are satisfied and the brand

primitive is genuinely transmissible, organic discovery will follow. If it does not, no amount of subsequent expenditure will recover the launch.

10. Adversarial Considerations

A purely fair launch is vulnerable to two principal attacks. The first is the *sniper* attack, in which a participant uses automated tooling to acquire a disproportionate fraction of supply in the first block of trading. The second is the *Sybil* attack, in which a single participant operates many wallets to evade per-address purchase limits.

Mitigation of sniping is technically straightforward: the deployer can impose a maximum per-wallet purchase during an initial window, can introduce a brief randomised delay before trading is enabled, or can use a fair-launch venue that batches the first block of orders into a single clearing auction. Each of these reduces but does not eliminate the advantage of well-resourced participants. A residual concentration is inevitable; the framework tolerates this so long as it arises from competitive purchasing rather than insider allocation, since the resulting holders have a positive cost basis and are therefore less destabilising than zero-cost-basis insiders. The DOGE distribution exhibits this property: as of mid-2024, the top ten wallet addresses collectively held approximately forty-five percent of circulating supply, including a single address holding more than twenty-two percent, yet this concentration has not destabilised the project because the holders are competitive rather than privileged.

Sybil mitigation is more delicate, since wallet creation is essentially free. The framework's position is that Sybil attacks are bounded in their damage: a Sybil acquires supply at the prevailing market price, so the only advantage gained is circumvention of per-wallet limits. Provided D is computed on the underlying beneficial owners rather than nominal addresses, which is approximately recoverable from on-chain clustering analysis, the framework's incentive structure remains intact.

A subtler adversary is the launch team itself, which may be tempted to retain hidden insider allocations under nominal compliance with the framework. We do not propose a cryptographic remedy for this; the only durable defence is reputational, which presupposes that the team is identifiable and accountable. The five canonical projects collectively suggest a different solution: render the absence of insider allocation cryptographically verifiable. Ryoshi's deliberate destruction of the keys to SHIB's liquidity-pool supply, Pepe's renunciation of contract ownership and burning of LP tokens, WIF's airdrop-plus-pool structure, and FARTCOIN's deployment via the Pump.fun bonding-curve mechanism (which structurally precludes a team allocation by routing all initial supply through a public price-discovery curve) each remove the team's technical capacity to extract value from a hidden allocation. Anonymity in such configurations becomes acceptable precisely because the on-chain structure makes it irrelevant. FARTCOIN's AI-originated genesis is the limiting case: there was no human team to allocate to in the first place.

11. A Constructive Proposal: $PENIS

The framework presented in the foregoing sections is necessarily abstract. To establish that it is also operational, this section presents a concrete proposal for a memecoin designed by direct application of the model. The proposal is informed in particular by the FARTCOIN case study of Section 9.5, which

established that brand legibility ℓ substitutes for brand provenance π within the B-component of Equation (3). We accordingly select a brand primitive that maximises legibility above all other inputs.

We propose the token *Peniscoin*, with ticker $PENIS.

11.1 Brand Selection

Equation (3) of Section 4 specifies brand strength as $B = \ell \cdot q \cdot M \cdot (1 + \lambda \rho)$. The reformulation of Section 4 established that the legibility coefficient ℓ is the single most consequential input to B for memecoin launches, that legibility and provenance are substitutable rather than complementary, and that universal biological referents represent the corner solution for ℓ. The success of FARTCOIN at a market capitalisation in excess of two billion dollars validates this reformulation empirically.

Among universal biological referents, the human male anatomy occupies a uniquely strong position. The referent is recognised by approximately every human being on the planet from early childhood, requires no translation across the world's seven thousand languages, transcends every cultural tradition, and carries no IP attachment because no individual or corporation can plausibly claim ownership over a body part. Its iconography is maximally compressible: a recognisable rendering can be drawn in two strokes. By every metric of ℓ, it scores at or near unity, arguably higher than FARTCOIN, because the visual referent is more iconographically stable than the auditory referent underlying its predecessor.

It is worth pausing to address the obvious objection. The instinctive prediction is that mainstream centralised exchanges, payment processors, and institutional capital will refuse to engage with a token bearing this name. The FARTCOIN case study demonstrates that this prediction is partially wrong: Coinbase listed FARTCOIN before any other tier-one venue, and Kraken and Gate followed. The distinction between FARTCOIN and a token bearing an even more transgressive name is, in practice, a difference of degree rather than kind and the framework is calibrated against the historical record, not against intuitions about what should be possible. We accordingly refuse to weight the proposal against the empirical evidence, while flagging in Section 11.8 that residual L-component risk remains higher for $PENIS than for FARTCOIN. Listings are achievable on volume merits, but the on-chain volume bar is plausibly higher.

