

Crypto Airdrop Pricing

Yunxia Bai

Tongji University

Xiaohan Liu

Tongji University

Bofu Zhang

bofu20131@163.com

Tongji University

Research Article

Keywords: cryptocurrency airdrops, token finance, option pricing, platform governance, Hyperliquid

Posted Date: September 3rd, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10240730/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Crypto Airdrop Pricing

April 21, 2026

Abstract

This paper develops a valuation framework for option-like claims on cryptocurrency airdrops. Standard option-pricing models are not designed for this setting because the payoff to an airdrop-linked claim depends not only on token price risk, but also on qualification uncertainty, wallet-level allocation uncertainty, vesting constraints, and event-driven dilution. I model an airdrop-linked claim as a compound contingent claim whose terminal payoff is the maximum of a strike-adjusted product of four primitives: an eligibility indicator, the token allocation, the tradable share of that allocation, and the post-event token price.

The paper embeds these institutional features in a reduced-form continuous-time environment under a pricing measure. Token prices follow a jump-diffusion process with stochastic volatility, while the airdrop event may induce a separate dilution shock at the distribution date. Eligibility is modeled as a Bernoulli random variable with a conditional qualification probability, allocation follows either a continuous or tiered distribution, and vesting is summarized by a tradable-share process. The general claim value is therefore the discounted expectation of a nonlinear payoff that jointly integrates market risk and project-design risk.

The paper contributes in three ways. First, it provides a tractable theory of airdrop-linked valuation that nests points-based qualification, variable token allocations, and lock-up schedules. Second, it shows that the most relevant comparative statics in airdrop markets concern eligibility, allocation, tradability, and dilution rather than only standard delta or vega. Third, it develops a calibration and robustness exercise built around the Hyperliquid genesis distribution. Using public information on the points program, token supply, launch valuation, and recipient scale, I show how the framework can be operationalized with strike grids, maturity grids, and event-dilution scenarios. The results indicate that airdrop-linked claim values are highly nonlinear in expected allocation and are materially compressed by dilution and vesting frictions. The framework therefore offers both a theoretical benchmark and a practical template for empirical work on token distributions, platform finance, and airdrop markets.

JEL classification: G12, G13, G32, O16

Keywords: cryptocurrency airdrops; token finance; option pricing; platform governance; Hyperliquid

Funding statement - Not Available

1 Introduction and Related Literature

Cryptocurrency airdrops have become an important financing and distribution technology in tokenized markets. Projects use airdrops to bootstrap user adoption, decentralize governance rights, reward early participation, and create ownership structures that differ sharply from venture-backed token launches [2, 3, 41]. The rise of decentralized finance more broadly has created an ecosystem in which token-based mechanisms substitute for many traditional financial intermediation functions [10, 12, 28, 51]. In practice, however, an airdrop creates a security-like contingent payoff at the wallet level. An address is valuable not only because of the token it may eventually receive, but because it embeds a bundle of contingent states: the address may or may not qualify, the number of tokens distributed may differ across users, those tokens may be locked or only partially tradable, and the event itself may affect market value through dilution and sell pressure.

These institutional details make airdrop valuation a natural corporate-finance problem in tokenized environments. A large literature studies token issuance, token-based platform financing, and governance design [1, 17, 18, 24, 25, 31, 36]. Another strand studies blockchain governance and the organizational implications of tokenized ownership [4, 15, 37, 54]. Yet the valuation of airdrop-linked claims has received far less formal treatment, even though market participants routinely price “airdrop rights” through over-the-counter negotiations, points markets, account purchases, strategic farming, and informal internal hurdle rates [22, 39].

Standard derivative-pricing models do not map neatly into this environment. A vanilla Black–Scholes contract assumes a tradable underlying, a well-defined strike, and a payoff that depends only on the terminal price of the underlying asset [9]. By contrast, the payoff from an airdrop-linked claim is jointly shaped by qualification risk, allocation uncertainty, vesting, and event-driven price impact. Relative to conventional option pricing, the relevant underlying is therefore a scaled and state-contingent object rather than the spot token alone. Relative to the token-finance literature, the problem is distinct because the goal is not to explain issuance or platform design, but to price a wallet-level contingent claim conditional on those design choices.

This paper proposes a parsimonious but extensible framework for this problem. The central object is an airdrop-linked option with payoff

$$X_T = \max(I_T N_T G_T S_T^{(a)} - K, 0), \quad (1)$$

where I_T is an eligibility indicator, N_T is the token allocation, G_T is the tradable share of the allocation at liquidation, $S_T^{(a)}$ is the token price after accounting for the airdrop event, and K is the strike or acquisition cost of the claim. This representation captures the economic structure of airdrop markets while allowing the pricing problem to be embedded in a standard risk-neutral environment.

The paper contributes in three ways. First, it gives a unified payoff representation that nests common real-world airdrop designs, including points-based qualification, rank-based allocations, and lock-up schedules. Second, it identifies the economically central valuation sensitivities: eligibility probability, expected allocation, tradable share, strike, maturity, and expected event dilution. Third, it provides a fully reproducible calibration and robustness exercise based on Hyperliquid’s genesis distribution. That exercise uses public information on token supply, points-based qualification, launch valuation, and recipient counts to show how the framework can be taken to data through wallet comparisons, strike grids, maturity grids, and event-dilution scenarios.

The paper is intended as a theory paper with a calibration component rather than a reduced-form event study. In that sense, it sits between the platform-finance literature, which studies why tokens are issued and how they support adoption [18, 25, 36], and the entrepreneurial-finance literature on token sales, which studies fundraising, signaling, and post-issuance outcomes [24, 31, 45]. The calibration exercise is designed to demonstrate economic magnitudes and

comparative statics, not to claim a full ex post reconstruction of every address-level Hyperliquid allocation.

The rest of the paper proceeds as follows. Section 1 continues with a review of the related literature on airdrops, derivative pricing, token finance, and platform governance, and explains why Hyperliquid provides a useful calibration illustration. Section 2 describes the economic environment, the airdrop-linked claim, the state dynamics for token prices, eligibility, allocation, and vesting, and interprets the model primitives as design choices. Section 3 derives the general pricing equation, tractable approximations, and the main comparative-static results. Section 4 explains how the model maps into public disclosures and empirical implementation. Section 5 presents the Hyperliquid case study, including baseline calibration, numerical results, robustness exercises, and design counterfactuals. Section 6 discusses extensions including incomplete markets, endogenous qualification, and general-equilibrium effects. Section 7 concludes.

1.1 Airdrops as financing and ownership-allocation mechanisms

An airdrop is often described colloquially as a marketing event or a retroactive user reward. That description is incomplete from a finance perspective. Airdrops allocate claims on future cash-flow rights, governance rights, or speculative resale value, and they do so through rules that map past user behavior into contingent ownership shares. In that sense, an airdrop is a financing arrangement and an ownership-allocation mechanism at the same time. Projects use it to allocate part of the economic surplus of a protocol to users rather than to outside investors, to compensate early adopters for supplying liquidity or attention, and to seed a governance constituency whose incentives are aligned with the network’s future growth [2, 3, 41].

This institutional interpretation connects airdrops directly to corporate-finance questions. In a conventional firm, ownership claims are distributed through equity issuance, option grants, acquisitions, or internal compensation plans [46, 47]. In tokenized markets, ownership is often distributed through smart-contract-based issuance, vesting, and retroactive rewards. The choice between a sale, a restricted allocation, a foundation reserve, or a user-facing airdrop affects both the project’s financing structure and the distribution of residual claims. The relevant question is therefore not only why a project chooses to airdrop, but what the resulting claims are worth to different recipients at the moment those claims are created.

The valuation problem is distinct from the one studied in the early ICO literature. ICO papers primarily analyze fundraising, project quality, information asymmetry, and post-issuance performance [24, 31, 45]. Airdrops instead distribute claims to users who may not contribute capital contemporaneously and who often face qualification uncertainty until close to the distribution date. The resulting object is neither straight token ownership nor a plain financing contract. It is a contingent claim on a future state-contingent token allocation.

1.2 Relation to the derivative-pricing literature

The model in this paper borrows from derivative pricing because the core object is a nonlinear payoff. At a formal level, the payoff looks like an option written on an endogenous underlying. Yet the mapping to standard option pricing is imperfect. Black–Scholes-style settings assume a tradable underlying and a contract whose payoff depends on the terminal realization of that underlying alone [9]. By contrast, an airdrop-linked claim depends on the token price only after it has been filtered through qualification, allocation, and tradability rules. Even if the token price process were identical to the one in a standard option model, the contract would remain institutionally different because part of the effective underlying is produced by project design rather than by market trading.

This distinction is important for two reasons. First, some sources of uncertainty are not spanned by liquid markets. Qualification risk, sybil filtering, rank-based allocations, and vesting

enforcement may be largely idiosyncratic to a protocol and a wallet. Second, the underlying token may itself be affected by the issuance event through dilution, liquidity shocks, or endogenous information revelation. In conventional option theory, issuance and payout rules are typically not modeled as part of the underlying’s law of motion. In token markets they frequently are.

Jump-diffusion and stochastic-volatility methods nevertheless remain natural tools because token prices around snapshots, announcements, and unlock events display the same kind of discontinuities that motivated the classic extensions of Black–Scholes in the first place [8, 14, 21, 30, 35, 43]. Empirical work on alternative option pricing models has shown that jump and stochastic-volatility extensions materially improve pricing performance relative to the plain Black–Scholes benchmark [6]. The contribution here is therefore not to reject option theory, but to adapt it to an institutional setting in which the relevant scale of the underlying is produced jointly by the market and by a distribution mechanism. Relatedly, the compound nature of the airdrop claim—where the effective underlying is itself a product of several random variables—connects to the literature on options on multiple assets and exchange options [20, 26, 42].

1.3 Relation to token finance and platform governance

The paper also sits inside the growing literature on token-based platform finance. Recent work shows that tokens can coordinate platform participation, support ecosystem growth, and distribute control rights in ways that differ from standard corporate securities [1, 16–18, 25, 36]. Governance papers emphasize that blockchains and tokens reshape the allocation of control and monitoring rights [4, 15, 37, 54]. Airdrops are one of the main operational mechanisms through which those rights are dispersed in practice.

What is missing in that literature is a valuation language for the transitional object that exists before the token is fully held and traded. A user who has accumulated points, trading volume, or protocol activity has not yet received the tokens, but has acquired an economically meaningful chance of receiving them. The *ex ante* value of that chance can affect user behavior, account acquisition, secondary OTC pricing, and even platform competition. A theory of airdrop pricing therefore complements the token-finance literature by pricing the contingent rights created by a token-distribution rule. This connects to the broader literature on two-sided platform economics, where distribution and participation rules shape equilibrium platform value [13, 48, 50].

This perspective also helps organize the boundary between governance and finance. Governance researchers often care about who ends up controlling a protocol and how distribution rules shape the legitimacy of that control. Finance researchers care about who bears risk, who receives upside, and how those claims should be priced. Airdrops sit at that boundary because they are simultaneously a governance design instrument and a contingent financial claim. The model below is intentionally reduced form, but its variables map directly into governance features: eligibility corresponds to screening and participation rules, allocation to the sharing rule, tradability to vesting and lock-up design, and event adjustment to supply overhang and market impact.

1.4 Why airdrop rights are a distinct asset-pricing object

The most useful way to frame the contribution is to note that an airdrop right is neither a spot token position nor a standard derivative. It is not a spot position because the address does not necessarily receive the tokens, does not know with certainty how many it will receive, and may be unable to trade them immediately. It is not a standard derivative because the effective underlying contains non-market primitives whose law is determined by protocol rules and wallet behavior. The claim is therefore best understood as a compound contingent claim on a protocol-designed scale factor.

