Fair Unlock and Secured Emissions model (FUSE)

Aligning early backers and community

V0.9, November 2025

Peter Bienek, Stefan Iliev, Tudor Malene, Cais Manai, Rosie Sargsian, Gavin Thomas

Introduction

We present a novel **token issuance mechanism that combines principal protection, yield generation, and deflationary incentives** into a structured token distribution model.

Traditional token launches often expose early supporters to extreme volatility and significant unlock events. Investors and teams are forced to choose between long, inflexible vesting schedules or risk the reputational damage associated with dumping tokens into thin markets.

The **FUSE model offers a more elegant solution by combining principal protection, yield generation,** and deflationary incentives into a single issuance framework, and it eliminates the need for **traditional vesting.**

Participants who purchase a new protocol's token via a **FUSE round** are able to request the return of their principal (downside protection) while retaining full upside to the token's price. At the same time, the capital raised is deployed into low-risk yield strategies. A portion of the resulting yield is distributed to protected token holders, and another portion is earmarked to buy back and burn tokens, creating a self-reinforcing scarcity effect.

In the initial configuration, participation in new issuance rounds is **token-gated** – prospective buyers must hold or stake a minimum amount of the token to qualify. This requirement ties demand in the secondary market to access to future protected rounds, closing the loop between primary and secondary markets and aligning incentives across all stakeholder groups.

This whitepaper presents the design of the FUSE model, analyses its payoff structure and game-theoretic incentives, compares it with comparable financial instruments, and explains why it eliminates the need for traditional vesting.

Mechanism Design

Primary Issuance

In a FUSE issuance event, users buy newly issued tokens directly from the protocol's issuance pool using stablecoins (e.g., USDC). Rather than treating these stablecoins as immediate revenue, the protocol encumbers them and deploys them into short-term, low-risk yield strategies. This capital serves two purposes:

- **Generating yield** that can later be shared with the protected token buyers and used for buybacks.
- **Preserving the principal** to back each buyer's redemption guarantee.

Buyers receive a **Convertible Staking Token (CST)** – a non-transferable, soul-bound NFT representing the protected position. The CST encapsulates three rights:

1. **Redemption right (principal floor)** – The right to return any portion of the covered tokens at any time and recover the original purchase price in stablecoins. This provides a perpetual floor to the investment.
2. **Conversion right** – The right to convert the CST into freely transferable tokens. Conversion gives up the redemption floor and the yield entitlement for the converted amount.
3. **Yield right** – While the CST remains unconverted, the holder accrues a share of the yield generated on the encumbered reserves, proportional to their contribution.

These rights together make the protected token behave like a **perpetual principal-protected note**: holders have unlimited upside if the token's price rallies, a guaranteed redemption floor if it falls, and a yield stream while they remain protected.

Convertible Staking Token (CST)

The **Convertible Staking Token (CST)** is a **soulbound, non-transferable NFT** minted to each primary buyer at the moment of a FUSE purchase. It is the on-chain container for that buyer's rights and accounting.

Each CST records immutable parameters (purchase price, quantity, issuance round, timestamp, policy snapshot) and mutable balances (remaining protected amount,

unclaimed yield, remaining redeemable principal, status flags).

CST Key lifecycle actions		
Action	**Description**	**Effect on rights**
Convert	The holder swaps some portion of the CST for freely transferable tokens.	Cancels the floor and yield on the converted amount.
Redeem	The holder returns tokens to the issuance contract and receives the original stablecoin principal.	Removes those tokens from circulation; rights are consumed.
Hold	The holder keeps the CST intact.	Continues to accrue yield and retains the redemption floor.

This design ensures that selling protected tokens is not frictionless: to sell, a holder must first convert, giving up the redemption guarantee and yield. Consequently, holders will typically sell only after meaningful price appreciation.

Conceptually, one CST = one “tranche” of protected exposure (often many CSTs per wallet over time).

CST – Payoff Intuition

At time t with market price S_t:

- **Hold:** payoff increases by yield dY_t.
- **Convert** ΔQ**:** receive ΔQ token; lose floor on ΔQ.
- **Redeem** ΔQ**:** receive $\Delta Q \cdot P_0$ stablecoins; remove ΔQ token from supply.

For the still-protected balance Q_{prot}, the instantaneous value to the holder is:

$$V_t = Y_{\text{acc}} + Q_{\text{prot}} \cdot \max(S_t, P_0)$$

(Holding retains the embedded put at P_0 and continues to accrue yield.)

CST – Yield Mechanics

- **Source:** Low-risk, short-duration strategies on the encumbered reserves backing P_0.

- **Accrual:** Each CST accrues **pro rata to its time-weighted protected principal**, e.g.,

  $$dY_t = \alpha_t \cdot (P_0 \cdot Q_{\text{prot}}) \cdot dt$$

  where α_t is the net distribution rate after fees/retentions.

- **Distribution policy:** Claimable as stablecoins by default. Governance may permit **auto-recycle** (holder opt-in) to use Y_{acc} for buybacks (burn) or to boost future allocations.
- **Extinguishment:** Yield accrual ceases on any converted/redeemed portion immediately at the block of the action.

How CST Drives the Flywheel

1. **Stickiness:** Selling tokens requires **Convert** (giving up floor + yield). Rational holders delay conversion unless upside clearly outweighs forgone protection.

2. **Quiet downside exits:** Bear phases trigger **Redeem**, which **shrinks supply** (tokens removed from immediate circulation) without order-book sell pressure.

3. **Deflation on exits:** When holders convert to free-floating tokens, the reserves corresponding to the **now-unprotected** amount can be partially deployed for **buyback & burn** under PUMP policy.

4. **Gating:** CST ownership can be used as a **multiplier or entrance** for future protected allocations (e.g., balance- and tenure-based boosts), reinforcing demand.

CST – Gating and Token Reflexivity

The design of the **gating mechanism –** the rule determining which participants may access new FUSE rounds **–** has significant implications for both market reflexivity and long-term value formation within the token ecosystem. The Convertible Staking Token (CST) serves not only as a container of principal and yield rights, but also as a potential eligibility instrument for subsequent rounds. The choice of gating model directly shapes liquidity behaviour, speculative feedback loops, and capital retention dynamics.

1. Gating by Secondary-Market tokens only

In this model, eligibility for each new issuance requires participants to hold **free-floating tokens** acquired on the open market. Existing CSTs are not considered valid collateral for participation.

This structure introduces a **mechanical demand loop** into every issuance cycle: before each sale, participants must purchase tokens to qualify, creating predictable upward pressure on the secondary-market price. This generates **strong reflexivity**, where rising prices reinforce demand and vice versa. Liquidity becomes concentrated in the transferable token, improving depth and price discovery.

2. Gating by Tokens or Existing CSTs

In this alternative design, holders may satisfy eligibility using **either** free-floating tokens **or** existing CST balances. CSTs thus act as reusable access credentials for future protected rounds.

This model reduces the immediate buy pressure in secondary markets but greatly increases **holder retention**. Participants who already possess CSTs have an incentive to remain within the system rather than exiting to liquidity. The result is a larger proportion of encumbered, long-term-aligned capital and a steadier monetary base.

3. Hybrid Weighted Gating

In early phases, the protocol **will use token-only gating** to maximise open-market demand, stimulate reflexivity, and establish deep liquidity and price discovery. This configuration ensures that every issuance injects direct buying pressure into the circulating supply, anchoring market confidence and accelerating capitalisation of the system.

As the network matures, the parameters **will gradually shift toward CST-inclusive gating** to favour capital recycling, stability, and long-term retention. At this stage, reflexivity becomes less important than continuity and sustainability. The CST transforms from a pure protection instrument into an **access passport** for recurring participation.

Redemption Rules

Redemption rights are intentionally designed to be **non-transferable, buyer-specific, and Soulbound**. Only the original purchaser of a given issuance tranche can exercise the right to redeem those protected tokens back into stablecoins, and only up to the amount they originally bought. If those tokens are sold at any point, the rights attached to them do **not** travel with the tokens and are **permanently extinguished**.

