xPrivFi: Minimal Fixed-Supply Digital Cash

A Layer-1 Proof-of-Work Chain with Equal-Chance Mining

BF
xPrivFi Project

Whitepaper v3 (Short Edition)
December 2025

Status: Informational – Non-contractual, Experimental Software

Abstract

XPrivFi is a minimal Layer-1 blockchain implementing a predictable, fixed-supply monetary system secured by a CPU-oriented proof-of-work algorithm derived from the RandomHash / interactive Proof-of-Work (iPoW) family. The protocol is paired with HexGrid, an optional Layer-2 mining system that coordinates equal-chance mining rounds without altering Layer-1 consensus rules. XPrivFi aims to minimize attack surface, remain fully auditable and transparent from genesis, and provide a clear path toward optional future privacy via the CP-Shield framework. This document describes the design philosophy, monetary policy, consensus rules, ledger model, mining architecture, security assumptions, limitations, and non-binding roadmap for the protocol.

Contents

1 Introduction

Modern cryptocurrency systems often suffer from two persistent problems: (1) unnecessary technical complexity and (2) economic and governance opacity. Layer-1 protocols accumulate large scripting engines, arbitrary tokenomics, upgrade levers, and governance mechanisms that are difficult for users to audit or fully understand. In practice, this tends to centralize power in the hands of large infrastructure operators, developers, and capital.

XPrivFi takes the opposite approach. The goals are:

- **Strict scarcity:** a fixed and easily verifiable total supply of 1,000,000 XPF.
- **Deterministic rules:** consensus, supply, and validation logic are intentionally minimal and static.
- **Human-scale mining:** equal-chance mining rounds coordinated at Layer-2 where hardware advantage does not determine success.
- **Transparency first:** a fully transparent, auditable ledger at launch, with privacy introduced later as an optional layer.

The protocol consists of:

- A Layer-1 (L1) blockchain with RandomHash/iPoW proof-of-work, deterministic monetary policy and an account-based ledger.
- A Layer-2 (L2) coordination layer, HexGrid, which provides equal-chance, round-based mining that is accessible from browsers and mobile devices.
- A forward-compatible privacy framework, CP-Shield, which can introduce shielded balances and private transfers in the future without compromising total supply verifiability.

The remainder of this paper presents the design philosophy, the monetary and consensus rules, the L2 mining model, and the security assumptions and limitations of the system.

2 Design Philosophy

XPrivFi is guided by three high-level principles.

2.1 Minimalism

Consensus rules must remain small enough that a single engineer can understand, audit, and re-implement the entire system. This excludes:

- General-purpose smart contract virtual machines.
- Arbitrary on-chain governance or voting logic.
- Dynamic monetary levers or administrative minting.

The base layer implements only the minimum logic required to maintain a consistent, fixed-supply ledger with proof-of-work security.

2.2 Transparency Before Privacy

XPrivFi launches as a fully transparent chain. All balances and transfers are publicly visible, and total emission can be recomputed independently by any full node. Privacy is treated as an additive, carefully designed layer (CP-Shield) that must not introduce inflation or weaken supply auditability.

In other words, transparency is a *default*, and privacy is an *option* that must be earned by correct cryptography and careful engineering.

2.3 Human-Scale Mining

Traditional proof-of-work mining rewards the largest concentration of hardware. Over time, this leads to industrial mining, ASIC centralization, and dependence on electricity budgets rather than broadly distributed users.

XPrivFi introduces HexGrid, a Layer-2 system where the mining experience is:

- **Equal-chance:** in each round, each participant has probability $1/N$ of winning, independent of device power.
- **Round-based:** mining occurs in structured time windows (rounds) with clear start and end points.
- **Human-focused:** participation happens in browsers and mobile devices, with simple visual interfaces.

HexGrid does not change consensus rules or monetary policy, but it reshapes the user experience of mining around human participation rather than machines.

3 Monetary Policy

3.1 Total Supply

XPrivFi defines a single native asset, XPF, with a fixed and immutable total supply of:

$$S_{\text{total}} = 1{,}000{,}000 \text{ XPF}.$$

This supply is split into two components:

- 500,000 XPF allocated at genesis (documented and transparent, for development, security, ecosystem bootstrap and long-term alignment).
- 500,000 XPF emitted over time as mining rewards.

There is no administrative minting or dynamic inflation mechanism. All coins must arise either from the explicit genesis allocation or from block rewards governed by consensus rules.

3.2 Emission via Mining Rounds

Mining emission is conceptually defined in terms of *rounds*. Each successful reward round corresponds to exactly 1 XPF being assigned to a winner and settled on Layer-1.

- Number of reward rounds: 500,000.
- Reward per round: 1 XPF.
- Total mined supply: 500,000 XPF.

