

Capturing and Distributing Cryptographic Luck: From NFTs and Meme Coins with Inherent Value to Future SuperCurrencies

Tristan Badface
tristan@0xbadface.xyz
<https://0xbadface.xyz>
0xbadface.eth

(First Draft, August 3, 2025)

Abstract

Cryptocurrencies, non-fungible tokens (NFTs), and more recently meme coins have emerged as significant assets in the digital economy. Yet their value is often, rightly or understandably, regarded as purely speculative. We initially propose and demonstrate an extension of the traditional Proof-of-Work (PoW) mechanism, employing two strictly sequential, highly correlated, perceptually significant PoWs to establish reproduction-cost-based inherent value in native, tokenizable blockchain assets.

A proof of concept is implemented on the Ethereum blockchain as an NFT collection `0xBadFace` with contract address

`0xBadFace1EaDA67F194fd202E48E5B2Ca2203c3e2`

created by a smart contract with address

`0xBadFace1AA3AbB02cb2Ce3c11F4EA3B9d428F676`.

The current (as of August 2025) estimated inherent value of such an asset is between 0.5 and 2.8 USD. The estimated required energy input is about 2.8 kWh.

We explore how this concept may be extended to the creation and initial distribution of a new class of digital currencies, discovered through a global, coordinated, yet decentralized search for high-value public cryptographic artifacts — tokenized, exceptionally rare, culturally significant memes embedded in PoWs on blockchains.

We caution all potential investors that the inherent value of such assets is a theoretical peculiarity — an unforeseen, paradoxical artifact of our digital society. This inherent value, independent of the market value, will diminish over time due to technological progress, similar to many physical assets. The inherent value of any such initial NFT collections or fungible tokens mined now or in the foreseeable future may decline very quickly. Investments in these assets should currently be regarded as purely speculative.

Contents

1	Introduction	3
2	Motivations, Previous Work, and Key Innovations	5
2.1	Initial Motivation: An NFT You Cannot Copy	5
2.2	A Greater Motivation: A Meme Coin You Cannot Create (Alone)	7

2.3	Previous Work and Proposed Innovation	8
2.4	Organization of the Paper and Call for Action	9
3	Tokenized Meme-Proof-of-Work	11
3.1	Proof-of-Intent: Adam and Eve	11
3.2	Inherent Value	12
3.3	0xBadFace NFT Collection: Digital Ownership Experiment	15
3.4	First Meme-PoW NFTs and Meme Coins	16
4	Collaborative Tokenized Meme-Proof-of-Work	18
4.1	Deployment Paths and Artifact Creation Models	18
4.2	Effort-and-Luck-Based Distribution	20
4.3	Mining Infrastructure	21
4.4	Need for Low-Monetary Use	22
4.5	Meme Races and The Look-Elsewhere Effect	22
4.6	World of Meme-PoW Currencies	25
5	Conclusion	27
	References	28
A	Mining 0xBadFace Adam and Eve on Ethereum	29
A.1	Mining Adam on Ethereum	29
A.2	Mining Eve on Ethereum	29
A.3	Mining Wallet Address for Etherscan	30
B	Meme Mining Extras	30
B.1	Parallel Meme Mining and GPU Logic	30
B.2	Address Encoding Schemes and Memetic Bandwidth	31
C	Examples of Low-Monetary Use and Utility	32
C.1	Examples of Non-Directly-Financial Uses	32
C.2	Discount Tokens as Practical, Low-Risk Entry Points	32
D	Possible Role of the Lottery Industry	33
D.1	Issues with Modern Digital Lottery	34
D.2	Trustless Digital Casino: A Verifiable Game of Chance	34
D.3	Auditable Game Sessions, Replayable Fairness, and Continued Mining	35
E	Economic and Monetary Considerations	36
E.1	Market Value and Look-Elsewhere Effect Calibration	36
E.2	New Money Without Debt?	36
E.3	Can Fixed-Supply Digital Gold Solve Anything?	37
E.4	Economics of Wasting Money in Luck-Driven Systems	38
E.5	Resolving the Bitcoin Paradox	39
F	Environmental Concerns	40
G	Generalized Lucky Genesis and Multipurpose PoW	41

1 Introduction

Since the genesis of Bitcoin^[1] in 2009, blockchain-based assets have grown into a significant component of the digital economy. Yet over time, the ecosystem has arguably diverged from its original ideals. Instead of serving primarily as decentralized mediums of exchange, many cryptocurrencies function today as speculative investment vehicles — volatile assets without fundamentals, traded in the hopes of extraordinary returns. The landscape has become saturated with hype cycles, insider advantages, and projects that fall short of their decentralization promises or turn out to be outright scams.

At the same time, traditional fiat currencies are undergoing both perceptual and structural shifts. In many societies, physical cash is rapidly vanishing as digital transactions and app-based payments become the norm. Central banks have increasingly relied on monetary policies involving expansive, debt-based money creation. While intended to promote economic stability, these interventions contribute to inflationary pressures and asset bubbles. Rising public debt and growing disparities in wealth distribution have led to increased public skepticism toward central banks and governments, undermining the very foundations of traditional currencies — trust in the issuing authorities and belief in their ability to manage monetary systems in the best public interest.

Modern fiat currencies can be seen as government-issued meme coins — narratives backed by institutional power, yet with no intrinsic or inherent value and without any fixed cap on supply. This institutional power also acts as a gatekeeper, preventing the emergence of new currencies as legal tenders. Meme coins, by contrast, can now be created and traded by anyone, thanks to permissionless blockchain infrastructure. However, blockchains and smart contracts enable creators to fix a token’s total supply immutably, something fiat systems do not technically or politically guarantee. In fact, meme coins are in many ways surprisingly close to the theoretical definition of ideal money: they are divisible, durable, portable, fungible, and impossible to counterfeit. They can be created with fixed supply, require no centralized issuer, and have no intrinsic utility beyond serving as a medium of exchange or store of value¹. Yet despite these properties, they lack authority, legitimacy, and coordination. Anyone can release a meme coin; there are no fundamentals, no central hierarchy, and no constraints on initial distribution.

But even for meme coins with a fixed supply, the challenge is their distribution — and the temptation to retain large allocations or manipulate early market dynamics is hard to resist. Bitcoin succeeded in part because its distribution mechanism was transparent and gradual, tied to proof-of-work mining. But it was also thanks to the apparent altruism of Satoshi Nakamoto, who vanished without claiming any special privilege or early stake.² Few projects since have managed to replicate that level of neutrality, especially after Bitcoin’s price explosion revealed just how valuable early ownership could be.

In this paper, we explore how blockchain technology can be used not only to record value but also to create it through shared discovery. We propose a method for generating and distributing digital assets whose inherent value is grounded in provable computational effort and cultural resonance — backed by large mining communities and game-theoretic principles — rather than driven by speculative hype. Our approach initially extends traditional Proof-of-Work (PoW) to a two-step process searching for digital artifacts with human-recognizable patterns, such as

¹Note that they even cannot be used for example to pay transaction fees on the network, unlike typical cryptocurrencies.

²A more definitive step would have been to provably burn the coins — ensuring they could never be claimed. By simply disappearing without moving the funds, the possibility of Satoshi’s reemergence and future influence remains, however remote.

names or memes embedded directly into cryptographic addresses.

We begin with a small-scale demonstration: the NFT collection 0xBadFace, deployed on the Ethereum blockchain, whose contract and creator (contract) address were both mined to include the same meaningful prefix. This collection, though symbolic in value, serves as a working example of how computational effort can create tokenized artifacts that are rare, verifiable, and non-replicable — attributes typically associated with physical collectibles or historical objects, but now achieved in a purely digital medium.

From this foundation, we outline a potential path toward a new class of digital currencies — currencies that are not issued, but discovered. In this model, value emerges not from a central authority, but from a global, decentralized lottery of effort and luck. Millions of participants may contribute to uncover ultra-rare, culturally significant blockchain artifacts, whose discovery could signal the genesis of new monetary systems — and the opportunity to distribute the newly created wealth at an unprecedented scale. A natural byproduct of such a search would be the emergence of an entire hierarchy of digital artifacts — varying in rarity and significance — most of which could be obtained by the general population not by buying them directly, but by contributing computing power at minimal marginal cost. The market dynamics of such assets could differ radically from those of traditional cryptocurrencies.

The proposed system can be viewed as a generalization of the Bitcoin PoW mining concept. Here, the mined artifacts are not part of a continuous operational blockchain, but rather potential seeds of standalone, newly discovered currencies — with their significance determined by the difficulty and recognizability of the meme embedded in the PoW. Unlike traditional mining, where rewards are claimed by individual discoverers, this system aims to distribute tokenized ownership of each artifact after its genesis. It should do so such that all contributors are rewarded based on their verifiable share of the computational effort, similar to mining pools.

This concept does not promise fast and easy profits or guaranteed early success. It is a proposal for a collective experiment in utilizing the authority of mathematics and cryptography to create new currencies. Currencies that are grounded in computational effort and energy but emerge from a system inherently driven by pure randomness and luck — a long-term, open-ended process through which energy, attention, and computation are gradually transformed into provably rare digital artifacts. These artifacts, shaped by both human intention and chance, can serve as a new foundation for value: one that transcends the limitations of fiat currency by being anchored in verifiable effort, cultural meaning and provable luck.

The proposed concept is not about discovering a single ultimate currency, but about establishing a system ready to recognize and capture such moments whenever they emerge — an infinite game of search, discovery, recognition, and redistribution. In this regard, let us stress that if Bitcoin (or any other existing cryptocurrency) is truly the final answer to digital money, then the opportunity to meaningfully participate — to gain from early adoption — would have occurred only once, over a decade ago. Such a model would imply that future generations will forever be priced out of foundational economic participation in crypto. We argue instead for an ongoing, decentralized process: a system capable of continuously generating value without debt, by discovering new currencies through shared effort and chance. This doesn't invalidate existing systems like Bitcoin or even fiat currencies, but rather suggests they should not be the final word. Economic evolution demands the possibility of rare but recurring transitions — new currencies and wealth emerging when shared effort and luck align.

2 Motivations, Previous Work, and Key Innovations

This section outlines the motivation and core ideas behind the proposal, its relation to earlier cryptoeconomic ideas, and the key innovations that distinguish it from prior work. We begin by identifying limitations in how NFTs currently encode ownership and uniqueness, then introduce the concept of collaboratively discovered tokenized PoW artifacts as an alternative foundation for value. We position this within the lineage of pre-Bitcoin proposals, highlight the cultural and collaborative aspects of our approach, and conclude with a roadmap for the rest of the paper and call for action.

2.1 Initial Motivation: An NFT You Cannot Copy

Despite widespread enthusiasm for tokenized ownership, NFTs often remain disconnected from actual control, rights, or permanence. Two fundamental issues underpin this critique:

- **Off-chain fragility:** The asset an NFT represents is typically stored off-chain. On-chain data usually includes only a reference (such as a URL or IPFS^[2] link)³ to an external resource such as an image or metadata. If that resource disappears or becomes inaccessible, the NFT becomes an ownership record pointing to nothing. This problem also affects NFTs that convey access, membership rights, or real-world asset ownership via centralized platforms or classical legal frameworks — thereby reintroducing trust-based failure modes.
- **Lack of inherent uniqueness:** Public digital content is trivially reproducible. Even when content is stored fully on-chain — as in some procedurally generated art — the data itself can be copied and reissued as a contract with identical or equivalent logic. The only truly unique component is the ownership record inscribed in the blockchain’s ledger.

This raises a fundamental question: what, if anything, truly cannot be duplicated on the blockchain?

On Ethereum, an NFT collection is uniquely identified by its contract address. Each token within the collection has a token ID, which acts as a key in an associative mapping that stores ownership and metadata. Crucially, the contract address itself is globally unique and immutable once deployed — it is deterministically derived from the creator’s address and deployment parameters.

In this sense, the most irreproducible feature of an NFT is not its artwork or metadata, but its *location* in Ethereum’s address space. Much like domain names in the web ecosystem, Ethereum exposes a fixed-length public address space — allocated on a first-come, first-served basis. But unlike domain names, Ethereum addresses are not selected by the user; they are cryptographically derived using a hash function and are effectively pseudo-random.

Standard NFT contracts are deployed to such pseudo-random addresses. But what if the address itself could carry meaning — not merely serve as a pointer, but become part of the message, or even part of the artwork?

This question becomes even more compelling when we shift our focus from NFTs to fungible tokens, especially meme coins. These lack any per-token metadata or images. In such cases, what if the contract address *is* the meme or its textual representation?

³While IPFS ensures that identical content always maps to the same hash, it does not guarantee availability or permanence unless the content is actively pinned by one or more nodes — introducing potential fragility in NFT references.

An Ethereum contract address is a 20-byte (160-bit) value, typically rendered as a 40-character hexadecimal string. For readability and basic error detection, addresses are often shown using the EIP-55 checksum format^[3].

For instance, one of the most iconic NFT collections, CryptoPunks^[4], was deployed at:

`0xb47e3cd837dDF8e4c57F05d70Ab865de6e193BBB`

The hexadecimal format is quite limited in terms of symbols and thus memes it can encode, which paradoxically makes such meaningful patterns more rare.⁴

As a simple example, imagine an address that begins with:

`0xBadFace1...`

Such an address could anchor a themed NFT collection — for example, profile pictures (PFPs) that are visually or conceptually linked to the idea of a “bad face.” This could encompass defiant, awkward, angry, unsettling, humorous, or subversive portrait styles — all exploring the fixed theme.

Obtaining such an address is not trivial. It requires mining. Mining means repeatedly generating candidate addresses until one matches the desired prefix. This is a form of *Proof-of-Work* (PoW): computational effort spent to find an output of a hash function with specific properties. Bitcoin uses PoW based on the SHA-256 hash function and is now mined with specialized hardware (ASICs). Ethereum addresses are instead derived using the Keccak-256 hash function⁵, and Ethereum no longer uses PoW for consensus as of 2022^[5;6]. However, no ASICs currently exist for Keccak-256 address generation, so the search for such addresses still requires general-purpose compute.