This eliminates every form of redemption arbitrage. A protected holder who sells their tokens forfeits both the redemption guarantee and the associated yield share. The secondary-market buyer receives **regular tokens** with no protection and no yield entitlements. As a result, no one can buy the token near the floor and attempt to redeem at a guaranteed higher payout, because only the original buyer can

redeem, and selling destroys that privilege. The protocol's liability is therefore always fixed and predictable.

Functionally, the redemption right behaves like a **perpetual convertible floor** at the issuance price: if the token ever trades below the purchase level, an original buyer can simply redeem rather than sell at a loss. This removes potential sell pressure during downturns. Instead of dumping into the market, protected holders redeem quietly in the background, removing tokens from circulation without affecting the price.

When the market trades above the issuance price, the original buyer may sell to realise upside. But doing so means giving up both the downside protection and the ongoing yield share. This creates a natural holding incentive: as long as the upside is not large enough to justify burning the protection, most buyers prefer to continue holding. The system therefore promotes **sticky, long-term alignment**, with protection and yield rewarding patience, and exit only becoming attractive once genuine upward momentum exists.

Together, these mechanics ensure that redemption rights cannot be abused, cannot be transferred, cannot create arbitrage, and continuously reinforce price stability and holder alignment.

Post-Sale Buyback and Burn Dynamics

When an original purchaser converts their CST to obtain self-custodied, unprotected tokens (thus voiding the attached rights on those tokens), the stablecoin principal corresponding to the sold tokens becomes **unencumbered** – the issuer (e.g. a Foundation) now has discretion over those freed reserves. Importantly, the mechanism dictates that a significant portion of that freed-up stablecoin capital is used for **buybacks and burns** of the token on the open market. In practice, once tokens lose their protection (because the original holder sold them), the protocol takes, say, a percentage, X% of the stablecoins that had been backing those tokens and immediately buys tokens from the market and **permanently burns** them (sends them to a burn address or otherwise removes them from circulation). The remaining (100–X)% of the stablecoins might be kept in reserve or used for other purposes (e.g. funding development, future yield, paid to insiders, etc.), according to agreed governance.

This **deflationary buyback** on every redemption-right expiry ensures that any increase in circulating supply from the original sale is counteracted by a reduction in supply later. Essentially, tokens that were initially issued under protection either *stay out of the market* (if holders redeem and the tokens are taken back) or, if they *enter the market* (via resale), they trigger a burn of other tokens using the stable reserves. Either outcome supports the token's value: redemptions directly shrink supply by

taking tokens out of circulation (in exchange for returning capital), whereas market sales indirectly shrink supply by funding equivalent burns.

After such a resale, the previously protected tokens **lose their special status** and continue circulating as ordinary tokens without any redemption privilege or yield entitlement. From that point on, they behave like normal free-floating supply. On the other side, the original buyer, having sold, keeps whatever upside they have captured from selling above the entry price, but no longer has any claim on the principal or yield. The issuer ends up with some stablecoins (after buyback allocation) and has fewer tokens in circulation due to the burn. Overall, this dynamic creates a **self-correcting supply mechanism**: any time an early buyer exits, the system uses that opportunity to reduce supply and bolster the market.

Payoff Profile and Rights Extinguishment (Mathematical Model)

To formalise the incentives for an original purchaser, consider the issuance price per token to be P_0 (paid in stablecoins). Upon buying one token at P_0, the investor's position can be viewed as:

- A long position in 1 token (market price P at a given time),
- **Plus** a redemption right to sell that token back at P_0 (effectively a put option with strike P_0),
- **Plus** entitlement to continuous yield Y from the deployed P_0 capital.

Ignoring time value, the *floor value* of the protected token to the original buyer is P_0 (since they can redeem at P_0 any time). The *upside value* is unbounded (since if market price P exceeds P_0, they could sell at P). At any time t, if the buyer still holds the token, their total payoff (including yield earned up to that time) is:

$$\Pi_{\text{hold}}(t) = Y(t) + \max(P(t),\, P_0)$$

where $P(t)$ is the market price of the token at time t and $Y(t)$ is the cumulative yield distributed to the holder up to time t. The $\max(P, P_0)$ term reflects that the holder will choose the better of redeeming at P_0 or selling at market price P. If $P(t) < P_0$, the rational action is to redeem for P_0 (realising $\Pi = P_0 + Y$). If $P(t) > P_0$, the rational action is to sell at market (realising $\Pi = P(t) + Y$), since selling generates more than the guaranteed P_0 redemption. In either case, $Y(t)$ – the yield collected during the holding period – is a bonus that improves the payoff. Thus, the original buyer's position is essentially **principal + call option**: they are guaranteed at least P_0 back and can capture upside from any P above P_0.

Importantly, upon conversion of the CST to free tokens, the holder receives the accumulated yield plus tokens valued at the current market price, thereby relinquishing the principal protection floor. The floor protection only exists while holding the CST; once converted to self-custodied tokens, the holder assumes full market risk but gains complete control and transferability of the tokens.

This resembles a **protective put** strategy in traditional markets (long asset + long put for downside protection), combined with an income stream. As long as they hold, the buyer continues to earn yield and keep the put. Once they **sell the token**, however, these additional rights are **extinguished**. Formally, if the original buyer transfers the token at time t, *for* $t>t$ their entitlement goes to zero and their payoff is fixed at $\Pi(t^) = Y(t^) + P(t^*)$ (the yield earned plus sale price). The new owner's payoff is just $P(t)$ (the market value) with no floor, since they hold a regular token. The extinguishment of rights upon sale is a crucial part of the mechanism's design: it prevents any transfer of the put option or yield entitlement.

From the perspective of the issuing authority, this mechanism creates a **contingent liability** that is fully funded by the stablecoin reserves. For each protected token issued, the protocol holds P_0 in reserve. If the token is later redeemed, the protocol pays out P_0 from reserves and receives the token back (which can be permanently burned or held by the protocol. Only burned tokens are truly removed from supply; held tokens may be reissued in future issuance events, subject to governance). If the token is sold instead, the protocol retains the P_0 in reserve (since the buyer didn't redeem) and is then free to deploy some of it for buybacks. In effect, the protocol is selling an embedded put option to the buyer financed by the buyer's own capital. Because the arrangement is **fully collateralised** (the stablecoins backing each guarantee), there is no danger of defaulting on redemptions, unlike an under-collateralised guarantee. The only cost to the protocol is the opportunity cost of locking the capital and the usage of yield for buybacks, but these actions directly enhance the project's tokenomics by boosting price support and scarcity.

Completing the Flywheel: Token-Gated Participation

A critical element that amplifies the efficacy of the FUSE model is the **Token-gating requirement** for participation in primary issuances. This design choice ensures that the **primary market and secondary market are coupled**, creating a demand flywheel that continuously reinforces the token's value and utility.

Participation Requirement

Access to the protected token sales is **restricted to users who already hold (or stake) a minimum amount of tokens**. In practice, before an address can contribute stablecoins to buy new protected tokens from the issuance pool, that address must prove it owns a certain threshold of free tokens (which must be locked in a staking contract). The exact threshold is determined by the protocol (for example, one might need to stake 15% in free tokens to qualify for the next sale). This creates an exclusive gateway: only existing token holders (or those willing to become holders) can benefit from the principal-protected issuance.

Effects of the Gating Mechanism

This requirement has two immediate effects on the ecosystem:

- **Structural Demand Driver:** Prospective buyers who hear about the attractive protected issuance can't directly just show up with stablecoins; they must **first acquire tokens on the open market** to meet the eligibility criteria. This inevitably generates **buying pressure in the secondary market** ahead of each issuance. Essentially, every new participant has to become a buyer of the tokens twice: once in the market to qualify, and once in the issuance itself. This is a powerful force for price appreciation and liquidity in the period leading up to an issuance event. It also means that even those who missed past issuances will drive up demand in order to not miss the next – a continuous incentive for market participation.

