

Decentralization: the Search for a Legal Definition

Salvatore Luciano Furnari¹

¹ University of Rome, Tor Vergata
salvatore.furnari@leplex.it

Abstract. Decentralization is the defining feature of decentralized finance, yet no legally workable definition of the concept exists. Regulatory instruments - most notably the MiCA Regulation - invoke decentralization as a threshold criterion for exemption from the regulatory perimeter, without specifying its content. Academic literature has developed sophisticated measurement tools, but these are inherently static and probabilistic, and cannot produce the binary determinations that legal qualification requires.

This paper argues that the definitional gap is not a secondary problem: it is the root cause of the persistent uncertainty surrounding the regulatory treatment of decentralized protocols. Without a precise and verifiable definition, it is impossible to determine, in any concrete case, whether a given system qualifies as decentralized - and therefore whether the associated exemptions apply.

To fill that gap, the paper proposes a legal definition built on three cumulative and verifiable criteria. The first is the presence of at least three independent decision-making centers - the minimum number that allows collective governance without degenerating into unilateral control or mutual veto. The second is the structural interdependence of participants within a protocol-governed framework, which distinguishes decentralized systems from traditional intermediated relationships. The third is the non-custodial nature of the infrastructure, understood as the structural renunciation of control over users' assets by any single entity.

These three criteria are non-mathematical but operationally verifiable. They are cumulative: the absence of any one of them is sufficient to bring a system within the scope of traditional regulatory categories. Together, they provide a minimum normative threshold that is both theoretically grounded and practically applicable by regulatory authorities and courts.

Keywords: Decentralized finance (DeFi), Legal definition of decentralization, MiCA Regulation.

1 The importance of defining decentralization.

According to a recent working paper of the European Central Bank there is “*a significant gap between the theoretical ideal of decentralisation and its practical implementation. The identified concentration of governance tokens among protocols’ treasuries, founders, and centralized exchanges reveals substantial managerial control, indicating that most DeFi platforms would not meet the criteria for full*

decentralisation.”[1]

This paper argues that such a gap is not merely the result of poor implementation. Rather, it reflects the absence of a clear and workable definition of decentralization. Without a precise definition, it is impossible to measure the distance between theory and practice. This paper therefore seeks to fill that gap by proposing a definition of decentralization that is both theoretically sound and practically useful - one that allows to determine, in a concrete case, whether a given protocol can be considered decentralized or not.

The analysis focuses on decentralization as a property of protocols - that is, of the software layers through which DeFi services are delivered - rather than on the decentralization of other systems and sectors where this principle may also be relevant.

2 DAOs and the Community-based Decision Making Procedures

In developing a legal definition of decentralization for protocols, the first observation worth making is that a protocol is typically governed by a DAO. DAOs and their decision making procedure are therefore the natural starting point for any analysis on decentralization.

The simplest definition of DAO is the following: a DAO is a ‘decentralized organization’ managed by an ‘autonomous agent’.[2] In other words, it is an organization composed of people and assets, structured without a central authority (i.e. a ‘decentralized’ organization) and managed by a software that operates independently of both participants and programmers (i.e. autonomous agent).

A *decentralized organization* has been described by its earliest theorists [2] as a non-hierarchical organization, characterized by the lack of a “central body” that direct its operations. To compensate for the lack of hierarchical and central control, any possible actions within these organizations are governed by *ex ante* established rules shared by participants or incentivized through reward mechanisms.

It is here that autonomous agents (i.e. ‘informatic protocols’ or just ‘protocols’) enter the scene. Autonomous agents are software whose existence depends on humans only in their creation: once created, their operation is indifferent to the will of their programmers. [2] While the first example of autonomous agents are computer viruses, modern technologies make possible to realize these systems using *smart contracts* [3]. This makes autonomous agents even more independent from their developer.

Decisions within a DAO are typically adopted through governance procedures that revolve around the submission of a voting proposal (in industry terminology simply referred to as a “proposal”). A proposal may concern virtually any aspect of the DAO’s operations and often addresses highly technical matters aimed at modifying the underlying protocol on which the DAO is built. In most governance frameworks, the right to submit proposals – and subsequently to vote on them – is restricted to holders of governance tokens issued by the DAO itself.