Mining such an address consumes energy and time. The longer and more specific the prefix, the harder it becomes to find — and the greater the expected cost. This allows us to define the *inherent value* of such an address as its reproduction cost: the estimated amount of computational effort required to regenerate a given pattern.

As a concrete demonstration, our prototype NFT collection is deployed at:

`0xBadFace1EaDA67F194fd202E48E5B2Ca2203c3e2`

By minting, say, 10,000 tokens from this contract, the address itself becomes tokenized. No other actor can reuse this address. If we treat the address as the asset — rather than metadata or visuals — then it truly cannot be duplicated.

Of course, others may mine addresses that also begin with `0xBadFace1E`, but most of the remaining characters will differ. Each such address must still be mined independently, at a cost that can be estimated from its difficulty. No one gets this pattern “for free.”

⁴Alternative address formats like Base58 allow for a broader range of characters and more expressive encoding. However, extending the character set also increases the difficulty of matching a specific pattern.

⁵`keccak256` is Ethereum’s standard 256-bit cryptographic hash function, based on the Keccak algorithm (winner of the SHA-3 competition). Note: Ethereum’s version differs slightly from NIST’s SHA3-256. See <https://keccak.team/keccak.html> for details.

2.2 A Greater Motivation: A Meme Coin You Cannot Create (Alone)

The situation becomes even more interesting when we move from NFTs to meme coins.

Today, anyone can create a meme coin with any name. While this openness is technically impressive, it has led to a chaotic landscape saturated with opportunism — coins launched in minutes, with no constraints and no accountability.

But now consider a meme coin whose name is not just branding — it is *written into the contract address itself*.

For short names, such addresses can still be discovered by individuals with access to GPU power or time. But as the name grows longer or more specific, the difficulty — and therefore the cost — increases exponentially.

At a certain point, such as an address like:

0xBadFaceOfBeef1...

the cost to reproduce this artifact with current hardware reaches into the millions of dollars (see Section 3.2). At that scale, it becomes economically irrational for any single person — or even a small to medium team — to attempt to create such an address on their own.

This leads to a natural question: could such rare addresses be discovered through collaborative effort?

We propose a model of collaborative mining. Instead of one actor investing millions in compute, millions of individuals could each contribute a small amount — say, \$1 worth of compute. As in Bitcoin mining pools, the contributions of individual miners could be registered and rewarded proportionally.

Importantly, this could be implemented entirely on-chain. Once a desired address is discovered, a smart contract can be deployed at that location. But unlike NFTs, which often represent unique artworks, these high-difficulty addresses are ideal for launching fungible tokens — meme coins.

Such a token could be created with a fixed supply,⁶ then distributed automatically to all contributors based on their recorded mining shares. This entire process could be coordinated via smart contracts — fully automated, transparent, and without privileged actors.

We believe the rarest of these meme coins — those discovered either through vast collaborative effort or extraordinary luck, and distributed across broad participant bases — could function as truly decentralized currencies.

In some cases, the actual cost of discovery may be far below the expected reproduction cost, due to statistical outliers or a wider search space in the search process. This discrepancy is not a flaw but a feature: it represents the system’s core element of *luck*, which plays a central role throughout this paper. It allows valuable artifacts to appear unexpectedly. In this way, computational/cryptographic luck becomes a genuine part of the value creation (and distribution) process and rarity perception.

Unlike standard meme coins, these cannot be created at will. And in cases of a very high reproduction cost, they may remain unique for years — or even decades. And since these currencies would be distributed automatically — with no pre-mines, pre-sales, or central authority — they would avoid many of the exploitative practices typical of meme coin launches.

Until a competing address with the same prefix is found, the memetic part of the address itself could serve as a recognizable identifier — both symbolically and practically. It is easier to remember `0xBadFaceOfDecade` than a random hexadecimal string, and safer to verify.

⁶The actual number of tokens is arbitrary. For intuitive comparison to Bitcoin, one might default to 21 million.

2.3 Previous Work and Proposed Innovation

This work extends several foundational ideas that predate Bitcoin, including Wei Dai’s *b-money*^[7], Nick Szabo’s *BitGold*^[8], and Hal Finney’s *Reusable Proof-of-Work* (RPOW)^[9].

Like *b-money*, it envisions decentralized value creation through verifiable computational effort; like *BitGold*, it treats PoW artifacts as inherently valuable; and like RPOW, it explores how such proofs can become transferable assets.

However, building on the emergence of smart contract blockchains — most notably Ethereum^[10] — this work goes significantly further. It enables the *tokenization of PoW itself*, not as a personal achievement or mining reward, but as a collaborative and culturally meaningful discovery.

Importantly, many of the typical concerns that shaped earlier PoW-based currency proposals are not the focus here. These core infrastructure problems are already solved by the underlying blockchain. Our system operates entirely within such an environment, and can be implemented in a fully trustless way using smart contracts. The base-layer blockchain provides consensus, finality, and verifiability — allowing this work to focus instead on how value can be *created*, not merely transferred.

Here, the PoW artifact is not a block or a coin, but a cryptographic address imbued with memetic significance — a pattern that is human-recognizable, computationally rare, and verifiably the result of costly effort. The resulting tokens derive their value not only from scarcity, but from the symbolic structure of the address, encoded directly in the artifact. These can be minted and distributed trustlessly using smart contracts — embedding both cultural and economic value into a native blockchain object.

It’s worth acknowledging that vanity addresses already exist. These are cryptographic addresses — particularly on blockchains that use Base58 encoding — intentionally mined to contain short human-readable patterns such as project names or token symbols. However, nearly all such addresses are trivial in difficulty: typically three to five characters, easily discovered on consumer-grade hardware.

This is understandable — no rational actor would spend tens or hundreds of thousands of dollars to mine a single vanity address, especially when the only result is symbolic.

What distinguishes our approach is the proposal to mine such addresses collaboratively at high or even extreme levels of difficulty. In this setting, no single actor must bear the full cost. Instead, the mining burden is distributed across a large population of participants — each contributing only a small share of effort. Smart contracts can be used to register these contributions, much like mining pools in traditional cryptocurrency ecosystems. Once a qualifying address is found, the tokenized asset linked to it — whether an NFT collection or a meme coin — can be automatically issued and fairly distributed to contributors.

This framing transforms vanity address mining into a structured, economically grounded, and potentially large-scale system for collaborative digital asset creation.

There is another crucial difference between this work and earlier PoW systems: in most traditional PoW protocols, the “challenge” is predefined and rigid — for example, finding a hash below a certain numeric threshold, often expressed as leading zeros. In our model, the PoW challenge is perceptual and flexible: we search for *some* culturally resonant pattern within the address space. It does not need to be a specific string; it only needs to be meaningful enough to be recognized and valued once found. This flexibility dramatically expands the space of potential artifacts and increases the probability that at least one such very rare discovery will emerge from any sufficiently large mining campaign. However, at the same time, it creates another issue: First, many low to medium difficulty-artifacts are potentially too easy to obtain

at once speculatively. Second, the miners could also attempt to search for too many "imperfect" artifacts eventually with some advanced methods. These concerns arise unless there are at least some constraints replacing the notion of a fixed PoW challenge for the final deployed artifact. These constraints can have different forms or could be ignored. In our demonstration, we decided to include the challenge in another PoW, which is in turn an input of the final contract address generation. This makes the challenge (almost) fixed for the second PoW, resembling more tightly the originally proposed pre-Bitcoin concepts. And in our particular Ethereum implementation, these two PoWs are tightly integrated and can be verified easily by users.

The ideas in this paper are also based on the recent evolution of the meme culture and the resulting emergence of meme coins. One important point of clarification concerns the term "meme." In internet culture, memes are typically visual — images, or animations. But the meme coin trend has shown that, for tokens intended to function as currencies, the meme ultimately reduces to a name or short phrase: "DOGE," "Dog with Hat," or "Shiba Inu," for example. This is the label users see on wallets, trading interfaces, and charts. In that sense, the meme is not an image — it is a compact, memorable linguistic marker.

In our context, then, a meme is best understood as a *name* or *phrase* — one that evokes recognition, cultural associations, or other narrative. It must be sufficiently short to be achievable at least in principle (but not too short - to not be too easy to obtain), sufficiently unique to stand out, and sufficiently resonant to be worth remembering. This framing allows for a direct analogy to fiat currencies: today's money is also just named tokens — "dollars," "euros," "yen" — backed by institutional authority.

Of course, over time, other addresses may be discovered that begin with the same memetic prefix, gradually reducing the perceived rarity of the original. In such cases, the full address will still serve as a full, unique identifier — just as the "United States Dollar" remains distinct from other "dollars." The shorter name might be sufficient in practice, but the full name (address) exists for disambiguation.

Similarly to difficulties of valuation of the PoW attempts in the earlier concepts, a non-trivial question is how users will perceive difficulty when evaluating discovered artifacts. The longer the prefix, the harder it is to mine — a one-character increase in a hexadecimal prefix increases difficulty by a factor of 16; in Base58, by a factor of 58. However, a longer meme may not necessarily be more recognizable or resonant.

This introduces an open question: will communities value raw difficulty, perceptual clarity, or cultural resonance more? How will these factors balance when evaluating new tokens?

The answer remains to be seen — and will likely emerge through real-world experiments with early such PoW meme coins. Difficulty, in this sense, is not a proxy for value, but one input into a more complex social evaluation. However, contrary to pre-Bitcoin ideas, here the PoW is supposed to be a "single rare object" (later tokenized) of a very large difficulty and with its own name. Thus market evaluation is significantly easier than for many small personal PoW attempts.

2.4 Organization of the Paper and Call for Action

The remainder of this paper is structured in two main phases: what we have already demonstrated, and what we hope to enable through community engagement and further development.

The first phase (Section 3) presents a concrete implementation of tokenized, perceptually meaningful Proof-of-Work artifacts on Ethereum. We show how both contract and creator

addresses can be intentionally mined to include memetic prefixes, and how their discovery cost can be estimated and used to define a notion of *inherent value*.

This is formalized through a two-stage mining process in which two sequential, correlated PoWs are used to establish rarity and intent. We demonstrate the concept using a symbolic NFT collection, `0xBadFace`, deployed on Ethereum. All artifacts in this phase were created centrally — by us — without collaborative mining or share registration. This serves as a minimal working proof that such artifacts can exist, be tokenized, and carry reproducible, computationally measurable value.

The second phase (Section 4) explores a more ambitious goal: *collaborative discovery and distribution* of high-difficulty meme artifacts, potentially involving millions of participants. The idea is to enable individuals to contribute small amounts of computing power toward the joint discovery of culturally meaningful blockchain artifacts.

However, the infrastructure for such collaboration does not yet exist. There is no working system for share registration, no reward-tracking contracts, and no finalized protocol for distribution. These components remain to be designed, tested, and built — and doing so will require contributions from developers, researchers, and communities with experience in mining infrastructure, smart contract tooling, and open-source coordination.

We therefore present this paper not only as a technical proposal, but as a call to action.

If this idea resonates, then the next steps must be community-driven. We hope this work will raise awareness, inspire experimentation, and help shape practical implementations: mining share verification, fair distribution schemes, incentive structures, user interfaces, and decentralized coordination models.

The final part of Section 4 speculates on what might follow if this model is widely adopted and eventually extremely rare artifacts emerge, completing a hierarchy of many other such meme coins, currencies, and eventually possibly SuperCurrencies. Could collaboratively mined meme coins evolve into a new class of digital currency — one not issued by central authority, but discovered through randomness and effort? Could such tokens become recognizable not by institutional branding, but by their cultural (and energy) footprint and the story of their genesis?

While these questions remain open, we believe they point to a deeper inquiry: how can currencies be created, distributed, and circulated in a decentralized digital world — and who should participate in their birth and benefit from it?

3 Tokenized Meme-Proof-of-Work

This section introduces our core implementation framework: mining verifiably rare, perceptually meaningful artifacts on Ethereum, and using them as the basis for token issuance. We begin by presenting a novel mining scheme for sequential two-stage artifacts — which we call *Adam and Eve* — designed to embed both effort and intent into the discovered artifacts. We then define a notion of inherent value based on reproduction cost, provide a live demonstration through the 0xBadFace NFT collection, and discuss broader implications for early-stage adoption. Technical details of contract address generation are deferred to Appendix A.

3.1 Proof-of-Intent: Adam and Eve

First of all, let us emphasize that in a strict sense, a hash represents a PoW only if the "challenge" is defined before the mining, which is greatly relaxed in the meme mining concepts.

The idea of mining blockchain artifacts with human-recognizable patterns remains surprisingly underexplored. Vanity addresses — for both wallets⁷ and contracts — do exist, and some are very likely intentional.⁸ Most of them are actually not used on Ethereum, but on blockchains using Base58 to represent addresses. However, they tend to be low in difficulty, rarely central to the associated project, and sometimes indistinguishable from chance.

This raises an important question: can one *prove* that a given artifact was mined deliberately even if its difficulty is relatively low? For such a proof, we need something like pre-defining the PoW challenge we are about to attempt and then doing the PoW as usual.

The Ethereum ecosystem provides a useful feature for this purpose. When viewing NFTs on most blockchain explorers, users not only see the contract address, but also the *creator address* — the address (of potentially another smart contract) from which the contract was deployed. This offers a simple, standardized way to associate a second visible artifact with each contract. This "creator artifact" must be mined first.

Unfortunately, the possibly most well-known Ethereum blockchain explorer, Etherscan^[11], does not conform to this standard. Instead, Etherscan shows always as the "creator" the account which initiated the deployment transaction, not the direct creator address, which might be a smart contract address. With such a setting, the "creator" address⁹ could be mined *after* the final (NFT or token) contract address is mined.

