Wrapper Architecture as Coordination Device: Multiple Equilibria in Autonomous-Agent Organizations

Ian Staley*

May 2, 2026

Abstract

Decentralized autonomous organizations (DAOs) and AI-agent systems combine cryptographic execution with blockchain-based governance, yet observed organizations almost universally combine these mechanisms with a conventional legal entity—a foundation, statutory DAO form, or limited liability wrapper. I develop a stylized model in which token-holders jointly determine wrapper choice and governance concentration, generating multiple equilibria: an inefficient *trap* in which the wrapper coalition cannot form because no holder will absorb the front-loaded fixed cost alone, and an efficient *wrapper equilibrium* in which the coalition reaches scale and amortizes fixed costs effectively. The trap is an empirically grounded coordination problem rather than an analytical artifact, and a global-games selection argument identifies the threshold at which institutional design tips the system between equilibria. The framework reframes the *CFTC v. Ooki DAO* ruling, the Wyoming DAO LLC and DUNA statutes, and AI legal personhood debates as questions of equilibrium selection rather than of substantive cost allocation, and bounds the "Coasean singularity" claim that AI agents dramatically reduce transaction frictions.

Keywords: decentralized autonomous organizations; AI agents; coordination; multiple equilibria; organizational design; blockchain finance
JEL classification: D23; L22; K22; O33; G23

*Independent Researcher, USA. Email: `ian.t.staley@gmail.com`. ORCID: 0009-0000-8592-3186.

1 Introduction

Thousands of decentralized autonomous organizations (DAOs) operate on-chain, governing protocol decisions and treasury allocation through token-holder voting. Yet Appel and Grennan [1] document that three or fewer entities control most decisions across 151 of the largest DAOs. Han et al. [2] and Fritsch et al. [3] confirm extreme voting-power concentration across major protocols. Sun et al. [4] show that in MakerDAO, one of the most institutionally mature DAOs, governance outcomes track concentrated whale preferences rather than median-voter views.

Despite this concentration, organizations operating DAO and autonomous-agent mechanisms exhibit striking institutional variety. The same protocols that route capital through smart contracts also organize themselves as Cayman Islands foundations, Swiss *Stiftungs*, Wyoming statutory DAO LLCs, Wyoming DUNAs, or—in some cases—defend themselves in federal court as unincorporated associations [5–7]. The *CFTC v. Ooki DAO* ruling, in which a Northern California district court found a DAO subject to CFTC jurisdiction as an unincorporated association and entered a default judgment of $643,542 plus a permanent trading ban (*CFTC v. Ooki DAO*, No. 22-cv-5416, June 8, 2023, Orrick, J.), is one extreme of this variety; the Wyoming DAO LLC Supplement and the Wyoming DUNA statute are the other.

What explains the co-occurrence of governance concentration with institutional diversity? The dominant economic framing treats wrapper choice as a cost-comparison problem: cryptographic execution is cheap but cryptographic recourse is expensive, so wrappers combining cryptographic execution with legal recourse should dominate [8, 9]. But this framing leaves unexplained why some protocols select wrapper architectures while others apparently identical do not, and why *Ooki DAO*-style enforcement actions occur at all if wrappers are unambiguously efficient.

This paper offers a stylized account in which the answer involves *equilibrium selection* rather than cost minimization in isolation. I develop a model in which token-holders' wrapper-support choices jointly determine equilibrium governance concentration, with externality probability π depending on both. The model yields multiple equilibria: a stable *trap* in which no individual holder will absorb the front-loaded fixed cost of forming a legal wrapper, and a stable *wrapper equilibrium* in which a coalition has reached scale and amortizes fixed costs effectively. An unstable interior threshold separates the two basins. Following Carlsson and van Damme [10] and Goldstein and Pauzner [11], I argue that under noisy observation of cost fundamentals, a global-games argument selects between equilibria via a threshold θ^* whose comparative statics characterize how institutional design choices tip the system between trap and wrapper outcomes.

Three observations follow. First, institutional designs that lower the fixed coordination cost (DUNA standardizes legal-form establishment; Wyoming DAO LLC pre-resolves partnership-liability questions) shrink the trap basin without affecting the underlying cost ordering—they operate as equilibrium-selection devices, not cost-reduction interventions. Second, the *Ooki DAO* ruling can be read not as an anomaly in the cost-comparison framework but as a regulatory shock that shifts the equilibrium-selection threshold by raising the realized recourse cost on the cryptographic side. Third, the Shahidi et al. [12] "Coasean singularity" claim—that AI agents dramatically reduce frictions in matching, search, and contracting—survives only on the execution dimension; on the coordination and recourse dimensions, the multiple-equilibria structure persists regardless of reductions in execution cost.

The model is formally about DAOs (token-holder governance plus smart-contract execution); the application to autonomous-agent organizations is by analogy. Where AI agents serve as token-holders' delegates or as automated oracles, the wrapper-coordination problem applies directly; where AI agents are themselves the relevant decision-makers, the model's predictions extend by treating the AI agent as a token-holder with deterministic preferences.