That distinction matters practically. When users discuss whether it is worth “farming” a protocol, purchasing a points-bearing account, or holding a token through a snapshot, they are

implicitly pricing a future airdrop right. When an OTC desk discounts a wallet’s expected airdrop, it is assigning value to a contingent claim with an endogenous scale, not simply to a token price. When a project chooses between full unlock and staged vesting, or between broad distribution and concentrated rewards, it is changing the parameters of that claim directly. The language of airdrop pricing therefore helps reconcile user behavior, project design, and secondary-market valuation.

Table 1: Institutional channels in crypto airdrops

Institutional feature	fea- ture	Model ob- ject	Direct valuation channel	Typical empirical proxy
Qualification or sybil filter	or	p_t	Changes the probability that the claim has any positive scale at all	Score thresholds, wallet labels, rule announcements
Allocation rule or points schedule	or	N_T	Alters the size of the effective underlying and can move a wallet across moneyiness regions	Points balances, trading volume, tiers, snapshots
Lock-up or vesting design		G_T	Compresses immediate realizable value and delays monetization even when allocation is large	Tokenomics tables, unlock calendars, vesting contracts
Distribution-induced pressure	sell	δ_a	Lowers the event-adjusted token price and reduces in-the-money states	Distributed share of supply, recipient concentration, float expansion
Market volatility and jump exposure		$S_T^{(a)}$	Governs convexity value and determines the dispersion of the monetizable component	Historical returns, announcement windows, listing events

Notes. The table emphasizes that the main drivers of airdrop valuation are partly market variables and partly protocol-design choices. The purpose of the model is to place these channels inside a single pricing object rather than treating them in isolation.

1.5 Why Hyperliquid is a useful illustration

The Hyperliquid case is especially useful because the institutional features that motivate the model are unusually visible. The points system created a forward-looking probability distribution over qualification and allocation; the genesis documentation pinned down aggregate supply, recipient scale, and unlock status; and the launch produced a public market price that can anchor the event-adjusted value of the claim [32, 33, 44, 49]. Hyperliquid is not unique, but it is disciplined enough to illustrate how the model can be populated with public information while still leaving genuine uncertainty at the wallet level.

The case is also conceptually attractive because it exposes the interaction between platform growth and contingent claims. Users accumulated points through activity on a trading venue, anticipating that those points might convert into future token ownership. That dynamic is exactly the bridge between platform finance and derivative-style valuation that the paper aims to formalize. The airdrop was not merely a post hoc transfer; it was part of the incentive environment under which the platform grew.

2 Model Setup and State Dynamics

2.1 Timing

Consider a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{Q})$, where $\mathbb{Q}$ denotes the pricing measure. Time is continuous. At current date t , an address or investor values a future airdrop-linked claim. The project announces, snapshots, or allocates the airdrop at date T_a . Tokens become partially or fully tradable at date T , where $T \geq T_a$. In the simplest version of the model, the claim is liquidated at T .

This timing structure separates three economically distinct moments. Before T_a , the market forms beliefs about qualification and allocation. At T_a , information arrives about the distribution event and may move prices through dilution or sentiment. Between T_a and T , vesting or lock-up restrictions affect what fraction of the distribution can actually be sold.

2.2 Primitive random variables

The terminal airdrop-linked payoff depends on four primitives:

- $I \in \{0, 1\}$: an indicator for whether the address is eligible;
- $N_T \geq 0$: the number of tokens allocated conditional on the project rules;
- $G_T \in [0, 1]$: the fraction of allocated tokens that is unlocked and tradable at date T ;
- $S_T^{(a)} > 0$: the token price at date T after incorporating airdrop-event effects.

Let $K \geq 0$ denote the strike or acquisition cost of the claim. If one interprets the contract as the right to acquire an airdrop-linked cash flow, then K is the exercise price. If instead one interprets the contract as an OTC agreement to buy an address's future airdrop rights, then K can be read as the contractual transfer price.

Table 2: Core notation

Symbol	Interpretation
I	Eligibility indicator for the address
N_T	Token allocation received by the address
G_T	Tradable fraction of allocated tokens at liquidation
$S_T^{(a)}$	Event-adjusted token price at liquidation
K	Strike or acquisition cost of the claim
p_t	Conditional qualification probability
δ_a	Event-dilution or sell-pressure parameter
A_t	Discounted value of the airdrop entitlement
C_t	Value of the call-like airdrop-linked claim

2.3 Payoff

The baseline airdrop-linked call claim has terminal payoff

$$X_T = \max(I_T N_T G_T S_T^{(a)} - K, 0). \quad (2)$$

This payoff differs from a standard call in two important ways. First, the effective underlying is not just the token price; it is the product $I_T N_T G_T S_T^{(a)}$. Second, the first three multiplicative components are institution-specific variables that reflect platform rules rather than only financial market dynamics.

2.4 Multi-date cash-flow interpretation

Many airdrops do not unlock all tokens at once. Suppose there are M future unlock dates $\tau_1, \dots, \tau_M$ with shares $g_1, \dots, g_M$ such that $\sum_{m=1}^M g_m = 1$. Then the time- t fundamental value of the received airdrop can be written as

$$V_t^{\text{airdrop}} = \mathbb{E}_t^{\mathbb{Q}} \left[I_T \sum_{m=1}^M e^{-r(\tau_m - t)} N_T g_m S_{\tau_m} \right]. \quad (3)$$

The single-date payoff X_T is therefore a reduced-form representation of a more general stream of discounted unlock cash flows.

2.5 Token-price dynamics

Token prices around airdrop events commonly display jumps, changing volatility, and event-driven order-flow shocks. To reflect these features, let the pre-event token price follow a jump-diffusion process with stochastic volatility:

$$\frac{dS_t}{S_t} = (r - q - \lambda_J \kappa_J) dt + \sqrt{v_t} dW_t + (J_t - 1) dN_t, \quad (4)$$

$$dv_t = \kappa_v (\theta_v - v_t) dt + \xi_v \sqrt{v_t} dB_t. \quad (5)$$

Here, r is the risk-free rate, q is a convenience-yield-type adjustment that can also absorb dilution or holding costs, N_t is a Poisson process with intensity λ_J , and J_t is the random multiplicative jump size. The variance process v_t mean reverts to θ_v with speed κ_v . The Brownian motions W_t and B_t may be correlated. The jump-diffusion component follows the logic of Merton [43], while stochastic volatility follows Heston [30]. Affine jump-diffusion methods also provide a broader analytical background for this structure [21]. Empirical evidence from crypto markets supports the presence of jumps and time-varying volatility in token returns [23, 38, 40], and the broader option-pricing literature documents the importance of these features across asset classes [6, 8, 35].

2.6 Airdrop event jump

The airdrop itself may move the token price because of expected dilution, unlock-induced sell pressure, or an information effect from the project's announcement. Let the price at the event obey

$$S_{T_a} = S_{T_a^-} (1 - \delta_a + \varepsilon_a), \quad (6)$$

where $\delta_a \geq 0$ denotes the predictable component of dilution and selling pressure, and ε_a is an idiosyncratic residual shock. A simple reduced-form parametrization is

$$\delta_a = \eta \cdot \frac{\text{airdrop supply}}{\text{circulating supply}}. \quad (7)$$

This term isolates a feature that is central in token markets but absent from standard option models: the contingent claim can itself affect the value of the asset on which it is written.

2.7 Eligibility

Eligibility is modeled as a Bernoulli random variable,

$$I \sim \text{Bernoulli}(p_t), \quad (8)$$

where

$$p_t = \mathbb{P}(I = 1 \mid \mathcal{F}_t). \quad (9)$$

For applications, p_t can be represented as a reduced-form function of observed activity variables:

$$p_t = \frac{1}{1 + \exp[-(\beta_0 + \beta^\top X_t)]}. \quad (10)$$

The vector X_t may include wallet age, active days, transaction count, bridge volume, gas spending, staking participation, KYC status, and sybil-detection features.

2.8 Allocation

Conditional on eligibility, the number of tokens allocated is uncertain. A flexible continuous specification is

$$\log N_T \sim \mathcal{N}(\mu_N, \sigma_N^2). \quad (11)$$

In cases where allocation is based on tiers, it is more natural to use a discrete distribution:

$$N_T \in \{n_1, n_2, \dots, n_m\}, \quad \mathbb{P}(N_T = n_j) = \pi_j, \quad j = 1, \dots, m. \quad (12)$$

More generally, one can write

$$N_T = f(Z_T) \varepsilon_T, \quad (13)$$

where Z_T captures publicly described scoring variables and ε_T summarizes residual discretion, rule changes, and implementation noise.

2.9 Vesting and tradability

Let $G_T \in [0, 1]$ denote the fraction of the allocated tokens that is unlocked and tradable at T . For a fully unlocked airdrop, $G_T = 1$. For a cliff-plus-linear-vesting schedule, one may write

$$G_T = g_0 + (1 - g_0) \frac{T - T_a}{L}, \quad T_a \leq T \leq T_a + L, \quad (14)$$

where g_0 is the immediately unlocked share and L is the vesting horizon.

2.10 Qualification as a screening technology

The qualification variable is not merely a statistical probability. Economically, it summarizes a screening technology chosen by the project. Protocols decide which actions count, how much discretion remains after the formal rules are published, and how aggressively they filter sybil behavior. Those design choices affect both the distribution of ownership and the perceived legitimacy of the distribution. In the model, all of that institutional richness is summarized by p_t , but the interpretation is broader: a higher p_t means that participation is more likely to translate into monetizable ownership.

This view matters for both comparative statics and empirical work. A protocol with permissive qualification may spread token ownership widely, but it may also encourage opportunistic activity and weaken the information content of points or historical usage. A protocol with strict qualification can increase expected value per qualifying wallet, but it also creates more binary risk and more demand for off-platform information about what behavior “counts.” The price of airdrop rights should therefore react not only to the scale of promised distribution, but to how credible and predictable the project’s screening technology appears to be.

2.11 Allocation as a contingent sharing rule

The allocation variable N_T plays the role that a payout ratio or sharing rule would play in a more traditional finance setting. Two wallets may both be eligible and face the same token price, yet still have sharply different values because the distribution mechanism converts historical behavior into rewards nonlinearly. In practical terms, N_T encodes a contract between the project and the user base: the contract says how much ownership follows from past contribution, liquidity provision, trading activity, social participation, or capital commitment.

Thinking of N_T as a sharing rule clarifies why the allocation distribution deserves separate treatment from token price risk. A project may widen the overall airdrop pool without changing the token’s fundamental prospects, or it may compress allocations through tiering and caps even if the token price outlook is unchanged. These are different economic decisions. One changes how much claim is distributed per qualifying wallet; the other changes the value of each unit of claim in secondary markets.

For the same reason, expected allocation is closely related to airdrop moneyness. A high strike can make a medium wallet effectively speculative, while the same strike can be almost irrelevant for a top-tier wallet with a much larger implied allocation. This is why wallet-level heterogeneity is central in the model rather than a secondary detail.

2.12 Tradability as dynamic supply control

Tradability is the most direct bridge between tokenomics and valuation. A vesting schedule does not eliminate value, but it changes both timing and marketability. From the holder’s perspective, a lock-up reduces the immediately monetizable component of the claim. From the market’s perspective, staggered unlocks smooth secondary-market supply, alter expected sell pressure, and may support the token price relative to a fully unlocked counterfactual. The variable G_T compresses these effects into a tradable-share term, while the multi-date cash-flow formulation in Equation (3) allows one to model richer unlock structures when needed.

This design channel is especially important because airdrop discussions often treat “amount received” and “amount effectively monetizable” as if they were the same object. They are not. A wallet that receives a large nominal allocation under a twelve-month vesting schedule may have less near-term value than a smaller wallet receiving immediately unlocked tokens. The model makes that comparison explicit.