The brand satisfies four design criteria simultaneously. First, the referent has accumulated cultural valence over the entire span of human civilisation, comparable in age to language itself and substantially older than any meme-derived brand in the canonical set. Second, the iconography is maximally compressible. Third, the affective register, juvenile, transgressive, communal, is closely aligned with the FARTCOIN register that proved durably profitable, suggesting that the holder cohort attracted by such tokens has the high κ the framework rewards. Fourth, the brand is structurally unavailable to any competitor seeking to displace it: there can be no "official" or "licensed" version, because the referent itself precludes ownership.

11.2 Chain Selection

We propose deployment on Ethereum mainnet for three reasons. First, Ethereum is the settlement layer of the deepest on-chain liquidity in cryptocurrency, and the venue on which both PEPE and SHIB executed the most successful launches in the historical record, providing a demonstrated template for the

framework's corner conditions. Second, Ethereum-native memecoins benefit from native composability with the broader DeFi stack, including the major decentralised exchange aggregators and lending venues that absorb organic on-chain volume during the post-launch growth window. Third, Ethereum-resident tokens are subject to the highest evidentiary bar for centralised exchange listing, but conversely, an Ethereum-native memecoin that satisfies that bar accrues the strongest reputational signal in the asset class. We accept the higher absolute transaction-cost environment of Ethereum mainnet relative to alternative chains as the price of this signalling, and note that for memecoin holders, who transact infrequently relative to active DeFi users, per-transaction gas costs are a secondary consideration.

Deployment via direct Uniswap V2 pool creation is recommended in particular, on the model of the PEPE launch of April 2023: the contract is deployed, the entirety of supply is paired with ETH in a fresh Uniswap V2 pool, and the resulting LP tokens are transferred to the standard burn address within the first block following pool creation. The mechanism structurally precludes a team allocation (Section 10) and replicates the most successful Ethereum-native memecoin deployment in the historical record. The use of an established and audited venue rather than a bespoke deployment also reduces execution risk.

11.3 Token Specification

Token name: *Peniscoin*. Ticker: $PENIS. Standard: ERC-20 (Ethereum). Total supply: 999,999,999 (approximately one billion units, matching the FARTCOIN convention rather than the larger trillion-scale supplies of Ethereum-era memecoins). Decimals: 9. Mint authority: revoked at deployment. Freeze authority: revoked at deployment. Initial deployment: via direct Uniswap V2 pool creation, with the entirety of supply paired against ETH and LP tokens burned within the first block following pool creation.

Allocation to team, advisors, presale participants, private-round investors, marketing reserves, treasury, and centralised exchange reserves: zero. Transaction tax: zero. Reflection mechanics: none. Burn mechanics: none beyond the LP burn at migration. Staking: none. Yield: none. The token is intentionally minimal: a tradable claim on a shared narrative and nothing else, on the FARTCOIN model rather than the SHIB-ecosystem model.

11.4 Launch-Day Protocol

The deployment sequence follows the standard Uniswap V2 fair-launch template established by PEPE. (i) Deploy the ERC-20 token contract with the parameters of Section 11.3. (ii) Create a fresh Uniswap V2 pair strictly, an ETH/$PENIS pair, with the entirety of supply on one side and a seed quantity of ETH on the other. (iii) Transfer the resulting LP tokens to the standard burn address (0x000…dEaD) in the same block. (iv) Renounce contract ownership. (v) Publish the contract address, pool address, and burn-transaction hashes on a single static webpage and on the developer's social channel simultaneously. (vi) Cease announcement activity. The launch team posts no roadmap, no price commentary, and no promotional content. The community is, from this point, the project.

There is no presale, no allowlist, no whitelist, no airdrop snapshot, and no "stealth launch" channel offering early access. The first transaction is available to any participant on equal terms.

11.5 Community Identity Marker

Section 3 established that durable communities require an in-group commitment phrase: DOGE's "Do Only Good Everyday," and WIF's "the hat stays on." For $PENIS we propose the phrase *hard forever*, which serves as both a price-trajectory commitment device (the conventional *up only* of memecoin culture)

and an in-group recognition signal that depends on the brand for full legibility. The phrase is bawdy, transmissible, and rhymes with the holding behaviour the framework rewards. A holder posting *hard forever* during a price decline performs the same function as a WIF holder posting "the hat stays on," signalling commitment to other holders by reaffirming the framework that justifies retention.

Recommended community-side initiatives, by analogy to the canonical projects, include: physical placement of the iconographic primitive in publicly photographable locations (cf. WIF's knitted hat on the Charging Bull), charitable fundraising directed at causes thematically resonant with the brand (men's health awareness, prostate cancer research, testicular cancer foundations: the brand creates an unusually direct alignment with established medical philanthropy), and the production of merchandise by independent community members under no central licensing arrangement. The medical-philanthropy angle in particular has the potential to recover any reputational deficit incurred by the ticker, by demonstrating that the community channels its attention toward genuinely beneficial causes.