The contribution is modest. The model does not predict which equilibrium any specific DAO will select; it characterizes the parameter region in which each equilibrium can be sustained and identifies how institutional design choices shift the boundary. The framework is consistent with observed institutional variety rather than a test of any particular wrapper architecture. The empirical implications are testable but not tested here.

The paper proceeds as follows. Section 2 develops the model and proves the multi-equilibrium and selection results. Section 3 applies the framework to three institutional debates—*Ooki DAO*, statutory DAO forms, and AI legal personhood—and bounds the Coasean-singularity claim. Section 4 documents pattern consistency with the existing empirical DAO governance literature without claiming validation. Section 5 discusses limitations and extensions. Section 6 concludes.

2 Model

2.1 Preliminaries

A *decentralized autonomous organization* (DAO) is an on-chain entity whose decisions are executed by smart contracts under token-holder governance. A *wrapper* is a conventional legal entity (LLC, foundation, statutory DAO form) placed around the DAO to interface with courts, regulators, and counterparties. *Cryptographic control* denotes authority exercised through smart-contract code; *legal control* denotes authority exercised through the wrapper's recognized legal capacities.

2.2 Setup

Consider a DAO with N token-holders, indexed $i = 1, \dots, N$, each performing a sequence of identical tasks. Each task has outcome $y \in \{I, O\}$ where I denotes in-distribution execution and O denotes an out-of-distribution outcome that triggers external recourse—counterparty dispute, oracle failure, regulatory action, or consumer harm.

Two endogenous variables determine outcomes. The first is *wrapper choice* $w \in \{C, W\}$, where C is cryptographic-only and W is wrapper architecture. The second is *governance concentration* $h \in [0, 1]$, defined as the share of voting power held by the top-k token-holders for some fixed small k, following the empirical concentration measures in Appel and Grennan [1], Han et al. [2], and Fritsch et al. [3].

Endogenous externality probability. The per-task probability of an out-of-distribution outcome is

$$\pi(h, w) = \pi_0 + \alpha h - \delta_w, \qquad \delta_C = 0, \quad 0 < \delta_W \leq \pi_0 + \alpha(h_0 - \delta_h), \tag{1}$$

where $\pi_0 \in (0, 1)$ is the baseline probability, $\alpha > 0$ captures the *concentration channel* (concentrated decision-makers miss out-of-distribution events that distributed monitoring would catch), and $\delta_W > 0$ captures the *accountability discount* specific to wrapper architectures (legal recourse and regulatory disclosure independently reduce out-of-distribution probability by changing the set of actions that are profitable to attempt). The upper bound on δ_W ensures $\pi(h, w) \geq 0$ for all admissible (h, w) since $h_W(s) \geq h_0 - \delta_h$ across $s \in [0, 1]$. The decomposition into concentration and accountability channels is the central modeling departure from prior work treating π as exogenous [9, 12].

Per-task costs. Following the cost structure motivated by the empirical literature on smart-contract execution [13, 14] and legal scholarship on DAO recourse [5, 6]:

$$\tau_C(h) = c_C^E + \pi(h, C) \cdot c_C^R, \tag{2}$$

$$\tau_W(h) = c_C^E + \kappa + \pi(h, W) \cdot (c_L^R + \lambda), \tag{3}$$

where c_C^E is cryptographic execution cost, c_C^R is cryptographic recourse cost, c_L^R is legal recourse cost, and $\kappa, \lambda \geq 0$ are wrapper integration and partitioning costs. We retain the cost asymmetry assumed in prior literature: $c_C^R > c_L^R$ (cryptographic recourse is expensive when an exploit, governance attack, or regulatory action requires off-chain action; documented quantitatively in the DeFi-attacks literature [15, 16]) with κ, λ small relative to recourse cost differences.

2.3 Token-holder choice

Each token-holder i chooses $a_i \in \{0, 1\}$, where $a_i = 1$ means *support wrapper adoption* and $a_i = 0$ means *stay cryptographic-only*. Let $s = (1/N) \sum_{i=1}^{N} a_i$ denote the fraction supporting wrapper. Two channels determine each holder's payoff.

Channel 1 — fixed-cost coordination. Wrapper infrastructure has a fixed cost $F > 0$ comprising registered-agent fees, legal counsel retainer, audit, regulatory disclosure compliance, and oracle subscriptions. Foundation-style and statutory wrappers face nontrivial annual operational expenses largely independent of coalition size, as the legal scholarship on entity formation and DAO statutes documents [5, 6]. This fixed cost is shared among supporting holders, generating the per-capita coordination cost

$$\psi(s) = \frac{F}{sN + 1} + \varepsilon, \tag{4}$$

where $\varepsilon > 0$ is an idiosyncratic individual cost. Note $\psi'(s) < 0$: the more holders support wrapper, the cheaper individual support becomes.