2.13 Event adjustment as an equilibrium feedback

The event-adjusted price $S_T^{(a)}$ captures the idea that issuance and valuation are jointly determined. In standard option settings, the contract’s scale is typically small relative to the underlying asset and does not materially alter the law of motion of the underlying. In token distributions, the event itself can change float, expectations, and resale pressure enough to move market price. That feedback is represented by δ_a and the residual event shock, but its deeper interpretation is equilibrium based: distribution design changes market outcomes, and those market outcomes feed back into the value of the distributed claim.

This feedback channel is also where the model is closest to platform-finance arguments. If a project distributes tokens broadly to active users, the distribution may deepen engagement and strengthen governance legitimacy [48, 50]. But it may also generate immediate resale pressure, especially if recipients are financially constrained or do not share long-term beliefs about the project. The sign and magnitude of the net price effect is empirical, but the fact that it should be modeled directly is conceptual.

2.14 The scale factor and effective exposure

It is useful to group qualification, allocation, and tradability into a single scale factor

$$M_T = I_T N_T G_T. \quad (15)$$

The payoff can then be written compactly as

$$X_T = \max(M_T S_T^{(a)} - K, 0). \quad (16)$$

This representation emphasizes that an airdrop claim exposes the holder to two conceptually distinct objects: a market price process and a protocol-designed scale factor. Much of the paper’s logic can be summarized as a study of how those two objects interact.

The scale-factor perspective also clarifies why the model is useful outside the precise case of token airdrops. Any allocation mechanism that grants users uncertain future tokens, rebates, or governance shares based on current activity can be treated similarly. The particular institutional language differs across settings, but the core economics remain the same: expected value depends on a designed scale factor multiplying a risky future market value.

Table 3: From model primitive to empirical object

Primitive	Empirical interpretation	Typical observable inputs	Main measurement challenge
p_t	Probability that a wallet survives rule evaluation and sybil filtering	Participation histories, scoreboards, public rule changes, labels	Qualification is only partially observable before the official distribution list
N_T	Reward size conditional on qualification	Points tallies, volume rankings, tier descriptions, total distributed supply	Mapping from public scores to wallet-specific allocations is often incomplete
G_T	Fraction of nominal rewards immediately monetizable	Tokenomics tables, vesting contracts, unlock schedules	True realizable liquidity may differ from contractual unlock due to market depth
δ_a	Event-specific price discount from issuance and resale pressure	Airdrop share of supply, recipient concentration, float expansion	Dilution and information effects are hard to separate empirically
$S_T^{(a)}$	Post-event market value of each unit of the claim	Exchange prices, OTC indications, realized volatility, jump windows	Liquidity conditions and announcement timing can contaminate simple price measures

3 Pricing Theory and Comparative Statics

3.1 General pricing formula

Under the risk-neutral measure $\mathbb{Q}$, the time- t price of the airdrop-linked call is

$$C_t = \mathbb{E}_t^{\mathbb{Q}} \left[e^{-r(T-t)} \max(I_T N_T G_T S_T^{(a)} - K, 0) \right]. \quad (17)$$

This equation is the central object of the paper. Pricing airdrop-linked claims requires integrating over market uncertainty, qualification risk, allocation uncertainty, and vesting restrictions at the same time.

3.2 Fundamental airdrop value

Before considering the option strike K , it is helpful to isolate the expected discounted value of the airdrop itself:

$$A_t = \mathbb{E}_t^{\mathbb{Q}} \left[e^{-r(T-t)} I_T N_T G_T S_T^{(a)} \right]. \quad (18)$$

This object can be interpreted as the shadow value of the address's future airdrop entitlement. The call price C_t is then the option value associated with this entitlement when it must exceed the strike K to be worth exercising.

3.3 Conditional independence approximation

The general pricing formula may be difficult to solve in closed form. A tractable approximation arises under the assumption that, conditional on $\mathcal{F}_t$, the variables I , N_T , and G_T are independent of $S_T^{(a)}$. Then

$$C_t \approx p_t \mathbb{E}_t^{\mathbb{Q}} \left[e^{-r(T-t)} \max(N_T G_T S_T^{(a)} - K, 0) \right]. \quad (19)$$

This approximation has a natural interpretation: eligibility rescales the value of the conditional claim.

3.4 Mean-scaling approximation

A further simplification replaces allocation and tradability by their conditional means:

$$\bar{N}_t = \mathbb{E}_t[N_T], \quad \bar{G}_t = \mathbb{E}_t[G_T]. \quad (20)$$

Under this mean-scaling approximation,

$$C_t \approx p_t \mathbb{E}_t^{\mathbb{Q}} \left[e^{-r(T-t)} \max(\bar{N}_t \bar{G}_t S_T^{(a)} - K, 0) \right]. \quad (21)$$

Equivalently, if one defines an effective scaled underlying

$$S_t^* = \bar{N}_t \bar{G}_t S_t, \quad (22)$$

the problem becomes close to pricing a jump-diffusion call on S_t^* , multiplied by the probability of qualification. This structure echoes compound-option and exchange-option settings in which the effective underlying is a product or ratio of traded and non-traded quantities [20, 26, 42].

3.5 Scale convexity and effective moneyness

The mean-scaling representation makes it useful to isolate a simple one-dimensional object:

$$\phi(a, K) = \mathbb{E}_t^{\mathbb{Q}} \left[e^{-r(T-t)} \max(a S_T^{(a)} - K, 0) \right], \quad a \geq 0. \quad (23)$$

This function prices a call with scale factor a . The airdrop problem differs from a vanilla option problem because the scale factor is itself endogenous, but once a is fixed the usual convexity logic applies.

Lemma 1 (Monotonicity and convexity in scale). *For fixed K , the function $\phi(a, K)$ is nondecreasing and convex in a . For fixed a , it is nonincreasing in K .*

Lemma 1 is conceptually central. It says that the airdrop claim does not merely rise proportionally with expected scale. Because the payoff is truncated at zero, larger scale also changes the share of states in which the claim finishes in the money. The wallet-level importance of qualification, allocation, and tradability therefore comes from both a first-moment effect and a convexity effect.

It is also useful to define an effective moneyiness ratio

$$\mathcal{M}_t = \frac{A_t}{K} = \frac{\mathbb{E}_t^{\mathbb{Q}} \left[e^{-r(T-t)} I_T N_T G_T S_T^{(a)} \right]}{K}. \quad (24)$$

When $\mathcal{M}_t$ is low, the contract behaves like a speculative option on qualification and allocation upside. When $\mathcal{M}_t$ is high, the claim behaves more like a discounted entitlement with limited truncation. This ratio is not sufficient for pricing, but it provides a compact way to compare wallets and contract designs across the same strike.

Assumption 1 (Conditional separability). Conditional on $\mathcal{F}_t$, the variables I , N_T , and G_T are independent of the event-adjusted price $S_T^{(a)}$.

Theorem 1 (Scaled-claim representation). *Under Assumption 1, the value of the airdrop-linked call satisfies*

$$C_t = p_t \mathbb{E}_t^{\mathbb{Q}} \left[e^{-r(T-t)} \max(N_T G_T S_T^{(a)} - K, 0) \right]. \quad (25)$$

If, in addition, allocation and tradability are replaced by their conditional means $(\bar{N}_t, \bar{G}_t)$, then the claim admits the approximation

$$C_t \approx p_t \mathbb{E}_t^{\mathbb{Q}} \left[e^{-r(T-t)} \max(\bar{N}_t \bar{G}_t S_T^{(a)} - K, 0) \right], \quad (26)$$

so that the airdrop-linked claim can be interpreted as a scaled call on an endogenous underlying.

Remark 1. Theorem 1 is useful because it cleanly separates market risk from distribution-design risk. The first enters through the law of $S_T^{(a)}$, while the second enters through the scaling terms $(p_t, \bar{N}_t, \bar{G}_t)$ and the event-dilution parameter δ_a .

3.6 Risk-neutral valuation versus subjective valuation

Equation (17) is written under a pricing measure because the paper's main objective is to characterize a benchmark no-arbitrage-style value in a reduced-form environment. In practice, however, users may not evaluate airdrops under a common pricing kernel. Retail recipients may use subjective probabilities, venture-backed traders may apply hurdle rates, and OTC counterparties may add discounts for unverifiable qualification or for operational frictions associated with account transfer. Behavioral considerations such as mental accounting [53], prospect-theory-style probability weighting [34], and extrapolative beliefs [7, 52] may further distort subjective valuations away from the risk-neutral benchmark.

This distinction does not overturn the usefulness of the risk-neutral benchmark. Instead, it clarifies the role of each object. The risk-neutral price is a clean theoretical quantity that separates market risk from design risk. Subjective or negotiated values can then be interpreted as departures from that benchmark due to heterogeneous beliefs, incomplete hedging opportunities, or transaction costs [19]. The same decomposition is often useful in corporate finance more broadly: a model first provides a frictionless benchmark, and empirical or institutional analysis then explains persistent wedges around it.

3.7 Decomposition of claim value

The model also permits a simple accounting decomposition. Write the option value as

$$C_t = A_t - D_t, \quad (27)$$

where

$$D_t = \mathbb{E}_t^{\mathbb{Q}} \left[e^{-r(T-t)} \min(I_T N_T G_T S_T^{(a)}, K) \right] - K \mathbb{Q} \left(I_T N_T G_T S_T^{(a)} \geq K \right). \quad (28)$$

The term D_t is the moneyness discount: it captures the part of the shadow airdrop entitlement that does not survive the strike truncation. This decomposition is useful in calibration because it separates a wallet's total entitlement value from the fraction of that value that is actually monetizable under a given contract structure.

The same logic explains why similar-looking airdrop programs can generate very different OTC prices. Two protocols may have comparable expected aggregate distribution values, but if one produces many small, out-of-the-money wallet entitlements while the other produces fewer large, in-the-money entitlements, the values of their wallet-level claims can be markedly different even before any difference in token price prospects is considered.

3.8 Comparative statics

This subsection summarizes the model's main economic implications. Formal proofs are collected in Appendix A.

Proposition 1 (Eligibility effect). *Holding the conditional distribution of N_T , G_T , and $S_T^{(a)}$ fixed, the option value C_t is increasing in the eligibility probability p_t .*

Proposition 2 (Allocation effect). *Holding all other primitives fixed, C_t is increasing in expected allocation. Under the mean-scaling approximation,*

$$\frac{\partial C_t}{\partial \bar{N}_t} > 0. \quad (29)$$

Proposition 3 (Vesting effect). *Holding all else constant, C_t is increasing in the expected tradable share:*

$$\frac{\partial C_t}{\partial \bar{G}_t} > 0. \quad (30)$$

Proposition 4 (Dilution effect). *Holding the distribution of all other variables fixed, C_t is decreasing in the expected event-dilution parameter δ_a :*

$$\frac{\partial C_t}{\partial \delta_a} < 0. \quad (31)$$

Proposition 5 (Volatility and jump risk). *Under standard convexity conditions, greater return volatility or jump intensity tends to raise the value of the call-like claim, though the magnitude of the effect depends on the interaction between jump risk and event dilution.*

Proposition 6 (Strike effect). *Holding all other primitives fixed, the option value C_t is decreasing in the strike K .*

Proposition 7 (Maturity under convex-order expansion). *Suppose a longer liquidation horizon preserves the conditional mean of the event-adjusted token price while increasing its conditional distribution in convex order. Then the airdrop-linked call value is weakly increasing in maturity.*

Proposition 8 (Wallet ordering). *Under the mean-scaling approximation, if two wallets share the same strike, horizon, and token-price process, and if*

$$p_t^A \bar{N}_t^A \bar{G}_t^A \geq p_t^B \bar{N}_t^B \bar{G}_t^B, \quad (32)$$

then wallet A weakly dominates wallet B in claim value.