- **Supply Sink (Locked Tokens):** The tokens that participants hold to qualify is required to be **locked or staked** for a duration (through the issuance period). A portion of circulating supply becomes effectively **illiquid** during these periods, reducing float. The gating causes a **temporary supply reduction** in the market – people holding (and staking) tokens to remain eligible won't be selling, which tightens supply and supports the price.

- **Unified Incentives:** Both primary and secondary market participants end up being largely the **same group of people** or at least overlapping groups. This alignment means whales or insiders cannot exploit one side of the market without affecting themselves on the other. Everyone who wants to benefit from the "safe" issuance must contribute to and believe in the token's market value. It fosters a community of holders who are all mutually invested in the token's success, blurring the line between "pre-sale investor" and "public token holder." In fact, the mechanism **converts secondary buyers into primary buyers**, linking their fate to the protocol's.

The Flywheel in Action

The interplay of the gating mechanism with the rest of the FUSE design creates a **flywheel effect** – a positive feedback loop where each phase feeds into the next:

1. **Acquire tokens on the Market:** Buyers who want the principal-protected, yield-bearing issuance must first *buy the token in the secondary market* to meet the eligibility. This pushes up demand and price in the market ahead of the sale.

2. **Participate in Protected Issuance:** Those holders then use their stablecoins to purchase new tokens from the protocol under the FUSE terms. They now hold more of the token, with protection and yield rights, reinforcing their commitment to the ecosystem. The stablecoins collected go into the protocol's reserve, ready to generate yield and back the redemption guarantees.

3. **Yield Generation and Distribution:** The protocol deploys the stablecoins into yield-bearing instruments. Over time, **yield flows back** to the protected token holders (as periodic interest) and to the protocol (which earmarks a portion for buybacks). This yield distribution gives holders ongoing returns, and the protocol accumulates funds for strategic market operations.

4. **Buyback and Burn on Exits:** If some of those holders later decide to sell their tokens (e.g., when price has gone up), their special rights expire and corresponding reserves are freed. The protocol uses those freed stablecoins to **buy tokens on the market and burn them**, reducing supply. This action supports the token price, benefiting all remaining holders (including new entrants who bought on the market).

5. **Rinse and Repeat:** The burned tokens and supportive price action create a scarcity effect, and the promise of ongoing yield and safety continues to attract more buyers. Because of the success of the last issuance, **more participants want to join the next issuance**, restarting the cycle at step 1 –

which means *again* buying tokens on the open market to qualify, and so on.

Through this cycle, the protocol's model ensures that **each stage reinforces the others**. The primary issuance drives secondary demand; secondary demand makes the primary issuance more successful (higher price, more funds); exits from primary issuance trigger burns that reward secondary holders; those secondary holders then want to join the next primary issuance. It is a tightly interlocked system that, if executed properly, can accelerate token adoption and value. Early adopters are strongly incentivised – they get the best of both worlds: unlimited upside with no downside and yield, a combination that naturally triggers **FOMO** among those who have not yet participated. This **FOMO-driven growth** can lead to rapid initial uptake, faster price discovery, and a surging market capitalisation, all while a safety net is in place to **anchor the downside risk** for participants.

Moreover, this mechanism ensures that growth is **organic and liquidity-first**. Because it emphasises a strong on-chain market (people must buy on-chain to get in), it encourages deep liquidity in decentralised venues before any centralised exchange (CEX) listings.

In summary, the gating mechanism is the catalyst that turns the FUSE features into a powerful **flywheel**, aligning incentives and actions across both primary sales and secondary market trading. It transforms the issuance model from a one-time token distribution event into an **ongoing, self-reinforcing ecosystem of demand**. Every participant's rational move – buying to qualify, holding to earn yield, or selling to take profit – ends up feeding back positively into the system: driving demand, reducing supply, or distributing value. This is a stark departure from traditional token release models where primary sales and secondary trading are disconnected and often at odds. Here, they are symbiotic, each one propelling the other forward.

Game-Theoretic Dynamics

The FUSE mechanism creates a strategic interplay between:

- **Primary buyers** – participants buying tokens directly from the issuance platform
- **Secondary market participants** – participants buying and selling tokens on CEXes and DEXes
- **Protocol treasury** – the collective tokens held by early backers.

By design, each actor's incentives are structured so that rational behaviour contributes to the overall health of the ecosystem. We analyse the motivations and likely strategies of each group:

Original Buyers (Primary Issuance Participants)

Original buyers are those who acquire the token directly from the protocol's issuance under the protected terms. Their incentives are engineered to favour **holding over selling** in most scenarios. While they hold, they enjoy **downside protection + yield + potential upside**:

- **Hold incentives:** As long as they hold their tokens, original buyers face very limited downside (thanks to the redemption floor) and continuously accrue yield rewards. They also remain fully exposed to upside if the token's price increases. This combination means their position is low-risk and high-upside – an attractive payoff profile. They effectively hold a principal-protected investment with unlimited upside, which encourages them to **remain invested for longer** durations than a typical token sale participant.

- **Sell incentives:** The primary reason an original buyer would sell on the open market is if **market price exceeds their entry price**. In that case, locking in upside by selling might outweigh the loss of yield and protection. Essentially, the protected holder will consider selling when the *opportunity cost* of not selling (i.e. missing out on a large price gain) becomes higher than the value of continued protection and yield. Even then, many might sell only a portion of their holdings to realise some upside while keeping some protected position. It's important to note that because selling forfeits the redemption right, an original buyer will **not** sell lightly or in panic – selling is attractive primarily in *very bullish* conditions where upside has been realised.

- **Redemption behaviour:** In bearish scenarios or market downturns, we expect original buyers to either **hold or redeem rather than sell**. Since selling when the price is below the entry would realise a loss (and also give up

the redemption guarantee), a rational holder will instead invoke their redemption right to quietly withdraw their principal from the reserve or hold with the expectation that future issuances will be higher. This has two beneficial effects: (a) it **removes tokens from circulation** (those redeemed tokens go back to the issuer and can be burned or withheld for future issuances, rather than being dumped on the market), and (b) it avoids slippage or further price pressure that a market sale would cause. Thus, protected buyers act as a stabilising force: when some might otherwise add to selling pressure, these participants *exit via redemption off-market*, helping to form a price floor near the issuance price. In summary, original buyers' dominant strategies are to **hold in neutral or modest market conditions (earning yield)**, **redeem in heavy downturns (protecting principal)**, and **sell in strong upturns (taking upside)**. All three actions are non-detrimental or net-positive for the ecosystem (holding is benign, redeeming and selling both trigger supply reduction as described).

Secondary Buyers (Open Market Participants)

Secondary buyers are those who purchase the token on the open market (e.g. on DEXes or CEXes) from circulating supply. These participants do **not** receive any principal protection or yield directly – they are buying regular tokens. Why would they be motivated to buy, given others have protection? There are two main game-theoretic incentives for secondary market participants outside of the general utility for any particular token:

- **Anticipation of Buyback & Burn Cycles:** Secondary holders know that the protocol is continuously using yield and freed reserves to buy back tokens and burn them. This creates a reflexive expectation: secondary buyers can be confident that as original protected holders eventually exit (via selling or redeeming), a portion of the capital will flow into **deflationary buybacks**, reducing supply and potentially boosting the token's price. In essence, secondary buyers are riding on the coattails of the mechanism's design – they expect that even though they lack direct protection, the protocol's actions (funded by primary buyers' capital and yield) will support the market price of the token over time, benefiting all holders. The prospect of these periodic **buyback-and-burn events** gives secondary buyers a reason to hold the token and even accumulate more when prices dip, expecting the protocol to intervene and provide a price floor through redemptions and buybacks.