Neglecting the specifics of the Etherscan explorer, this leads us to a two-stage Proof-of-Work structure. We mine both the *creator (contract) address* and the *contract address* to match a desired prefix — for example, both beginning with 0xBadFace1. When these two entities share a recognizable pattern, the resulting artifact becomes perceptually distinct, cryptographically credible, and provably intentional.

The prefix 0xBadFace1 includes 8 hex characters after 0x. Matching this prefix requires 32 bits of difficulty. Because 7 of these characters are alphabetic and we want them to match capitalization under the EIP-55 checksum, the total effective difficulty becomes $32 + 7 = 39$ bits. Mining the creator address (Adam) with this prefix requires roughly $2^{39} \approx 550 \times 10^9$

⁷Wallets tied to private keys cannot be tokenized or transferred in a trustless way — the key remains forever bound to the original creator.

⁸We have identified three Ethereum contract addresses beginning with 0xbadface, but none include the digit 1, nor do they use the intended capitalization. They also lack any special creator address.

⁹Naming both options as "contract creator" is probably confusing. Ideally, Etherscan should show both the creator and the deployer, where deployer will always be an externally-owned-account with an associated private key, currently shown as "creator".

attempts. Once Adam artifact is found, the search for the Eve artifact can begin. This second PoW process has the same difficulty in our case — bringing the total expected effort to about $2^{39} + 2^{39} \approx 1 \times 10^{12}$ hash attempts.

The chance of observing both artifacts in sequence by accident is $1/2^{78} \approx 1/3 \times 10^{23}$ — a probability of flipping 78 heads in a row. Such a double structure serves not only as a scarce digital object, but as a *proof-of-intent* — a clear, verifiable signal that the resulting artifact was generated deliberately.

We refer to these paired artifacts as *Adam* and *Eve*. Adam is the creator — typically a smart contract factory. Eve is the final NFT or token contract, deployed deterministically using Ethereum’s `CREATE2` opcode. Full derivation paths on Ethereum are described in Appendix A. Together, they form a traceable, tamper-proof record of effort and authorship. A single PoW result might appear random or arbitrary, but when a second address in the same structure shares the same symbolic pattern, intent becomes “undeniable”.

This design also introduces a valuable constraint: in a single-stage PoW system, a miner could scan broadly for any pattern of interest, producing many low-difficulty artifacts in parallel. But in our two-stage model, a second, independent PoW is required for each Eve artifact, creating a bottleneck that defends against uncontrolled inflation of perceived value. Even if a miner discovers many interesting Adam candidates, they must still invest proportional effort to complete their Eve counterparts.

Finally, this dual-PoW setup enables a new mode of decentralized competition. Anyone can search for interesting Adam artifacts — with no predefined notion of what counts as “valuable.” Communities can scan broadly, flag candidate memes or visual patterns, and launch public searches to complete their corresponding Eve contracts. In this way, the system becomes a platform not just for artifact generation, but for cultural exploration and symbolic coordination.

3.2 Inherent Value

We define the *inherent value* of a mined artifact as the expected cost of reproducing an equivalent result under current conditions (with current hardware and algorithms). This includes verifiable inputs such as computational efficiency, energy consumption, and hardware amortization. The reproduction cost works as a kind of SPAM protection — once some artifact is obtained, sometimes very cheaply, reproduction of its exact prefix requires the full inherent value investment on average.

In our case, because computing resources are widely accessible, any individual could in principle remine an artifact with a matching prefix. However, doing so may take a very long time and cost a lot of money, depending on the difficulty and available hashrate. This makes (re)production of such artifacts fully democratic, yet inherently constrained.

Calculation

To estimate the reproduction cost depending on the artifact difficulty, let D denote the total difficulty in bits (i.e., the average number of hashes needed is 2^D). We define the following input parameters:

- H — hashrate of the mining machine (hashes per second)
- P — power consumption (watts)
- c_e — electricity cost (USD per kWh)

- c_h — hardware cost of the mining rig (USD)
- L — total assumed lifespan of the hardware (in hours)

The expected time T (in hours) required to mine an artifact of difficulty D is:

$$T = \frac{2^D}{H \cdot 3600}$$

The estimated inherent value V (in USD) is then calculated as:

$$V = T \cdot \left(\frac{P}{1000} \cdot c_e + \frac{c_h}{L} \right)$$

The first term represents energy consumption, and the second captures hardware amortization — the fraction of the device’s value used up during the mining process.

Estimation of Inherent Value on Ethereum

For practical estimation, we used an implementation based on NVIDIA CUDA running on a consumer-grade notebook NVIDIA GPU and consider the following measured or estimated input values for our setup or their ranges:

- Power draw: approximately 80W
- Hashrate: 200 MH/s (Adam), 500 MH/s (Eve), 400 MH/s (“Eve creator address”) see Appendix A for the mining paths
- Hardware cost: \$200–\$400
- Lifespan: 2–4 years
- Electricity price: \$0.05–\$0.40 per kWh

Note the different hashrates for Adam and Eve, resulting from different derivation paths (see Appendix A). With these inputs, mining a 44-bit PoW for 0xBadFace1A and 0xBadFace1E (9-character hex prefix with 8 checksum-sensitive letters) takes approximately:

$$T_{\text{Adam}} \approx 24 \text{ hours}, \quad T_{\text{Eve}} \approx 10 \text{ hours}$$

leading to an estimated cost of \$0.4–\$2 for the Adam artifact and \$0.1–\$0.8 for the Eve artifact, that is \$0.5–\$2.8 in total. Regarding the energy consumption, one would expect to need about 2 kWh for Adam and 0.8 kWh for Eve on average for our hashing efficiency, giving total expected energy input of 2.8 kWh.

The “Eve deployer address” is a standard wallet address used to deploy our final NFT contract, see also next Section 3.3. We mine it such that the Etherscan explorer shows a memetic address starting with 0xBadFace1B as the “creator”, which is slightly non-standard and only introduced for full user experience in our project. The associated mining cost are similar to the Eve derivation path, but we neglect them in all our total estimations, as this is more a visual quirk than a cryptographically relevant element.

Using a mid-range estimate of \$0.5 for Eve as a 44-bit artifact using the Ethereum `CREATE2` opcode (see Appendix A.2), which will be probably used preferably due to its higher mining efficiency, we extrapolate reproduction costs (and expected energy inputs) for longer, more difficult prefixes:

- 0xBadFaceOfBeef (64 bits): approximately \$0.5 million (840 MWh)
- 0xBadFaceOfBeef1 (68 bits): approximately \$8 million (13 GWh)
- 0xBadFaceOfBeef1E (73 bits): approximately \$270 million (430 GWh)

These estimates assume a single artifact, unchanged hardware efficiency and no parallel scanning advantage - that is one has to target this specific artifact. Due to the stochastic nature of PoW, any specific mining attempt may succeed earlier or later than average.

Production vs Reproduction Cost

It is also important to distinguish between the cost of producing an artifact and the cost of reproducing it.

The *reproduction cost* — as used in our definition of inherent value — refers to the expected energy and hardware expenditure required to discover an equivalent artifact (with the same prefix) again through brute-force mining. This cost can be estimated objectively using public mining benchmarks and hardware and electricity prices.

However, the *production cost* — the actual resources consumed in creating a specific artifact — may differ significantly. While some aspects of production effort may be inferred from mining logs or time stamps, the actual cost can be much more variable and difficult to estimate. In addition, it stays constant, unlike the reproduction cost, which decreases (read below).

One important nuance arises in the Adam phase. Because a single PoW search can scan for many potential memes simultaneously, the discoverer may find multiple viable artifacts from the same computational run at no increased cost. In such cases, we risk "overcounting" the value of each result. In cases where both Adam and Eve share the same prefix, the system naturally avoids such kind of "inflation", see also Appendix E.1 for short discussion of market calibration of this aspect. Eves appear roughly in proportion to total effort, while Adams emerge at a higher rate due to broad scanning.

If another actor wants to obtain the same artifact class, they must invest the full reproduction cost — which will almost certainly require brute-force search from scratch. This, in turn, incentivizes wide, open-ended searches over narrow brute-force replication. Rare artifacts cannot be easily spoofed or cloned — scammers have to expect to pay full price.

Evolution of Inherent Value

The inherent value of these digital artifacts mirrors something rare in the cryptoeconomic world: like physical objects, their inherent value depreciates with time — not through decay, but through the technological progress.

It should be emphasized that our cost estimates are only a first-order approximation. They are based on a single GPU model and a straightforward CUDA implementation, which may not be fully optimized. Future optimizations — even modest improvements in mining efficiency — could shift this estimate significantly. If the idea gains traction, it's plausible that specialized hardware such as FPGAs or even ASICs will be developed to further significantly accelerate this specific mining path by several orders of magnitude.

However, such developments do not break the system — they simply shift the ground truth. If a single actor attempts to use private newly developed highly efficient hardware to hoard rare artifacts, their overproduction will undermine the perceived rarity, and the community can respond by adjusting its notion of value accordingly. In effect, the cost function must be

re-evaluated in light of the new technological baseline. For the system to remain credible, new hardware must be made accessible (likely by renting it) to the broader miner base that sustains the search. In this way, the decentralized ethos can be preserved even in the face of accelerating mining efficiency.

Any such algorithmic or technological progress will reduce the inherent value — even for artifacts that have already been discovered. This is not necessarily a flaw. Since the market value is influenced not only by fundamentals but also by scarcity and demand, a declining inherent value for a currency with a fixed supply may actually help stabilize its market value and promote circulation. This stands in contrast to the typical investment mindset — where assets are hoarded in hopes of appreciation — a behavior fundamentally at odds with any asset aspiring to function as a medium of exchange.

Inflation-like spending pressure is thus instead formally achieved through the persistent possibility of appearance of new, more valuable currencies—not by increasing the supply of existing ones. Unlike for the “eternal” Bitcoin narrative, here we expect “better” currencies will eventually appear and none of them will dominate forever.

3.3 0xBadFace NFT Collection: Digital Ownership Experiment

As mentioned before, we selected for our demonstration one of the simplest and most flexible meme concepts — one with broad potential for interpretations. The Adam artifact is mined as the first contract address deployed from a new wallet address. The resulting contract address (creator address)

`0xBadFace1AA3AbB02cb2Ce3c11F4EA3B9d428F676`

was used to deploy a factory contract: This contract was then used to deploy the final NFT collection based on the ERC-721 standard (using the `CREATE2` opcode) at the Eve artifact address

`0xBadFace1EaDA67F194fd202E48E5B2Ca2203c3e2`

To show a memetic address as a “creator” also on the Etherscan explorer, we mined also the wallet address used to submit the Eve deployment transaction to the contract at the Adam artifact address. This “Eve creator address” is

`0xBadFace1B0E5d06F4BF9d8fa92B75794fd955781`

The NFT collection at the address `0xBadFace1EaDA67F194fd202E48E5B2Ca2203c3e2` is created with a maximum supply of 10,000 tokens. These will be minted by the creator or another entity with minting rights and sold in batches via auctions to their new owners.

But, what kind of ownership is it, if the only actions available to the holder are transfer or burn? In line with our broader exploration of what truly makes NFTs unique, we chose to challenge a common assumption: that an NFT asset must always be immutable.

Instead, the 0xBadFace collection is envisioned as a shared, public, evolving digital gallery. Each NFT represents a frame in this gallery, initially populated with an image and meta-data inspired by the “bad face” theme. While each owner can propose an updated version, new submissions are expected to remain consistent with the spirit of the original or previous version.¹⁰

¹⁰For example, if the original artwork depicted a pixel-style eagle with a “bad face,” a new proposal might offer an improved pixel version of the same or another eagle, or reinterpret it in a different visual style — as long as the theme and subject are preserved. Exceptions could be made in rare cases, with community approval, which should ultimately become the final arbiter.

This collection also more fully exploits a property often underutilized in NFTs: that the asset is not merely represented on-chain — it *is* on-chain. The contract address itself is a mined artifact, and the list of token owners exists directly in the blockchain state under this address. Ownership here is not symbolic; it is technical and literal.

Building on this, we introduced a user-editable metadata field for each token — a form of personal inscription. Each token owner can modify their field freely using a designated `setUserURI` function. This field is not simply a metadata proposal, but a persistent on-chain message — a personal extension of ownership. It is not shown as the default `tokenURI`, but remains publicly accessible via Etherscan and other block explorers.

Crucially, only the current owner of a token can propose a change to its main metadata. The designated *approver* — initially the creator — can only accept or reject that proposal. This separation of roles is encoded in the contract. It ensures that the power to initiate changes lies exclusively with the holder, while the approver acts merely as a confirmer, not a controller.

Over time, we expect the approval process to be handed over to the community — for example, through token-based voting or DAO-based governance. Even if the metadata becomes frozen by community decision, individual holders can continue updating their user-defined fields. These always remain editable, giving every holder a personal blockchain space tied directly to their token.

The result is a hybrid form of NFT ownership: partially personal, partially communal, and inherently dynamic. Rather than freezing meaning at the moment of minting, the 0xBadFace collection allows ongoing reinterpretation — a living memetic artifact shaped by those who hold it.

Because we did not want to use generative art, but rather allow individual artists to contribute most of the initial images for the 0xBadFace NFT collection, we decided to cap the collection at 10,000 tokens. If the project gains popularity, the initial release through auctions may make our original demonstration artifact inaccessible to many early followers. Fortunately, the concept is fully democratic: anyone can have their own 0xBadFace—just not 0xBadFace1EaDA67F194fd202E48E5B2Ca2203c3e2.

3.4 First Meme-PoW NFTs and Meme Coins

Anyone with a GPU can generate similar artifacts like our 0xBadFace demonstration within hours or days. With a CPU, it may take longer—possibly weeks or a month—but still remains within reach. As a result, we expect a wave of similar low-difficulty artifacts created by gamers and others with access to basic GPU power. These may be used to release new NFT collections or meme coins—often in a more centralized manner. Such easily created digital assets can be later distributed freely, experimentally, or according to the creator’s intentions.