Corollary 1 (Fully unlocked benchmark). *If qualification is deterministic and tokens are immediately tradable, so that $I = 1$ and $G_T = 1$ almost surely, the model collapses to an allocation-scaled option on the token price:*

$$X_T = \max(N_T S_T^{(a)} - K, 0). \quad (33)$$

Proposition 9 (Unlock-dilution interaction). *If a lower tradable share also lowers expected sell pressure, then the net effect of tighter vesting on option value is ambiguous ex ante. The direct payoff-compression effect works through G_T , while the indirect price-support effect works through a lower δ_a .*

Taken together, these propositions show that the most important sensitivities in an airdrop market are not only delta and vega. The more distinctive state sensitivities are

$$\frac{\partial C_t}{\partial p_t}, \quad \frac{\partial C_t}{\partial \bar{N}_t}, \quad \frac{\partial C_t}{\partial \bar{G}_t}, \quad \frac{\partial C_t}{\partial \delta_a}. \quad (34)$$

Several economic interpretations follow immediately. First, strike sensitivity matters differently in airdrop markets than in standard options because expected allocation can change the effective moneyness of otherwise similar wallets. Second, maturity has value not only because of diffusion volatility, but because longer horizons allow more time for vesting to unlock and for uncertainty about qualification or dilution to resolve. Third, wallet ordering provides a practical ranking rule for OTC traders and for empirical researchers who want to compare ex ante expected value across observable wallet types.

The unlock-dilution interaction deserves particular emphasis. Practitioners sometimes speak as if tighter vesting is always bad for recipients because it lowers immediate liquidity. That is true in partial equilibrium for any fixed token price process. But in general equilibrium the same vesting schedule may reduce expected sell pressure enough to support the token price. The model keeps those effects conceptually separate. This is useful not only for pricing but also for project design, because it clarifies the tradeoff between immediate monetizability and market stabilization.

Finally, the propositions explain why airdrop pricing naturally produces coarse wallet taxonomies. In a standard option market, fine-grained differences in account history are usually irrelevant because the contract is publicly standardized. In an airdrop market, small differences in wallet activity can affect p_t , $\bar{N}_t$, or both. Those changes then feed through a convex payoff map. The result is a market in which “low activity,” “average recipient,” and “high activity” are not merely descriptive labels but economically distinct regions of the claim-value distribution.

4 Calibration Framework

4.1 Why calibration matters in this setting

The theory developed above is intentionally general. That generality is useful conceptually, but a valuation paper must also show how the primitives can be disciplined in practice. Airdrop markets pose a special challenge because no single database contains all relevant information. Projects disclose some tokenomics variables publicly, on-chain data reveal some dimensions of user activity, and market prices reveal the realized value of the token once it begins trading. The calibration problem is therefore one of partial observability: the researcher must map incomplete public information into the model’s latent variables in a disciplined way.

The central strategy of this paper is to separate directly observed objects from inferred objects. Directly observed quantities include the announced token supply, the distributed share of supply, vesting rules, launch dates, and benchmark token prices after trading begins. Inferred quantities include the wallet-level eligibility probability, the conditional allocation distribution, and the equilibrium event-dilution parameter. This division makes the calibration transparent. It also clarifies where the main identification assumptions sit.

4.2 From public disclosures to model primitives

The first step is to map aggregate protocol disclosures into the model’s wallet-level variables. If a project discloses total genesis allocation N^{agg} and the approximate number of recipient wallets M^{rec} , then a natural benchmark for the conditional mean allocation is

$$\mathbb{E}[N_T \mid I = 1] \approx \frac{N^{\text{agg}}}{M^{\text{rec}}}. \quad (35)$$

This is not a full characterization of the allocation distribution, but it anchors the scale of a representative eligible wallet. If scoreboards or tier descriptions are available, the researcher can then impose additional cross-sectional structure around that aggregate anchor.

Qualification is more subtle. In some airdrops, qualification is almost deterministic once activity thresholds are met. In others, the formal rules are underspecified and ex post filtering remains important. A practical approach is to model

$$p_t = \mathbb{P}(\text{meets public rules} \mid \mathcal{F}_t) \times \mathbb{P}(\text{survives discretionary filtering} \mid \mathcal{F}_t). \quad (36)$$

The first term can often be approximated from observable activity; the second term reflects sybil screening, exclusions, or last-minute rule updates and is harder to pin down without historical examples from the same project family.

Tradability is typically easier. Whitepapers, tokenomics dashboards, and foundation announcements often state whether distributions are fully unlocked, partially unlocked, or cliff vested. Those disclosures can be translated directly into G_T or into the multi-date unlock sequence in Equation (3).

Event dilution is the most delicate object because it mixes supply expansion with a market-response component. The empirical strategy in this paper is therefore to treat δ_a as a scenario parameter disciplined by the airdrop’s scale relative to float rather than as a fully estimated reduced-form coefficient.

Table 4: Parameter mapping from public information to model inputs

Model input	Primary source	Mapping logic	Residual assumption
S_t and $S_T^{(a)}$	Exchange or first-trade price	Use observed launch price as a post-event anchor and back out pre-event value from the dilution scenario	Requires a maintained event-adjustment convention
$\bar{N}_t$	Announced aggregate distribution and recipient count	Convert public aggregate allocation into representative wallet scale	Requires an assumed cross-sectional wallet taxonomy
p_t	Public rules and wallet behavior	Map observed activity into qualification likelihood using reduced-form screening logic	Filtering and discretionary exclusions remain latent
$\bar{G}_t$	Tokenomics and unlock schedule	Translate vesting into immediate tradable share or dated cash-flow shares	Market depth may still constrain realizability
δ_a	Distributed share of supply and market context	Use scenario analysis based on expected float expansion and sell pressure	Not point identified without richer event-study evidence
$\sigma, \lambda_J, \mu_J, \sigma_J$	Market microstructure and comparable crypto launches	Choose high-volatility benchmark values appropriate for short-horizon crypto trading	Calibration is illustrative rather than fully structural

Notes. The table summarizes the empirical philosophy of the paper. Some inputs are observed, some are inferred from public protocol design, and some are varied in scenario form because the available data do not identify them sharply.

4.3 Wallet taxonomies as a bridge between theory and data

Because address-level allocations are usually unobserved before the event, it is useful to work with wallet taxonomies. A taxonomy groups addresses into economically interpretable bins such as low activity, average recipient, medium activity, and high activity. Each bin is assigned a representative qualification probability and expected allocation. The purpose is not to claim that all wallets within a bin are identical. Rather, it is to approximate the cross section with a set of economically meaningful representative states.

This approach mirrors standard calibration practice in corporate finance and macro-finance, where heterogeneous firms or households are often represented by a small number of bins or types. In the present setting it is especially attractive because airdrop rules are themselves frequently tiered. Even when a project does not disclose an explicit tier structure, community expectations often organize around informal categories such as “power users,” “average farmers,” and “tourists.” Those categories are useful because they line up with different regions of effective moneyiness.

Table 5: Representative wallet taxonomy used in the calibration

Wallet type	p_t	$\bar{N}_t$	Economic interpretation
Low activity	0.55	1,200	Sparse interaction history; meaningful chance of exclusion; small expected allocation even if included
Implied average recipient	0.70	3,298	Representative beneficiary implied by aggregate recipient count and total distribution size
Medium activity	0.75	4,000	Consistent participation and moderate confidence of qualification; used as baseline valuation case
High activity	0.90	9,000	Heavy participation or top-tier standing; large expected scale and high in-the-money likelihood

4.4 What the calibration does and does not claim

The calibration in this paper is deliberately modest in its ambitions. It is designed to show that the model can be brought to public data, to identify economically meaningful orders of magnitude, and to illustrate which primitives matter most. It is not designed to recover the entire ex post distribution of Hyperliquid wallet allocations, to infer the precise internal scoring algorithm, or to estimate equilibrium token demand around launch with structural precision.

That limited claim is a feature rather than a weakness for a theory paper. The point is to separate what the model says from what the public record reveals. When a project provides unusually rich disclosures, the same framework could be estimated more structurally. When disclosures are thinner, the model still provides disciplined scenario analysis. In either case, the mapping from narrative institutional detail to formal valuation is the main contribution.

4.5 Identification limits and sources of model risk

Three sources of model risk are especially important. First, qualification risk is only partially observable before the official distribution list is released. Second, allocation rules may contain discontinuities or discretionary adjustments that are difficult to infer from public information alone. Third, event dilution is an equilibrium object that depends on market microstructure, recipient behavior, and the broader crypto environment at launch. Scenario analysis is therefore not a convenience; it is a necessary part of disciplined valuation.

These limitations are also a guide for future empirical work. Researchers with wallet-level outcome data could estimate the mapping from points and user behavior into realized qualification and allocation. Researchers with detailed transaction-level data could study how unlocks and float expansion affect post-airdrop liquidity and price impact. The current paper is best viewed as providing the pricing architecture into which those richer empirical measurements can later be inserted.

5 Hyperliquid Case Study

5.1 Why Hyperliquid is a disciplined calibration case

Hyperliquid is a useful demonstration case because the public record identifies the main institutional objects required by the model while still leaving address-level uncertainty. According to the official points documentation, the protocol distributed one million points weekly beginning on November 1, 2023 for six months, followed by an “L1 phase” starting on May 29, 2024 with 700,000 points distributed weekly through November 2024 [33]. The same documentation emphasizes that criteria could be updated and prior distributions revised, which makes qualification and wallet-level allocation fundamentally stochastic from the ex ante perspective.

The official Hyper Foundation genesis announcement states that the maximum HYPE supply is one billion tokens and that 310 million tokens, or 31% of supply, were allocated to the genesis distribution [32]. Those genesis tokens were fully unlocked at launch, whereas contributor tokens were subject to materially longer lockups. Contemporaneous market coverage reported a launch valuation of approximately \$3.90 on November 29, 2024 and described the genesis distribution as reaching slightly more than 94,000 users [44, 49]. These facts make Hyperliquid attractive for calibration because they pin down aggregate distribution mass, the qualification mechanism, the unlock regime, and a benchmark market price.

This section uses Hyperliquid as a calibration exercise rather than as a forensic reconstruction. The exact mapping from points to wallet-specific HYPE allocations is not public in full, so the exercise treats wallet-level allocation and qualification as latent primitives disciplined by public aggregate targets.

5.2 Calibration targets and baseline parameterization

The calibration values the claim immediately before the genesis event and assumes liquidation 30 days later. Because the genesis distribution was fully unlocked, the baseline tradable share is set to $G_T = 1$. The reported launch valuation of \$3.90 is treated as the post-event benchmark. With baseline event dilution $\delta_a = 0.10$, the corresponding pre-event reference price is

$$S_t = \frac{3.90}{1 - \delta_a} \approx 4.33. \quad (37)$$

The aggregate figures also imply a natural wallet-level anchor. If 310 million HYPE were distributed across approximately 94,000 recipient wallets, the implied mean allocation was

$$\frac{310,000,000}{94,000} \approx 3,297.87 \quad (38)$$

HYPE per recipient. I therefore use 3,298 HYPE as an implied-average wallet benchmark and 4,000 HYPE as the representative medium-activity wallet. This creates a clean mapping between the public aggregate distribution and the wallet-level objects required by the model.