Channel 2 — belief coordination. Holder i's evaluation of wrapper versus cryptographic-only depends on her expectations of equilibrium concentration on each side. Concentration responds to coalition size:

$$h_C(s) = h_0 + \eta s, \tag{5}$$

$$h_W(s) = h_0 - \delta_h \cdot s, \tag{6}$$

where $\eta \geq 0$ captures *residual concentration* (holders flowing to the wrapper side leave smaller-stake distributed monitors behind, raising concentration on the cryptographic side), $\delta_h \in (0, h_0)$ captures the *accountability-scale channel* (wrapper accountability infrastructure scales with coalition size; a thousand-member DUNA enforces accountability more effectively than a three-member foundation), and h_0 is the baseline concentration prevailing at $s = 0$. The strategic complementarity in beliefs is direct: a holder who expects others to support wrapper expects h_W to be low (lower out-of-distribution probability on her side) and h_C to be high (worse outcomes for those who stay), making wrapper relatively more attractive.

Holder i's problem. Holder i chooses a_i to minimize expected per-task cost. Define

$$\Phi(s) \equiv \tau_C(h_C(s)) - \tau_W(h_W(s)) - \psi(s). \tag{7}$$

Under perfect-foresight equilibrium beliefs, an equilibrium fraction s^{eq} satisfies one of three conditions: $s^{eq} = 0$ with $\Phi(0) \leq 0$; $s^{eq} = 1$ with $\Phi(1) \geq 0$; or $s^{eq} \in (0, 1)$ with $\Phi(s^{eq}) = 0$.

2.4 Multiple equilibria

Proposition 1 (Multiple equilibria)**.** *Under assumptions*

(A1) $\pi(h, w) = \pi_0 + \alpha h - \delta_w$ *with* $\alpha > 0$ *and* $0 < \delta_W \leq \pi_0 + \alpha(h_0 - \delta_h)$*;*

(A2) $\psi(s) = F/(sN + 1) + \varepsilon$ *with* $F > 0$ *and* $\varepsilon > 0$*;*

(A3) $h_C(s) = h_0 + \eta s$ *and* $h_W(s) = h_0 - \delta_h s$ *with* $\delta_h \in (0, h_0)$ *and* $\eta \in [0, 1 - h_0]$*;*

(A4) $c_C^R > c_L^R + \lambda$ *(cryptographic recourse is more expensive than legal recourse net of partitioning cost);*

(A5) $\Phi(0) < 0$ *and* $\Phi(1) > 0$*;*

the equilibrium correspondence has at least three Nash equilibria: a stable trap equilibrium $s_L = 0$*, an unstable interior separator* $s_M \in (0, 1)$*, and a stable wrapper equilibrium* $s_H = 1$*.*

Proof. Differentiating Φ:

$$\Phi'(s) = \alpha\eta c_C^R + \alpha\delta_h(c_L^R + \lambda) + \frac{FN}{(sN + 1)^2}.$$

All three terms are strictly positive under (A1), (A3), (A4), so $\Phi'(s) > 0$ for all $s \in [0,1]$. By (A5) and the intermediate value theorem, Φ has a unique interior zero $s_M \in (0,1)$. This is unstable in the Nash sense: at s slightly above s_M, $\Phi(s) > 0$ implies wrapper is preferred, so additional holders prefer to join and s rises; at s slightly below s_M, $\Phi(s) < 0$ implies cryptographic-only is preferred, so supporters prefer to defect and s falls. The boundary equilibria $s = 0$ and $s = 1$ are stable Nash equilibria: at $s = 0$, $\Phi(0) < 0$ ensures no individual holder profits by unilaterally absorbing the unamortized fixed cost; at $s = 1$, $\Phi(1) > 0$ ensures no individual holder profits by defecting. Coordinated coalition-level deviations are not refinement-eliminated under Nash; the trap equilibrium therefore reflects a coordination failure rather than a Pareto-efficient outcome. □

The mechanism is a *fixed-cost coordination trap* rather than a textbook strategic-substitutes structure. The first holder to support wrapper adoption faces $\psi(0) = F + \varepsilon$, an empirically non-trivial cost that no individual holder has incentive to bear alone. The trap exists because wrapper economics has front-loaded fixed costs that are realistic in practice: a foundation requires legal counsel before supporters exist, and a DUNA requires statutory filings before a coalition has formed. The result is non-trivial relative to prior cost-comparison frameworks: $\tau_W(h_0) < \tau_C(h_0)$ at baseline does not imply wrapper adoption, because the equilibrium correspondence depends on whether the coalition can clear the fixed-cost threshold.

2.5 Equilibrium selection

The multi-equilibrium structure raises the standard selection question. Following the global-games tradition of Carlsson and van Damme [10], applied to finance by Morris and Shin [17], Goldstein and Pauzner [11], and Frankel et al. [18], I introduce noise in token-holders' observations of the underlying cost asymmetry $\theta \equiv c_C^R - (c_L^R + \lambda)$. Each holder i observes

$$x_i = \theta + \sigma\xi_i, \qquad \xi_i \overset{iid}{\sim} \mathcal{N}(0,1),$$

with precision $\sigma > 0$. Under threshold strategies (support wrapper iff $x_i \geq x^*$), the standard global-games argument selects a unique equilibrium as $\sigma \to 0$. The argument proceeds as in Goldstein and Pauzner [11]: the action ordering supports threshold strategies because $\Phi(s)$ is monotone in s (Proposition 1); the marginal type's indifference condition pins down a critical fundamental θ^*; and iterated elimination of dominated strategies converges to the threshold equilibrium.