Focusing on NFTs in particular, many of these lower-difficulty tokens could be deployed on Ethereum Layer-2 solutions or other low-fee blockchains. This would allow them to reach a wider audience—including people without GPU access or technical experience¹¹. Such tokens could still find a market, especially if offered at mint for a low price. Interestingly, this sets up a new kind of economic opportunity: if a miner spends hundreds or thousands of dollars in computing resources to mine such an artifact, they can still profit by distributing many low-cost NFTs to many buyers. The buyers would be people who like the meme and value owning something provably scarce and effort-backed. Such buyers also might not have access to GPU

¹¹Over time, we expect most technical barriers to disappear through user-friendly apps and interfaces, making it even easier to participate.

compute or do not want to wait for the address to appear on a CPU. Or they might want to save the environment by not producing their own artifact - so they buy a tokenized version of an existing one. Such a small "business model" can operate without needing to "exploit" buyers through hype or artificial scarcity.

While this kind of centralization is not necessarily problematic in the case of low-value NFTs—where artistic input or creative authorship adds unique value to the meme—for memecoins, the situation is quite different. If memecoins are minted or sold directly for real money (e.g., ETH), we risk recreating the same exploitative dynamics seen in today's speculative markets. But a meme coin is something more—it is, in essence, a precursor to a currency. And if we want to shift away from viewing these tokens as investment assets and instead treat them as mediums of exchange, we cannot rely on purchasing as the main method of acquisition. Just like most people earn their national currency through work, these tokens must be earned—through effort—if they are to function fairly and meaningfully.

The situation changes dramatically, however, when we begin to consider artifacts whose inherent value could reach millions of dollars or more. In such cases, centralized creation becomes risky—since a single entity may invest enormous resources into producing a high-value artifact, only to face catastrophic losses if no one accepts or desires "their" object.

4 Collaborative Tokenized Meme-Proof-of-Work

In this section, we outline steps necessary to move from small-scale, single-player artifact discovery to large-scale, collaborative mining of high-difficulty meme-based digital assets, focusing on the Ethereum ecosystem. We first examine improved deployment paths — avoiding mined private keys in favor of deterministic contract factories using `CREATE2` — which allow for secure artifact creation while enabling future token issuance from either Adam or Eve contracts. We then discuss effort-based share registration, enabling millions of participants to contribute compute power toward a shared discovery goal and receive proportional rewards if a valuable artifact emerges. We discuss the potential mining infrastructure, possible uses for many of these newly created tokens, probabilistic and luck-based aspects of open searches and finally try to give a generic overview of the potential landscape and hierarchy of these new currencies.

4.1 Deployment Paths and Artifact Creation Models

The mining approach used in the 0xBadFace demonstration — where a private key is mined to produce a specific first contract address (Adam) — is conceptually simple but not optimal. Similar concerns arise for mining of wallet addresses (our "Eve creator address"), as we have used to have a memetic "creator" address also on the Ethrescan explorer. While sufficient for small-scale, symbolic artifacts, these methods become problematic in high-value settings.

First, these methods require miners to generate and manage private keys directly, which poses a security risk in e.g. cloud environment. The mined address must deploy the artifact contract and may also receive the discoverer's token allocation. Second, the computation required to mine such keys is less efficient than using `CREATE2`-based deployment. The derivation from private key to address involves public key recovery and (up to two sequential) Keccak-256 hash(es), resulting in roughly $1.25\text{--}2.5\times$ the cost per hash attempt compared to `CREATE2`-based mining.

To address these issues, we propose three alternative deployment architectures that avoid private key mining and allow more flexible, secure, and deterministic artifact creation. All rely on `CREATE2`. In all models, the deployed bytecode — for Adam and optionally Eve — should *not* be supplied dynamically by users but is instead *pre-loaded into a deployment factory* as immutable code. Only constructor parameters can be variable, and these become part of the `CREATE2` inputs (contract constructor parameters) along the salt (see Appendix A.2) — ensuring that every mining input (including the miner wallet address which prevents any potential front-running submission attacks) is cryptographically bound to the resulting artifact and verifiable by anyone.

1. Predefined Nested Factories with Optional Delay and Eve Integration

A general-purpose deployment system can begin with a top-level factory contract (Level 1) deployed at a fixed address. This contract logs submitted mining shares and deterministically deploys second-level factories (Level 2), each corresponding to a mined Adam artifact.

Each Level-2 factory is initialized with a standard token contract bytecode — for example, a dedicated ERC-20-based implementation — and minimal constructor arguments such as miner wallet.

Crucially, the Level-2 Adam contract may allow to delay actual token deployment. Instead, it may allow to wait for a qualifying Eve — mined with the same or similar prefix

— to be discovered. If such an Eve artifact appears, a Level-3 contract can be deployed from it (using another fixed bytecode). Tokens may then be issued from the Eve address, with ownership or share attribution handled via pre-registered contributions in the Level-1 contract.

If no suitable Eve is found, the Adam contract might allow to fall back to a direct token deployment using its own address. This design supports a wide range of release strategies, while maintaining a unified registry and common infrastructure. This would probably not be automatic, but controlled by the Adam creator, who decides whether to deploy a sole Adam artifact or wait for its (matching) Eve.

2. Semi-Centralized Adam, Community-Mined Eve

Another practical approach is for a well-known figure or organization—such as a public authority, celebrity, or influencer—to mine a low-difficulty Adam artifact and deploy the deterministic deploy factory. The factory includes a fixed contract bytecode, and expects constructor parameters for the Eve discoverer address and a `CREATE2` salt. This setup allows some Eve artifact (possibly in a meme race, see Section 4.5), which can be of significantly higher difficulty, to be mined openly by the community.

This method offers a compromise: the Adam phase is relatively easy and can be done by someone publicly trusted, while the Eve phase—more difficult and prestigious—is left to an open global race, without complications of collaborative two-stage full-difficulty artifact.

3. Single-Stage Artifacts

Finally, the simplest option is to ignore the Adam+Eve dual artifacts and simply deploy a generic mining share registration and deployment factory. Anyone could register shares or deploy any resulting address as a meme coin. Similarly to other schemes, these new meme coins would not actually automatically appear in miners wallets, but instead would require some kind of claim mechanism allowing for share validation using aggregated share registry state fingerprints in the smart contract. So arguably people would only try to claim "reasonable" meme coins automatically issued to them via such mechanism. This social filter could then remove most of the "imperfect" or "random" submissions to only keep such currencies whose names (significant part of their contract addresses) can be remembered easily and resonate culturally.

Such a mechanism might also be more viable for blockchains not compatible with Ethereum (but possibly using Base58 for address encoding), where the multi-stage artifacts could be more difficult to implement or would not be so easily seen in blockchain explorers. But in case of very efficient advanced search methods and such open contracts, the number of appearing "imperfect meme currencies" to "socially filter" could be quite overwhelming. Initially, these could be limited to predefined sets of memes implementing the proposed meme races. But allowing any discovery is of course attractive - what if?

Ultimately, multiple deployment paths will likely coexist, shaped by evolving standards, available tooling, and collective experimentation. A downside is that the search might become quite fragmented, making truly wide distribution complicated, unless people unite under some popular generic contracts on specific blockchains, possibly allowing more expressiveness than hexadecimal addresses, for example using Base58 for the memes.

4.2 Effort-and-Luck-Based Distribution

Blockchains are often celebrated for their trustlessness — their ability to enable open, verifiable interactions without the need for trusted intermediaries. But while the infrastructure is transparent and decentralized, the *distribution* of value within these systems has often been anything but fair. Early adopters, insiders, and technically advantaged actors frequently capture disproportionate rewards. Even so-called “fair launches” tend to favor those with faster access, deeper knowledge, or superior infrastructure. This system aims to bridge this gap.¹²

If these new meme coins are to evolve into meaningful digital currencies — rather than remain speculative tokens — their initial distribution must be fair and based on measurable effort. In particular, if we want to break free from the financial logic of “buying in early” to extract gains from later participants, we cannot rely on direct purchase as the primary method of acquisition. Instead, meme coins should be *earned* — through demonstrable contribution of computing power.

Much like how national currencies are distributed through labor or service, we propose that meme coins should be earned based on recorded computational work. Participants contribute hashpower to a shared mining effort, often at negligible personal cost — a few cents or dollars worth of electricity or device wear — but with the potential for upside if the community succeeds in discovering a valuable meme artifact. Of course, the discoverer(s) should get a special reward of some fraction of the total supply of the new currency.

Unlike traditional mining pools that forward rewards to a centralized operator, this system would use smart contracts to register and verify work transparently and trustlessly. Tools and infrastructure from decentralized mining research, in particular SmartPool^[12] — show that such systems are not only feasible, but have been field-tested. In our case, the requirements are even simpler: there is no need to solve blocks or track chain consensus, only to verify that valid shares were submitted. In addition, at least on Ethereum, one can easily make the miner wallet as a smart contract constructor parameter and in turn input of the PoW attempt, linking it immutably to a potential discovery, or any of the (much more frequent) mining shares.

To make this process efficient, each participant could self-select a difficulty level according to their hashrate. When they discover a valid share — such as constructor inputs resulting in contract address beginning with a sufficient number of zeros (e.g., 0x0000...) — they submit it to a (possibly Layer-2) reward-tracking smart contract. Unlike standard shares used in mining pools, which have the same weight if over some small difficulty threshold, here the shares may also be weighted by their difficulty, giving more value to less likely discoveries. This gamifies the mining process and rewards both effort and luck.

Because on-chain storage and gas fees are non-trivial, the frequency of submissions must be carefully managed. Classical mining pools aim to reduce statistical variance by submitting shares frequently. Here, by contrast, each submission incurs real-world cost — so participants are incentivized to reduce their submission rate.

Importantly, there is no need to record the specific pattern a miner was targeting — the effort can be generic. This aligns well with meme races where the winning artifact is not known in advance or there is a predefined set. Once a meme coin is discovered — that is, a sufficiently rare contract address is mined and successfully deployed — the token supply can be minted and distributed proportionally to all registered miners. A small fraction of the total supply should be allocated to the discoverer and encoded in the distribution contract ahead of time.

This creates a system where effort is recorded and rewarded, but no one knows in advance

¹²See Appendix E.4 for economic implications of mining-based rewards and its potential effects on this potential new wealth distribution.

which pattern will win and when. Miners compete and cooperate simultaneously, trusting that if something rare is discovered, their contributions will be honored. Even participants with minimal hardware — such as smartphones — can take part, submitting one valid share every day or week. While their expected return is small, the upside is real, and the broadness of distribution increases the cultural legitimacy and viral potential of the resulting token.

Importantly, the system reduces the typical advantages of early participants. As hardware and algorithms improve, the cost of computing each share drops. With new miners joining, earlier shares become diluted. To maintain their proportional stake in an ongoing search, users should periodically mine new shares under current efficiency conditions. Participants are not rewarded indefinitely for being early, but for not leaving the game.

Over time, we expect ecosystems of mining tools, UIs, and simplified participation apps to emerge, lowering barriers further. With sufficient scale and optimization, the cost of running such systems — in terms of smart contract fees and share verification — can remain well below the expected value of discovered artifacts. The key challenge is scaling: batching, off-chain aggregation (as in SmartPool), and Layer-2 optimization are essential, but the fundamental model seems viable.

4.3 Mining Infrastructure

While participants may use their own hardware — smartphones, laptops and PCs (with GPUs) — the combined compute power of casual users is inherently limited. Large-scale mining will likely be performed by commercial infrastructure: data centers, cloud GPU providers, or specialized operators offering time-rented hashpower. If dedicated ASICs are developed, they should again be mostly rented to be profitable. But unlike traditional mining pools, the goal here is not to secure a blockchain or earn predictable block rewards. Instead, users are searching for perceptually meaningful, culturally resonant, and computationally rare artifacts — which may or may not be discovered.

The key economic difference is that this system does not require decentralization of compute itself — only decentralization of *production cost*. Each miner, whether an individual or an infrastructure provider, bears their own cost of attempting discovery. And since ownership is linked to the wallet that requested the computation (and is input of the PoW), not to the machine performing it, compute can be rented, shared, or pooled with minimal trust assumptions.

This shifts mining dynamics toward open, permissionless access: users can rent an hour of GPU time for a dollar, search for a meme artifact or submit shares, and disconnect. Mining may resemble the casual, burst-based participation seen in Bitcoin’s early days — not continuous hashrate competition, but periodic contribution.

As such mining becomes more widespread, this infrastructure could be monetized similarly to renting hash power or even online gambling. In this way, meme mining could evolve into a large-scale system securing online gambling as a by-product — with public randomness, auditable outcomes, and community-wide rewards triggered by successful discoveries, beyond standard rewards financed by increased fees compared to non-lottery providers. Appendices D–D.3 explore how this model may naturally intersect with the lottery and online gambling industry — potentially reframing such activities as provably fair, community-driven events with real digital value.

4.4 Need for Low-Monetary Use

For cryptocurrencies or token-based systems to evolve into meaningful mediums of exchange, they must find real, practical uses beyond speculation. The most significant of those then could start working as universal mediums of exchange and global currencies. But in our proposal, we assume anyone can create a currency, just likely not a significant one. The hierarchy of such low and medium market cap tokens would prepare the ground for potential discoveries of exceptional artifacts, created and distributed globally.

One of the core requirements for widespread adoption and circulation is the development of *non-financial* or *low-financial* utilities — applications where tokens serve a purpose other than storing or transferring monetary value, or where the value involved is generally negligible.¹³ As a result, typical issues like market volatility do not pose significant problems for circulation for such low and medium-value tokens.