Table 6: Public Hyperliquid calibration targets

Target	Value	Role in calibration
Maximum HYPE supply	1.0 billion	Supply denominator for genesis share
Genesis allocation share	31.0%	Aggregate airdrop mass
Genesis tokens distributed	310 million	Wallet-level allocation anchor
Recipient wallets	$\approx 94,000$	Cross-sectional scaling target
Implied average recipient allocation	3,297.87 HYPE	Average-wallet benchmark
Launch price	\$3.90	Post-event price anchor
Genesis unlock status	Fully unlocked	Baseline tradable share

Notes. The table combines official Hyperliquid documentation on the points program and genesis distribution [32, 33] with contemporaneous reporting on launch valuation and recipient scale [44, 49]. The implied average recipient allocation is computed as 310 million divided by 94,000.

Table 7: Baseline model parameters for the Hyperliquid calibration

Parameter	Baseline value
Risk-free rate	$r = 0.04$
Horizon	$T - t = 30/365$ years
Pre-event reference price	$S_t = \$4.33$
Event dilution parameter	$\delta_a = 0.10$
Eligibility probability	$p_t = 0.80$
Expected allocation	$\bar{N}_t = 4,000$ HYPE
Tradable fraction	$\bar{G}_t = 1.00$
Diffusive volatility	$\sigma = 1.10$
Jump intensity	$\lambda_J = 4.0$ per year
Mean jump size	$\mu_J = -0.08$
Jump volatility	$\sigma_J = 0.18$
Strike	$K = \$10,000$

Notes. Institutional inputs are anchored to Hyperliquid public disclosures. The financial parameters are chosen to reflect high-volatility crypto trading conditions over a short horizon and are used consistently across the robustness exercises unless otherwise stated.

5.3 Simulation design

For tractability, the calibration uses the jump-diffusion version of the model rather than the full stochastic-volatility system. Each Monte Carlo path draws an eligibility indicator, a jump-diffusion terminal token value, and the discounted airdrop and option payoffs. The simulation uses 50,000 paths and fixed seeds to keep the calibration and robustness outputs reproducible.¹ The simulation follows standard Monte Carlo principles for derivative pricing as described in Glasserman [27], with variance-reduction considerations informed by Avramidis et al. [5] and Broadie and Glasserman [11].

¹The full replication files are `hyperliquid_case.py`, which produces the numerical outputs and L^AT_EX-ready table rows, and `make_figures.py`, which produces the figure files used below.

5.4 Baseline valuation and wallet heterogeneity

Under the baseline calibration, the expected discounted value of the representative wallet’s airdrop entitlement is \$12,442.84, while the corresponding call-like claim with strike \$10,000 is worth \$4,618.10. The wedge between these objects is economically meaningful: the first is the shadow value of the full airdrop entitlement, whereas the second is the option value after truncating states in which the realized claim ends below the strike.

Table 8 shows how quickly values diverge across wallet types. A low-activity wallet is effectively out of the money. An implied-average recipient wallet is worth materially more, but remains far below the representative high-activity case. Once both qualification probability and expected allocation rise, the option value scales sharply.

Table 8: Representative Hyperliquid wallet comparisons

Wallet type	p_t	$\bar{N}_t$ (HYPE)	Airdrop value (USD)	Option value C_t (USD)
Low activity	0.55	1,200	2,558.83	4.56
Implied average recipient	0.70	3,298	9,084.54	2,397.77
Medium activity	0.75	4,000	11,650.45	4,334.86
High activity	0.90	9,000	31,612.26	22,634.28




Figure 1: Option values across representative Hyperliquid wallet types. The steep gradient is the practical manifestation of the model’s convexity: once expected allocation becomes large relative to the strike, the wallet moves from near out-of-the-money to deep in-the-money.

5.5 Baseline decomposition and vesting counterfactuals

It is useful to decompose the baseline calibration into layers. The first layer is the full discounted entitlement value, which asks what the wallet’s expected future token receipt is worth before any contractual strike truncation. The second layer is the strike-sensitive option value. The wedge between the two is the cost of money: a large share of the representative wallet’s economic value lives in states that are valuable in expectation but not always above the strike.

Table 9: Baseline value decomposition for the representative Hyperliquid wallet

Object	Value (USD)	Interpretation
Discounted airdrop entitlement A_t	12,442.84	Expected discounted value of the full wallet-level claim
Call-like value C_t	4,618.10	Value after strike truncation at \$10,000
Moneyness discount $A_t - C_t$	7,824.74	Value lost because many states finish below the strike
Option-to-entitlement ratio	37.11%	Share of total entitlement value that survives the strike

This decomposition is helpful for interpreting OTC trading in airdrop rights. A buyer of the address’s future claim may care primarily about the option value because the realized payoff is contractually truncated or because the buyer is effectively underwriting only upside states. By contrast, a protocol designer or a welfare analysis may care more about the entitlement value because it measures the total scale of value being distributed to users.

Vesting counterfactuals further illustrate why tokenomics design belongs inside the pricing problem. Hyperliquid’s genesis distribution was fully unlocked, which makes it a high-tradability benchmark. Table 10 shows how dramatically the value of the representative wallet would fall under alternative unlock structures if all other parameters were held fixed.

Table 10: Vesting counterfactuals around the Hyperliquid baseline

Immediate tradable share $\bar{G}_t$	Option value (USD)	Airdrop value (USD)	Interpretation
0.25	1.07	3,125.00	Deeply delayed monetization; claim is almost entirely out of the money
0.50	310.84	6,221.42	Partial unlock restores some value but leaves strong strike truncation
0.75	1,948.85	9,332.13	Moderate vesting still materially compresses the claim relative to full unlock
1.00	4,699.70	12,442.84	Fully unlocked benchmark closest to the realized genesis structure

The table underscores a simple but important point: a nominal token amount is not a sufficient statistic for value. The same expected allocation can correspond to a near-zero option value or to a large in-the-money claim depending on the unlock regime. That is why pure “headline airdrop size” comparisons across projects can be misleading.

5.6 Strike and maturity grids

Journal-style calibration exercises usually ask whether the core result survives basic contract-design perturbations. Tables 11 and 12 therefore vary the strike and maturity while holding the Hyperliquid baseline fixed.

Table 11: Strike-grid robustness

Strike	Option value (USD)	Airdrop value (USD)	Option-to-airdrop ratio
\$5,000	8,457.65	12,446.54	68.0%
\$7,500	6,516.15	12,489.68	52.2%
\$10,000	4,628.85	12,453.33	37.2%
\$12,500	3,034.40	12,437.89	24.4%
\$15,000	1,879.72	12,448.58	15.1%

Table 12: Maturity-grid robustness

Days to liquidation	Option value (USD)	Airdrop value (USD)	Monte Carlo s.e.
7	4,502.71	12,504.52	14.17
14	4,525.88	12,508.09	17.32
30	4,632.43	12,486.21	22.83
60	4,986.75	12,525.05	31.12
90	5,165.86	12,387.89	37.45
180	6,100.94	12,563.83	59.08

Two patterns stand out. First, the strike grid behaves as expected for a call-like claim: option value falls from \$8,457.65 at a \$5,000 strike to \$1,879.72 at a \$15,000 strike, and the option-to-airdrop ratio falls from 68.0% to 15.1%. Second, the maturity grid shows that the Hyperliquid claim retains more upside over longer liquidation horizons. The expected airdrop value remains comparatively stable because discounting and drift effects are modest over the horizons studied, while option value rises from \$4,502.71 at 7 days to \$6,100.94 at 180 days as distributional dispersion increases.

5.7 Event-dilution scenarios and one-way robustness

The most distinctive uncertainty in token-distribution settings is that the issuance event can move the value of the token itself. I therefore study a four-state event-scenario grid that jointly perturbs dilution, diffusive volatility, and jump intensity. The optimistic state sets dilution to 5% and reduces jump severity; the stress state sets dilution to 20% and raises return dispersion and negative jump exposure.

Table 13: Event-dilution scenarios

Scenario	Dilution δ_a	Volatility σ	Jump intensity λ_J	Option value (USD)
Optimistic	5.0%	0.95	3.0	5,242.96
Baseline	10.0%	1.10	4.0	4,657.95
Sell-off	15.0%	1.20	5.0	4,076.49
Stress	20.0%	1.35	6.0	3,642.07

Table 14: Selected one-way comparative statics around the Hyperliquid baseline

Parameter	Value	Option value C_t (USD)
Eligibility probability p_t	0.60	3,486.17
Eligibility probability p_t	0.80	4,643.26
Eligibility probability p_t	0.90	5,187.93
Expected allocation $\bar{N}_t$	3,000	1,968.81
Expected allocation $\bar{N}_t$	4,000	4,629.40
Expected allocation $\bar{N}_t$	6,000	10,739.62
Tradable share $\bar{G}_t$	0.50	310.84
Tradable share $\bar{G}_t$	1.00	4,699.70
Event dilution δ_a	0.05	5,281.82
Event dilution δ_a	0.20	3,379.98

The scenario table shows that event-dilution risk is first order. Option value declines from \$5,242.96 in the optimistic state to \$3,642.07 in the stress state, a drop of roughly 31%. Table 14 sharpens the same point with one-way perturbations. The claim is nearly linear in qualification probability over the relevant range, strongly convex in expected allocation, and extremely sensitive to vesting and dilution. For example, holding all else fixed, moving the tradable share from 0.50 to 1.00 raises option value from \$310.84 to \$4,699.70.




Figure 2: Strike, maturity, dilution, and unlock robustness in the Hyperliquid calibration. The four panels confirm that the model's predictions are not an artifact of a single parameter choice.




Figure 3: Option values under event-dilution scenarios. The main forces remain the same across contract specifications and market conditions.

5.8 Interpreting the Hyperliquid calibration through the lens of effective moneyness

One of the most useful lessons from the calibration is that wallet-level value is best organized around effective moneyness rather than around token price alone. In the representative baseline, the post-event token price anchor is the same for all wallets. What changes from one wallet to another is the scale factor. Once the scale factor is small enough, the claim behaves like a lottery ticket on upside states. Once the scale factor is large enough, the same contract behaves much more like a monetizable entitlement.

This observation helps explain why points markets and account-purchase markets tend to form coarse pricing tiers. Market participants do not need a full structural model to see that some wallets are effectively out of the money and others are not. The model formalizes that intuition. It also clarifies why public disclosures about eligibility thresholds, anti-sybil enforcement, or unlock design can move the value of airdrop rights even before any change in the token's own secondary-market price.

5.9 Interpretation

The Hyperliquid exercise delivers two broader lessons. First, airdrop-linked claims are best understood as claims on a scaled underlying rather than as vanilla token options. The wallet comparison and strike grid show that small differences in expected allocation can move the claim across moneyness regions and create large value differences. Second, contract design and market design interact tightly. Because issuance affects both the quantity distributed and the value of the token itself, event dilution and unlock design belong directly inside the pricing problem rather than in a separate tokenomics discussion.

5.10 Broad versus concentrated distributions

The model provides a simple way to think about one of the most important strategic choices facing token issuers: whether to distribute broadly across many users or concentrate rewards in a smaller set of highly active accounts. A broad distribution tends to lower the expected allocation of the median wallet while potentially increasing the legitimacy and reach of the protocol. A concentrated distribution tends to generate larger in-the-money claims for a smaller number of recipients and may strengthen short-run incentives for heavy activity. In model terms, this is a choice about the cross-sectional distribution of N_T , and in some cases about the joint distribution of p_t and N_T .

From a valuation standpoint, concentration increases option value for top-tier wallets but may reduce the mass of moderately in-the-money claims across the broader user base. From a market-design standpoint, the same choice may affect post-event selling pressure and governance concentration. A protocol that concentrates distribution heavily may create a smaller number of sophisticated recipients who internalize governance value, but it may also generate a more fragile coalition if excluded users feel that participation was under-rewarded. The pricing model does not resolve that tradeoff normatively, but it does make clear how the financial consequences appear at the wallet level. This tradeoff is closely related to the “chicken-and-egg” problem in two-sided platform design [13, 50] and to broader questions about innovation and governance in platform ecosystems [48].