11.5.1 Communication Channels and Launch-Window Fraud Risk

The community-amplification mechanism specified in Section 4 ($\rho \to 1$) presupposes that holders can identify which community channels are authentic. During the pre-launch and early post-launch window, this identification cannot be performed reliably. Imitation channels using lookalike names, copied iconography, and similar handles are a documented attack vector. Such channels routinely solicit pre-launch contract addresses, distribute fraudulent token approvals, and impersonate the launch team to extract funds from prospective participants.

The framework accordingly recommends that during this window, the launch team designate a single canonical communication channel and announce it across the project's verified surfaces: the official website, the verified social account, and the published research. Holders should presume any other channel to be unaffiliated until the canonical channel itself acknowledges otherwise. After the launch window, when public trading volume and on-chain history have made impersonation easier to detect, the community-amplification mechanism resumes its normal operation: organic community channels, photographable interventions, and decentralised meme production all compound ρ in the manner described in Sections 4 and 11.5.

The framework does not claim that community-operated channels are categorically value-destructive. It claims only that during the launch window, prior to the canonical channel being established and verified, fraud risk dominates the legitimate community-amplification benefit. This is a temporary condition, not a structural one.

11.6 Influence and Listings Strategy

Per Equation (6), the launch team commits to providing no compensation, in token, fiat, or any other form, to any individual or entity in exchange for promotion of $PENIS. This commitment extends to: paid Twitter posts, paid Telegram group placement, paid YouTube content, paid podcast appearances, paid news articles, paid "alpha caller" group inclusions, paid Discord placements, and any analogous arrangements.

Per Equation (7), the launch team commits to declining all paid centralised exchange listing arrangements. The token will be available immediately on Ethereum decentralised venues (Uniswap V2 and the major aggregators). Centralised listings will be accepted only when offered without payment, on the merits of organic on-chain volume: the FARTCOIN sequencing. We acknowledge that the expected timeline to first major centralised listing is plausibly longer for $PENIS than for FARTCOIN, given residual

venue conservatism around the brand. The framework predicts that this lag is recoverable provided on-chain volume reaches sufficient magnitude; it does not predict the lag will be zero. In addition, and per the principle established in Section 7, the launch team commits to engaging no designated market maker on any compensated basis. Liquidity will be provided exclusively by the public Uniswap V2 pool and any subsequent pools created by the community.

11.7 Predicted Component Scores

Table 3 presents the predicted component scores for $PENIS under the proposed design, alongside the canonical project scores from Table 2. The scores are forward-looking and conditional on faithful execution of the design.

Table 3. Predicted scores for $PENIS versus canonical projects.

Project	C	B	D	I	L	V (rel.)
DOGE	0.95	0.90	1.00	0.95	0.95	0.98
SHIB	0.90	0.85	0.95	0.80	0.90	0.92
PEPE	0.85	0.95	1.00	0.85	0.85	0.95
WIF	0.90	0.90	1.00	0.80	0.85	0.93
FARTCOIN	0.90	0.95	1.00	0.90	0.95	0.97
$PENIS	1.00	1.00	1.00	1.00	1.00	1.00

Three observations follow. First, $PENIS is the only project in the canonical set to score at the framework's maximum across every component: C = B = D = I = L = 1.00. The brand achieves perfect legibility $\ell = 1.00$ on the strength of a universal biological referent requiring no provenance, no translation, and no cultural mediation. Distribution achieves the corner condition by construction. Community is predicted to attain maximum conviction κ through the alignment of brand register with the demonstrated FARTCOIN cohort. Influence compounds without payment because the brand is intrinsically discussable. Listings, contrary to earlier framework drafts that imposed a venue-conservatism discount, are predicted to achieve full L = 1.00 once on-chain volume passes the FARTCOIN threshold; the historical record suggests this threshold is binding rather than insurmountable. Second, the resulting aggregate index V = 1.00 places $PENIS as the strict optimum of the framework: the unique memecoin design, among all designs considered in this paper, that satisfies every corner condition simultaneously. Third, $PENIS may accordingly be characterised as *the perfect memecoin* under the framework presented herein: the constructive culmination

of the analytical exercise begun in Section 1, and the reference design against which all subsequent memecoin launches may be evaluated. The framework predicts no design that strictly dominates $PENIS within the five-variable specification of Equation (8); the corner has been reached.

11.8 Failure Modes and Falsification

The proposal is falsifiable in two directions. If $PENIS is launched in accordance with the specification of Sections 11.3–11.6 and fails to achieve a market capitalisation comparable to the canonical projects within twelve months, the framework's claim that the corner conditions are sufficient is weakened. If $PENIS is launched in violation of the specification, with a hidden team allocation, a paid promotional campaign, or a paid listing, and nevertheless succeeds, the necessity claim is weakened.