Proposition 2 (Selection threshold)**.** *Under assumptions (A1)–(A5) and the noisy-signal structure above, the global-games equilibrium selects the wrapper outcome for $\theta > \theta^*$ and the trap for $\theta < \theta^*$, where*

$$\theta^* = \frac{\kappa - \delta_W(c_L^R + \lambda) + \dfrac{F}{\bar{s}N + 1} + \varepsilon - \alpha\bar{s}(\eta + \delta_h)(c_L^R + \lambda)}{\pi_0 + \alpha h_0 + \alpha\bar{s}\eta}, \tag{8}$$

with $\bar{s} \in (0,1)$ the Laplacian threshold belief (typically $\bar{s} = 1/2$ in the symmetric case).

Proof. The marginal type $x^* = \theta$ is indifferent between supporting wrapper and staying cryptographic-only at the threshold belief $\bar{s}$:

$$\tau_C(h_C(\bar{s})) - \tau_W(h_W(\bar{s})) - \psi(\bar{s}) = 0.$$

Substituting (2)–(3) and (4), with $h_C(\bar{s}) = h_0 + \eta\bar{s}$ and $h_W(\bar{s}) = h_0 - \delta_h\bar{s}$:

$$(\pi_0 + \alpha h_0 + \alpha\bar{s}\eta)\, c_C^R - (\pi_0 + \alpha h_0 - \alpha\bar{s}\delta_h - \delta_W)(c_L^R + \lambda) = \kappa + \frac{F}{\bar{s}N + 1} + \varepsilon.$$

Using $\theta = c_C^R - (c_L^R + \lambda)$ to substitute $c_C^R = \theta + (c_L^R + \lambda)$ in the first term and collecting the θ-coefficients on the left:

$$(\pi_0 + \alpha h_0 + \alpha\bar{s}\eta)\, \theta = \kappa - \delta_W(c_L^R + \lambda) + \frac{F}{\bar{s}N + 1} + \varepsilon - \alpha\bar{s}(\eta + \delta_h)(c_L^R + \lambda).$$

Solving for θ at the threshold yields (8). The denominator $\pi_0 + \alpha h_0 + \alpha\bar{s}\eta = \pi(h_0 + \eta\bar{s}, C)$ is the cryptographic externality probability evaluated at the threshold belief, which is strictly positive under (A1). $\square$

The threshold formula has direct comparative statics:

Proposition 3 (Selection comparative statics)**.** *Under (A1)–(A5):*

(i) $\partial\theta^*/\partial F > 0$: a higher fixed coordination cost expands the trap basin.

(ii) $\partial\theta^*/\partial\delta_W < 0$: a stronger accountability discount expands the wrapper basin.

(iii) $\partial\theta^*/\partial\kappa > 0$: a higher integration cost expands the trap basin.

(iv) $\partial\theta^*/\partial\delta_h < 0$: a stronger accountability-scale channel expands the wrapper basin.

Proof. Direct partial differentiation of (8), treating $\bar{s}$ and the parameters of π as fixed. $\square$

The substantive content is that institutional design choices operate as *equilibrium-selection devices*, not as cost-reduction interventions for the inframarginal organization. A statute that lowers F shifts a DAO from trap to wrapper without changing the underlying cost ordering between architectures. DAOs sit on different sides of θ^* for reasons that depend on history, jurisdiction, and the institutional resources available at formation rather than on intrinsic cost differences across protocols.

2.6 Numerical illustration

To visualize the equilibrium correspondence, consider the calibration $N = 1000$, $\pi_0 = 0.05$, $\alpha = 0.2$, $\delta_W = 0.08$, $h_0 = 0.6$, $\delta_h = 0.4$, $\eta = 0.05$, $c_C^E = 1$, $c_C^R = 10$, $c_L^R = 3$, $\kappa = 0.5$, $\lambda = 0.4$, $F = 200$, $\varepsilon = 0.05$. Figure 1 plots $\Phi(s)$. The function is monotonically increasing, with $\Phi(0) \approx -199$ (deep