5.11 Strict versus permissive qualification rules

Qualification design introduces a related tradeoff. Permissive rules raise the probability that light users receive something, but they can reduce the informational value of points and exacerbate sybil behavior. Strict rules make qualification more valuable for serious users while increasing the option-like binary character of the claim. This is directly visible in the comparative statics: when p_t moves from modest values toward one, option value rises nearly proportionally for the representative wallet, but the qualitative interpretation also changes. Low- p_t wallets are paying for a chance to qualify, while high- p_t wallets are mostly exposed to allocation and market-value risk.

From the protocol’s perspective, qualification policy therefore determines whether the airdrop behaves more like a lottery, a rebate, or a deferred ownership grant. A lottery-like design can attract speculative attention and wide participation, but it may also encourage inefficient farming. A rebate-like design is easier to value and may better reward genuine usage. An ownership-grant design moves the program closer to a corporate-finance mechanism for allocating claims to quasi-insiders and committed users.

5.12 Unlock design and post-event market quality

Unlock design is one of the clearest places where partial-equilibrium and general-equilibrium reasoning can diverge. Holding token price fixed, tighter vesting lowers wallet value because less of the allocation is immediately monetizable. But once token price is allowed to react, tighter vesting can improve market quality by reducing instantaneous float expansion, smoothing sell pressure, and allowing the market more time to absorb new supply. The relevant question is therefore not whether vesting is “good” or “bad” in the abstract, but whether the price-support effect is strong enough to offset the direct monetization loss.

This question has practical relevance for projects deciding how to split the supply between user rewards, foundation reserves, and insider allocations. A project that fully unlocks user rewards but heavily vests insider allocations may create a very different market path than one that vests users and insiders symmetrically. The present model focuses on wallet-level user claims, but the same logic can be embedded in richer supply-side analyses that account for the

staggered release of other token buckets.

Table 15: Illustrative design regimes and predicted valuation effects

Design regime	Main primitive affected	Predicted direct effect on wallet value	Broader implication
Broad, lightly tiered retrodrop	$\bar{N}_t$ dispersion	Compresses top-wallet upside while raising the mass of moderate claims	Supports inclusion and may broaden governance base
Highly concentrated reward schedule	$\bar{N}_t$ upper tail	Raises the value of top-tier wallets sharply	Can increase perceived unfairness and recipient concentration
Strict sybil filtering	p_t	Reduces value for ambiguous wallets, raises value for confidently eligible users	Strengthens legitimacy but increases binary qualification risk
Immediate full unlock	$\bar{G}_t$	Maximizes immediate monetizable value	Can increase short-run sell pressure and float expansion
Staggered unlock	$\bar{G}_t$ and δ_a	Lowers immediate value but may support the token price	Trades off user liquidity against market stabilization

5.13 Secondary markets for points, accounts, and airdrop rights

Although the paper is written as a valuation model rather than a market-microstructure study, it also speaks to informal secondary markets that exist before an airdrop is distributed. Points-bearing accounts may trade privately, users may negotiate side agreements over future distributions, and funds may value participation in a protocol partly through the embedded option on a future token grant. The present framework rationalizes those practices. The object being traded is not a claim on current token price; it is a claim on a future protocol-designed scale factor multiplied by the future token price. Crypto market microstructure features, including segmentation and return predictability patterns documented in Makarov and Schoar [40] and common risk factors identified in Liu et al. [38], further shape the environment in which these secondary valuations occur.

This also implies that information in airdrop markets differs from information in ordinary token markets. News about eligibility thresholds, sybil enforcement, points snapshots, or vesting updates can move the value of airdrop rights even when the current token price is unchanged or when no spot token yet exists. An empirical literature on pre-token trading would therefore benefit from measuring announcement effects at the level of expected scale rather than only at the level of subsequent token returns.

5.14 Corporate-finance interpretation

Stepping back, the model can be read as a corporate-finance framework for user ownership. The project commits to a distribution rule that resembles a contingent compensation contract. Users invest attention, liquidity, order flow, or network participation, and they receive a stochastic future claim in return. Qualification risk acts like employment or performance uncertainty,

allocation acts like the compensation schedule, vesting acts like restricted stock or option vesting, and event dilution acts like the price impact of new issuance.

That analogy is useful because it places crypto airdrops inside a broader tradition of thinking about how organizations distribute contingent claims. The difference is that tokenized organizations can often commit to publicly observable rules and can distribute claims to a much larger and more heterogeneous user base than a conventional firm. The resulting claims are novel in form but familiar in economic substance: they are contingent compensation and ownership contracts that must be priced under uncertainty.

6 Extensions

6.1 Numerical valuation

The general model can be valued by Monte Carlo simulation. A practical implementation proceeds in four steps:

1. Simulate token-price paths under jump diffusion with stochastic volatility and the event jump at T_a .
2. Draw eligibility I from a Bernoulli distribution with parameter p_t .
3. Draw allocation N_T from either the chosen lognormal or discrete tier distribution.
4. Compute the payoff $\max(I N_T G_T S_T^{(a)} - K, 0)$ and average across simulated paths.

This procedure is straightforward, flexible, and naturally accommodates path-dependent vesting schedules or state-dependent dilution terms. Standard references on Monte Carlo methods for financial engineering provide the theoretical underpinning and variance-reduction techniques [5, 11, 27].

6.2 Endogenous qualification and strategic behavior

The present paper treats p_t as exogenous from the perspective of valuation at time t . In practice, wallet owners may strategically alter their behavior to improve qualification probability. A richer model would let p_t depend on an optimal control problem in which users trade off farming costs against expected airdrop value. That extension is important for market design but is separate from the pricing problem addressed here.

6.3 Incomplete markets and utility-based valuation

Airdrop claims are often idiosyncratic and not perfectly hedgeable. Qualification and allocation risk may be address-specific, and there may be no liquid secondary derivatives market on the eventual token. In such cases, the risk-neutral benchmark remains useful as a baseline but may not be fully sufficient. A natural extension is utility-indifference pricing. Let W_T^π denote terminal wealth under portfolio policy π , and let X_T denote the airdrop-linked payoff. With exponential utility $U(w) = -e^{-\gamma w}$, the indifference price π_t solves

$$\sup_{\pi} \mathbb{E}[-e^{-\gamma W_T^\pi}] = \sup_{\pi} \mathbb{E}[-e^{-\gamma(W_T^\pi + X_T - \pi_t)}]. \quad (39)$$

This extension is especially appealing when the claim's non-traded risks are large and investor-specific [29]. Transaction costs in crypto markets, which can be substantial, provide an additional wedge that incomplete-market frameworks are well suited to handle [19].

6.4 Empirical calibration

The model is well suited to empirical calibration. Eligibility probability can be estimated from historical qualification rules or labeled on-chain outcomes. Allocation distributions can be in-

ferred from points programs, scoreboards, token-distribution snapshots, or governance disclosures. Vesting terms and token-supply schedules are often directly observable in project documentation. Event dilution can be disciplined using tokenomics and post-distribution price reactions. The framework therefore provides a practical bridge between derivative-style valuation and the broader empirical literature on token issuance, platform growth, and governance [1, 16, 18, 36, 41, 55].

6.5 Learning and belief updating

A natural extension is to allow users to update beliefs over time as more information arrives about qualification rules and expected allocation. In many live airdrop programs, the market learns gradually from partial rule announcements, unofficial scoreboards, community speculation, and comparable prior launches. In such settings, p_t and the conditional distribution of N_T are themselves state variables that evolve with information. A dynamic learning model would make the pre-event value of the claim path dependent and could generate announcement effects even without any contemporaneous change in the token price process.

The current framework is compatible with such an extension. One can reinterpret p_t and the distribution of N_T as posterior beliefs rather than static parameters. The main valuation equation would remain intact, but the empirical design would focus more explicitly on how information shocks change those posteriors. This would be a promising bridge between asset-pricing analysis and the high-frequency study of protocol communications.

6.6 General-equilibrium token-demand effects

The present paper treats the token price process as reduced form. A richer framework would embed the airdrop inside a model of token demand, where users, speculators, and liquidity providers interact strategically around the launch. In that environment, the distribution mechanism could affect not only immediate float and sell pressure, but also the demand curve for the token by changing beliefs about governance legitimacy, future platform adoption, or the stickiness of the recipient base.

Such a general-equilibrium extension would be especially useful for comparing alternative launch architectures. For example, a project might prefer a broader but smaller user airdrop if the resulting governance legitimacy increases token demand enough to offset the higher distribution mass. Another project might prefer a concentrated distribution if it believes that sophisticated recipients are less likely to dump and more likely to provide liquidity or governance effort. The current model does not endogenize those responses, but it identifies the channels through which they matter. Platform-economics models of two-sided competition [13, 50] and token-based governance design [16, 17] offer natural starting points for this extension.

6.7 Cross-project comparisons and external validity

The framework is also portable across project types. In a decentralized exchange or derivatives venue, qualification may depend heavily on trading activity and liquidity provision. In a layer-1 or layer-2 ecosystem, qualification may instead depend on wallet activity, bridging, staking, or governance participation. In consumer or gaming applications, user activity and retention may be more important than direct financial contribution. These differences matter for empirical measurement but do not require a different pricing architecture; they simply change which observables are most informative about p_t , N_T , and G_T .

This portability suggests that airdrop pricing could become a useful common language across otherwise disparate token launches. Instead of comparing projects only on the basis of headline distribution size or first-day market capitalization, researchers and practitioners can compare them on a common set of objects: qualification uncertainty, allocation dispersion, tradability,

and event-adjusted market value. That is a more disciplined basis for both empirical work and market commentary. The broader literature on crypto market dynamics, including common risk factors [38], boom-bust cycles [23], and the digitalization of money more generally [12], provides useful context for such cross-project comparisons.

7 Conclusion

This paper provides a theoretical framework for pricing option-like claims on future cryptocurrency airdrops. The key insight is that an airdrop-linked payoff cannot be represented adequately as a standard option on a token price alone. Instead, the economically relevant underlying is a product of qualification, allocation, tradability, and market value. Once those components are modeled explicitly, the valuation problem becomes conceptually clear.

The main pricing object,

$$C_t = \mathbb{E}_t^{\mathbb{Q}} \left[e^{-r(T-t)} \max(I_T N_T G_T S_T^{(a)} - K, 0) \right], \quad (40)$$

captures the four defining institutional features of crypto airdrops: qualification risk, allocation uncertainty, vesting restrictions, and event-driven price impact. Under conditional independence and mean-scaling approximations, the model reduces to a tractable scaled jump-diffusion option. The Hyperliquid case study shows that this structure is not merely conceptual: the dominant drivers of value in practice are precisely the objects highlighted by the theory, especially qualification probability, expected allocation, and expected dilution.

The framework opens several directions for future work. A structural model of strategic farming could endogenize eligibility. A general equilibrium approach could study how anticipated airdrops affect user acquisition and token demand before distribution. An incomplete-market extension could quantify how non-hedgeable qualification risk alters valuation relative to the risk-neutral benchmark. Even in its current form, however, the model offers a useful language for organizing theory and empirics around the pricing of airdrop rights and for linking token-distribution design to more standard corporate-finance questions about contingent claims, financing structure, and market discipline.