These uses create social, cultural, or practical value, helping people understand and interact with blockchain-based assets without needing to engage in financial risk or speculation. Such systems naturally build trust, familiarity, and everyday presence in people’s lives — essential prerequisites for any emerging exchange medium.

Of course, these tokens will operate in the same space as existing meme coins or NFTs, and nothing can prevent people from trading them — not just for similar meme tokens, but also for other crypto-assets or fiat. Similarly, one cannot prevent exchanges from listing such tokens. Arguably, some of these tokens may become pegged to traditional markets, and their post-issuance dynamics and demand could drive their effective market value upward — potentially calibrating the value of other artifacts, where the inherent cost of reproduction becomes a reference point for comparing different meme coins (without relying on less rigorous valuation methods).

To support broader use, many such tokens may find application in everyday, low-risk scenarios such as thank-you tokens, personal or community badges, and school or team recognition coins, or as lightweight local currencies within events and interest groups. Discount tokens represent a particularly practical model — functioning as digital vouchers that offer minor benefits or slightly reduce the purchasing price of goods or services, without behaving like fully fungible currencies. These examples are discussed further in Appendices C.1–C.2.

4.5 Meme Races and The Look-Elsewhere Effect

Searching for a single, specific PoW pattern is generally inefficient. Because the mining logic itself involves computationally expensive hashing operations, each candidate requires thousands of machine instructions to process. Once a hash or address is produced, it must be (in the simplest case) compared against a list of candidate patterns to determine whether it qualifies as a potential “discovery.” Fortunately, checking for simple prefix matches (e.g., the first few bytes of the address) is a lightweight operation, often requiring only a couple of instructions per comparison. This makes it feasible to test each mining output against a large number of candidate prefixes, allowing for the simultaneous search for many meme patterns. We put a more detailed discussion of advanced search methods and utilization of different address encoding schemes into Appendices B.1 and B.2.

¹³Some projects could attract millions of miners and result in assets with multi-million dollar inherent value. In such cases, one can expect that, on average, each miner will be effectively rewarded with only a few dollars’ worth of tokenized inherent value — which could, however, be a typical amount held by individual participants.




Figure 1: The probability to mine a given pattern as a function of multiples of the expected search time.

Probability to Find a Specific Meme

Let T be the estimated time required to discover a specific meme artifact (see Sec. 3.2) at specific position (usually as a prefix), based on previous benchmarks or the mining capacity and mining path in question. Assuming the discovery follows a Poisson process, the actual time to discovery is exponentially distributed with mean time T . The probability that the meme is discovered within a given time t is:

$$P_{\text{found}}(t) = 1 - e^{-t/T} \quad (1)$$

For example, the chance of finding the meme within the expected time is:

$$P_{\text{found}}(T) = 1 - e^{-1} \approx 63.2\% \quad (2)$$

as can be seen from the plot of the probability in Figure 1, showing the probability to find the meme as a function of multiples of the expected time (assuming constant hashrate). To reach a higher confidence, such as 95%, one would need to mine longer:

$$t_{95\%} \approx 3T \quad (3)$$

So there is still a 5% probability the meme was not found, despite three times more hashes than expected were attempted.

Probability to Find Some Meme: The Look-Elsewhere Effect

In most mining scenarios, participants will arguably not focus on a single predefined meme. Instead, they will typically search for a broad set of candidate patterns simultaneously, possibly hundreds while still avoiding a significant slowdown of the mining. This increases the chance of a discovery and reflects the open-ended nature of cultural value: we set up the stage for the memes to compete and let pure randomness and luck to choose among them.

Let T again denote the expected discovery time for a single specific meme. Assuming that we are (effectively) mining for n independent meme patterns in parallel and each meme has a roughly equal expected discovery time T , the probability of finding at least one of them during time t becomes:

$$P_{\text{found}}(t; n) = 1 - (e^{-t/T})^n = 1 - e^{-nt/T} \quad (4)$$

In other words, searching for n distinct memes in parallel is statistically equivalent to searching for one meme, but with n times the hash rate—or, alternatively, n times the amount of time compressed.

This formulation captures the essence of what is known in statistics and physics as the *look-elsewhere effect*, albeit for a fixed list, in a much constrained scenario. When searching across many possibilities, even highly improbable outcomes become likely—simply because of the vast number of chances we give them to occur. As more memes are included in the mining process, the probability of discovering *some* match rises accordingly, even though the odds for any individual meme remain unchanged.

This dynamic transforms meme mining into a kind of probabilistic cultural competition: we do not know in advance which pattern will be chosen by the system, but we know that with enough participants and time, *some* pattern will almost certainly emerge earlier than expected from the difficulty—sometimes surprisingly early. While this effect enhances the odds of discovery, it also complicates the interpretation of the *luck factor* associated with any individual artifact. To convince observers that a meme has “won” the race by chance and not by manipulation (e.g., a concentrated GPU effort targeting a single prefix), the artifact must be sufficiently expensive and difficult to reproduce. Only then can the underlying randomness be trusted on game-theoretic grounds, and the outcome be accepted as genuine proof of computational luck.

It is worth emphasizing how economically unreasonable such attempts to artificially manipulate luck would be. If the eventual distribution of the resulting meme-based tokens will be tied directly to recorded and publicly verifiable mining efforts, any secretive mining operation aimed at artificially boosting a particular meme would remain invisible to the network—and thus unrewarded. The “attacker” could only benefit through the (uncertain) discoverer’s reward, which should typically represent only a small fraction of the total token supply. Therefore, attempting such manipulation would require a substantial expenditure on computational resources, all without any proportionate financial return.

Meme Races

The simultaneous search for multiple memes creates a unique economic situation. Suppose a community wishes to discover an artifact with an inherent value of approximately 1 million USD. Instead of searching for a single pattern, the community could compile a list of, for example, 100 candidate memes—each of similar difficulty—and begin mining across the entire set. In this broader task, the effort required is substantially reduced: rather than needing one million participants each contributing \$1 of compute, only 10,000 miners contributing the same amount are needed to “ensure” that at least one of the memes is likely to be found. The result could be an artifact with a \$1 million reproduction cost but only \$10,000 actual production cost.

The search may stop after the first match or continue to produce an ordered list of “winners”—from the luckiest discovery down to the most brute-forced patterns. Importantly, the inherent value of an artifact is not diminished by the fact that it was discovered quickly. In our framework, value is defined by the reproduction cost, not the actual cost of discovery. The

difference between the two becomes a quantitative measure of the luck embedded into the artifact. Publicly recorded mining efforts further validate the credibility of the search process. In this system, cheap production through luck doesn't undermine value—because reproduction is still expensive.

For inexpensive artifacts like `0xBadFace`, there will be many similar-looking artifacts with matching prefixes. To interact with the correct smart contract, one must often remember/check additional characters from the address—typically a random-looking suffix. Over time, we can expect intentionally created “spam” artifacts to appear, attempting to mimic the extended parts of the addresses of well-known or valuable memes. Of course, there will be also unintentional such discoveries, arguably growing with time if the search is sustained (but moves to more difficult artifacts in general).

However, for higher-difficulty artifacts, the situation changes. When an artifact is sufficiently rare—relative to the realistically available mining power—it becomes likely that only a single matching address exists for an extended period. In such cases, the meme itself becomes a unique identifier for the artifact, at least for some time. Because of the look-elsewhere effect, repeatedly mining over the same set of memes will almost certainly lead to discovering a different meme.

4.6 World of Meme-PoW Currencies

We have only outlined a general framework for discovering and distributing digital artifacts through perceptually significant Proof-of-Work. While many technical proposals have been made, the broader picture remains open-ended and fundamentally exploratory. What, then, is the lifecycle of such artifacts? How might they evolve from small experiments to widely accepted currencies? And could these currencies have actual function in the economy?

We suggest that the progression unfolds along a continuum—from *TestCoins*, through *Community Coins*, and eventually toward *Currencies* and even potential *SuperCurrencies*. This spectrum reflects not only technical and economic maturity, but also levels of cultural coordination, memetic resonance, and infrastructural support.

Early efforts will likely take the form of lightweight meme tokens discovered by individuals or small teams with popular prefixes existing thousands of times. These *TestCoins* may target amusing or low-difficulty prefixes. Given their likely high numbers, two-stage artifacts might be more credible for non-collaborative addresses. But we think these would be mostly used for themed NFT collections, managed by their creators. Some of the initial ones, possibly using a given low to medium difficulty meme first, could eventually gain historical recognition and become collectibles. In a form of memecoins, these might be distributed centrally (and mostly freely) by their creators for example to their fans or supporters¹⁴

As infrastructure matures, artifact discovery will shift toward collaborative protocols using deterministic contract deployment and public mining registries. The first meme races could produce the first *Community Coins* with inherent value exceeding millions of dollars and many thousands to possibly millions of contributors. Naturally, the landscape will be fragmented—at first. Competing share registries or contract templates will coexist across Layer1 and Layer2 blockchains. Encoding formats (hex, base58) will vary. And new platforms may emerge with dedicated support for artifact-first discovery systems. The Ethereum ecosystem remains a strong candidate due to its tooling and composability, but other or future systems may offer more flexible implementations.

¹⁴For low-monetary, symbolic, or social applications of such tokens, see Appendices C.1–C.2.

These efforts may employ meme races, variants of two-stage mining (Adam–Eve), or other experimental paths. Successful Community Coins should demonstrate:

- A medium to high-difficulty memetic pattern with aesthetic or symbolic appeal,
- A contract design that transparently embeds fairness or redistribution,
- And a community effort that anchors legitimacy through shared discovery.

Some Community Coins may achieve wide enough recognition and economic footprint to function as true *Currencies*. These are likely to emerge from particularly improbable discoveries coupled with a publicly recorded search history and wide cooperation. Here, value stems not just from mining cost, but from the legitimacy and coordination surrounding the artifact. Importantly, the supply of each such currency is fixed at discovery, but the overall ecosystem can expand as new artifacts are found. This enables the emergence of a growing family of rare fixed-supply assets. As discussed in Appendix E, these currencies—unlike mined tokens used for network fees—could be seen as autonomous monetary primitives: not backed by debt, but emerging from cryptographic improbability and shared authorship and backed by energy, similar to Bitcoin.

Their use as mediums of exchange, stores of value, or collectibles will ultimately depend not on design, but on chance, perception, social uptake, and potential legal recognition.

Should the system reach global scale and become efficient enough, a few artifacts may attain memetic and economic dominance sufficient to function as *SuperCurrencies* — global universal mediums of exchange. However, it is expected that more valuable artifacts will be discovered eventually, overtaking the (likely growing) market share of all crypto. The search becomes an infinite game: not to win, but to keep playing, occasionally winning a jackpot as a civilization and being repaid our investments. Such a never-ending scheme, especially considering energy consumption of the PoW and its environmental impact (see Appendix F), only makes sense when the whole civilization could profit from it.

The proposed scheme utilizing existing blockchains might be too fragmented for such a global purpose, and might require generalization of the concept to achieve a truly efficient search. We discuss briefly some aspects of such generalizations in Appendix G. The basic idea is that the meme PoW can be performed directly on block hashes of a "non-existing" blockchain, eventually arriving at a genesis "block sequence" resulting in a block with a memetic hash of extreme difficulty. True SuperCurrencies might require either such generalizations or global convergence of the effort under possibly a single contract and blockchain. A true SuperCurrency would likely, beyond eventual market cap possibly even in trillions of dollars, exhibit:

- Extremely high PoW difficulty - possibly being a true far outlier,
- A crisp, compelling meme or phrase recognized globally,
- Extremely broad token distribution through share registration,
- And long-term persistence, unlikely to be surpassed for years or decades

Importantly, the concept does not target any particular or even a "final" SuperCurrency. Thanks to the look-elsewhere effect and its potential future generalization to multi-stage searches (see Appendix G), the aim is to simply scan for *some* culturally resonant artifact among many potential candidates. This allows the reach of the search to exceed naive expectations and in very unlikely but possible cases produce artifacts whose PoW reproduction cost (or even

effective difficulty) and extremely broad and fair distribution could rival or even exceed that of Bitcoin (for some time). There is no inherent reason to rely on a single, centralized narrative of scarcity (see also Appendix E.5).

5 Conclusion

We have proposed a new class of digital artifacts utilized as currencies whose value is derived not from issuance by a central authority, but from discovery through provable computational effort. These artifacts, at least initially (ideally) constructed through sequential correlated Meme-Proof-of-Work, offer a transparent, verifiable link between energy, intention, and rarity. However, the concept is flexible and once it transitions to community mining, the integrated additional PoW can be simply replaced by public knowledge of the intention. So in practice, community consensus and later solid market evaluation may emerge even around single-stage artifacts, particularly in the early phases of experimentation, until a potential "inflation" of many "imperfect" artifacts would occur. By that time, some of the first high difficulty artifacts could already gain significant adoption, overshadowing many existing meme coin projects.

This collaborative approach marks a significant shift. It moves value creation away from arbitrary issuance and speculative hype toward a decentralized currency genesis process anchored in measurable effort and randomness.

By aligning incentives around discovery, contribution, and recognition rather than control or early access, this model opens a path to democratized value creation. It decentralizes not only currency issuance, but also the very authority to define what counts as valuable. The memetic nature of these artifacts reinforces their symbolic power while enabling broader community involvement in their emergence and adoption. We believe, the greatest barrier to crypto adoption is not technological. It is the suspicion that the game is rigged. By designing a system where no one—no miner, no investor, no institution—can hold an early advantage, we remove that suspicion.

Of course, the system is not without challenges. The environmental cost of PoW, the speculative tendencies of crypto markets, likely fragmentation of the search, and the technological limits of current blockchains—such as high fees and scalability constraints—remain significant concerns. We note, however, that many of these issues are being actively addressed across the blockchain ecosystem. If scalable and low-cost infrastructures mature, the proposed discovery mechanisms could be run at global scale with minimal resource requirements per user.