- **Token-Gated Issuance Eligibility:** The other major incentive is that holding the token in the secondary market is the only way to become eligible for future protected issuance rounds. The protocol requires participants to already hold or stake a minimum amount of the token to access the next issuance (Section

6 will detail this). This **demand-gated entry** creates a powerful motivation for secondary buyers: by accumulating tokens now (even without protection), they position themselves to participate in the next round of primary sales with principal protection and yield. In game theory terms, this encourages a **race to accumulate** – buyers buy in secondary markets not just for speculative upside, but to obtain a ticket into the principal-protected issuance. This dynamic can generate **FOMO** (fear of missing out) in the market, as those who are not yet in rush to buy the token so they don't miss the next protected opportunity. It also "locks in" secondary buyers as future primary buyers, feeding the cycle of demand.

Overall, secondary buyers are betting on the **positive feedback loop** created by the mechanism. They may not enjoy guaranteed downside protection themselves, but they understand that the **system's design benefits anyone holding the token** via price support and exclusive access. Thus, even unprotected tokens have enhanced attractiveness in this ecosystem compared to a standard token scenario.

Protocol Foundation

The protocol's governing body is the issuer of tokens. Its incentives and strategies align with long-term ecosystem health, and the FUSE mechanism grants it new tools for active supply management:

- **Flexible Reserve Management:** Through the redemption flows, the association gains a form of **automatic stabiliser** for reserves. In bull markets or hype cycles, lots of participants want to buy into the protected issuance (and secondary buyers accumulate tokens to qualify), meaning the association receives more stablecoins to manage. In bear markets, some participants redeem, meaning the association returns stablecoins but also withholds tokens, contracting supply. This ebb and flow allows the association to strategically manage how much reserve it holds and when to deploy it. The reserves are only invested in short-duration, high-quality instruments, ensuring the association can always meet redemption demands (liquidity management risk is minimal by design). The association can also decide what fraction of freed funds to put into buybacks versus keep as treasury; this discretionary element can be adjusted based on market conditions to maximise impact (for example, a higher burn fraction during severe downturns to strongly support the price).

- **Reflexive Cycle Benefits:** The association benefits from the **reflexivity** of the system. In exuberant times, increased demand for the token (partly driven by the gating requirement) raises the token's price and market cap, which in turn can attract more participants and capital to the protocol's ecosystem. This allows the association to sell tokens (from the insider pools or treasury) at

higher prices under protected terms, bringing in more reserves. In downturns, the mechanism releases reserves to buy back tokens, helping arrest price declines and maintain confidence. Essentially, the association has a more direct influence on supply and demand dynamics than in a traditional model, and it can **ride the market cycle** in a way that strengthens the token's fundamentals rather than undermining them. The result is a more **market-responsive monetary policy** for the token – the association isn't dumping tokens arbitrarily; tokens only enter circulation when there's active demand (people buying them with stablecoins), and exits are cushioned.

- **Alignment with Holders:** Importantly, the association's interests are fully aligned with token holders in this mechanism. Because any insider token unlock also goes through the same protected issuance (as detailed in Section 9), the association cannot simply flood the market with tokens to raise funds without simultaneously committing to the principal protection scheme that forces responsible use of those funds (i.e. yield generation and buybacks). The association's treasury grows primarily when new buyers join (increasing reserves), and it "pays out" when primary buyers leave (redemptions). This symmetry means the association is motivated to foster genuine growth and adoption so that more demand-side capital flows in, rather than relying on scheduled token releases that could damage market trust. In game-theoretic terms, the association has *skin in the game* – its ability to utilise reserves or unlock insider tokens is directly tied to **market demand signals**.

Through these dynamics, the FUSE mechanism aligns the strategies of all participants toward a common goal: strengthening the token ecosystem. Primary buyers are protected and loyal, secondary buyers anticipate and backstop the market, and the protocol uses its reserves to reinforce price stability and scarcity. The equilibrium outcome is that **all actors benefit from positive behaviour** (holding, gradual exiting through the mechanism) and are disincentivised from actions that harm the system (like dumping tokens). This is a case where **game theory in tokenomics design** is applied to encourage cooperative behaviour: the rules are set so that the individually rational decisions (seeking yield, avoiding losses, qualifying for access) also lead to collectively optimal outcomes (price support, gradual supply release). The result should be a more resilient token economy that can grow sustainably.

Benefits and Advantages

The FUSE mechanism yields a number of important benefits for the token project and its stakeholders, addressing many pain points of conventional tokenomics:

- **Principal Protection for Investors:** By guaranteeing redemption at the purchase price, the model **lowers the risk** for new investors to enter the ecosystem. This can attract a broader range of participants, including more risk-averse investors or institutions that would normally avoid volatile token sales. In effect, it provides a safety net that can significantly increase confidence in participating in the token launch or subsequent sales. Importantly, this is achieved without any external insurance or downside pooling – it's inherent in the token's design.

- **Reduced Sell Pressure and Volatility:** The presence of redemption rights means that during market downturns, protected holders will **redeem instead of market-selling**, as discussed, which curtails the usual cascade of panic selling. Furthermore, the continuous buyback and burn funded by yields and exited positions creates a constant **deflationary pressure** that supports the token price. Together, these effects stabilise the market: price floors are established by rational redemption behaviour, and upside volatility is more controlled because even when people sell, the protocol is shrinking supply. The mechanism thus dampens the extremes of price swings, creating a smoother growth trajectory.

- **Deflationary Token Dynamics:** Unlike most token distributions which introduce inflation (release of locked tokens, mining rewards, etc.), tokens launched through FUSE are actively **deflationary** in net effect. Every time someone converts from a protected position, tokens are permanently burned. Also, yield that could have been simply given out to holders is partly used to remove tokens from circulation. This means over time, if the platform succeeds in attracting participants, the circulating supply could **decrease relative to what it would have been**, concentrating value in the remaining tokens. Deflationary tokens are appealing to investors as they imply increasing scarcity. The token's supply becomes more scarce as adoption increases – an opposite of typical models where more adoption often means more tokens unlocked (inflation).

- **Longer-Term Holding ("Sticky" Holders):** The combination of earning yield and having a guaranteed way out if needed encourages investors to **hold for longer periods** than they would in an unprotected scenario. There's less fear of being left holding the bag, since the bag has a floor value. Yield rewards

patience and penalises short-term flips (since selling forfeits that yield). Consequently, the holder base becomes more **aligned with long-term success**, reducing the rapid turnover or immediate dump dynamics seen after many ICOs or IDOs. A stickier holder base is beneficial for governance, community, and price stability.

- **Market-Driven Token Release:** The mechanism effectively **paces token release according to market demand**. If demand is high, more tokens enter circulation (through more issuance sales), which is fine because the market can absorb them. If demand is low, issuances will either not sell out or people will redeem, meaning circulation doesn't swell without support. This **market-responsiveness** prevents the scenario of large token unlocks hitting a weak market and crashing the price. It replaces predetermined vesting timelines with an **adaptive process**: tokens come out when there are willing buyers to take them. This is a fundamentally more efficient way to distribute tokens, akin to how commodity supply might increase only when there's demand, rather than dumping product irrespective of need.

- **No Ongoing Subsidy Needed:** The yield that fuels the system is generated from the participants' own contributed capital (deployed in safe yields). The buybacks are funded either by that yield or by the participants' money when they choose to exit. Unlike some token models that rely on continuous inflation or rewards from the protocol treasury to incentivise behaviour, this mechanism is largely **self-sustaining**. It doesn't require the foundation to constantly spend its own assets to prop up the market; the participants and the mechanism's design create their own incentives. This makes it more **sustainable in the long run**, as the model can keep running as long as there is a yield on stablecoins and some trading activity.

- **Enhanced Price Discovery and Early Participation:** The urgency to buy tokens to enter the protected sale and the powerful incentives for early adopters (full upside, no downside) create a rapid initial **price discovery** phase. The token is likely to reach a fair market valuation faster than usual because so many are trying to acquire it for its utility in the issuance. This can accelerate the project's growth and visibility. Moreover, it **rewards early believers** – those who jump in first get the most benefit (they get to participate at presumably lower initial prices with protection). This fairness can turn early users into strong evangelists for the project, driving adoption through network effects.