A Proofs

A.1 Proof of Theorem 1

Proof. Under Assumption 1, one may condition first on the eligibility indicator:

$$C_t = \mathbb{E}_t^{\mathbb{Q}} \left[e^{-r(T-t)} \max(I_T N_T G_T S_T^{(a)} - K, 0) \right]. \quad (41)$$

Since $I \in \{0, 1\}$ with conditional mean p_t , the expectation decomposes into eligible and ineligible states. Ineligible states contribute zero to the payoff, so the claim reduces to

$$C_t = p_t \mathbb{E}_t^{\mathbb{Q}} \left[e^{-r(T-t)} \max(N_T G_T S_T^{(a)} - K, 0) \right]. \quad (42)$$

Replacing N_T and G_T by their conditional expectations yields the mean-scaling approximation stated in the theorem. The resulting expression is a call-like claim on the scaled object $\bar{N}_t \bar{G}_t S_T^{(a)}$, multiplied by the qualification probability p_t . $\square$

A.2 Proof of Proposition 1

Proof. Under the conditional independence approximation,

$$C_t = p_t Y_t, \quad (43)$$

where

$$Y_t = \mathbb{E}_t^{\mathbb{Q}} \left[e^{-r(T-t)} \max(N_T G_T S_T^{(a)} - K, 0) \right] \geq 0. \quad (44)$$

Since Y_t is nonnegative and does not depend on p_t , the option value is affine and increasing in p_t . Hence $\partial C_t / \partial p_t = Y_t \geq 0$, with strict inequality whenever the conditional claim has strictly positive value. $\square$

A.3 Proof of Proposition 2

Proof. Under the mean-scaling approximation in Equation (21), define

$$\phi(a) = \mathbb{E}_t^{\mathbb{Q}} \left[e^{-r(T-t)} \max(a S_T^{(a)} - K, 0) \right] \quad (45)$$

for $a \geq 0$. For every realization of $S_T^{(a)}$, the map $a \mapsto \max(a S_T^{(a)} - K, 0)$ is nondecreasing because $S_T^{(a)} > 0$. Therefore $\phi(a)$ is nondecreasing in a . Since $a = p_t \bar{G}_t \bar{N}_t$ up to a positive multiplicative constant, C_t is nondecreasing in $\bar{N}_t$ and strictly increasing whenever the claim has positive probability of finishing in the money. $\square$

A.4 Proof of Proposition 3

Proof. The proof is identical to that of Proposition 2. Under Equation (21), the payoff is nondecreasing in the scale factor $\bar{N}_t \bar{G}_t$. Since $\bar{N}_t > 0$, the option value is nondecreasing in $\bar{G}_t$ and strictly increasing whenever positive-probability in-the-money states exist. $\square$

A.5 Proof of Proposition 4

Proof. Suppose the event-adjusted token price can be written as

$$S_T^{(a)} = (1 - \delta_a) \tilde{S}_T, \quad (46)$$

where $\tilde{S}_T$ collects all non-dilution uncertainty and is independent of δ_a over the comparative-static exercise. Then for each realization of $(I, N_T, G_T, \tilde{S}_T)$, the map

$$\delta_a \mapsto \max(I_T N_T G_T (1 - \delta_a) \tilde{S}_T - K, 0) \quad (47)$$

is nonincreasing on $[0, 1]$. Taking discounted expectations preserves this ordering, so C_t is nonincreasing in δ_a . Strict monotonicity follows whenever states with positive intrinsic value occur with positive probability. $\square$

A.6 Proof sketch of Proposition 5

Proof. Fix the scale factor $\alpha = I_T N_T G_T \geq 0$. The terminal payoff can be written as $f(S_T^{(a)}) = \max(\alpha S_T^{(a)} - K, 0)$, which is convex in $S_T^{(a)}$. Therefore any parameter change that raises the terminal distribution of $S_T^{(a)}$ in convex order raises the expected payoff. In jump-diffusion environments this includes standard cases in which higher diffusion volatility or more dispersion-inducing jump risk increases terminal price dispersion without lowering the mean enough to overturn convexity gains. The exact sign can be parameter-dependent when jump changes primarily load on large negative event shocks, which is why the proposition is stated under standard convexity conditions rather than as a global unconditional result. $\square$

A.7 Proof of Lemma 1

Proof. Fix a realization of $S_T^{(a)}$. For any $a \geq 0$, the pathwise payoff

$$g(a; S_T^{(a)}) = \max(a S_T^{(a)} - K, 0) \quad (48)$$

is nondecreasing in a because $S_T^{(a)} > 0$. It is also convex in a because it is the maximum of two affine functions, $a S_T^{(a)} - K$ and 0. Discounting by the positive scalar $e^{-r(T-t)}$ preserves both properties, and taking expectations preserves them as well. Hence $\phi(a, K)$ is nondecreasing and convex in a .

For fixed a , the same pathwise payoff is nonincreasing in K because a larger strike can only reduce the intrinsic value. Discounting and taking expectations therefore imply that $\phi(a, K)$ is nonincreasing in K . $\square$

A.8 Proof of Proposition 6

Proof. Fix the random tuple $(I, N_T, G_T, S_T^{(a)})$. The map

$$K \mapsto \max(I_T N_T G_T S_T^{(a)} - K, 0) \quad (49)$$

is nonincreasing for every realization because raising the strike weakly lowers the payoff state by state. Multiplying by the positive discount factor and taking conditional expectations preserves that ordering. Hence C_t is nonincreasing in K . $\square$

A.9 Proof of Proposition 7

Proof. Under the stated hypothesis, the longer-maturity event-adjusted price is a mean-preserving convex-order spread of the shorter-maturity price. The payoff function

$$s \mapsto \max(\alpha s - K, 0) \quad (50)$$

is convex in s for any fixed $\alpha \geq 0$. Therefore the expected payoff weakly increases under a convex-order expansion of $S_T^{(a)}$ that preserves its conditional mean. Integrating over the nonnegative random scale factor $\alpha = I_T N_T G_T$ preserves the ordering. Consequently the airdrop-linked claim value is weakly increasing in maturity under the maintained convex-order condition. $\square$

A.10 Proof of Proposition 8

Proof. Under the mean-scaling approximation,

$$C_t^j \approx \phi(p_t^j \bar{N}_t^j \bar{G}_t^j, K) \quad (51)$$

for wallet $j \in \{A, B\}$, where $\phi(\cdot, K)$ is nondecreasing by Lemma 1. If

$$p_t^A \bar{N}_t^A \bar{G}_t^A \geq p_t^B \bar{N}_t^B \bar{G}_t^B, \quad (52)$$

then monotonicity implies $C_t^A \geq C_t^B$. $\square$

A.11 Proof of Corollary 1

Proof. If $I = 1$ and $G_T = 1$ almost surely, the general payoff simplifies mechanically:

$$X_T = \max(I_T N_T G_T S_T^{(a)} - K, 0) = \max(N_T S_T^{(a)} - K, 0). \quad (53)$$

No further argument is required. $\square$

A.12 Proof sketch of Proposition 9

Proof. The direct effect of tighter vesting is to lower the payoff scale through G_T , which weakly decreases claim value. The indirect effect operates through the token-price channel: if lower immediate tradability reduces expected sell pressure and therefore lowers δ_a , then since lower dilution raises $S_T^{(a)}$, this second channel increases claim value. The net sign is therefore ambiguous without additional structure on the magnitude of the price-support effect relative to the direct loss of immediacy. $\square$

B Additional Hyperliquid Calibration Notes

This appendix records the main numerical choices used in the Hyperliquid illustration.

1. **Qualification environment.** Hyperliquid’s points system ran in repeated weekly rounds, with one million points distributed weekly from November 1, 2023 through May 1, 2024 and 700,000 weekly points in the L1 phase beginning on May 29, 2024 [33]. This motivates treating qualification as probabilistic rather than deterministic from the ex ante perspective.
2. **Aggregate airdrop mass.** The genesis event allocated 310 million HYPE, corresponding to 31% of the one billion token maximum supply, and genesis tokens were fully unlocked [32].
3. **Address-count anchor.** Market coverage reported that the launch reached just over 94,000 users [44]. This implies a crude average of approximately 3,298 HYPE per recipient, which anchors the medium-wallet calibration.
4. **Price anchor.** The launch-day reference price of \$3.90 is taken from contemporaneous reporting [49]. The pre-event benchmark of \$4.33 is obtained by backing out a baseline dilution discount of 10%.
5. **Counterfactual vesting.** The Hyperliquid genesis airdrop was fully unlocked, but the model also studies lower values of G_T to show how strongly vesting can compress option value in alternative token launches.

C Measurement Issues and External Validity

This appendix records several practical issues that arise when the model is taken to protocol data.

C.1 Observability problems

Even when tokenomics are public, the researcher usually does not observe all primitives with equal clarity. Qualification rules are often announced gradually, allocation formulas may only be partially disclosed, and the relevant event window for measuring dilution can be blurred by listing effects and broader market moves. The calibration strategy in the paper therefore treats the public record as informative but incomplete. This is why the numerical exercise emphasizes ranges, wallet taxonomies, and scenario analysis rather than a claim of point identification for every primitive.

The observability problem is particularly acute for anti-sybil filtering. Projects routinely describe high-level principles without revealing every enforcement detail. As a result, a wallet can appear highly active in public data and still face material exclusion risk. In a valuation setting, that uncertainty belongs directly in p_t rather than being treated as a purely qualitative caveat.

C.2 Interpretation across project types

While Hyperliquid is used as the central illustration, the framework is meant to travel. The economic meaning of the primitives is stable across projects even when the empirical proxies change. On a decentralized exchange, points may depend on trading volume or liquidity provision. On a consumer protocol, they may depend on engagement or referral activity. On a proof-of-stake network, staking or validator behavior may be more important. The same pricing equation remains useful as long as the researcher can form disciplined beliefs about qualification, allocation, tradability, and event-adjusted value.

Table 16: External validity across token-launch environments

Project environment	Most informative observable	Dominant primitive	Typical valuation concern
Decentralized exchange	Trading activity and liquidity provision	N_T and p_t	Large reward dispersion across active and inactive wallets
Layer-1 or layer-2 network	Bridging, wallet activity, staking	p_t and G_T	Qualification ambiguity plus future unlock overhang
Governance retrodrop	Proposal voting, delegation, historical usage	p_t	Hard-to-measure quality of participation and sybil filtering
Consumer or social protocol	Retention, referrals, engagement streaks	N_T	Weak mapping between usage metrics and token monetization
Gaming or in-app economy	Item ownership, play history, spending	N_T and δ_a	Illiquidity and post-launch float shocks

C.3 Threats to identification

Three threats are especially important in future empirical work. First, public rules may differ from effective rules if projects retain discretion. Second, wallet-level activity can be endogenous to anticipated airdrop value, which complicates causal interpretation of observed correlations between activity and realized allocation. Third, post-event token prices are affected by broad crypto conditions, making it difficult to isolate pure airdrop-induced dilution without narrow event windows or richer control designs.

These threats do not invalidate the framework; they indicate where supplementary empirical structure is most valuable. If a future study obtains wallet-level labels, it can estimate the screening function more directly. If it combines that information with high-frequency launch data, it can isolate event dilution more cleanly. The present paper’s contribution is to make clear exactly which economic primitives those additional data would identify.