We emphasize the importance of the bootstrapping phase. Only by experimenting—through community-driven deployments, token distributions, and open mining efforts—can we begin to understand how such a system behaves economically and socially. The idea carries long-term implications, potentially reshaping how currencies emerge and circulate. But its trajectory will depend on countless interactions, technical iterations, and emergent behaviors. Any predictions beyond this are speculative.

Given the foundational questions this concept touches—about value, fairness, and coordination—we invite researchers, developers, economists, and broader communities to engage. This is not a polished product, but a proposal for a distributed experiment in creating new forms of value—grounded in energy, effort, chance, and cultural meaning. Whether this approach becomes a serious component of future economies or remains a thought-provoking experiment of the present will depend on what people choose to do next.

References

- [1] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. <https://bitcoin.org/bitcoin.pdf>, 2008. White paper, October 2008; accessed 28 Jun 2025.
- [2] Protocol Labs. Interplanetary file system (ipfs). <https://ipfs.tech/>, 2025. Official website; accessed 18 Jul 2025.
- [3] Vitalik Buterin and Alex Van de Sande. Erc-55: Mixed-case checksum address encoding. <https://eips.ethereum.org/EIPS/eip-55>, 2016. Ethereum Improvement Proposal No. 55, January 2016; accessed 28 Jun 2025.
- [4] Larva Labs. Cryptopunks. <https://cryptopunks.app>, 2025. Official homepage of the CryptoPunks NFT collection; accessed 18 Jul 2025.
- [5] Ethereum Foundation. Ethereum merge upgrade. <https://ethereum.org/en/upgrades/merge/>, 2022. Documentation of Ethereum’s transition to Proof-of-Stake; accessed 18 Jul 2025.
- [6] Ethereum Foundation. Proof-of-stake (pos). <https://ethereum.org/en/developers/docs/consensus-mechanisms/pos/>, September 2024. accessed 11 Jul 2025.
- [7] Wei Dai. b-money. <http://www.weidai.com/bmoney.txt>, 1998. Posted to the Cypherpunks mailing list, November 1998; accessed 28 Jun 2025.
- [8] Nick Szabo. Bit gold. <https://unenumerated.blogspot.com/2005/12/bit-gold.html>, 2005. Blog post, December 2005; accessed 28 Jun 2025.
- [9] Hal Finney. Reusable proofs of work (rpow). <https://nakamotoinstitute.org/finney/rpow/>, 2004. Project white paper and code release, August 2004; accessed 28 Jun 2025.
- [10] Vitalik Buterin. Ethereum: A next-generation smart contract and decentralized application platform. https://www.weusecoins.com/assets/pdf/library/Ethereum_white_paper-a_next_generation_smart_contract_and_decentralized_application_platform-vitalik-buterin.pdf, 2014. White paper, originally published November 2013 – revised 2014; accessed 28 Jun 2025.
- [11] Etherscan. Ethereum block explorer. <https://etherscan.io>, 2025. Block explorer for Ethereum blockchain; accessed 18 Jul 2025.
- [12] Loi Luu, Yaron Velner, Jason Teutsch, and Prateek Saxena. Smartpool: Practical decentralized pooled mining. In *26th USENIX Security Symposium (USENIX Security 17)*, pages 1409–1426, 2017.
- [13] Kristóf Poduszló. Introduction to provably fair gaming algorithms (5th draft). <https://cryptogambling.org/whitepapers/provably-fair-algorithms.pdf>, July 2017. White paper (draft), Cryptogambling.org; accessed 28 Jun 2025.

A Mining 0xBadFace Adam and Eve on Ethereum

This appendix describes the technical process used to mine the 0xBadFace Adam and Eve artifacts on Ethereum, including how contract addresses can be deterministically controlled.

A.1 Mining Adam on Ethereum

To avoid a scheme where the total difficulty is multiplicative for our two PoWs, we rely on the `CREATE2` opcode in the second step — when deploying Eve. However, `CREATE2` deployments must be initiated by a *smart contract*, not a regular externally owned account (EOA).¹⁵ Therefore, if we want to mine the *creator address* — that is, the Adam artifact — the simplest process begins at the level of a private key. Specifically, we search for a private key such that the *first* contract deployed from its associated account address becomes the Adam artifact, which will then be used to deploy Eve via `CREATE2`.

In Ethereum’s standard `CREATE` path, the address of the n -th contract deployed from a wallet is computed using the Keccak-256 hash function as ¹⁶ :

$$\text{Adam} = \text{keccak256}(\text{rlp_encode}([\text{wallet_address}, n]))[-20:]$$

where $[-20:]$ means the last 20 bytes (last 40 characters in hex). For the first contract (where $n = 0$), this address is entirely determined by the wallet address. The wallet address itself is derived from the private key as:

$$\text{wallet_address} = \text{keccak256}(\text{public_key}(\text{private_key}))[-20:]$$

By varying the private key, one can search for EOAs whose *first deployed contract* (via `CREATE`) ends up at a desirable address.¹⁷ That contract can then be used as the deployer for subsequent `CREATE2` deployments.

However, as can be seen above, this mining path requires not only public key generation from a generated private key, but also twice hashing with Keccak-256. This makes this particular path more computationally difficult than the one using `CREATE2` as used below. Fortunately, it is easy to simply add another layer of smart contracts and deploy a deterministic deployment factory at a normal pseudo-random address and only from this factory deploy the original deterministic deploy factory, only now already at a memetic address. This way, one can entirely avoid mining private keys directly, which would not be secure on a remote infrastructure. As this approach is slightly more advanced, we leave it for future implementations, but generally recommend it.

A.2 Mining Eve on Ethereum

The `CREATE2` opcode in Ethereum enables deterministic smart contract deployment: the resulting contract address is independent of the transaction count and instead uses the contract

¹⁵An externally owned account (EOA) is a standard Ethereum account controlled by a private key, typically used by individuals to sign transactions and hold assets. Unlike smart contracts, EOAs contain no code and cannot execute logic on their own.

¹⁶`rlp_encode` refers to Recursive Length Prefix encoding, the serialization method used by Ethereum for lists and data structures. See the Ethereum Yellow Paper or <https://eth.wiki/fundamentals/rlp> for details.

¹⁷Private key search should always use high-quality entropy sources and refresh entropy regularly to avoid biased or predictable key material. Deterministic or low-entropy generation methods pose serious security risks and can lead to loss of funds or unintended key collisions.

bytecode as one of its inputs. The address generated by `CREATE2` is defined as:

$$\text{Eve} = \text{keccak256}(\text{0xff} \parallel \text{deployer} \parallel \text{salt} \parallel \text{keccak256}(\text{init_code}))[-20:]$$

where `0xff` is a constant prefix byte, `deployer` is the address of the contract performing the deployment (i.e., `Adam`), `salt` is a user-specified 32-byte value, and `init_code` is the contract’s creation bytecode. Changing the salt from some initial random sequence¹⁸ makes it possible to *mine* contract addresses.

Importantly, in this scheme, the contract code must be fixed before mining Eve. Once Eve is discovered, the resulting address is bound to that exact contract bytecode — it cannot be changed. While Ethereum supports upgradeable patterns like proxy contracts, this immutability may in fact be *desirable*, as it limits the potential influence of the discoverers. In such cases, only a small set of predefined parameters may remain adjustable, as contract constructor parameters or settable by the discoverers of Adam and Eve just after deployment and prior to full activation and token distribution.

Regarding the influence of the discoverers, there is a significant asymmetry in our most simple scheme. Once Adam is discovered, the owner of the EOA has full control over what is deployed at the Adam address — they are not required to follow any predefined rules, nor to wait for Eve’s discovery under contract conditions they may not control. This is not a problem for individually mined artifacts, but for a collaborative effort, the community might expect rewards for participation to be guaranteed by the code and not entirely rely on the good will (and possibly abilities) of the discoverer, despite being chosen in a perfectly decentralized way.

A.3 Mining Wallet Address for Etherscan

As we mentioned in the main text, the possibly most popular Ethereum blockchain explorer, Etherscan^[11], shows as the “creator” the wallet address initiating contract deployment, not the creating address of a possibly another smart contract. So specifically for this explorer, we have mined another artifact - but in this case not a tokenizable one. The wallet can only be used by the author and can never be shared. The process to mine a memetic wallet address on Ethereum is basically a simplified Adam generation described in Appendix A.1, we look for a private key such that

$$\text{wallet_address} = \text{keccak256}(\text{public_key}(\text{private_key}))[-20:]$$

starts with the desired prefix.

B Meme Mining Extras

B.1 Parallel Meme Mining and GPU Logic

In practice, mining for a single meme prefix is inefficient. Our mining of the `0xBadFace1A` artifact did not target only that exact string, but searched simultaneously for several variants such as `0xDadFace1A`, `0xBabeFace1A`, and others. Each mining attempt was checked against a predefined list of such candidates, hardcoded into the mining logic for performance. This

¹⁸A sufficiently high-entropy random input is required to ensure that multiple miners do not inadvertently explore the same portion of the address space. Then such a high-entropy input can be simply increased by one in each hashing attempt as usual.

hardcoding enabled compiler-level optimizations, reducing branching and improving memory layout — important factors for high-performance GPU mining.

The address generation process involves computationally expensive step of hashing. But prefix checks are typically just a few instructions and can be compiled efficiently for GPU execution. As a result, hundreds or possibly even thousands of candidate patterns can be evaluated per hash without significantly affecting the mining rate.

Our implementation did not check for EIP-55 checksum casing on the GPU level (which would require another round of Keccak-256 hashing). Instead, any match in raw characters (case-insensitive) was flagged and sent to the CPU for postprocessing, where capitalization was verified. This architecture — GPU for raw scan, CPU for final selection — improves the search efficiency.

As a side note, scanning the full address or hash for matches (i.e., looking for a meme anywhere in the byte string, not just at the start) would require significantly more processing, especially on GPU architectures. In most cases, these “misaligned” discoveries would be missed unless specifically scanned for, even though they might still hold value—especially if extremely rare. In practice, however, it is probably more efficient to prioritize the discovery of “perfect” artifacts—those that begin with a meaningful prefix.

But this search can be even more efficient than the primitive version looking for a fixed list of candidates. For example, one could implement a multi-stage trigger system, similar to those used in high-energy physics experiments: fast, low-level filters running directly on mining GPUs could flag candidate addresses, which are then passed to slower, high-level evaluation — possibly using machine-learning techniques. It is plausible advanced methods will be developed, producing many “imperfect” candidates, for example instead of `FightFightFight...` meme address, candidates like `FightFX6htFight...`, `Fight9ightFight...` and many others could still be found and of course appear much more often than perfect phrases. At such point, this would require heavy social filtering of all possible candidate currencies, taking into account the full perceptual rarity and difficulty. Similar problem will likely already appear with the early hexadecimal meme coins, especially variants with “wrong capitalization”, like `0xBadFAce0fBeef...` instead of `0xBadFace0fBeef...` Unless some other (human) filter would exist, the market would have to deal with all such variants and weight their “value” in a broader context - but hopefully “perfect” artifacts, if discovered, would stand out and get recognized appropriately.

B.2 Address Encoding Schemes and Memetic Bandwidth

Most of this paper focuses on Ethereum and its Layer 2 systems, which use hexadecimal address representations. Hex encoding limits expressivity to 16 characters (0–9, a–f), reducing the set of recognizable or culturally resonant memes.

Alternative encodings offer greater expressivity, but also higher difficulty per character:

- **Base58:** Used in Bitcoin and some altcoins. Provides a richer character set and avoids visually confusing characters. However, encoding is complex and GPU-unfriendly due to modular arithmetic, padding, and Base58Check checksums.
- **Base64:** Offers high expressivity, using all uppercase and lowercase letters, digits, and two special symbols (+, /). It allows fast prefix comparisons and case-sensitive matches.

Selecting an encoding scheme involves trade-offs between expressivity (memetic bandwidth) and mining efficiency. For systems that value cultural and symbolic density, Base64 may represent a sweet spot between computational feasibility and linguistic freedom.

However, in this work we emphasize the advantage of implementing PoW meme coins on existing blockchains. So we might need to just use whatever can be implemented. On the other hand, representing Ethereum addresses in Base58 or Base64 is in principle only a detail of the user interface - in the end, all these addresses are simply big numbers and the encodings are just their shorter (in number of characters) representations.

C Examples of Low-Monetary Use and Utility

This appendix expands on the examples briefly discussed in the main text, illustrating how individually or collaboratively mined tokens can find meaningful, low-risk use cases beyond speculation. These include symbolic, social, and community-driven applications that encourage circulation and engagement without requiring significant monetary value.

This is of course mainly important because there will many low to medium difficulty meme coins eventually, many re-using the same (popular) memes, so they would not likely work as universal global currencies and could have only small market caps. However, because people will not buy them for money, but also will not get them "for free" (had to mine with their GPU and pay electricity bill, or paid a small fee to some cloud GPU or ASICs provider), their circulation will likely differ from typical meme coins distributed via pre-sales. Given that typical miner could receive only a couple of dollars of effective value for many smaller and medium coins, these would be best used for various micro-transactions.

C.1 Examples of Non-Directly-Financial Uses

Tokens can act as proofs of participation, recognition, or shared culture. Some example applications include:

- **Thank-you tokens** — digital gestures of appreciation or gratitude, possibly extending the notion of likes to socially distributed low- and medium-cap meme coins.
- **Friend or family coins** — personal tokens used privately or socially in small groups
- **School or team tokens** — issued to reward a small voluntary work, mark achievements or participation, etc. in larger groups
- **Local or community currencies** — usable within specific events, spaces, or interest groups, possibly with global reach

Importantly, these types of tokens are not primarily bought or sold in a speculative market — they are earned, gifted, or traded for similar tokens in playful or meaningful contexts. This greatly reduces their appeal to bad actors and makes the ecosystem much less exploitative than the current state of the memecoin market.