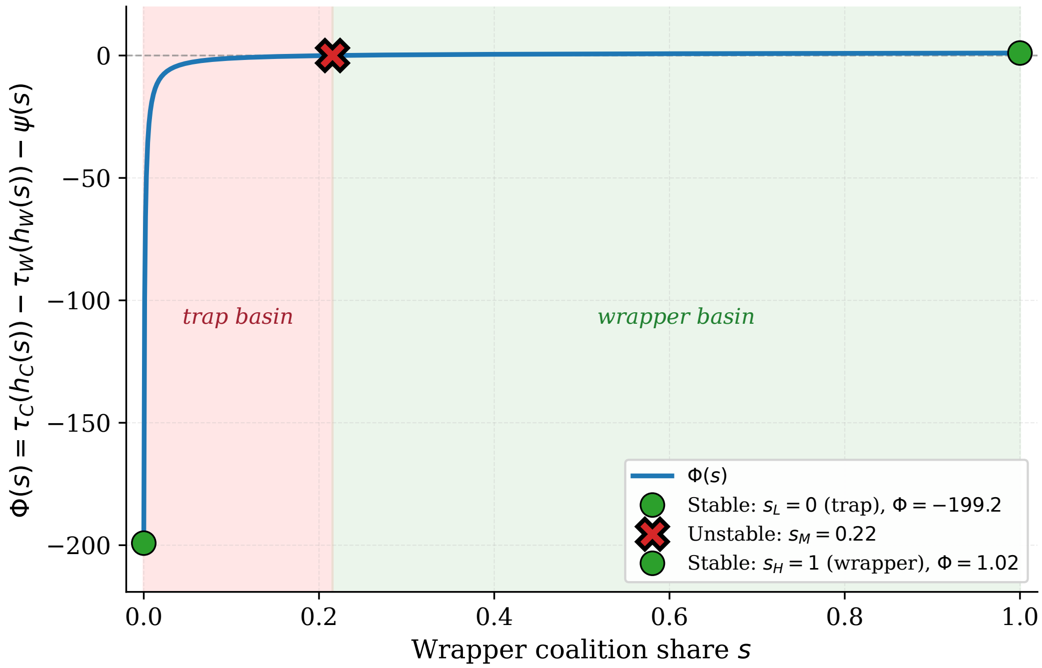

Figure 1: Indifference function $\Phi(s) = \tau_C(h_C(s)) - \tau_W(h_W(s)) - \psi(s)$ under the calibration $N = 1000$, $\pi_0 = 0.05$, $\alpha = 0.2$, $\delta_W = 0.08$, $h_0 = 0.6$, $\delta_h = 0.4$, $\eta = 0.05$, $c_C^E = 1$, $c_C^R = 10$, $c_L^R = 3$, $\kappa = 0.5$, $\lambda = 0.4$, $F = 200$, $\varepsilon = 0.05$. The function is monotonically increasing with $\Phi(0) \approx -199$, $\Phi(1) \approx 1.02$, and unique interior zero $s_M \approx 0.22$. Three Nash equilibria coexist: a stable trap at $s_L = 0$, an unstable separator at s_M, and a stable wrapper equilibrium at $s_H = 1$. The trap basin $[0, s_M)$ collects coalition initial conditions that fall back to the unwrapped outcome; the wrapper basin $(s_M, 1]$ collects those that converge to full wrapper adoption.

negative due to unamortized fixed cost), $\Phi(1) \approx 1.02$ (wrapper preferred at full coalition), and unstable separator $s_M \approx 0.22$. The basins of attraction are: trap $[0, s_M)$ and wrapper $(s_M, 1]$.

Figure 2 traces s_M as the fixed coordination cost F varies, mapping observed institutional designs to points on the trap-basin boundary. DUNAs (low F) sit at $s_M \approx 0.06$, an easily reachable basin transition. Wyoming DAO LLCs (moderate F) require $s_M \approx 0.22$ coalition support. Foundation-style wrappers (higher F) require $s_M \approx 0.49$. Unwrapped multisig arrangements (F effectively unbounded) require near-universal coordination—a basin transition that observed unwrapped DAOs systematically fail to make, which the model treats as predicted, not anomalous.

Table 1 summarizes how each institutional design maps to the model's parameters and the resulting basin geometry. The qualitative rankings reflect the substantive content of the legal scholarship cited in Section 3; the numerical values are illustrative under the calibration of Figures 1–2.

3 Implications

The model recasts three institutional debates as equilibrium-selection questions.