Once the voting process is concluded, the technical implementation of the approved

proposal may occur either automatically or manually. The manual implementation by developers appointed by the DAO is required in all cases where third-party services are used to conduct the vote (i.e., off-chain voting). Although third-party voting platforms (e.g., Snapshot) offer a more user-friendly experience, their use does not guarantee that the approved proposal will be technically executed. In such cases, what is recorded on-chain is merely the voting outcome, not the operational effects that the vote is intended to produce. [4]

The foregoing analysis makes it evident that DAOs incorporate governance mechanisms that diverge significantly from those of traditional organizations, most notably through the degree of automatism with which decisions, once adopted, may be implemented.

However, a deeper understanding of the organizational dynamics of a DAO requires careful attention to the procedures - largely unconstrained in form - that precede the formal voting stage. Behind each DAO there is typically a community that engages in continuous, often informal, discussion through social platforms such as Discord channels or Telegram groups, which constitute the principal *fora* for preliminary deliberation. Within these communities, smaller sub-communities may emerge, occasionally functioning as *de facto* voting syndicates.

The above observations underscore the need to critically assess the meaning attributed to the concept of decentralization. The question that could arise regards the possibility of discussing the decentralization requirement when beyond the DAO there are informal coordination structures made by socially influential actors. Because a DAO remains a “*sociotechnical*” structure involving human participation, the natural tendencies of individuals toward aggregation and cooperative behaviour unavoidably manifest themselves. It is therefore entirely possible - even predictable - that within a DAO, concentrations of influence may arise in the hands of groups of token holders who coordinate around a shared objective.

The key question is the following: does the presence of influential actors or voting syndicates affect the degree of decentralization of a protocol?

3 Decentralization according to European law and authorities

Before engaging with that question directly, it is necessary to survey the existing legal and technical frameworks that have sought to give meaning to the concept of decentralization. This survey will provide the analytical foundation on which the definition proposed in this paper rests.

The most legally significant reference to decentralization is found in the MiCA Regulation. Although MiCA does not provide a specific or autonomous definition of decentralization, the concept is expressly acknowledged and given decisive importance, to the point that the Regulation itself, in Recital 22, states that, where “*crypto-asset services are provided in a fully decentralised manner without any intermediary, they should not fall within the scope of this Regulation.*”

It must be noted, however, that the concept of decentralization is embedded exclusively within a recital of MiCAR [5], and finds no specific operative counterpart within the body of this Regulation. Indeed, no express mention is made within the articles of the regulation with respect to protocols, nor with regard to the offering of decentralized services.

This is itself symptomatic of a restrictive reading of the "non-definition" contained in Recital 22, as further confirmed by the positions adopted by European supervisory authorities on the matter.

Recently, for instance, the EBA and ESMA have expressly stated that MiCAR applies even where only part of the relevant crypto-asset activities or services is carried out in a decentralised manner, whereas services provided in a "*fully decentralised*" manner "*without any intermediary*" fall outside its scope; yet, in the same passage, they immediately acknowledge that MiCAR does not specify how references to "*fully decentralised*" should be interpreted [6].

ESMA had already anticipated this difficulty in its 2023 consultation paper, where, in relation to DEXs, it recognised that Recital 22 excludes services provided in a "*fully decentralised manner without any intermediary*", but also remarked that the exact scope of that exemption remains uncertain and that each system should therefore be assessed on a case-by-case basis in light of its specific features [7]. The same restrictive approach was subsequently confirmed by ESMA in its final report on the second MiCA package, where the Authority stated that the exemption under Recital 22 is only available for services provided in a "*fully decentralised manner without any intermediary*", and that, precisely because of the restrictive nature of that exemption, the draft RTS should also cover DEXs and AMMs that do not meet that "undefined" threshold [8]. At the same time, however, ESMA also expressly admitted in its final report on the first MiCA package that determining whether a service is provided in a partially decentralised manner or in a fully decentralised manner "*is not straightforward*", thereby recognising that the distinction on which the exemption is based remains conceptually and operationally indeterminate [9]. A similar approach emerges from the French AMF, which, after recalling both limbs of Recital 22, stated that a legal analysis of projects must be conducted on a case-by-case basis in order to determine whether crypto-asset services are in fact performed "*in a fully decentralised manner without any intermediary*" [10].

A more structured attempt to give operative content to the concept of decentralization is found in the principles proposed by the Danish Financial Supervisory Authority for the assessment of decentralization under MiCAR. The Danish Authority identified three decisive issues: whether a user can identify a specific legal entity as counterparty; whether smart contracts operate autonomously and without embedded control mechanisms; and whether decision-making powers remain concentrated in a legal entity or are genuinely dispersed among participants. On that basis, the Danish Authority treats decentralization not as a formal label but as the outcome of a substantive, case-by-case assessment encompassing both technical and governance dimensions.