C.2 Discount Tokens as Practical, Low-Risk Entry Points

One particularly practical use case is that of **discount tokens**. These are tokens accepted by businesses, communities, or individuals — such as cafes or online shops — in exchange for a small discount on a purchase or service. For example, a coffee shop might support the token `0xCafe0fDecade...`, and offer a 5% discount to any customer who trades them this token at checkout.

The business then *keeps* the token. Over time, they may reuse it for loyalty purposes, gift it to regulars, or trade it for another token supported by a partnering business. In this way, the token circulates socially, helping spread awareness about the business and connecting customers through shared interests or memes. Note that during each transaction, the memetic address appears in the blockchain receipts, duplicating and propagating the meme over the blockchain.

This model might also solve a modern problem: today, each business runs its own loyalty or discount program, often requiring account registration, privacy compromises, or app downloads. In contrast, this system allows for spontaneous participation. If a customer doesn't hold the needed token, they can *quickly trade* it for another token — using smart contracts — without requiring trust or intermediaries. Many loyalty points are not used. If these could be "traded" for a very low cost or in exchange for another loyalty tokens, it could attract new customers and improve the efficiency of the loyalty programs, allowing to share excess points outside peoples' direct social circles - fully online in a trustless way. Note that such loyalty points already work like "unofficial currencies", but usually in closed proprietary or non-digital ecosystems.

A key insight supporting the "not-directly-financial" perception of these tokens is the *non-linear relationship between token quantity and value*. For example, some amount of a particular token may entitle the holder to a one-time 5% discount — but offering more tokens does not multiply the benefit. The discount does not stack. As such, the token functions more like a *ticket* or a *pass* than a divisible unit of money, while still allowing divisibility.

This non-linearity shifts how people perceive the token. It no longer behaves like traditional currency, but more like a symbol of shared culture, access, or community engagement. Its value is contextual, not entirely market-driven despite being available on the free market. Of course, eventually high value tokens could escape this paradigm and start really working as currencies, but we think it would be useful to have the full hierarchy of such tokens, while "true" currencies would be simply the most significant ones.

These low and medium value tokens offer a safe and engaging introduction to crypto for individuals and communities. They reward participation and cultural connection, not speculation. And they give businesses new tools for loyalty, visibility, and community interaction — all without needing to manage complex infrastructure.

On the path to discovering or developing new forms of money, *non-financial and not-directly-financial utility* may be the critical bridge between curiosity and real-world usage. It could bring crypto into people's lives as something useful, personal and to some extent familiar — not just something hyped on social networks to bet on with real money.

D Possible Role of the Lottery Industry

We explore how the meme mining can be used for an additional purpose: the verifiable proof-of-work mechanism could transform digital gambling and lotteries into transparent, auditable systems—with fairness enforced not by trust, but by computation. In addition, next to classical rewards, one could also discover a new currency during a compatible game of chance, distributing the otherwise individual luck to millions while getting one of the largest stakes¹⁹.

¹⁹The discoverer reward should be probably smaller than 1% of the total supply - larger rewards could be seen as unfair and the community acceptance is fundamental for the currency legitimacy.

D.1 Issues with Modern Digital Lottery

Most modern digital gambling systems—whether online casinos, app-based games of chance, or video lottery terminals—rely on pseudorandom number generators (PRNGs) to determine game outcomes. These systems typically begin with a secure entropy source, such as hardware randomness, user interaction timing, or system-level secure random APIs. In most cases, the initial seed can be considered a sufficiently strong source of randomness for cryptographic or gaming purposes.

However, the problem lies not in the seed itself, but in the deterministic nature of the PRNG progression that follows. Once the seed is set, the sequence of outcomes is fully determined and—importantly—can be simulated or predicted by the system operator (or any actor with access to the internal state) for free. This creates a significant asymmetry: the player sees only a stream of seemingly random results, while the operator can fully predict or simulate future outcomes at no cost.

This enables subtle forms of front-running or silent manipulation, such as:

- Silently re-seeding games to avoid large payouts.
- Ending or aborting games when a win is close.
- Offering bonus rounds only when the outcomes are known to be bad for the player.

In reality, none of this may matter to the average player, since there is no practical way to prove whether a result was genuinely random or dynamically adjusted. In regulated physical casinos or official lotteries, trust is placed in institutional frameworks—independent audits, legal accountability, and human oversight. But in online environments with little external regulation, players cannot easily verify fairness.

D.2 Trustless Digital Casino: A Verifiable Game of Chance

The concept of a trustless lottery or digital casino can be built directly on the foundations introduced in this paper. While the basics are not new^[13], the integration with meme-currency discovery extends the model far beyond any individual game. The core principle is simple: the outcome of each round is derived from a deterministic hashing procedure seeded by entropy contributions from both the player and the operator. This ensures unpredictability before the game, and full verifiability afterwards.

To initiate a round, the player submits a personal entropy seed and other mining inputs and selects the search algorithm used at the provider, possibly from a public repository. The casino commits to its own entropy via a hash, which is revealed only after the game ends.²⁰ Optionally, a third entropy source (e.g., from a public regulator) may also be included. A combined seed is then formed:

$$\text{seed} = \text{Hash}(\text{player_entropy} \parallel \text{casino_entropy} \parallel \text{regulator_entropy}) \quad (5)$$

This seed allows to define a deterministic hash sequence. To be specific, for Ethereum address mining of the Eve artifact:

$$\text{Eve} = \text{keccak256}(\text{0xff} \parallel \text{Adam} \parallel \text{salt}=\text{seed} + n \parallel \text{keccak256}(\text{init_code}))[-20 :]$$

²⁰The casino commits in advance by publishing a hash of its entropy.

where `Adam` is the address (potentially memetic) of some deployment factory. The `init_code` contains the Eve contract bytecode and ABI-encoded constructor parameters, in particular the miner wallet:

$$\text{init_code} = \text{Eve_creation_code} \parallel \text{ABI_encode}(\text{miner_wallet}) \quad (6)$$

The contract bytecode can be immutably set in the factory, expecting as a constructor argument e.g. only the miner wallet. For deployment transaction, this argument shall be provided together with "winning" seed. The deployed contract can be configured to accept for example setting the token name and symbol after deployment (by the owner of the "winning" wallet). The miner would only provide fixed inputs like the full hashed `init_code` and his part of the seed to be combined with casino and possibly regulator entropy and then incremented as usual.

Importantly, a win is not based on a single draw but on potentially trillions of hashes. The casino must calibrate its pricing to mining efficiency, using the PoW output to define winning probabilities per time frame or hash batch.²¹

If the casino attempts to suppress a win, anyone can recompute the sequence and verify that a reward condition was met. Conversely, the player cannot precompute outcomes, since the casino's entropy was unknown until after the game. This symmetry enforces fairness through design.

Here, Proof-of-Work is not used for consensus or difficulty adjustment, but to prevent low-cost prediction and manipulation. Traditional systems allow cheap forward simulation, enabling real-time strategic abuse. In contrast, this model requires the casino to compute each outcome honestly, with no shortcuts or pre-filtering. Any attempt to manipulate outcomes—such as offering to switch games if a win is near—requires the same costly computation as honest play.

Casinos fund or mass-rent infrastructure (e.g., GPUs, later possibly ASICs) and finance it through player fees. The reward structure is linked to the bet size via proportional hashing. Larger bets increase hash rate, not just payout size—ensuring fairness without exploitable asymmetries. Because all outcomes must be computed by PoW, selective manipulation becomes computationally and economically irrational, especially in a competitive market environment of such providers.

Once demonstrated and understood, this model could shift player expectations. Cryptographically provable fairness might become standard—just as encrypted payments did in e-commerce. Such a technical improvement could trigger cultural and regulatory evolution.

D.3 Auditable Game Sessions, Replayable Fairness, and Continued Mining

Every game session can be replayed deterministically using known inputs: the player's wallet address, player-supplied entropy, and the casino's revealed seed. Once fixed, the entire hash sequence is known, requiring no new randomness. This allows third-party verification of wins, payouts, and early exits.

Players can also explore "what if" simulations to see how a session might have continued. This creates a uniquely transparent gambling experience, where randomness is a journey, not a black box. While the original bet is over, players can extend the sequence to search for meme PoW artifacts, continuing to register mining shares anchored to the original session.²²

²¹When mining shares are submitted to meme mining contracts, they may be withheld until the game concludes to avoid revealing casino entropy.

²²We assume shares will be considered registered at time of submission, not generation

Unlike traditional games—where session end often coincides with a financial loss—this model lets users keep participating, even if at reduced reward potential. Note that users can always continue on their phone, with much lower hashrates, but only needing to charge it - the rewards would not be directly financial, but would have such potential. Such design may reduce harmful compulsive behavior while preserving long-term engagement.

Sessions can be independently verified by players or regulators, turning fairness enforcement into a public, transparent process. Every outcome is replayable. The wallet-based ID model allows artifact ownership without disclosing private keys—supporting secure, non-custodial mining.

For multi-player games or jackpots, additional entropy is needed. Without it, a malicious player could collude with the provider to replay known sequences. To prevent this, external certified randomness must be introduced.

E Economic and Monetary Considerations

E.1 Market Value and Look-Elsewhere Effect Calibration

We defined the inherent value of an artifact as the reproduction cost. However, for single-stage artifacts, it is possible to search for many address prefixes simultaneously. It is difficult to predict how efficient this search may become. With advanced multi-stage triggers, eventually one will be able to discover thousands of relevant artifacts during a single pass up to a certain difficulty threshold. Initially, this will probably not be a problem, as the artifacts will be mostly curiosities and their total market value would not be too significant. But eventually, we would see an “inflation” of single-stage artifacts due to the improved search efficiency (even at constant hashrate) and parallel mining attempts for already existing artifact prefixes. We believe the market will naturally calibrate this effect, simply from observing how many artifacts, with what difficulties (and hash functions or derivation paths), appear. The inherent value will likely be generally disconnected from the market value due to memetic and cultural potential factors and the level of adoption, but eventually, these currencies will be competing among each other for some fraction of the economy using them. It will not be possible to estimate their value in isolation from inherent-value-related fundamentals, but only in the broader market context.

Simply summing the reproduction costs would result in inflated or largely “imaginary” valuations—as reproduction and production costs can diverge significantly for single-stage artifacts. Around such time, the double-stage artifacts, where the Adam and Eve prefixes will be (ideally) identical, might gain more relevance, where only the inherent value associated with the Eve artifact might be “counted” (and reflected in the market capitalization). If even at the second-stage artifact we could mine for multiple artifacts simultaneously (using several different Adam artifacts as inputs, see Appendix G), the calibration of the resulting look-elsewhere effect might again be necessary. In this case, the number of simultaneously attempted prefixes could be publicly known from mining share registration, unlike in the case of potentially various mining software implementations looking for very different prefix patterns for the Adam artifacts, only minimally revealed in case of a published discovery.

E.2 New Money Without Debt?

In the modern economy, basically all money is created with a corresponding debt obligation, which in addition includes interest to be paid. This mechanism can have different forms im-

plemented by central banks or licensed commercial banking institutions utilizing the fractional reserve banking mechanism. Even without a growing economy, this system tends to create a “shortage” of money (though this is more complicated due to the velocity of money entering the considerations) which must be continuously created to cover the interest obligations. In such a system, a slight inflation can be (among other functions) seen as a mechanism allowing to counterbalance the burden of interest (and debt in general). Unlike many, we do not believe such a system is inherently flawed, but admit our worries about its long-term consequences. At the same time, we believe it is often not sufficiently recognized in the crypto community how often and on what scale market interventions and monetary operations need to be executed by central banks to keep the value of fiat currencies stable. It should be recognized that we really need “money printing” to keep up with the growing economy. If no new money were produced, the economy could soon collapse in a deflationary spiral. The problem is how to properly calibrate the amount of monetary interventions and that this system in general tends to increase wealth inequality by disproportionately affecting the low- and medium-income population relying mostly on depreciating fiat.

But couldn’t some fraction of money be created without any debt obligation? Hypothetically, if enough money could be created by discovering new currencies, these could offset the amount of fiat money supply needed to cover interest and reduce the need for continuous inflation of fiat. Note that unlike concepts like helicopter money or universal basic income, these currencies are only earned in exchange for a cryptographic proof of contribution (and proportionally to invested effort, with some luck factor), not for free. Cryptocurrencies like BTC, which are mined, were likely on the right path initially, but unfortunately, without being used and recognized as a medium of exchange, they do not really matter for the monetary system itself. If the newly discovered cryptocurrencies we propose are to help with covering some part of the money supply, we would need them to really work as money. This would, for example, require them to be directly taxable as income without conversion to national currency. For currencies which would be distributed upon discovery and not mined continuously, the initial issuance could be used by states to make reserves from a fraction of such currencies paid as taxes. In addition, states could choose to recognize only certain such currencies, possibly completely ignoring the low-cap currencies and exempting related microtransactions from tax obligations entirely. Arguments for a special classification of such assets as actual currencies could be that they can only be obtained initially by mining (there are no pre-sales), and unlike other mined cryptocurrencies, cannot be used to pay blockchain network fees—i.e. they are really “useless” beyond being a medium of exchange or a collectible.

In our view, the future monetary system could be a mix of fiat currencies and cryptocurrencies with various issuance models as well as the proposed PoW currencies, which could not only coexist, but actually complement the fiat system in a productive way. If recognized and used as mediums of exchange, we could have a range of money with different velocity. Of course, optimal management of a potential “basket” of one’s currencies might be difficult for a human, but algorithmic and machine-learning-based solutions could be used to, for example, pay for transactions (if agreed by both parties) with a mix of fiat and crypto. Note that less sound money might require some premium in the payments, taking into account possible risks of volatility or limited liquidity for the accepting party.

E.3 Can Fixed-Supply Digital Gold Solve Anything?