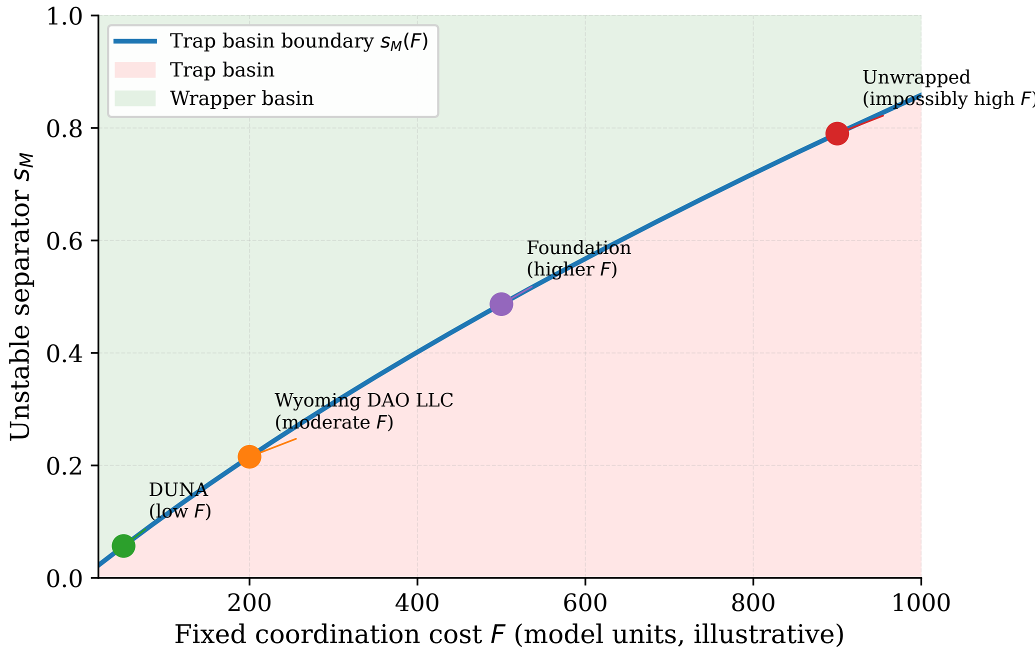

Figure 2: Trap-basin boundary s_M as a function of fixed coordination cost F, holding all other calibration parameters at the values of Figure 1. Lower F produces a smaller trap basin: a coalition need only reach $s_M \approx 0.06$ under DUNA-style infrastructure to escape the trap, but $s_M \approx 0.79$ under unwrapped multisig. Wyoming DAO LLC and foundation-style wrappers occupy intermediate positions. The mapping of institutional designs to F values is illustrative; absolute magnitudes depend on jurisdiction, complexity of operations, and the maturity of supporting infrastructure.

3.1 The *Ooki DAO* ruling as regulator-induced equilibrium switch

In *CFTC v. Ooki DAO*, Judge Orrick of the Northern District of California found that an unwrapped DAO operating leveraged-margin trading software constituted an unincorporated association subject to CFTC jurisdiction (*CFTC v. Ooki DAO*, No. 22-cv-5416, N.D. Cal., Dec. 20, 2022 (capacity to be sued), and June 8, 2023 (default judgment)). The court entered $643,542 in monetary penalties and a permanent trading ban. Grant et al. [5] and Allen [6] argue from doctrinal first principles that the ruling is consistent with general partnership and unincorporated association doctrine; the model offers an economic complement.

Ooki operated at $s = 0$ (no wrapper) facing a population $\theta = c_C^R - c_L^R - \lambda$ that was below θ^* ex ante: cryptographic recourse c_C^R was perceived as low because regulatory enforcement against an unincorporated DAO was assumed unlikely. The CFTC's enforcement action revised that perception. Ex post, c_C^R rose substantially—to at least $643,542 plus the cost of permanent business termination—which raises θ above θ^* for any DAO operating in similar regulatory territory. The ruling functions as an information shock that shifts perceived θ across the threshold, not as an anomalous decision in the cost-comparison framework.

The companion cases *Sarcuni v. bZx DAO*, 664 F. Supp. 3d 1100 (S.D. Cal. 2023), and *Samuels v. Lido DAO*, No. 23-cv-06492 (N.D. Cal. Nov. 18, 2024), extend the partnership-liability theory

Table 1: Institutional designs mapped to model parameters. Lower F and higher δ_W both shrink the trap basin. Specific numerical values are illustrative.

Design	F	δ_W	δ_h	**Trap basin** $[0, s_M)$
Wyoming DUNA	low	high	high	$[0, 0.06)$
Wyoming DAO LLC	moderate	moderate	moderate	$[0, 0.22)$
Foundation	high	moderate	low	$[0, 0.49)$
Unwrapped multisig	very high	none	none	$[0, 0.79)$

to other unwrapped DAOs and reinforce the same cross-DAO inference: the cryptographic-only recourse cost is higher than $\theta < \theta^*$ would imply, and unwrapped operation systematically fails to clear the wrapper threshold.

3.2 Statutory DAO forms as F-reduction devices

Wyoming's DAO LLC Supplement (W.S. §§17-31-101 to 17-31-115, enacted 2021) and Decentralized Unincorporated Nonprofit Association (DUNA) Act (enacted 2024) reduce the fixed cost F in different ways. The Wyoming DAO LLC anticipated the partnership-liability problem that *Ooki DAO* and *Sarcuni v. bZx DAO* subsequently illustrated for unwrapped DAOs: by registering as a DAO LLC, a protocol's contracting capacity is statutorily defined and member liability is capped, eliminating the bespoke legal-formation work otherwise required. Neither *Ooki* nor *Sarcuni* tested the Wyoming statute itself—both cases involved unwrapped organizations—but their dispositions clarified the magnitude of c_C^R for protocols that had not registered. The DUNA goes further than the DAO LLC by removing the for-profit-entity framing entirely and providing a default constitution that resolves governance procedure, voting threshold, and dispute resolution [5].

In the model, both statutes shrink the trap basin by lowering F. A DAO that would face $F = 200$ as an unwrapped multisig coordinator faces $F \approx 50$ as a DUNA. Proposition 3 translates this into a falling θ^*, expanding the parameter region where the wrapper equilibrium is selected. The statutes are, in this framing, equilibrium-selection devices rather than cost-reduction devices for the inframarginal organization—they affect which equilibrium is reached, not the comparative efficiency of the equilibria once reached.

3.3 AI legal personhood as equilibrium-selection capacity

Debates over whether AI agents should be granted legal personhood [19] and whether DAOs should be granted entity status [6] typically frame the question as whether the system has properties (consciousness, agency, moral status, organizational integrity) that warrant personhood treatment. The model reframes the question. Legal personhood for an AI-DAO complex is a wrapper feature that lowers F (legal-formation cost) and raises δ_W (accountability discount), both of which shift θ^* in the direction of wrapper selection. Whether AI agents are "persons" in some metaphysical sense is irrelevant; whether the legal apparatus has *capacity* to recognize and constrain them is a coordination input. Cohen et al. [20] and Chan et al. [21] make analogous points in the AI governance literature, treating personhood-adjacent legal infrastructure as instrumental rather than ontological.