The Danish approach is noteworthy precisely because it moves beyond the circular invocation of decentralization as a self-evident concept, and instead attempts to disaggregate it into verifiable criteria. It remains, however, an administrative guidance document with no binding force, and its criteria - while analytically useful - are not

accompanied by any threshold or metric capable of producing determinate outcomes in concrete cases [11].

4 Definitions given by international authorities

The European regulatory landscape thus presents a consistent pattern: decentralization is acknowledged as legally relevant, but left undefined. A survey of the approaches adopted by other international authorities reveals that this ambiguity is not peculiar to the European framework, but reflects a broader absence of consensus on the meaning of the concept.

At the international level, the simplest definition is provided by the International Organization for Standardization (ISO), which - in the context of blockchains - describes decentralization in very general terms as a “*distributed system wherein control is distributed among the persons or organizations participating in the operation of the system.*”. While this definition captures the core intuition behind the concept, it remains too abstract to generate operational criteria for legal or regulatory purposes.

Moving from standardization bodies to financial regulators, a markedly different approach is adopted by the Financial Action Task Force (FATF), which avoids an abstract definition entirely and adopts a substance-over-form approach: whether a project is considered “decentralized” depends not on self-description but on the actual absence of “a person with sufficient control or influence”; where such a person exists, AML/CFT obligations apply [12]. The FATF approach is significant because it shifts the inquiry from the structure of the system to the factual existence of a controlling individual - a criterion that is verifiable in practice, but that leaves open the question of what “sufficient control” means in the context of distributed governance.

A more analytically developed framework is offered by the International Organization of Securities Commissions (IOSCO), which observes that decentralization may refer to ownership of the service provider, governance rights, user asset control, network architecture, or off-chain infrastructure, and explicitly states that “*there is no agreed definition*” of what constitutes decentralization, understood as the absence of concentrated ownership, voting power, or control over the activity, the organization, or users’ assets [13]. IOSCO’s multidimensional account is notable for the idea of recognising the concept’s possible dimensions, thereby acknowledging that decentralization is not a unitary phenomenon but a composite one.

From a financial stability perspective, the Financial Stability Board (FSB) employs the term in a systemic risk context, describing DeFi as a set of services that aim to replicate financial system functions by “*disintermediating service provision and decentralising governance,*” replacing traditional intermediaries with self-executing smart contracts [14]. Similarly, the Bank for International Settlements (BIS) describes DeFi as the provision of financial services “*without central intermediaries*” while warning that this may give rise to a “*decentralisation illusion*” due to persistent governance centralization [15]. The BIS observation is particularly apposite: it signals that decentralization, as claimed by market participants, may not correspond to decentralization as it

actually operates - a gap that any workable definition must be equipped to detect.

Finally, both the Organisation for Economic Co-operation and Development (OECD) and the International Monetary Fund (IMF) emphasize the substitution of traditional financial institutions with blockchain-based protocols. The OECD defines DeFi as attempting to replicate financial services in a manner that is “*open, decentralised, permissionless and autonomous*” [16], and the IMF states that DeFi involves using public blockchains to execute transactions “*without relying on centralized service providers such as custodians, central clearinghouses, or escrow agents. Instead, these roles are assumed by so-called smart contracts*” [17]. Both institutions treat decentralization primarily as a functional characteristic - the absence of traditional intermediaries - rather than as a structural property of the system itself.

Overall, while international and European authorities converge on the idea that decentralization entails the marginalization or removal of traditional intermediaries, they diverge on whether it constitutes a binary legal threshold or a continuum. Frequently, its recognition is subordinated to a factual analysis of power distribution, counterparty identifiability, and the degree of effective control. What is striking, across all of these frameworks, is the absence of any shared criterion: each authority approaches the concept from the angle most relevant to its own mandate, without engaging with the definitional question as such.

5 Decentralization definitions from the technical point of view

The approaches surveyed above share a common limitation: they are elaborated within specific regulatory mandates and address decentralization from a functional or institutional angle, without engaging with its structural properties. A more systematic understanding of the concept requires turning to the technical and socio-organizational literature, which has developed its own analytical tools for measuring and qualifying the phenomenon independently of any regulatory purpose.