In summary, the FUSE mechanism **tightly aligns token distribution with project success** and user incentives. It addresses the classic issues of token sales (dumping, oversupply, misaligned insiders, volatility) with a structured, win-win

approach. The mechanism fosters a healthier market and community from day one: investors who come in are protected and committed, the protocol actively supports the token's value, and the tokenomics become an engine for scarcity rather than dilution. These advantages make a compelling case for this model as a next-generation standard in token issuance, which we will strengthen further by examining how it replaces traditional vesting and insider lockups.

Vesting Becomes Obsolete

Traditional token vesting schedules have long been used to prevent early investor or team token dumps and to align stakeholders with long-term success. Typically, team members and VCs get their tokens released gradually over a period (with cliffs, etc.), so they cannot all sell immediately at launch. Vesting is meant to **protect the market** from oversupply and signal commitment. However, vesting comes with its own problems: it creates looming **supply overhangs** and scheduled **supply shocks** that traders anticipate and often *pre-sell* into, harming the price even before the unlock occurs. In many cases, large unlocks have caused drastic price drops as soon as tokens hit circulation. The FUSE mechanism offers an alternative that can **replace time-based vesting with demand-based release**. Let's examine why vesting is largely redundant in this model and how the new mechanism deals with insider tokens:

Role of Traditional Vesting: Conventional vesting's goals are to (a) **prevent massive token dumps** in early phases, (b) **align team/investors with project long-term** (they only realise value as the project grows), and (c) **smooth out the introduction of new supply** so as not to shock the market. These are sensible objectives in an unstructured distribution.

Why Vesting Becomes Redundant under FUSE: In this new model, the same objectives are achieved through different means:

- **No Dump Incentive:** Insiders (or any primary recipient) under FUSE are not handed large liquid token amounts that they can dump. Instead, insiders have their allocation in a protected issuance pool (Section 9 details this) meaning they get **stablecoins when they "sell"** tokens, not a pile of freely tradable tokens upfront. Since any exit requires either selling through the mechanism or redemption, there is no scenario where insiders have nothing to lose by dumping – they either have to give up yield and protection to sell, or if they try to game the system, they still can't flood the market without consequences (because sells trigger burns). In short, the mechanism itself **prevents sudden dumps** by structuring how insiders liquidate.

- **Demand-Linked Release:** The expansion of circulating supply (be it from insiders or treasury) happens **only when there is buyer demand to absorb it**, which is functionally equivalent to having vesting but instead of time being the gate, *demand is the gate*. If demand is low, insiders simply *cannot* sell large quantities because either no one will buy or they'd have to redeem (which returns tokens to the project rather than harming the market). If demand is high, then selling is fine because the market can take it – plus the sale will even boost the market via burns. This accomplishes the same goal as vesting (preventing early excess supply) but in a way that is **automatically tuned to market conditions**.

- **"Sticky" Holders and Signaling:** Early investors and team now have yield, giving them reasons to hold rather than rush to sell as soon as a lock expires. It shows the public that insiders are literally playing by the same rules as everyone else; **confidence is higher** because the community sees that *even insiders can't freely dump tokens*. In a sense, the FUSE model itself is a strong **signal of commitment and discipline**, potentially even stronger than a vesting schedule, since it is an ongoing commitment to only exit via the public markets in a constructive manner. That said, some conservative investors might still expect some vesting-like structure as a reassurance of good behaviour (since vesting is a familiar concept). For those audiences, the project can reframe vesting in this context: effectively, **insider vesting is "as long as it takes for the market to buy your tokens"**. If needed, one could still impose a minimal lock period before insiders are allowed to use the mechanism, but fundamentally the design aligns incentives enough that time-based vesting is no longer the primary tool to ensure trust.

In conclusion, **time-based vesting is supplanted by a more dynamic, market-driven process** in this model. There is no cliff at which a large number of tokens suddenly hit the market irrespective of conditions. Instead, tokens *trickle out* when conditions warrant – akin to an "IPO lockup" that doesn't expire until buyers show up. This not only avoids the dreaded **supply shocks** (e.g., 100%+ increase in float causing 40-60% price crashes as seen in projects like Optimism, Sandbox, Axie, etc.), but it also **aligns insider incentives with the network's growth** far more tightly than arbitrary timelines. Insiders will get liquid only if the network succeeds in attracting genuine users and buyers, which is exactly when distributing more tokens is healthy. The next section delves into how insider allocations are handled in detail under this mechanism.

Insider Allocations Under Protected Issuance

Typically, teams, advisors, and early investors (VCs) receive allocations of tokens that are locked and vested over time. These insider tokens often constitute a large portion of total supply and represent a potential **sell pressure** once unlocked. Under the FUSE model, we **treat insider distributions the same way as any other issuance**, ensuring that insider liquidity events are aligned with market demand and even contribute positively to the token economy.

Traditional Insider Vesting vs Protected Issuance Approach

Traditional Model: In a conventional setup, insiders might have, say, 20% of tokens vesting monthly over 2 years after a 6-month cliff. This means after 6 months, some tokens unlock, then every month more tokens become liquid. Insiders could then sell those tokens on the market, potentially leading to large **unlock events that flood the market** with tokens. These events are often public knowledge and cause anticipatory price drops (as mentioned, supply shock phenomenon). Insiders are motivated to sell at least some portion on each unlock to realise returns, especially VCs with fund timelines, etc. This scenario often creates an adversarial dynamic: insiders want to cash out, community fears the dumps, and the market is wary of upcoming unlocks. There is also an inequality: insiders got their tokens cheap and can still profit even at prices that current holders would consider low, so they have different risk tolerances.

Protected Issuance for Insiders: The approach is to allocate team, foundation, and VC tokens into **segregated protected issuance pools** that function exactly like the public one. Rather than receiving freely tradable tokens, insiders effectively receive an *entitlement to tokens* that they can sell via the issuance mechanism. For example, suppose the team has 10 million tokens to distribute. Instead of a time lock, these 10M could be put into a Team Issuance Pool, from which any buyer (including possibly community members or even the team themselves if they wanted) can purchase tokens with stablecoins under the FUSE terms. The insiders in this case are on the other side of that trade: when someone buys from the Team Pool, the stablecoins are encumbered (just like in public sales) and held in yield strategies. These proceeds are only released to the insider (team) when the buyer converts their CST, thereby relinquishing refund rights. This ensures insiders cannot exit until buyers are satisfied **–** the community "eats first" and the buyer gets the tokens with rights. The insider’s tokens are thus “sold” not dumped, and importantly:

- **Insiders realise liquidity only when there is a buyer** willing to pay for their tokens. If no one wants those tokens yet, they remain effectively locked (since no one buys them). If there is high demand, they get to sell – which is fine because that means new people are joining.

The rules for insiders mirror those of the public pool: the stablecoins from sales are **encumbered until rights expire**. An insider token sale initially just moves stablecoins into reserve backing that token's redemption. Only when the buyer eventually converts and sells does the insider get the freed capital. In essence, insider liquidity is **vested by demand**: they "vest" (get paid) when a buyer decides to hold their token without redeeming..

Implications: This unified framework has profound implications:

- All players – team, investors, public – operate under **the same market rules**. Insiders don't have an unfair advantage or separate secret vesting; their ability to monetise their tokens is directly tied to user participation. This is an ultimate alignment, likely boosting community trust.

- **No large scheduled unlocks:** Circulating supply only increases when tokens are purchased. There is never a moment where a big tranche just "appears" on the market. If 10M team tokens are allocated, they may take *years* to all enter circulation, or maybe never fully circulate, depending on demand. This removes the periodic shock factor.

- **Deflationary offsets on insider sales:** When insiders do sell via the mechanism (i.e., an outsider buys from their pool), that triggers the same post-sale dynamic – a portion of those stables go into buybacks and burns. So an insider cash-out is not dumping on the community; it actually provides funds that reduce supply and potentially bolster price. In other words, *insider exits become accretive, not dilutive*. This is a remarkable turnaround from typical situations where insider selling is purely extractive. Here it's additive: the insider gets their money, and the remaining holders get a supply reduction benefit.