D Additional Appendix Tables

Table 17: Comparative-static summary

Primitive	Sign for C_t	Economic intuition
Eligibility probability p_t	Positive	More states survive the qualification filter and carry positive scale
Expected allocation $\bar{N}_t$	Positive	Larger scale raises both intrinsic value and the probability of finishing in the money
Tradable share $\bar{G}_t$	Positive	More of the nominal allocation becomes immediately monetizable
Strike K	Negative	More states are truncated at zero and surviving states lose intrinsic value
Dilution δ_a	Negative	Expected event-adjusted token price falls when issuance pressure intensifies
Volatility or jump dispersion	Usually positive	Convex payoff benefits from greater dispersion unless negative jump risk dominates
Maturity	Usually positive	More dispersion and, in some designs, more unlock time increase option value

Table 18: Illustrative empirical roadmap for future work

Research question	Required data	Model object	Example output
How predictable is qualification?	Wallet-level outcomes and pre-event activity	p_t	Estimated screening model or qualification forecast
How dispersed are wallet rewards?	Realized allocations and points histories	N_T	Allocation distribution and wallet taxonomy
How costly is vesting?	Unlock schedules and token trading data	G_T	Discount from nominal allocation to monetizable value
How large is event dilution?	High-frequency price and volume around launch	δ_a	Event-window price-impact estimate
How do users price pre-token rights?	OTC trades, account sales, points-market quotes	C_t	Comparison between model value and negotiated market price

Table 19: Interpretation of the main numerical outputs

Output object	Economic use	Interpretation
Airdrop entitlement A_t	Welfare or distributional analysis	Measures the full discounted scale of value assigned to a wallet before strike truncation
Call value C_t	OTC or derivative-style valuation	Measures monetizable upside under the specified contract and strike
Option-to-entitlement ratio	Cross-wallet comparison	Summarizes how much of the full entitlement survives moneyiness truncation
Sensitivity to p_t	Rule-change analysis	Shows how strongly ambiguous qualification states drive value
Sensitivity to $\bar{N}_t$	Tier-design analysis	Quantifies how quickly rewards scale into the money
Sensitivity to δ_a	Tokenomics analysis	Measures the cost of event-induced price pressure

E Numerical Procedure

The simulation algorithm for the Hyperliquid case is as follows.

1. Set $(p_t, \bar{N}_t, \bar{G}_t, S_t, K, r, \sigma, \lambda_J, \mu_J, \sigma_J, \delta_a)$.
2. Draw $I \sim \text{Bernoulli}(p_t)$.
3. Draw a Gaussian shock $Z \sim \mathcal{N}(0, 1)$ and a Poisson jump count $N_J \sim \text{Poisson}(\lambda_J(T - t))$.
4. Draw N_J jump magnitudes from $\mathcal{N}(\mu_J, \sigma_J^2)$ and sum them.

5. Construct the event-adjusted terminal price

$$S_T^{(a)} = S_t(1 - \delta_a) \exp \left([r - q - \lambda_J \kappa_J - \tfrac{1}{2} \sigma^2] (T - t) + \sigma \sqrt{T - t} Z + \sum_{i=1}^{N_J} Y_i \right). \quad (54)$$

6. Compute the discounted airdrop value and option payoff.

7. Average across paths.

The script `hyperliquid_case.py` implements exactly this procedure, while `make_figures.py` generates the chart assets used in the main text.

References

- [1] Jason Abadi and Markus Brunnermeier. Token-based platform governance. *Journal of Financial Economics*, 162:103951, 2024. doi: 10.1016/j.jfineco.2024.103951.
- [2] Douglas W. E. Allen. Crypto airdrops: An evolutionary approach. *Journal of Evolutionary Economics*, 34(5):849–872, 2024.
- [3] Douglas W. E. Allen, Chris Berg, and A. M. Lane. Why airdrop cryptocurrency tokens? *Journal of Business Research*, 163:113945, 2023.
- [4] Benito Arruñada and Nuno Garoupa. Blockchain: The birth of decentralized governance. *Economic Inquiry*, 56(4):2227–2242, 2018.
- [5] Avishai N. Avramidis, Anand Deshpande, and Carl Lenos. Simulation experiments and common random numbers. *Management Science*, 44(9):1273–1284, 1998.
- [6] Gurdip Bakshi, Charles Cao, and Zhiwu Chen. Empirical performance of alternative option pricing models. *The Journal of Finance*, 52(5):2003–2049, 1997.
- [7] Nicholas Barberis and Richard H. Thaler. A survey of behavioral finance. *Handbook of the Economics of Finance*, 1:1053–1128, 2003.
- [8] David S. Bates. Jumps and stochastic volatility: Exchange rate processes implicit in Deutsche mark options. *The Review of Financial Studies*, 9(1):69–107, 1996.
- [9] Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654, 1973.
- [10] Rainer Böhme, Nicolas Christin, Benjamin Edelman, and Tyler Moore. Bitcoin: Economics, technology, and governance. *Journal of Economic Perspectives*, 29(2):213–238, 2015.
- [11] Mark Broadie and Paul Glasserman. Pricing American-style securities using simulation. *Journal of Economic Dynamics and Control*, 21(8–9):1323–1352, 1997.
- [12] Markus K. Brunnermeier, Harold James, and Jean-Pierre Landau. The digitalization of money. *Journal of Economic Literature*, 61(3):824–863, 2023.
- [13] Bernard Caillaud and Bruno Jullien. Chicken & egg: Competition among intermediation service providers. *RAND Journal of Economics*, 34(2):309–328, 2003.
- [14] Peter Carr, Hélyette Geman, Dilip B. Madan, and Marc Yor. Stochastic volatility for Lévy processes. *Mathematical Finance*, 13(3):345–382, 2003.
- [15] Christian Catalini and Joshua Gans. Some simple economics of the blockchain. NBER Working Paper 22952, National Bureau of Economic Research, 2016.
- [16] Lin Chen, Lina Xu, and Nihar Shah. Tokenomics: Dynamic adoption and valuation. *The Review of Financial Studies*, 33(3):1105–1135, 2020.
- [17] Jiri Chod and Evgeny Lyandres. Tokenized platforms and the role of governance tokens. *The Review of Financial Studies*, 36(11):4503–4543, 2023.
- [18] Lin William Cong, Ye Li, and Neng Wang. Token-based platform finance. *Journal of Financial Economics*, 144(3):972–991, 2022. doi: 10.1016/j.jfineco.2021.10.002.
- [19] George M. Constantinides. Capital market equilibrium with transactions costs. *Journal of Political Economy*, 94(4):842–862, 1986.

- [20] John C. Cox and Stephen A. Ross. Equilibrium valuation of options on stocks and bonds. *Journal of Financial Economics*, 3(1–2):145–166, 1976.
- [21] Darrell Duffie, Jun Pan, and Kenneth Singleton. Transform analysis and asset pricing for affine jump-diffusions. *Econometrica*, 68(6):1343–1376, 2000.
- [22] Jim Euchner and Jonathan Chiu. Cryptocurrency and blockchain: A survey. *Research Technology Management*, 65(1):57–64, 2022.
- [23] Libing Fang, Elie Bouri, and David Roubaud. What explains the boom and bust of crypto? *Journal of Banking and Finance*, 160:107402, 2024.
- [24] Christian Fisch. Initial coin offerings (ICOs) to finance new ventures. *Journal of Business Venturing*, 34(1):1–22, 2019. doi: 10.1016/j.jbusvent.2018.09.007.
- [25] Jessie Gan, George Tsoukalas, and Serguei Netessine. Decentralized platforms: Governance, tokenomics, and ICO design. *Management Science*, 69(11):6667–6683, 2023. doi: 10.1287/mnsc.2021.02076.
- [26] Robert Geske. Valuing corporate securities: Some effects of bond indenture provisions. *The Journal of Finance*, 32(2):517–534, 1979.
- [27] Paul Glasserman. *Monte Carlo Methods in Financial Engineering*. Springer, 2003.
- [28] Campbell R. Harvey, Ashwin Ramachandran, and Joseph Santoro. DeFi and the future of finance. *John Wiley & Sons*, 2021.
- [29] Vicky Henderson. Valuation of claims on nontraded assets using utility maximization. *Mathematical Finance*, 12(4):351–373, 2002.
- [30] Steven L. Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The Review of Financial Studies*, 6(2):327–343, 1993.
- [31] Sabrina T. Howell, Marina Niessner, and David Yermack. Initial coin offerings: Financing growth with cryptocurrency token sales. *The Review of Financial Studies*, 33(9):3925–3974, 2020. doi: 10.1093/rfs/hhz131.
- [32] Hyper Foundation. HYPE genesis [official announcement, published november 28, 2024]. Retrieved April 21, 2026, from <https://hyperfnd.medium.com/hype-genesis-1830a4dc2e3f>, 2024.
- [33] Hyperliquid Docs. Points [official documentation page]. Retrieved April 21, 2026, from <https://hyperliquid.gitbook.io/hyperliquid-docs/points>, 2025.
- [34] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292, 1979.
- [35] S. G. Kou. A jump-diffusion model for option pricing. *Management Science*, 48(8):1086–1101, 2002.
- [36] Jerry Li and William Mann. Digital tokens and platform building. *The Review of Financial Studies*, 38(7):1921–1954, 2025. doi: 10.1093/rfs/hhaf015.
- [37] Evan Lipnowski and Martin Schanzenbach. Blockchain governance: An overview. *Journal of Economic Perspectives*, 36(4):183–204, 2022.

- [38] Yukun Liu, Aleh Tsyvinski, and Xi Wu. Common risk factors in cryptocurrency. *The Review of Financial Studies*, 35(3):1133–1177, 2022.
- [39] Joshua D. Lyon and Christopher Mayer. Decentralized finance: Markets, regulation, and systemic risk. *Annual Review of Financial Economics*, 16:289–318, 2024.
- [40] Igor Makarov and Antoinette Schoar. Pricing and return structure in crypto markets. *Journal of Financial Economics*, 137(1):1–22, 2020.
- [41] Christos A. Makridis, Michael Fröwis, Karthik Sridhar, and Rainer Böhme. The rise of decentralized cryptocurrency exchanges: Evaluating the role of airdrops and governance tokens. *Journal of Corporate Finance*, 79:102358, 2023.
- [42] William Margrabe. The value of an option to exchange one asset for another. *The Journal of Finance*, 33(1):177–186, 1978.
- [43] Robert C. Merton. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3(1–2):125–144, 1976.
- [44] Tom Mitchelhill. How Hyperliquid’s insanely lucrative airdrop launch changed the game [CoinTelegraph, published january 8, 2025]. Retrieved April 21, 2026, from <https://cointelegraph.com/news/how-hyperliquid-airdrop-changed-game>, 2025.
- [45] Paul P. Momtaz. Entrepreneurial finance and moral hazard: Evidence from token offerings. *Journal of Business Venturing*, 36(5):106001, 2021. doi: 10.1016/j.jbusvent.2020.106001.
- [46] Stewart C. Myers. Determinants of corporate borrowing. *Journal of Financial Economics*, 5(2):147–175, 1977.
- [47] Stewart C. Myers. The capital structure puzzle. *The Journal of Finance*, 39(3):574–592, 1984.
- [48] Geoffrey Parker, Marshall Van Alstyne, and Sangeet Paul Choudary. Innovation and governance in platform ecosystems. *Strategy Science*, 4(3):231–236, 2019.
- [49] Ana Paula Pereira. Hyperliquid’s HYPE token surges 60% after billion-dollar airdrop [CoinTelegraph, published november 29, 2024]. Retrieved April 21, 2026, from <https://cointelegraph.com/news/hyperliquid-hype-token-surges-after-billion-dollar-airdrop>, 2024.
- [50] Jean-Charles Rochet and Jean Tirole. Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4):990–1029, 2003.
- [51] Fabian Schär. Decentralized finance: On blockchain- and smart contract-based financial markets. *Federal Reserve Bank of St. Louis Review*, 103(2):153–174, 2021.
- [52] Robert J. Shiller. *Irrational Exuberance*. Princeton University Press, 3rd edition, 2015.
- [53] Richard H. Thaler. Mental accounting matters. *Journal of Behavioral Decision Making*, 12(3):183–206, 1999.
- [54] David Yermack. Corporate governance and blockchains. *Review of Finance*, 21(1):7–31, 2017. doi: 10.1093/rof/rfw074.
- [55] Yiming Zhang, Cong Liang, and Frank Weikai Liu. DeFi lending: Design and risks. *The Review of Financial Studies*, 37(4):1053–1090, 2024.