Early academic attempts to measure decentralization focused primarily on quantitative indicators. For instance, the Gini coefficient is a statistical measure of inequality within a distribution, ranging from 0 (perfect equality) to 1 (maximum inequality). Originally developed for income distribution analysis, it is widely used in blockchain studies to assess the concentration of token ownership or mining power. Similarly, the Nakamoto coefficient takes in consideration the minimum number of independent entities that must collude to control more than 50% of the resources in a distributed system, thereby compromising its security or operational integrity. A higher coefficient corresponds to a more decentralized and resilient network.

These two “early” metrics capture different aspects of concentration: the Gini coefficient measures how unequally a resource is distributed across all participants, while the Nakamoto coefficient identifies the minimum number of independent entities that would need to collude to compromise the security of a peer-to-peer network. Together, they provided the first systematic tools for comparing decentralization across different blockchain systems.

A more structured and analytically refined approach is advanced in Gochhayat *et al.*, [18] who propose a multilayered framework distinguishing among three different layers: the governance layer, the network layer, and the storage layer. In their view, a blockchain system cannot be evaluated as a single unit. Rather, it must be decomposed into its constituent layers, and the degree of power concentration or dispersion must be assessed separately for each one. A system may exhibit high concentration at one layer and broad distribution of control at another, so any overall assessment must account for this internal heterogeneity. Accordingly, the authors offer the following definition: *“Decentralized systems are a subset of distributed systems where multiple authorities control different components and no authority is fully trusted by all. Decentrality is a property related to the control over the system. Better decentralization means higher resistance against censorship and tampering.”*

A similar multidimensional analysis is proposed by Axelsen *et al.* [19] Rather than relying on a single metric, they identify five key dimensions through the TIGER framework: Token-weighted voting, Infrastructure, Governance, Escalation, and Reputation. Each dimension captures a distinct locus of control within the system. A key insight is that the degree of concentration or dispersion of power must be assessed separately across all five dimensions, as a system may score differently on each. To visualize this, the authors propose a "decentralization radar" - a spider chart allowing cross-system comparison across multiple axes simultaneously.

Even though the last two approaches allow for a more nuanced assessment than single metrics like the Gini or the Nakamoto coefficient, all definitions and methodological approaches examined above exhibit a predominantly “static” character. They provide conceptual tools for attributing a qualitative assessment to a phenomenon at the specific moment in which it is observed. Yet, in the life of a DAO, governance tokens are typically tradable assets, and their distribution can shift rapidly as tokens move across holders. As a result, the balance of control within the system may change abruptly, rendering any static assessment quickly outdated.

The dynamic dimension of the decentralization phenomenon is more effectively captured in the study by Papangelou *et al.*, [20] the proponents of the *Apokedro Decentralization Index* – a metric designed to evaluate decentralization by considering the probabilities of all possible subsets of nodes that could collectively centralize control. The underlying intuition is particularly effective: decentralization can be meaningfully assessed by quantifying the system’s “exposure” to the risk of centralization.

From a less mathematical and more socio-organizational perspective, the work of Vergne [21] is particularly valuable because it reframes decentralization not as a purely technical property of a system, but as an organizational and performative phenomenon. His work defines decentralization as the dispersion of communication and coordination processes across multiple independent actors, without reliance on a central authority. This definition already signals a shift away from infrastructure-centric approaches toward an analysis focused on how power is actually exercised within a system.

A key analytical distinction introduced by Vergne is that between decentralization and distribution. While the two terms are often used interchangeably in the blockchain

discourse, Vergne argues that they capture different dimensions: decentralization concerns the absence of a central point of control in communication and coordination, whereas distribution refers to the dispersion of decision-making authority among multiple agents. A system may therefore be distributed without being truly decentralized, if coordination ultimately converges toward a dominant actor.

Crucially, Vergne emphasizes the performative nature of decentralization. Rather than treating it as a static or binary attribute, he conceptualizes decentralization as something that must be continuously enacted through practices, governance mechanisms, and interactions among participants. In this sense, claims of decentralization should not be assessed solely on the basis of protocol design or formal rules, but through an empirical analysis of how the system operates in practice - who effectively exercises influence, how decisions are made, and whether power is genuinely dispersed.

This leads to an important methodological implication: decentralization should be evaluated empirically and contextually, by observing real-world dynamics rather than relying on formal or declarative criteria. For instance, even in systems designed to be decentralized, phenomena such as voting concentration, governance capture, or reliance on core developers may reintroduce centralization in practice.