Benefits to Insiders and VCs

One might wonder: why would insiders or VCs agree to these constraints? What is in it for them compared to traditional vesting? In fact, the FUSE model is attractive to insiders as well, offering a more flexible and potentially lucrative way to realise value, *if* the project does well. Some benefits include:

- **Liquidity Without Market Impact:** Insiders can liquidate positions **without crashing the market**, because every sale is matched with a real buyer and triggers buybacks. This means they can exit more **orderly and at better prices** than if they were dumping into thin markets post-vesting. They effectively avoid the "post-vesting dump discount" – in many projects, insiders rush to sell at any price, hurting themselves with slippage and crashing the price. Here, if an insider sells, the price is supported. It's akin to having a built-in block trade mechanism that soaks up the sell pressure. Over time, this

could allow insiders to actually realise *more* total value from their tokens than a scenario where their aggressive selling depresses the price for remaining allocations.

- **Continual Alignment with Growth:** Because insiders only get significant liquidity as the *external demand increases*, their fortunes are tied to the project's user adoption and success in real time. This is philosophically similar to equity: founders get rich if the stock goes up (company succeeds). In tokens, often founders could in theory get rich even if the project stagnates, as long as their vesting ends during a bull market or due to hype. The demand-linked model ensures that **value creation (network growth) is directly what enables insider monetisation**. For long-term oriented insiders like founders and committed investors, this is ideal – it guarantees they are rewarded proportionally to the project's impact, not just the passage of time. It turns their holdings into a sort of **"growth-linked bond"**: they can only cash in big if the network is big.

- **Reflexive Upside Participation:** Insiders under this model still hold a lot of tokens (those not yet sold) and also indirectly benefit from the deflationary mechanics. Each time someone converts (an insider completes a sale) it triggers a burn, the remaining insider-held tokens become more scarce and potentially more valuable. So insiders have an extra upside feedback loop: not only do they gain from price appreciation of the token, but the very act of others selling out increases the value of their remaining stake. This can make them **think twice about selling too fast** – there's a game theory element where holding longer could yield more if others exit first and burn supply. It's a bit like being a long-term shareholder who benefits when short-term holders sell out cheap.

- **Signaling and Reputation:** By agreeing to this mechanism, insiders send a strong **signal of confidence and alignment** to the community. It shows they are not afraid to subject themselves to market forces and that they believe tokens will be bought by real users (they're not relying on hype and then dumping on a schedule). This can greatly enhance the project's credibility. For VCs, it might even become a mark of prestige to participate in such fair structures – demonstrating they are good actors supportive of the ecosystem, not pump-and-dumpers. A better reputation can benefit these investors in the long run (community goodwill, better deal flow, etc.).

- **Capital Efficiency for Funds:** Venture funds often have fixed lifecycles (say 7-10 years) and need to return capital to LPs. Traditional vesting might lock them in longer than they'd prefer or force suboptimal selling (e.g., selling as soon as vesting ends even if market conditions are poor). The FUSE approach, acting like a **perpetual convertible instrument**, gives funds

flexibility. If the market is not ready, they simply don't sell (their clock isn't ticking on a vesting schedule). If the market booms, they can start selling earlier. They can also **choose the timing** of sales more granularly, selling in batches when demand is present, rather than being constrained by a vesting cliff. This can align better with fund return timing, essentially providing a more **liquid-yet-aligned exit path**. It's like holding a note that pays off when the company grows, which might align better with LP expectations than unpredictable unlock dates.

In sum, for insiders and VCs who are *truly aligned* with the project's success, this mechanism offers a way to **have their cake and eat it too**: they can eventually get liquidity, but only by growing the pie (the network value) so that their selling doesn't diminish the pie for others. It transforms what used to be a contentious process (insider unlocks) into a transparent, market-driven, and even **value-positive event** for the ecosystem. The mechanism can be seen as creating a new paradigm for insiders: their token holdings become more like an **investment vehicle tied to network performance**, rather than just locked tokens waiting for time to pass.

Conclusion

The FUSE mechanism represents a significant advancement in token issuance design, recasting tokens from a standard utility token into a **structured asset with embedded game-theoretic incentives**. By fusing principal protection, yield participation, and scarcity enforcement, the model creates a class of token holders who are **incentivised to stay, participate, and help stabilise the market**, a dynamic seldom achieved in previous token launches. The requirement for existing token ownership to access issuances ties every new participant's fate to the token's value, completing a feedback loop wherein each stage of growth reinforces the next. This design allows the protocol to actively **manage volatility** and liquidity in a way that aligns with market conditions, all while preserving the **upside potential** that attracts investors in the first place.

Crucially, extending this mechanism to insiders (team, VCs, foundation) **eliminates the need for traditional vesting schedules** and the uncertainties they bring. Liquidity for insiders becomes **demand-driven instead of time-driven**, meaning insiders can only exit in step with genuine user growth and market depth. When they do exit, instead of harming the ecosystem, their exit directly **benefits remaining stakeholders** through buybacks and strengthened scarcity. All stakeholders thus operate under the same fair rules, dissolving the usual tension between insiders and public investors. Everyone's incentives are unified: the token's success is the common denominator for realising value.

By implementing a fully **market-responsive issuance and distribution model**, we pioneer a new paradigm in crypto tokenomics. It shifts away from arbitrary time-based releases and toward a **continuous price discovery and supply adjustment process** guided by real demand. In doing so, it addresses the long-standing issues of early-stage token projects – volatility, speculation-driven dumping, misaligned unlocks – with a holistic solution that enhances trust and stability. The FUSE mechanism can be seen as **crypto's equivalent of structured finance innovation**, tailoring the payoff structures to suit all parties' needs and using mechanism design (in the spirit of game theory and incentive alignment) to create a more adaptive and resilient token economy.

Moving forward, if widely adopted, such a model could lead to an ecosystem where token launches are **safer for investors, more engaging for communities, and more disciplined for teams**. It bridges the gap between the expectations of traditional investors (principal safety, yield) and the realities of crypto markets (high growth potential, volatility) in a novel way. We can create token distribution mechanisms where **adoption, liquidity, and scarcity** are not at odds but instead *interwoven*, each enhancing the other to lay a strong foundation for long-term growth.

Appendix 1
Formal Modelling of the FUSE Mechanism

To rigorously capture the **FUSE Mechanism**, we introduce stochastic models for key dynamics: **reserve drawdown under redemptions** and **deflationary burn rate under yield & volatility**. We assume the token price $P(t)$ follows a stochastic process (e.g. a geometric Brownian motion with drift μ and volatility σ) to model market fluctuations and trigger conditions for redemptions or sales. In downturns ($P(t) < S$, with S the stablecoin-paid entry price), original protected holders are likely to exercise redemption rights quietly (reclaiming S per token) rather than sell into a weak market. Conversely, when $P(t)$ significantly exceeds S, holders are incentivised to sell on the secondary market for profit (forfeiting their protection and yield). These behaviour-based transitions can be modeled with state-dependent stochastic rates.

These parameters are empirical and will vary based on holder psychology, market conditions, time preferences, and individual risk tolerance, rather than following predetermined thresholds.