HexGrid coordination determines which human participant is associated with each reward, but Layer-1 still enforces that the cumulative mined supply never exceeds 500,000 XPF.

3.3 Time Horizon

A typical HexGrid round consists of:

- approximately 6 minutes of work,
- followed by a short transparency/reveal phase,
- followed by a lobby/next-round preparation phase.

In combination with a target Layer-1 block interval of approximately $T = 360$ seconds (6 minutes), this yields an expected distribution period of roughly 5.7–6.6 years to mine the full 500,000 XPF, allowing for variance and implementation details.

3.4 No Monetary Surprises

The monetary design obeys the following invariants:

- Total circulating XPF can never exceed 1,000,000.
- Mined supply can never exceed 500,000.
- There are no halvings, dynamic inflation regimes or burned-fee monetary feedback mechanisms.
- All coins can be accounted for by iterating from genesis.

4 Layer-1 Consensus

4.1 Overview

The Layer-1 chain is a conventional, linear proof-of-work blockchain with:

- a block header containing minimal consensus-relevant metadata,
- a list of transactions,
- a RandomHash-based PoW condition,
- a difficulty controller (CP-Diff) targeting 6-minute blocks.

4.2 Block Structure

Each block B_n consists of a header and a transaction list. A simplified header structure is:

```
version          : uint32
prev_block_hash  : bytes32
tx_root          : bytes32   // Merkle or similar commitment to tx list
timestamp        : uint64
difficulty       : uint64
nonce            : bytesN
```

The header is hashed using the RandomHash/iPoW engine. A block is valid from a proof-of-work perspective if:

$$\text{iPoW(header)} < \text{target(difficulty)}.$$

4.3 Interactive Proof-of-Work (RandomHash / iPoW)

RandomHash/iPoW is a memory-oriented PoW function that:

- derives a seed from the block header,
- fills an in-memory scratchpad with pseudo-random state,
- performs a series of data-dependent reads and mixes,
- collapses the final state to a 256-bit hash.

The design aims to:

- require non-trivial memory bandwidth,
- reduce the extreme advantage of highly specialized ASICs,
- keep verification relatively cheap for nodes.

Miners search over nonce space (and potentially other mutable header fields) until they find a header whose iPoW hash falls below the difficulty target.

4.4 Difficulty Adjustment: CP-Diff

The CP-Diff controller maintains the average block interval around the target $T = 360$ seconds. Let D_n be the difficulty at height n, and Δ_n be the measured time difference between blocks $n-1$ and n.

The raw update would be:

$$D_{n+1}^{\text{raw}} = D_n \cdot \frac{\Delta_n}{T}.$$

To bound the effect of outliers and timestamp manipulation, Δ_n is clamped into a range $[T/4,\, 4T]$, producing:

$$\Delta_n' = \min\big(\max(\Delta_n, T/4),\, 4T\big),$$

and the effective update becomes:

$$D_{n+1} = D_n \cdot \frac{\Delta_n'}{T}.$$

Nodes recompute the expected difficulty from history. A block with an incorrect difficulty is rejected.

4.5 Chain Selection

Among competing chains, nodes adopt the chain with the greatest accumulated work, which in practice corresponds to the sum of difficulties across all blocks. This rule provides eventual convergence under the assumption that honest hashpower remains non-trivial over time.

4.6 Block Validation

To accept a new block, a full node verifies at minimum:

- the parent block referenced by `prev_block_hash` is known,
- the timestamp is not too far in the future and is consistent with recent blocks,
- the encoded difficulty matches CP-Diff given prior history,
- the PoW hash is below the target,
- all transactions are valid and correctly signed,
- total rewards and fees obey the monetary rules.

Only blocks passing all such checks are appended to the active chain.

5 Ledger Model

5.1 Account-Based State

XPrivFi uses a simple account-based ledger. Each address has:

- a balance, measured in the smallest unit (atoms),

- a transaction nonce (a monotonically increasing counter).

Transactions reference a sender (`From`), a recipient (`To`), an amount, an optional fee, a nonce and a digital signature.

5.2 Transaction Validity

At a high level, a transaction is valid if:

- the sender and recipient addresses are well-formed,
- the amount is positive,
- the sender has a sufficient balance to cover amount + fee,
- the nonce equals the sender's expected nonce,
- the signature is valid and matches the sender's public key.

When a valid transaction is applied, the ledger state updates:

- the sender's balance decreases by amount + fee,
- the recipient's balance increases by amount,
- the fee accumulates in the current block's fee pool,
- the sender's nonce increases by 1.

There is no general-purpose token creation mechanism at Layer-1, and no scripting engine. The ledger tracks only balances of the native XPF asset.