6 The search for a legal definition. The Minimum Threshold of Independent Decision-Makers

The definitions and methodological approaches outlined above provide a valuable point of departure for our discussion, whose aim is to articulate a single normative criterion for the concept of decentralization. What is needed, in particular, is the identification of the minimum characteristics that a system must possess in order to be classified as decentralized. All definitions surveyed thus far focus on the degree to which a system may be controlled. They presuppose the existence of a centralized structure whose control is subsequently dispersed, and then they measure – or essentially calculate – the probability that such a system may revert to the control of a single individual.

It must be observed, first, that because a DAO is, at its core, a human system, it is highly unlikely that its governance structure will remain static over time. It is inherently impossible to predict how such a system may evolve: a single participant may one day acquire all governance tokens from the other holders, or multiple participants may decide to bind themselves through a legal agreement - functionally akin to a shareholders' agreement - in order to manage the DAO's governance in a coordinated and effectively "centralized" manner.

Drawing this conclusion has significant implications for the DAO paradigm. Recognizing the impossibility of ensuring - except through probabilistic methods - the sustained decentralization of a DAO's governance framework leads to a paradox similar to the famous Schrödinger's cat. Focusing solely on the probabilistic dimension means to accept that a DAO may be simultaneously centralized and decentralized at any given moment, given the inherent impossibility of ascertaining the true identity and

relationships of all token holders.

There are few avenues to avoid this paradox. The first is to accept and value the *potential* dimension of decentralization within a DAO, accepting the need to also recognize the socio-economic nature of this concept. In other words, this approach must reckon with a well-documented human tendency already seen in practice: individuals and groups naturally gravitate toward coordination and coalition-building. Even in systems designed to distribute power broadly, participants tend to aggregate, align interests, and act collectively - a dynamic that may progressively reconcentrate governance power regardless of how evenly tokens are distributed at the outset.

One might argue, in other words, that the mere tokenization of decision-making power - and the corollary that each token is, by definition, tradable and easily transferable - makes it possible for governance power to be widely redistributed. Decision-making authority may therefore be characterized as (*potentially*) *decentralized*, irrespective of its contingent concentration at a specific point in time.

Unfortunately, this “probabilistic” approach offers limited utility for legal analysis, as legal definitions necessarily rely on binary classifications: a system must be considered either centralized or decentralized. It is self-evident that legal norms cannot be applied on the basis of probabilities (nor can they be applied *probabilistically*).

From the standpoint of legal qualification, it is of little practical value to know that a system composed, for example, of three layers with ten participants per layer has only a 2% probability of being centralized. Legal application would instead require a concrete assessment of who the participants are and whether any relationships among them may give rise to conflicts of interest or situations of joint control.

If one seeks, therefore, to assign a numerical threshold suitable for normative purposes - while abandoning probabilistic reasoning - it becomes apparent that a system is decentralized whenever it is *not* centralized. If centralization occurs when decision-making power over the entire system can be traced back to a single locus of control, then a system must be deemed decentralized when multiple centers of power exist, provided that these centers do not collude. The minimum “normative” number for defining a system as decentralized is therefore two, assuming that the relevant actors are independent and possess equal decision-making authority.

Indeed, just as a system composed of one million independent actors can be described as decentralized (i.e. not centralized), the same conclusion may be reached when the system comprises ten thousand, one thousand, or even smaller numbers such as one hundred, fifty, or twenty-five. Repeating this reductive exercise reveals that the minimum number of actors required for a system not to be centralized is at least three - that is, the smallest number that allows a collective to form a majority whenever a decision must be taken.

A system may therefore be regarded as decentralized when, at least from a logical and governance-theoretical standpoint, it is composed of *no fewer than three independent actors*. Indeed, a system composed of only two decision-makers cannot be considered centralized, yet the mutual veto power inherent in a structure with just two parties forces both to reach identical decisions, failing which the system becomes paralyzed.

Thus, the presence of at least three independent decision-making centers should become the focus of analysis for regulatory authorities seeking to determine whether a DAO system satisfies the requirement of decentralization.

7 The Symbiotic Relationship within an Immutable Environment

A closer examination of commonly observed decentralized systems reveals an additional feature, which may be described as a form of “mutualistic symbiosis.” To appreciate this, the protocol offering a given service must be considered *holistically*, rather than solely from the perspective of the party providing the service. For instance, in assessing a decentralized exchange, one must not limit the inquiry to the developers of the exchange or to those capable of exercising control over it. Rather, it is necessary to account for the users of the system and the manner in which they interact with one another and with the protocol itself.