Stochastic Reserve Dynamics Under Redemptions and Resales

We define $N(t)$ as the number of **protected tokens** still held by original buyers (with active redemption rights) at time *t*. Let $M_r(t)$ be the cumulative count of tokens **redeemed** by time *t*, and $M_s(t)$ the cumulative count of tokens **sold** on the secondary market by time *t* (causing their protection to expire). The protocol's **stablecoin reserve** $R(t)$ (initially $R(0) = S \cdot N(0)$) earns yield at a continuous rate y from deployment in low-risk strategies. Each redemption event withdraws stablecoins from $R(t)$, while each resale event frees the corresponding stablecoins (since the protocol no longer owes that seller a redemption). We denote by β the fraction of these freed stablecoins that the protocol allocates immediately to **buyback-and-burn**, with the remainder $(1 - \beta)$ kept in reserves. Under these assumptions, the infinitesimal evolution of $N(t)$ and $R(t)$ can be described as:

$$dN(t) = -\,dM_r(t) \;-\; dM_s(t),$$

$$dR(t) = y\,R(t)\,dt \;-\; S\,dM_r(t) \;-\; \beta\,S\,dM_s(t)$$

Where:

- S is the stablecoin principal paid per token at issuance (the **protected redemption price**).
- $dM_r(t)$ = 1 if a redemption occurs in $[t, t+dt)$ (0 otherwise), representing a token returned for S stablecoins.
- $dM_s(t) = 1$ if an original token is **sold** in $[t, t+dt)$ (0 otherwise), causing redemption rights to extinguish.
- y is the instantaneous **yield rate** on reserves (assumed constant for modeling).
- β is the **buyback allocation fraction** from freed reserves on a sale (e.g. $\beta = 0.5$ means half of the freed stables are used for buybacks, half retained).

The first equation simply tracks the decrease in protected tokens: each redemption or sale removes one token from the protected pool $N(t)$. The second equation governs the **reserve balance** $R(t)$: reserves grow continuously by yield $yR(t)dt$, but drop by S for each redeemed token (payout to the redeemer) and by βS for each resale (funds diverted to buy back tokens when a holder sells). This stochastic differential equation encapsulates the **reserve drawdown under redemptions** (direct stablecoin outflows) and the **reserve recycling** on resale (since only a portion $(1-\beta)$ of freed capital remains in $R(t)$). It ensures reserves are *always sufficient* for redemptions because they're only invested in short-duration, low-risk instruments and only reduced when redemptions actually occur (guaranteeing the promised floor protection). If yield y falls over time, the growth term $yR(t)$ shrinks, slowing the pace of buybacks funded by yield – but the **floor redemption guarantee remains intact**.

Assuming the hazard (instantaneous probability) of a redemption or sale depends on price $P(t)$, one can further model the **probability of eventual redemption** for a given token. For example, let $\lambda_r(P)$ and $\lambda_s(P)$ be state-dependent intensities (rates) for redemption and sale when the token price is P. A simple case is $\lambda_r(P) = 0$ if $P > S$ and $\lambda_r(P) = \rho$ (some constant ρ) if $P \le S$ (redemption becomes likely when market price is at or below the protected price), and conversely $\lambda_s(P) = \sigma_0$ (some rate) when P is well above S. In such a framework, the probability that a given token is ultimately redeemed (rather than sold) can be expressed by a **competition of exponentials**. For instance, if we approximate λ_r, λ_s as constant over time for tractability, the fraction of tokens expected to end up redeemed is $\dfrac{\lambda_r}{\lambda_r + \lambda_s}$, whereas a more general time-varying intensity model yields:

$$\Pr\{\text{redeemed}\} \;=\; \int_0^{\infty} \lambda_r(t)\, \exp\!\Big(-\int_0^{t} [\lambda_r(u) + \lambda_s(u)]\, du \Big) dt$$

which is the standard formula for an absorbing **continuous-time Markov process** with two competing exit rates (redeem vs. sell). This formalism reflects that a protected token will eventually exit its protected state via whichever event occurs *first*. Although the exact functional form of $\lambda_r(t), \lambda_s(t)$ depends on market conditions and holder behaviour, the model qualitatively captures the **game-theoretic choice**: holders will delay action while $P(t)$ is in a comfortable range, but if a **price surge** occurs, the sell rate λ_s spikes (many holders take profit), whereas in a **downturn** the redemption rate λ_r rises as holders opt for principal protection.

Expected Deflationary Burn Rate with Yield and Volatility

A core objective of the FUSE mechanism is to create a **deflationary token supply dynamic** through **buybacks and burns funded by yield and resale events**. We now model the **token burn process** $B(t)$, i.e. the cumulative number of tokens removed from circulation (burned) by time t via two channels: (1) direct redemptions (tokens returned are taken out of circulation), and (2) open-market buybacks (purchases funded by yield or freed reserves, with tokens subsequently burned). Using the same notation as above, the **infinitesimal burn** $dB(t)$ can be expressed as:

$$dB(t) \;=\; dM_r(t) \;+\; \frac{\beta\, S}{P(t)}\, dM_s(t) \;+\; \frac{(1-c)\, y\, R(t)}{P(t)}\, dt\,.$$

Where:

- The first term $dM_r(t)$ accounts for **one token removed from immediate circulation per redemption**
- The second term $\frac{\beta S}{P(t)}, dM_s(t)$ represents **tokens bought back and burned at a sale event**. When an original holder sells and forfeits rights, βS stablecoins are deployed to buy the token on the open market; if the price at that moment is $P(t)$, this buys $\beta S/P(t)$ tokens for burning. For example, if $\beta = 1$ and a holder sells at $P = 2S$ (price doubled from entry), the protocol spends S to buy back $S/2S = 0.5$ tokens to burn (burning half the amount sold). If instead the price crashes to $P = 0.5S$ and a sale occurs, S in reserves can buy $S/(0.5S) = 2$ tokens, **burning more tokens** than were sold – a stabilising, deflationary effect.

  Note: This model assumes efficient markets with negligible price impact. In practice, large buybacks may experience slippage and should be executed

strategically (e.g., via TWAP or gradual accumulation).

- The third term $\frac{(1-c), y, R(t)}{P(t)}, dt$ captures **continuous burn from yield recycling**. Here $0 < c < 1$ is the fraction of yield **distributed to original holders** as an incentive, while $(1-c)$ is the fraction of yield plowed into buybacks. Thus over an infinitesimal dt, the protocol uses $(1-c), y, R(t), dt$ stablecoins from interest to repurchase tokens; at price $P(t)$ this corresponds to $(1-c), y, R(t), dt/P(t)$ tokens burned. This term embodies the **yield-driven buyback pressure**: even in the absence of any redemption or sell events, the growing reserves continually fund token burns, causing a gradual decay in circulating supply.

Integrating the above $dB(t)$ expression over a long horizon (e.g. $t \to \infty$) gives the **total expected token burn**. In particular, if $N(0)$ tokens were originally issued under protection, their eventual fate (assuming the system runs until all those tokens either redeem or sell) will reflect in $B(\infty)$. In expectation, we have:

- Approximately $E[M_r(\infty)] = N(0), \Pr \text{redeemed}$ tokens burned via redemptions (since redeemed tokens are removed from supply).

- Plus an expected $E\Big[\frac{\beta S}{P(\tau_s)}\Big], E[M_s(\infty)] = N(0), (1 - \Pr \text{redeemed}), E\Big[\frac{\beta S}{P(\tau_s)}\Big]$ tokens burned from buybacks triggered by sales, where τ_s is the sale time and $P(\tau_s)$ the price at sale.

- Plus additional tokens burned from ongoing yield-funded purchases over time (the integral of the third term). In a steady-state or long-run average sense, the **burn rate** from yield is $(1-c)y\frac{R(t)}{P(t)}$ tokens per unit time, which one can integrate or average over the distribution of $P(t)$ and $R(t)$ dynamics.

Importantly, **price volatility** σ plays a role in the burn dynamics. Higher volatility increases the likelihood of hitting extreme price levels, which in our model accelerates state transitions: more frequent **up-swings** induce sells (triggering buyback burns), and more frequent **down-swings** induce redemptions or cheap-buyback opportunities. As seen in the term $\beta S/P(t)$, a low price $P(t)$ (often arising from volatile downcycles) magnifies the number of tokens repurchasable per unit of stablecoin, **intensifying the burn** when it is most needed to support the

market. Conversely, during price rallies $P(t)$ is high and fewer tokens are bought back per dollar – but rallies also mean fewer redemptions and a healthier market, so the mechanism naturally **leans in** during bear phases and **eases off** in bull phases. This reflexive behaviour aligns with the protocol's design goal of reinforcing long-term value: **yield-driven buybacks and resale-triggered burns continuously reduce the circulating supply**, offsetting new issuance and ensuring that insider or early-backer exits *strengthen* scarcity rather than cause excess supply. Each parameter in the above formulas can be adjusted or estimated based on empirical behaviours (e.g. calibrating λ_r, λ_s or β, c, y) to quantitatively project the token supply trajectory and reserve longevity under various market scenarios, thereby adding a layer of formal clarity to the protocol's economic design.