3.4 Bounding the Coasean singularity

Shahidi et al. [12] argue that AI agents may dramatically reduce transaction costs across search, matching, and contracting, producing a "Coasean singularity." The model bounds this claim. AI agents primarily affect c_C^E (cryptographic execution cost) through faster matching and contract drafting; on the execution dimension, the singularity claim is plausible. But c_C^R (recourse cost) does not collapse with execution efficiency, and the fixed cost F is determined by legal infrastructure, not execution speed. A genuine Coasean singularity requires $\theta = 0$ and $F = 0$ jointly, which is implausible: even at $c_C^E = 0$, the model's multi-equilibrium structure persists, and equilibrium selection still depends on θ and F. AI agents make execution frictionless on the in-distribution side but cannot make coordination frictionless on the wrapper side.

4 Pattern consistency

The framework is consistent with three patterns documented in the empirical DAO governance literature, though it does not test them rigorously.

Pattern 1: Concentration coexists with institutional variety. Appel and Grennan [1], Han et al. [2], and Fritsch et al. [3] document extreme governance concentration across major DAOs. The model treats this as the steady-state outcome at any equilibrium: h remains near h_0 regardless of w, modulated by the accountability-scale channel δ_h. Institutional variety arises through wrapper choice w, which depends on whether the DAO has cleared the trap-basin threshold s_M.

Pattern 2: Whales matter more for outcomes than coalition size. Sun et al. [4] document for MakerDAO that protocol decisions track concentrated whale preferences rather than median-voter views. The model is consistent with this through the $h_C(s) = h_0 + \eta s$ relation: marginal coalition movement does not fundamentally redistribute control among whales, who occupy the high-h end of the distribution irrespective of s.

Pattern 3: Smart-contract exploits map to high-π states. The systematization-of-knowledge literature on DeFi attacks documents that exploits cluster in protocols with weak governance oversight [15, 16, 22]. The model treats exploits as out-of-distribution events whose probability rises with h and falls with δ_W. Pattern consistency would predict that DUNA-wrapped or Wyoming DAO LLC-wrapped protocols experience proportionally fewer exploits per capital unit at risk than unwrapped protocols, an empirical claim the existing literature does not directly test but which is observable in principle.

The framing here is deliberately modest. Pattern consistency is not validation; the patterns above are necessary but not sufficient conditions for the model's empirical adequacy. A direct test would require a structural estimation of F, δ_W, δ_h, and θ across DAO populations, which is beyond the scope of this paper.

5 Discussion

The model has three principal limitations.

First, π is treated as a probability of out-of-distribution events without specifying the dynamic structure that generates such events. A more complete model would derive π from the interaction between governance decisions and external state realizations, possibly drawing on the principal-agent literature on capability-accumulation [23] or the security economics literature on incentive-compatible audit [24]. The current treatment is a reduced form.

Second, h and w are determined simultaneously rather than dynamically. In practice, DAO formation precedes coalition formation, and wrapper decisions are made by founders well before token-holder governance becomes meaningful. A dynamic version of the model would treat early-stage wrapper choice as commitment that subsequent token-holders inherit, possibly producing path dependence in equilibrium selection that the static model does not capture.

Third, the model treats institutional design (F, δ_W, δ_h) as exogenous parameters. In practice, statutes are themselves outcomes of political-economy processes that respond to observed DAO behavior. The Wyoming DAO LLC and DUNA were legislative responses to existing DAO operational difficulties; *Ooki DAO* was a regulatory response to perceived enforcement gaps. Endogenizing institutional design would require modeling regulators and legislators as players in a meta-game above the DAO-level game studied here.

6 Conclusion

This paper offers a stylized account of why legal wrappers often accompany autonomous-agent systems. The account treats wrapper choice and governance concentration as jointly determined, generating multiple equilibria with distinct institutional implications. A trap equilibrium exists in which the wrapper coalition cannot form because no individual holder will absorb the front-loaded fixed cost; an efficient wrapper equilibrium exists in which the coalition has reached scale and amortizes fixed costs effectively. Global-games selection identifies a threshold determined by the cost asymmetry, the fixed-cost wedge, and the strength of accountability discounts.

The framework reframes three institutional debates. The *CFTC v. Ooki DAO* ruling is consistent with a regulator-induced equilibrium switch rather than a doctrinal anomaly. Wyoming's DAO LLC and DUNA statutes are equilibrium-selection devices that shrink the trap basin without changing underlying cost ordering. AI legal personhood is a coordination capacity, not a metaphysical question. The Shahidi et al. [12] Coasean singularity claim survives only on the execution dimension; coordination friction persists.

The contribution is modest. The paper does not predict which equilibrium any specific DAO will select but characterizes the parameter region in which each equilibrium is sustainable. Future work could endogenize π through richer governance dynamics, model the dynamic timing of wrapper formation, or treat institutional design itself as endogenous to observed DAO behavior. The framework offers a coordination-game lens through which the institutional variety of autonomous-agent organizations becomes legible as predicted equilibrium-selection outcomes rather than regulatory happenstance.