Focusing on user interaction highlights that the decentralization of a protocol cannot be reduced to the issue of control alone. Even where control over the protocol is, *de facto*, attributable to a single entity, the protocol cannot operate without the participation of a plurality of actors (i.e., at least more than one). This observation diminishes the normative relevance of probabilistic assessments of control.

If a protocol requires, for any meaningful functionality, the participation of at least more than one actor performing distinct roles, then one may already observe the presence of a minimal degree of decentralization within that system.

To this, it could be objected that many services require multiple actors in order to operate (i.e. one that sells a service and one that buys it). But here, what distinguishes decentralized finance is the protocol itself – an immutable set of rules to which all participants must conform and within its boundaries they move and act.

It is true that in traditional sectors services also require multiple actors (traditional securities markets depend on a variety of intermediaries – such as clearing houses and brokers – to function). Yet the most decisive role is occupied by a central and centralized authority. In decentralized systems, the decisive role is covered by the protocol (a mere software) whose immutability favor the creation of symbiotic relations between the parties involved and whose basic actions should not be subject to the *stable* control of any particular entity.

To summarise, the criterion of mutualistic symbiosis identifies two features that a system must exhibit in order to qualify as decentralized: (i) a multiplicity of interactions among participants, none of whom can operate the system unilaterally; and (ii) the presence of a protocol that functions as the shared and neutral framework within which all participants coordinate their activity, without being subject to the stable control of any single entity.

8 The Non-Custodial Infrastructure

The analysis conducted in the preceding sections has approached decentralization primarily as a question of control over decision-making: who governs the protocol, how many independent centers of power exist, and whether governance is genuinely dispersed among participants. This dimension, however, does not exhaust the legal significance of decentralization. Alongside the governance criteria developed above, a structurally distinct and equally fundamental dimension must be considered: the non-custodial nature of the infrastructure through which decentralized services are delivered.

This third criterion finds a normative anchor in Recital 22 of MiCAR. As noted above, that recital contains two elements that point to the same underlying idea. While attention has largely focused on the requirement that services be offered in a "fully decentralised manner", a second and equally important element is the specification that this must occur "without any intermediary." This latter characteristic should not be understood as an additional requirement distinct from the former, but rather as a clarification of its meaning. Disintermediation, in this sense, is inherent in the very notion of decentralization - and it is precisely this dimension that the third criterion seeks to capture.

Decentralization, understood in this sense, denotes the renunciation of control over users' assets. In a non-custodial system, the service provider - whether a person, an entity, or a protocol - never acquires possession, custody, or effective control over the assets that users deploy within the system. Users retain direct, unmediated access to their own assets at all times, through self-custodied cryptographic keys. No commingling occurs between the patrimony of the service provider and that of the user. This stands in sharp contrast to traditional financial intermediation, where the defining feature of the relationship between intermediary and client is precisely the transfer - temporary, fiduciary, or contractual - of asset control from the latter to the former.

This structural distinction carries decisive regulatory implications. The regulatory architecture developed for traditional financial intermediaries is grounded, at its core, in the risks that arise from that transfer of control. Prudential requirements - capital adequacy rules, liquidity buffers, leverage constraints, asset segregation obligations - are designed to protect users against the risk that the intermediary, having obtained control over their assets, may become insolvent, misappropriate those assets, or otherwise fail to return them [22][23]. Similarly, conduct-of-business requirements - best execution, conflicts of interest, suitability assessments - presuppose an ongoing relationship in which the intermediary exercises discretion over assets it holds on behalf of the client.

Where the infrastructure is genuinely non-custodial, these risks do not arise in the same form. An entity that never holds user assets cannot become insolvent in the sense that threatens their recovery; it cannot commingle or misappropriate what it has never possessed; it cannot generate systemic risk through balance-sheet interconnections that do not exist. The regulatory rationale for imposing entity-level prudential requirements therefore dissolves, or at least transforms fundamentally. Applying to non-custodial protocols the same capital and organizational requirements designed for custodial intermediaries would not only be technically inappropriate - because the regulated entity

lacks the structural features that those requirements presuppose - but would also be practically ineffective, since the protocol would continue to operate regardless of whether any legal person associated with it satisfies such requirements [24][25].

This observation is confirmed, rather than contradicted, by the existing regulatory framework - not only under MiCAR, but across the broader architecture of financial regulation. The entire body of rules governing financial intermediaries - from banking prudential requirements to investment firm conduct obligations, from payment service provider licensing to securities custody rules - is premised on a single structural assumption: that a regulated entity obtains, holds, or exercises discretion over the assets or interests of its clients. Capital adequacy requirements, liquidity buffers, asset segregation obligations, best execution duties, and suitability assessments all presuppose this custodial or quasi-custodial relationship. They are designed, at their core, to protect users against the risks that arise when another party acquires control over what belongs to them.