5.3 Future Shielded Balances (CP-Shield)

CP-Shield is an optional future upgrade that introduces a shielded representation of XPF. Conceptually:

- transparent balances reside in the account-based state,
- shielded balances are represented by commitments in a Merkle tree,
- nullifiers prevent double-spending,
- zero-knowledge proofs allow private transfers.

A key invariant is that the sum of transparent and shielded value continues to match total emission. CP-Shield is out of scope for the initial mainnet launch and remains research-stage at this time.

6 HexGrid Layer-2 Mining

6.1 Motivation

In classical proof-of-work networks, mining rewards concentrate where the most specialized hardware and cheapest electricity reside. HexGrid aims to:

- equalize winning probabilities across human participants,
- decouple mining from industrial-scale hardware races,
- provide a clear, game-like user interface for mining.

6.2 Equal-Chance Rounds

HexGrid operates in rounds with the following properties:

- A round begins when users join and register as participants.
- The system requires N participants, where $200 \leq N \leq 890$.
- Each participant performs a deterministic, equalized CPU workload in their browser or device.
- At the end of the work phase, exactly one winner is selected with probability $1/N$ among the participants.
- The winner is paid 1 XPF via a transaction on the Layer-1 chain.

Visually, the process can be summarized as:

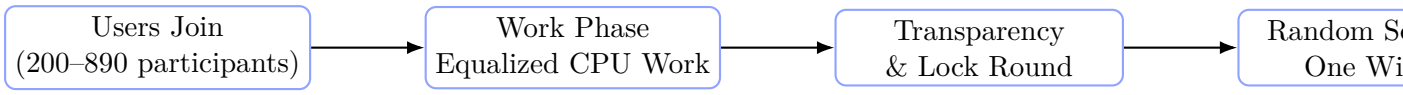

Figure 1: HexGrid Round Flow (conceptual).

6.3 Layer Separation

HexGrid is strictly a Layer-2 coordination mechanism. It:

- does *not* define consensus rules,
- does *not* mint coins directly,
- cannot change the total supply or difficulty,
- cannot bypass Layer-1 validation.

The worst-case failure of HexGrid is misallocation of a given round's reward; it cannot produce fake blocks or inflation. All rewards must still obey L1 monetary rules and pass standard validation.

6.4 Participation Fees

Mining via HexGrid is free at launch. No fee is required to join or participate in mining rounds.

If XPF later becomes listed on a decentralized exchange, the HexGrid Layer-2 engine may optionally enable a very small participation fee. This fee:

- exists only at Layer-2 and is not part of monetary policy,
- does not affect supply, difficulty or consensus rules,
- is intended solely as an anti-spam and ecosystem support mechanism,
- is paid in XPF by participants who elect to join rounds.

Participation in L2 mining remains entirely optional, and the core protocol (Layer-1) remains free to use.

7 Security Model

7.1 Adversary Capabilities

The security model assumes that an adversary A may:

- control significant hashpower, including potentially a majority,
- run many Sybil nodes in the peer-to-peer network,
- delay, drop or reorder network messages,
- observe network-level metadata and mempool contents,
- attempt to manipulate HexGrid participation and randomness.

We do **not** assume:

- honest majority of node operators,
- trusted committees or governance,
- trusted hardware or secure enclaves,
- honest Layer-2 infrastructure.

We only assume that widely used cryptographic primitives (hashes, signatures, commitments) behave as expected and that, over time, honest hashpower remains sufficiently large to make long-range attacks economically difficult.

7.2 Security Objectives

High-level security goals are:

- **Consensus safety:** two honest nodes should not permanently disagree on the canonical chain.
- **Monetary integrity:** total XPF in existence must never exceed 1,000,000, and emission must follow the specified schedule.
- **Double-spend resistance:** transactions (and future shielded notes) should not be spendable more than once.
- **Equal-chance L2 mining:** each participant in a round should have probability $1/N$ of winning, independent of device power, within the assumptions of the HexGrid implementation.

7.3 Layer-2 Risk Boundaries

HexGrid is explicitly *non-consensus*. Its failure modes include:

- unfair or biased selection of winners,
- downtime or instability of the coordination service,

- incorrect accounting of entry fees or rewards at L2.

However, HexGrid cannot:

- create new XPF beyond the consensus cap,
- alter difficulty adjustment,
- force acceptance of invalid blocks.

All L2 outcomes are ultimately encoded as standard Layer-1 transactions and must pass normal validation.

8 Privacy Roadmap (CP-Shield)

CP-Shield is a planned privacy extension that introduces shielded balances and private transfers while preserving the ability to verify total supply. The design goals include:

- hiding transaction amounts,
- hiding links between senders and recipients,
- preventing double-spends via nullifiers,
- maintaining a global supply invariant.