Appendix 2
Traditional Analogues and Precedents

Although the FUSE mechanism is an innovative design in the crypto space, it draws inspiration from several **established financial instruments** in traditional finance (TradFi) and also shares themes with some previous crypto fundraising experiments. Understanding these parallels helps contextualise the mechanism's payoff structure and its novelty. We compare it with **principal-protected notes, convertible bonds, covered call strategies,** and **reversible ICOs**, highlighting similarities and differences in each case.

Principal-Protected Notes (PPNs)

Similarity: Principal-protected notes are investment products that guarantee the return of the initial principal to buyers at maturity, while providing some participation in upside of a reference asset. For example, a bank might sell a 5-year note that guarantees 100% of your capital back, plus say 50% of any stock index gains. This scheme's issuance is conceptually similar: original buyers are guaranteed their **principal back** (via on-chain redemption) and have **unlimited upside** to the token's price appreciation. In both cases, the investor's downside is limited to zero loss of principal (if held to maturity or redeemed) and upside is positive if the underlying asset performs well. This structure appeals to investors who want to **invest in a risky asset without risking capital**, a direct parallel between PPNs and this scheme's protected tokens.

Difference: Traditional PPNs are typically time-bound (you get your principal back at a set maturity date) and rely on the issuer's creditworthiness for the guarantee. If the issuer (usually a bank) goes bust, your guarantee might be worthless. In our case, the redemption right is **perpetual** (no fixed end date) and is executed via smart contract with funds already reserved on-chain, removing **issuer default risk**. Another difference is that PPNs often cap the upside or have a participation rate less than 100%, whereas ours gives full upside (no cap) to the token holder. Finally, PPNs don't inherently include a **deflationary buyback mechanism** – they are passive instruments. Our mechanism goes further by using yield and exit events to actively support the token price (burns), which has no direct analog in a PPN. Nonetheless, in terms of **payoff profile (principal + possible gain)**, Our protected issuance is essentially a crypto-native, real-time PPN.

Convertible Bonds

Similarity: A convertible bond is a debt instrument that pays regular interest (yield) and can be converted into equity (stock) if the stock price rises above a certain level.

Investors in convertibles thus enjoy **fixed income plus the option to share in upside** if the company's stock does well. This is akin to this scheme's protected buyers who receive **yield from stablecoins plus the ability to profit from the token's price upside**. In both cases, there is an element of **downside buffer** (the bond principal in one case, the redemption right in the other) and an upside link to equity/token performance. Convertible bonds are popular as a way for investors to have a safer claim (as creditors) while not missing out on growth if the company succeeds, which is conceptually what FUSE offers to token buyers (safer entry with growth participation).

Difference: When convertible bonds are converted to equity, they **dilute** existing shareholders – new shares are issued, increasing supply. This is often a trade-off: the company gets to borrow at lower interest due to the conversion option, but if conversion happens, more stock is out there. In our model, something almost opposite happens on "conversion" (which in our case is an original buyer exiting): instead of diluting supply, an exit triggers **buybacks and burns**, resulting in a *deflationary outcome*. In other words, convertible bonds turning into stock can put downward pressure on stock price (more shares in circulation), whereas protected tokens turning into circulating supply results in upward pressure (tokens taken out of circulation via burns). Thus, this flips the dilution issue on its head – *exits reduce supply*. Another difference is that convertibles have a fixed maturity and interest rate, while this scheme's yield is variable and not time-limited. Nonetheless, it's fair to say the FUSE issuance behaves like a **"perpetual convertible note"** where conversion (selling) results in deflation rather than dilution, a novel twist on the traditional convertible structure.

Reversible ICOs (rICOs) – Crypto Precedent

Similarity: The Reversible ICO (rICO) concept, proposed by Fabian Vogelsteller in 2018–2020, was a fundraising model where investors could **withdraw their funding from a project's ICO gradually over time** if they lost confidence. In an rICO, participants reserve tokens and their contributed funds are released to the project slowly; investors always have the ability to "reverse" any remaining commitment by returning tokens to get back the unspent portion of their money. This provides a form of **downside protection** – if the project disappoints, investors reclaim their capital (at least the part not yet used) rather than being stuck with worthless tokens. FUSE issuance echoes this principle: buyers have an individual right to **reclaim their contributed capital** by returning tokens at any time, protecting them from losses beyond the opportunity cost of time. Both mechanisms enforce a kind of real-time accountability: the project/issuer cannot simply take all the money and run; funds are effectively **conditional on investor satisfaction**, and investors can pull the plug (rICO) or redeem (the token) if things are going south.

Difference: rICOs were primarily about **fund release over time** and did not incorporate any yield or token buyback dynamics. If an investor withdrew in a rICO, they received their remaining ETH back, and the corresponding tokens were cancelled – there was no concept of using funds to buy tokens on the market or rewarding remaining holders. In contrast, our mechanism **adds yield sharing and deflationary mechanics** to the equation, creating a more elaborate "flywheel" of incentives. Another difference is that rICO investments could be reclaimed only for a set period (e.g., within the ICO duration or a fixed window), making it more like a progressive commitment. The token's redemption right is *perpetual* for original buyers, not just during a funding period. Additionally, rICOs lack a gating mechanism – anyone could invest and withdraw; whereas our approach requires holding the token to participate in issuance, creating a structural demand driver. Finally, rICOs aimed to solve ICO trust issues, whereas FUSE aims to solve both trust and **market liquidity/vesting issues**. In summary, the FUSE Issuance Mechanism extends the rICO idea of **reversible commitments** by integrating it with tokenomics: not only can buyers reverse their purchase, but doing so (or selling) triggers **value-accretive actions (burns)** for the ecosystem. It's a more comprehensive solution marrying **investor protection with token value management**, which rICO alone did not address.

By examining these analogues, we see that the FUSE mechanism stands at an intersection of structured finance and crypto innovation. It offers the **capital protection of PPNs, the yield+upside mix of convertibles, the income strategy flavour of options selling (without capping upside), and the reversible commitment of rICOs** – all within one unified model. This interdisciplinary heritage underscores its robustness: each component of the mechanism has a rationale rooted in financial theory or prior practice, but their combination in this manner is novel. Next, we explore how all these pieces combine into a powerful feedback loop for the token's economy.

Flying Tulip – Perpetual PUT options

Similarity: The FUSE model shares a strong conceptual lineage with the Perpetual PUT structure used by Flying Tulip, both of which create a **principal-protected note** for contributors. This shared core ensures downside protection without a traditional vesting schedule.

Difference:

- **Primary Goal & Alignment:** Flying Tulip's scheme is primarily a **fundraising mechanism**. In contrast, FUSE is strategically designed to **replace traditional VC/insider vesting** by requiring key stakeholders to use the same FUSE terms, thereby aligning them with the community.

- **Yield Utility:** FUSE distributes yield from locked reserves directly to the CST holder as a continuous **Yield Right and Buyback & Burn**. Flying Tulip dedicates that yield as a secondary source for open-market **Buyback & Burn**.
- **Deflationary Exit:** FUSE's **Redeem** action removes tokens from circulation (allowing for immediate burn), directly contracting supply off-market. Flying Tulip's equivalent Divest action is purely a return of principal and does not inherently trigger a token burn.

Despite these differences, both mechanisms represent a significant innovation over traditional token sales by using collateralised principal protection to encourage long-term holding and introduce a powerful, non-dilutive, and market-responsive supply mechanism.