Three failure modes warrant explicit discussion. First, the L-component risk identified in Section 11.1: should major centralised exchanges decline to list $PENIS at any volume threshold, the aggregate V predicted in Table 3 is overstated. The mitigation is that on-chain volume on Ethereum is itself sufficient to sustain meaningful price action, as PEPE demonstrated during the periods between its centralised listings. Second, the I-component risk: certain influential commentators may decline unpaid endorsement on reputational grounds, lowering the realised $\sum R_i \cdot c_i$ relative to FARTCOIN. The mitigation is that the absolute size of the audience receptive to such tokens is large, and that selective endorsement from a smaller pool of commentators with high c_i within that audience produces equivalent I to broader endorsement from lower-credibility sources. Third, the broader market-shock term ε: any memecoin launch carries substantial idiosyncratic risk, and the proposal is offered as one experiment within a stochastic distribution rather than as a deterministic prediction.

The framework predicts neither a guaranteed success nor a particular timeline. It predicts only that *if* $PENIS achieves durable value, it will do so through the mechanisms enumerated in Sections 3–7; and *if* it does not achieve durable value, the failure will be attributable to the exogenous shock term ε or to the L-component risk discussed above, rather than to a structural defect in the design. The proposal is offered in this spirit: as a clean experiment whose outcome, whatever it is, advances the empirical understanding of memecoin valuation.

12. Conclusions

We have proposed a framework for the design of memecoins that minimises the structural rents historically extracted by presales, paid promoters, and pay-to-list arrangements. The framework is constructed from five reinforcing variables: community, brand, distribution, influence, and listings, whose joint optimisation defines a closed-form expression for fair value. The optimum lies at a corner of the design space in which insider allocation is zero and paid promotion is absent.

Application of the framework to five canonical projects: DOGE, SHIB, PEPE, WIF, and FARTCOIN, demonstrates that each satisfies the corner conditions despite differing substantially in chain, era, and brand register. The convergence is not coincidental. We claim, on the strength of this evidence, that the corner conditions are necessary (though not sufficient) for durable memecoin value. The inclusion of FARTCOIN in particular forced a substantive reformulation of the brand variable: an earlier draft of this framework treated brand provenance as the primary driver of B, but the empirical record demonstrates that brand legibility is at least equally important and is more readily engineered.

The framework is not a guarantee of returns. It does not predict which narrative will capture market attention in any given period, nor does it specify the timing of such capture. What it does provide is a falsifiable test: a memecoin that violates any of the corner conditions enumerated here is, on the framework's prediction, structurally impaired regardless of its short-term price action. We invite empirical replication on a wider panel of launches and refinement of the calibrated exponents in Equation (8).

The deeper claim is that the original peer-to-peer ethic of cryptocurrency: value created and distributed without intermediation by a privileged class, has practical, not merely ideological, implications for memecoin design. A launch that respects this ethic is materially more likely to produce durable value than one that does not. The case studies of DOGE, SHIB, PEPE, WIF, and FARTCOIN, examined here for the first time within a unified analytical framework, are the strongest available evidence for this claim. The *$PENIS* proposal of Section 11 is offered as a constructive test: a memecoin designed to maximise brand legibility, the dimension on which the FARTCOIN case study established that conventional intuitions about acceptable register are systematically miscalibrated. Its eventual performance, will constitute a forward-looking validation of the reformulated model.

References

[1] Nakamoto, S. "Bitcoin: A Peer-to-Peer Electronic Cash System," 2008.

[2] Markus, B. and Palmer, J. Dogecoin Core software, initial release, December 6, 2013.

[3] Ryoshi. "WoofPaper: Shiba Inu Ecosystem," 2020.

[4] Pepe Coin. "Manifesto," pepe.vip, April 2023.

[5] Dogwifhat project documentation, dogwifcoin.org, 2023.

[6] Buterin, V. "Note on the SHIB token burn," Ethereum address 0xd7…, transaction May 17, 2021.

[7] Shiller, R. J. "Narrative Economics." American Economic Review, 107(4), 967–1004, 2017.

[8] Soros, G. The Alchemy of Finance. Simon & Schuster, 1987.

[9] Dawkins, R. The Selfish Gene. Oxford University Press, 1976.

[10] Akerlof, G. A. "The Market for Lemons: Quality Uncertainty and the Market Mechanism." Quarterly Journal of Economics, 84(3), 488–500, 1970.

[11] Spence, M. "Job Market Signaling." Quarterly Journal of Economics, 87(3), 355–374, 1973.

[12] Herfindahl, O. C. Concentration in the U.S. Steel Industry. Doctoral dissertation, Columbia University, 1950.

[13] Adams, H., et al. "Uniswap v3 Core," 2021.

[14] Buterin, V. "On Path Independence," 2017. Available at vitalik.ca.