At a high level, CP-Shield would use:

- a commitment scheme to encode notes,
- a Merkle tree to allow efficient membership proofs,
- nullifiers derived from secret keys to block re-use of notes,
- zero-knowledge proofs to show correctness of transfers without revealing sensitive details.

CP-Shield is **not active at mainnet launch** and remains in the research and design phase. Any deployment would require additional audits and review.

9 Limitations and Future Work

9.1 Limitations

XPrivFi is experimental software with several limitations:

- Layer-1 security depends on sufficient honest hashpower over time.
- HexGrid fairness depends on randomness quality and correct implementation; it is possible for a malicious operator to bias selection if users rely on a centralized service.
- Privacy is absent at launch; all transfers are fully transparent.

- Browser-based mining depends on the integrity and security of user devices and execution environments.
- Network and peer-to-peer behavior may be vulnerable to as-yet unknown attacks or implementation errors.

9.2 Future Work

Potential areas of future work include:

- development and evaluation of CP-Shield, including concrete ZK constructions,
- a Dandelion-like (or similar) anonymity layer for transaction relay,
- additional independent node implementations (e.g., in Rust or C),
- formal analysis of HexGrid fairness and randomness,
- improved wallet and explorer tooling,
- external audits of the codebase and cryptographic design.

Roadmap items are non-binding and may change, be delayed, or never be delivered. No development outcome, listing, or future integration is guaranteed.

10 Conclusion

XPrivFi is a minimal, fixed-supply proof-of-work chain with an emphasis on simplicity, transparency and human-scale mining. By separating concerns across layers—a transparent, deterministic Layer-1, an equal-chance Layer-2 mining system, and a research-stage privacy layer—the protocol attempts to remain auditable and understandable while still offering a path toward meaningful privacy and fairer access to mining rewards.

Users and developers should treat the system as experimental and carefully assess the risks before running nodes, mining or holding XPF.

Acknowledgements. Community testers, reviewers and contributors who provide feedback, report issues, and stress-test early releases strengthen the protocol over time.

Appendix A: Architecture Overview Diagram

Appendix B: Conceptual Emission Shape

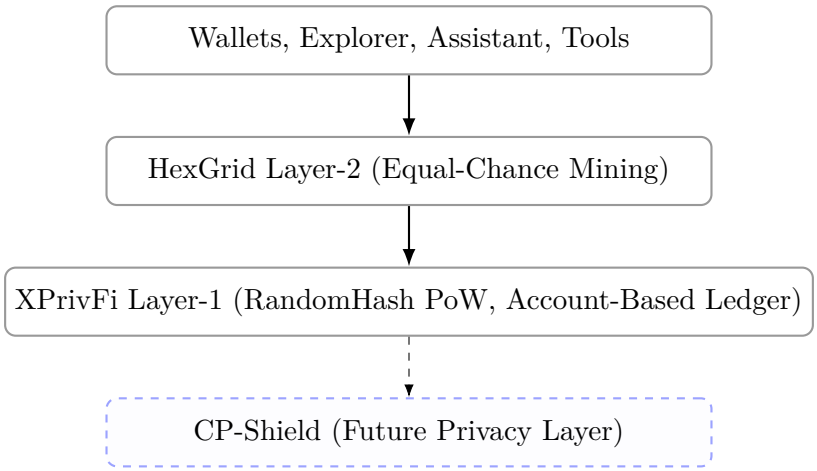

Figure 2: High-level architecture: tools and wallets on top of HexGrid L2 and the XPrivFi L1 chain, with CP-Shield as a future privacy layer.

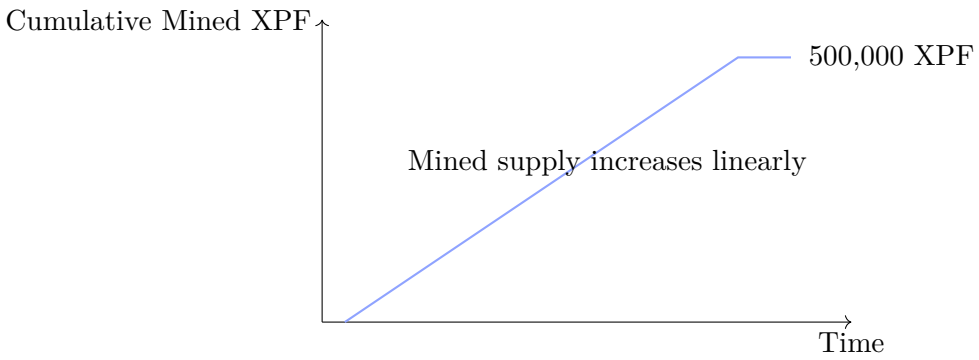

Figure 3: Conceptual emission curve for the mined portion of supply (not to scale).