Where that structural assumption does not hold - where no entity ever acquires possession, custody, or effective control over user assets - the regulatory rationale for these requirements is correspondingly diminished. The risks of insolvency, misappropriation, and balance-sheet contagion presuppose the existence of a custodial relationship. In its absence, those risks do not arise in the same form, and the imposition of entity-level requirements designed to manage them would be not only technically inappropriate but also practically ineffective: the protocol would continue to operate regardless of whether any associated legal person satisfies such requirements.

This judgment is implicitly reflected in the exemption established by Recital 22 of MiCAR for services provided in a "fully decentralised manner without any intermediary." That exemption is not an oversight or a drafting gap: it reflects a substantive recognition that where no entity holds or controls user assets, the principal dangers that financial regulation is designed to address do not present themselves in their traditional form. The existing regulatory framework thus already contains, in embryonic form, an acknowledgment that non-custodial systems occupy a structurally different position - one that calls for a different regulatory response, or, in the most genuinely decentralized cases, for the absence of entity-level regulation altogether.

This does not mean that non-custodial systems are risk-free. The risks that persist in genuinely decentralized infrastructures are, however, partially different in nature: they concern the integrity and transparency of the code itself, the conditions under which it may be modified, and the exposure of users to smart contract vulnerabilities [26][27]. These are real concerns, but they are concerns of a different kind - and the existing framework, by exempting fully decentralized services from its scope, implicitly acknowledges that distinction. What it fails to do is articulate the criteria by which that exemption can be concretely applied - a gap that, as argued throughout this paper, the three criteria proposed here seek to fill.

To summarise, decentralization in the custodial dimension provides a normatively coherent basis for distinguishing decentralized protocols from traditional intermediaries, and for justifying the non-application of existing regulatory response.

9 Conclusions and further research

The analysis conducted in the preceding sections reveals a fundamental inadequacy in the existing approaches to decentralization. Regulatory authorities have treated the concept as self-evident, invoking it as a threshold criterion without defining its content. Technical and academic literature has developed sophisticated measurement tools - from the Gini and Nakamoto coefficients to multidimensional frameworks such as TIGER - but these tools are inherently static and probabilistic. They capture a snapshot of a system at a given moment; they cannot account for the dynamic nature of human systems, nor do they produce the binary determinations that legal qualification requires.

This paper has argued that a legally workable definition of decentralization must be built on a different foundation. Rather than measuring the statistical distribution of tokens or computing the probability of centralization, it must identify a set of minimum structural criteria that are verifiable in practice and stable enough to support normative application. Three such criteria have been proposed.

The first is the presence of at least three independent decision-making centers. This criterion follows from a reductive logical argument: a system is decentralized when it is not centralized, and the minimum number of actors that prevents centralization while allowing majority formation is three. This threshold is not arbitrary - it is the lowest number at which collective governance becomes structurally possible without degenerating into mutual veto or de facto unilateral control.

The second criterion concerns the structural interdependence of participants within an immutable environment. A decentralized protocol is not merely a system with multiple controllers - it is a system in which the protocol itself, rather than any single actor, constitutes the shared framework within which all participants operate. The multiplicity of interactions, governed by immutable or collectively amendable rules, generates a form of mutualistic symbiosis that is structurally distinct from traditional intermediated relationships.

The third criterion is the non-custodial nature of the infrastructure. A system cannot be regarded as genuinely decentralized if a single entity retains - or may at any time acquire - effective control over users' assets. Non-custody is not a technical detail; it is the structural condition that removes the principal regulatory rationale for entity-level prudential requirements.

These three criteria are cumulative. The absence of any one of them is sufficient to bring a system within the orbit of traditional regulatory categories. Together, they define a minimum threshold below which a claim of decentralization cannot be sustained for legal purposes.

The implications for regulatory policy are significant. Existing frameworks - including MiCAR - invoke decentralization as an exemption condition without specifying its content, leaving both market participants and supervisory authorities without operational guidance. The criteria proposed here offer a basis for filling that gap. Looking forward, the analysis developed in this paper suggests that future regulatory instruments - whether legislative or issued by competent authorities in the form of technical standards or guidance - should articulate decentralization as a composite legal concept

defined by these three verifiable elements, rather than treating it as a binary label to be assessed on a purely case-by-case basis.