Because the number and value of the discovered PoW currencies is on average proportional to the invested computing resources and to the mining efficiency, more and more valuable

currencies can be discovered occasionally indefinitely, while keeping the total supply of each currency fixed. Such a system might be advantageous, or at least complementary, to the Bitcoin monetary model. If we were to truly adopt Bitcoin as the primary currency, or even just use it for central bank reserves in a similar way as gold is (or was) used, the fixed supply would probably cause significant issues in the economy due to deflationary pressures. Bitcoin is considered digital gold, but with an even stricter limit on supply²³. The fixed BTC supply might in fact be one of its most problematic aspects in an economy where Bitcoin would eventually truly dominate. There were good reasons to abandon the gold standard in the past. Adopting Bitcoin as a gold replacement would potentially reintroduce problems already “solved” by fiat systems or destabilize them in novel ways.

The proposed currencies offer some kind of compromise: the supply of each coin is fixed and all supply is emitted upon discovery, but new competing currencies can emerge over time (inflation of the number of currencies). To some level, this is similar to the evolving cryptocurrency landscape, but right now, anyone can release a currency with any name, while in the proposed system, the selection of which currency should emerge is constrained by the laws of nature and cryptography. To some extent, the emergence of these currencies could effectively look like occasional discoveries of gold treasures in sunken ships, allowing us to operatively increase the potential gold-backed money supply if needed²⁴. Unlike for such imaginary treasure discoveries, the rate and value of the PoW currency discoveries can be influenced by the total invested mining resources and, importantly, the amount of “treasures” to discover is unlimited.

E.4 Economics of Wasting Money in Luck-Driven Systems

We generally expect that mining will not be profitable unless some of the discovered currencies gain market value (through further adoption) truly reaching or exceeding their inherent value. However, occasionally, people will get a share (and later the corresponding coin) very quickly after starting mining, maybe even with a mobile phone and minimal hashrate. Some people will be lucky to do so repeatedly. Eventually, some will be “lottery” winners in the sense that their investments repay significantly. For most, the actual gains won’t be large, like in typical lotteries. But at the same time, for most, the actual investment would be very small. Some will have a PC with a GPU and could occasionally join mining of some currency to get some shares and, in very rare cases, actually discover a new currency and get some of its supply as a discoverer’s reward. Others could rent remote GPU compute (likely readily available in the future) for a short time, possibly for say a dollar per month. But unlike current cryptocurrency mining operations, the resulting shares would not have a value on their own, unless they are converted to a particular currency. Thus, running large-scale mining operations would not allow expenses to be continuously paid by selling the mined asset. Instead, the mining infrastructure would have to be rented for a short time to many occasional participants. For mining paths without highly efficient ASICs, GPU computing power would be readily available, and as the mining would not involve private key generation, the mining can be perfectly secure when executed, e.g., in a cloud environment.

²³Sometimes, people argue that while there can never be more than 21M BTC, one could, for example, find a new source of gold by mining in space, and that this is an advantage of BTC. This argument neglects the probably extreme cost of such gold mining operations. In contrast, any level of compute will not be able to mine any new BTC after 2140, unless the minimum unit (1 satoshi = 0.00000001 BTC) is redefined.

²⁴But in the case of these currencies, the supply would be emitted even if it was not needed by the economy. This is yet another reason to keep fiat with central bank control, which could operatively respond to temporary market instabilities in the (rare) events of discoveries of PoW currencies with extreme market valuation.

In such a lottery-like system, one would not expect that investments (bets) into mining scale linearly with wealth. In fact, more wealthy people will likely invest in such uncertain endeavors much less in proportion to their available resources than low- and mid-income population, for which a potential discovery or cheaply obtained set of mining shares (later rewarded) could become life-changing. Wealthier individuals will likely prefer investments not entirely driven by random processes and luck, as they can usually find better and more reliable means of increasing their wealth.

Thus, in the case of actual extreme currency discoveries distributed to a large fraction of the world population, this newly created wealth could be mostly initially distributed toward the less wealthy. At sufficient scales, such a new wealth distribution could mean a restart in the process of wealth extraction by “the rich.” Wealthy and “successful” people could then simply buy this currency (possibly by selling some of their assets) or earn it by accepting it for payments in their businesses. This would be similar to how they mostly missed the early Bitcoin years and only invested much later, as they likely found more certain economic opportunities in the past.

E.5 Resolving the Bitcoin Paradox

Bitcoin is widely understood as a fixed-supply digital commodity. With its 21 million coin cap and global inflationary environment, its purchasing power is expected to increase forever. If adopted at scale, this leads to a paradox: a passive, non-productive asset could outperform all productive investments indefinitely.

This implies:

- Bitcoin becomes more desirable than houses, businesses, or education.
- Early adopters are rewarded indefinitely, reinforcing inequality.
- Economic activity may be distorted by hoarding and underinvestment.

But Bitcoin’s scarcity is not physical—it is a product of consensus and design. If the perceived value of scarcity becomes socially or economically problematic, nothing prevents the creation of other provably scarce digital assets with a fixed supply. Let us stress the following:

There is no fixed supply of fixed-supply assets.

The system proposed here offers a way to mine new rare fixed-supply digital artifacts:

- backed by real and measurable computational effort,
- collaboratively discovered, not privately mined,
- and embedded with culturally legible meaning.

A plural ecosystem of such assets dilutes the need for any single “perfect” digital medium of exchange or (temporary) store of value. It becomes possible to:

- distribute speculative interest across multiple verifiably rare tokens,
- reduce the risk of runaway value concentration in one asset,
- and stabilize expectations about long-term digital wealth preservation.

By allowing society to co-create provably rare and culturally resonant digital assets, we move toward a healthier system where digital scarcity exists—but is not monopolized. Bitcoin was the first. It does not have to be the only.

F Environmental Concerns

To mine the 0xBadFace artifacts, we spent around 3 kWh of electricity. If we consider our various tests during development, our mining efforts consumed around 15 kWh in total. The 0xBadFace prefix is easily achievable by anyone with a GPU, but everyone making their own worthless artifact for about 1 kWh would be really a waste of energy. Instead, we should aim to join our forces to produce "objects" which would otherwise be inaccessible without "collective wasting of energy". Similarly to Bitcoin and in fact all money, a "sacrifice" (work/time/energy/damaged environment/...) or some "opportunity cost", aggregated over a large population demonstrates the commitment and "locks" the value in the monetary asset, which becomes representation of the collective effort of those receiving it. It is exactly the easiness of how fiat money is "created" today, which this system avoids, but of course, at the cost of the PoW. The effort put in (or what would need to be theoretically put in, but we have been lucky) is what would give these currencies the much desired fundamentals which fiat currencies (and of course also many cryptocurrencies as well meme coins) lack.

Note however, that for the initial experimentation, the scale of the energy input will be small compared to Bitcoin mining, which consumes over 150 TWh annually. Even if one billion people decide to invest an equivalent of 10 kWh of compute (worth at most a couple of dollars), the total required energy is around 10 TWh, and this assumes truly global and significant participation.

Many would see it as a step back when we formally re-introduce PoW to the Ethereum or other blockchains, albeit not for securing the transactions, but for artifact mining. There is however an important distinction to the Bitcoin-like dynamically adjusted difficulty – here more hashrate means access to more and more valuable artifacts, reflecting the invested mining resources.

Because the PoW does not need to be continuous here, but people can try to get mining shares occasionally, the mining can run in times of cheap (green) electricity. As far as the mining will only be possible on general purpose hardware, widespread GPU compute capacity of various providers can be used in otherwise underutilized time for such mining, allowing additional monetization for such infrastructure.

More importantly, we briefly discuss in Appendix G, we could potentially utilize a single global PoW mechanism using a particular chosen hash function, possibly post-quantum, for up to three distinct purposes simultaneously:

- To secure a single dominant (or multiple) PoW blockchain(s) using this hash function for consensus. This blockchain could be already designed to support creation of smart contracts from LuckChains (see Appendix G) in some way, and simultaneously supporting merged mining together with the artifact discovery headers - but this is not possible on existing blockchains
- To discover LuckChain artifacts, launch-able possibly on this dominant PoW blockchain, or as standalone blockchains
- To secure digital lotteries and various games of chance (see Appendix D), where the headers could contain also casino entropy, being revealed only after each game (sub)session for share registration or in case of artifact discovery. Of course if the casino, inserting also part of the header of a given mining pool the user is using (or this part is inserted by the user frequently enough to refresh the game inputs) finds a hash with difficulty

accepted by the PoW chain, it must publish it as soon as possible, maybe ending a game prematurely and starting a new one.

For the existing blockchains, on which we mine contract addresses and not block hashes, we can still at least combine the meme search with securing digital game of chance.

With such a potential multi-purpose usage of each hashing output, the resources utilized for the PoW could be used more efficiently, making the work of the PoW much more useful.

G Generalized Lucky Genesis and Multipurpose PoW

One could generalize the proposed concept to the idea of *LuckChains* — raw blockchains with no initial protocol with extreme-difficulty genesis, where rare artifacts are discovered directly in block hashes of initial blocks rather than smart contract addresses. These artifacts are found by hashing a genesis block or a short initial sequence of blocks until a hash with a perceptually meaningful pattern appears. The difficulty of reproducing such a hash (the hash function may be arbitrary) defines its *inherent value* as before. Once found, these artifacts can seed new blockchains or, if implemented, new currencies on newly established blockchains allowing artifact genesis from a raw blockchain, with its smart contract address being the discovered memetic final block hash.

In this appendix, we outline how this concept can be extended into a *generalized PoW coordination framework*, enabling more efficient, shared, and multipurpose mining.

Rather than searching for a single meme pattern or artifact at a time, future systems could support *merged mining across multiple Adam artifacts* — previously discovered hashes awaiting their matching Eve continuation. In this model, miners would construct a *multi-header block* containing:

- A `prev_hash` reference to own raw short blockchain starting from a genesis block, allowing to discover a potential Adam artifact
- several `prev_hash` references to previously discovered popular Adam artifacts, to which a matching Eve is searched for,
- optional headers of active blockchain(s) (for merged mining), if such blockchain natively supports this feature
- and optionally entropy seeds for lottery games, see Appendix D.2.

Every hash attempt over such a multi-header or a short block following it would simultaneously:

- look for candidate Adam artifacts in a completely open search,
- search for matching Eve artifacts matching any of the included previously discovered (popular) Adams,
- contribute to consensus of one (or more) standard PoW blockchains,
- and generate lottery-compatible randomness for fair games of chance.

This *multipurpose PoW architecture* transforms mining from a single-purpose effort into a shared computational process. However, it introduces challenges. Miners could attempt to include thousands of previously discovered "imperfect" Adam artifacts (e.g., `FightFi9ghtFight`, `FightFightFight1`, etc.), inflating their chance of some, even very "imperfect" discovery.

To mitigate this, we propose:

- **natural constraints:** First, search for exact matches might slow down a generic procedure. Second, the longer the multi-header, the more expensive the computations become (especially on-chain for share submission and verification),
- **social filtering:** communities, protocols and tooling may limit the size of the multi-header, discouraging search for imperfect artifacts.

Importantly, this structure also supports *passive background mining* of extremely rare or "impossible" artifacts. If all participants include a fixed list of public ultra-high-difficulty Adams in their merged searches, one such outlier may eventually be discovered — a genuine *proof-of-luck* event.

In sum, generalized LuckChain mining enables:

- simultaneous discovery across multiple meme artifacts,
- integration with live blockchains and public randomness,
- and efficient use of every hash attempt.

While speculative, this direction marks a critical design target: a system where massive distributed effort, recorded across many ecosystems, can be later *resolved into a single recognized artifact and fairly tokenized into a new currency*.

This would be the "holy grail" of generalized Proof-of-Work discovery: a decentralized architecture that not only captures rare value from collective effort without reliance on any single contract or chain, but also does so with *maximum efficiency* — ensuring that no good meme gets lost if it appears — and reuses each hash attempt across multiple purposes: securing PoW chains, mining new Adams or Eves, and anchoring fair lotteries in public randomness.

Distribution Challenges and the Need for Share Aggregation

While generalized LuckChain mining could allow broad, efficient discovery of new artifacts, it raises a difficult coordination problem: *where and how should mining shares be registered*, especially when no specific contract or chain is targeted? And how should the distribution of the tokenized discovery work?

In practice, share registration would still need to happen on existing blockchains — such as Ethereum or its Layer 2 networks — using predefined smart contracts that allow miners to submit generic shares, only referencing the selected submission contract (or possibly also chain ID, so they cannot be reused elsewhere).

But in the generalized setting described above, where miners include many artifact candidates in a single search header, there is no single destination contract under which all participants are organized. This complicates reward distribution once a discovery is made. To ensure fairness, some mechanism must aggregate these dispersed, cross-chain mining shares and associate them with the discovered artifact retroactively.

In the absence of such infrastructure, the discoverer — the entity (or entities) that finds the Adam and Eve block(s) — may hold significant responsibility. They must either:

- **cooperate with the community** to launch a new blockchain from the discovered artifact and import or verify share proofs from multiple registration systems, or
- **deploy the LuckChain as artifact on a dedicated blockchain** that supports native LuckChain genesis and connection of the distribution mechanism and share aggregation via a smart contract which can be deployed at the artifact address (hash of the last LuckChain block).

This introduces trust assumptions and logistical challenges that remain unresolved.

For high-value discoveries, the challenge is magnified. Ideally, such artifacts should be *distributed as broadly as possible*, reaching many millions of small participants who contributed to the global search effort — often with only a few low-difficulty shares. Achieving this would likely require:

- predefined aggregation protocols, including cryptographic proofs that can be submitted from many blockchains,
- interoperable registration formats, allowing shares to be cross-verified or batch-transferred into a newly launched chain,
- and possibly a canonical redistribution engine, governed by community or multi-sig consensus, that can operate across multiple networks.