References

[1] I. Appel and J. Grennan, "Control of decentralized autonomous organizations," *AEA Papers and Proceedings*, vol. 113, pp. 182–185, 2023.

[2] J. Han, J. Lee, and T. Li, "A review of DAO governance: Recent literature and emerging trends," *Journal of Corporate Finance*, vol. 91, p. 102696, 2025.

[3] R. Fritsch, M. Müller, and R. Wattenhofer, "Analyzing voting power in decentralized governance: Who controls DAOs?" *Blockchain: Research and Applications*, vol. 5, no. 3, p. 100208, 2024.

[4] X. Sun, C. Stasinakis, and G. Sermpinis, "Decentralization illusion in DeFi: Evidence from tokenized voting in MakerDAO polls," *Journal of Financial Stability*, vol. 73, p. 101286, 2024.

[5] D. M. Grant, E. M. Kirby, and S. Hawkins, "Decentralized autonomous organizations: To statutorily organize or not?" *Wyoming Law Review*, vol. 24, no. 1, pp. 59–103, 2024.

[6] J. G. Allen, "Bodies without organs: Law, economics, and decentralised governance," *Stanford Journal of Blockchain Law & Policy*, vol. 5, no. 1, 2022.

[7] P. De Filippi, M. Mannan, and W. Reijers, *Blockchain Governance.* MIT Press, 2022.

[8] L. W. Cong and Z. He, "Blockchain disruption and smart contracts," *Review of Financial Studies*, vol. 32, no. 5, pp. 1754–1797, 2019.

[9] S. Davidson, "The nature of the decentralised autonomous organisation," *Journal of Institutional Economics*, vol. 21, pp. e5, 1–14, 2025.

[10] H. Carlsson and E. van Damme, "Global games and equilibrium selection," *Econometrica*, vol. 61, no. 5, pp. 989–1018, 1993.

[11] I. Goldstein and A. Pauzner, "Demand–deposit contracts and the probability of bank runs," *Journal of Finance*, vol. 60, no. 3, pp. 1293–1327, 2005.

[12] P. Shahidi, G. Rusak, B. S. Manning, A. Fradkin, and J. J. Horton, "The Coasean singularity? Demand, supply, and market design with AI agents," National Bureau of Economic Research, NBER Working Paper 34468, 2025.

[13] K. John, L. Kogan, and F. Saleh, "Smart contracts and decentralized finance," *Annual Review of Financial Economics*, vol. 15, pp. 523–542, 2023.

[14] A. Lehar and C. A. Parlour, "Decentralized exchange: The Uniswap automated market maker," *Journal of Finance*, vol. 80, no. 1, pp. 321–374, 2025.

[15] L. Zhou, X. Xiong, J. Ernstberger, S. Chaliasos, Z. Wang, Y. Wang, K. Qin, R. Wattenhofer, D. Song, and A. Gervais, "SoK: Decentralized Finance (DeFi) attacks," in *Proceedings of the 2023 IEEE Symposium on Security and Privacy (S&P)*, 2023, pp. 2444–2461.

[16] S. Kitzler, B. Haslhofer, and P. Saggese, "The economic impact of DeFi crime events on decentralized autonomous organizations," *Journal of Finance and Data Science*, vol. 11, p. 100171, 2025.

[17] S. Morris and H. S. Shin, "Unique equilibrium in a model of self-fulfilling currency attacks," *American Economic Review*, vol. 88, no. 3, pp. 587–597, 1998.

[18] D. M. Frankel, S. Morris, and A. Pauzner, "Equilibrium selection in global games with strategic complementarities," *Journal of Economic Theory*, vol. 108, no. 1, pp. 1–44, 2003.

[19] C. Novelli, L. Floridi, and G. Sartor, "AI as legal persons: Past, patterns, and prospects," *Journal of Law and Society*, 2025.

[20] M. K. Cohen, N. Kolt, Y. Bengio, G. K. Hadfield, and S. Russell, "Regulating advanced artificial agents," *Science*, vol. 384, no. 6691, pp. 36–38, 2024.

[21] A. Chan *et al.*, "Infrastructure for AI agents," arXiv:2501.10114; forthcoming *Transactions on Machine Learning Research*, 2024.

[22] K. Qin, L. Zhou, B. Livshits, and A. Gervais, "Attacking the DeFi ecosystem with flash loans for fun and profit," in *Financial Cryptography and Data Security 2021*, ser. LNCS, vol. 12674. Springer, 2021, pp. 3–32.

[23] J. Chen, M. Elliott, and A. Koh, "Capability accumulation and conglomeratization in the information age," *Journal of Economic Theory*, vol. 210, p. 105653, 2023.

[24] P. Daian, S. Goldfeder, T. Kell, Y. Li, X. Zhao, I. Bentov, L. Breidenbach, and A. Juels, "Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability," in *Proceedings of the 2020 IEEE Symposium on Security and Privacy (S&P)*, 2020, pp. 910–927.