Such an approach would serve multiple objectives simultaneously: it would provide legal certainty for protocol developers and users; it would give supervisory authorities a structured framework for assessment; and it would resist the risk - already documented in practice - that the label of decentralization is invoked strategically to avoid regulatory obligations without satisfying any substantive criterion. A definition of decentralization is not a merely academic question. It is a precondition for any coherent regulatory engagement with decentralized finance.

References

1. Born, Z., Gati, Z., Lambert, C., Naeem, M., Pellicani, A.: Who to Regulate? Identifying Actors within DeFi's Governance. ECB Working Paper Series No. 3208, p. 33 (2026).
2. Buterin, V.: DAOs, DACs, DAs and More: An Incomplete Terminology Guide. Ethereum Blog (2014).
3. Villani, C.: Vulnerabilities of the DeFi Ecosystem. Legal Liability Challenges of New Autonomous Agents. SSRN (2025).
4. Baliotti, S., Saggese, P., Kitzler, S., Haslhofer, B.: Slaying the Dragon: The Quest for Democracy in Decentralized Autonomous Organizations (DAOs). SSRN (2025)
5. Regulation (EU) 2023/1114 of the European Parliament and of the Council of 31 May 2023 on markets in crypto-assets, and amending Regulations (EU) No 1093/2010 and (EU) No 1095/2010 and Directives 2013/36/EU and (EU) 2019/1937
6. EBA, ESMA: Joint Report on Recent Developments in Crypto-Assets (Article 142 of MiCAR), p. 4, 13 January 2025.
7. ESMA: Second Consultation Paper on MiCA, p. 29, 5 October 2023.
8. ESMA: Final Report - MiCA: Second Package, p. 31-32, 3 July 2024.
9. ESMA: Final Report on Draft Technical Standards Specifying Certain Requirements of the Markets in Crypto-Assets Regulation (MiCA) - First Package, p. 48, 25 March 2024.
10. AMF: Summary of Responses to the AMF Discussion Paper on Decentralised Finance (DeFi), p. 13, July 2024.
11. Danish Financial Supervisory Authority: Principles for the Assessment of Decentralisation in the Markets for Crypto-Assets, pp. 2–4, 24 June 2024
12. FATF: Updated Guidance for a Risk-Based Approach to Virtual Assets and VASPs, p. 23, (2021).
13. IOSCO: Policy Recommendations for Decentralized Finance (DeFi), p. 8, (2023).
14. Financial Stability Board: The Financial Stability Risks of Decentralised Finance (2023).
15. Bank for International Settlements: DeFi Risks and the Decentralisation Illusion, p. 4, (2021).

-
16. Organisation for Economic Co-operation and Development: Why Decentralised Finance (DeFi) Matters, p.6, (2022).
 17. International Monetary Fund: DeFi's Promise and Pitfalls, p. 24, (2022).
 18. Gochhayat, S.P., Shetty, S., Mukkamala, R., Foytik, P., Kamhoua, G.A., Njilla, L.: Measuring Decentrality in Blockchain Based Systems. *International Journal of Network Security*, p. 191 (2020)
 19. Axelsen, H., Jensen, J.R., Ross, O.: When is a DAO Decentralized? *Complex Systems Informatics and Modeling Quarterly* 31 (2022).
 20. Papangelou, S., Christodoulou, K., Inglezakis, A.: Apokedro: A Decentralization Index for DAOs and Beyond. *Blockchains* (2025).
 21. Vergne, J.-P.: Decentralization: The Future of Online Governance. De Boeck (2020).
 22. Basel Committee on Banking Supervision: Basel III: A Global Regulatory Framework for More Resilient Banks and Banking Systems, revised version (2011); FINANCIAL STABILITY BOARD, *Key Attributes of Effective Resolution Regimes for Financial Institutions*, 2014, p. 5
 23. Financial Stability Board: Key Attributes of Effective Resolution Regimes for Financial Institutions, p. 5 (2014).
 24. International Organization of Securities Commissions: Policy Recommendations for Decentralized Finance (DeFi), p. 21 (2023).
 25. Bank for International Settlements: DeFi Risks and the Decentralisation Illusion, p. 10 (2021).
 26. ESMA: Final Report - MiCA: Second Package, pp. 31-32 (2024)
ESMA, *Final Report - MiCA: Second Package*, 3 July 2024, pp. 31–32;
 27. IOSCO: Policy Recommendations for Decentralized Finance (DeFi), pp. 28-30 (2023)