Politecnico di Torino

Master's degree in engineering and Management

Ottobre 2025

The Impact of Web3 on Platform Economy: A Transaction Cost Economics Perspective

Tutors

prof.

PERBOLI GUIDO
VANDONI CHIARA

Candidate

ADEM ABD RABOU

ABSTSTACT

This thesis dives into how Web3 can create efficiencies compared to the Web2 model, all through the lens of Transaction Cost Economics (TCE) in the digital economy. Rather than just sticking to the traditional TCE framework, it expands on it to address the new complexities that come with decentralized systems. It takes into account factors like implementation and training costs, the hurdles of distributed governance, the necessity for secure code and infrastructure, as well as issues related to scalability and usability, not to mention the long-term resilience of protocols. To bring these ideas to life, the study conducts a comparative analysis between a Web2 benchmark solution (Oracle SCM) and a Web3 alternative (VeChain). The findings indicate that Web3 could potentially lower ex post costs those associated with monitoring and enforcement, largely due to automation, transparent shared ledgers, and the ability to embed contractual rules into smart contracts. However, the analysis also points out that Web3 might raise adaptation costs and introduce new challenges related to governance, security management, and user experience. In this context, achieving net efficiency isn't a given; it hinges on factors like the number and diversity of participants, the significance of end-to-end auditability, the level of institutional stability, the organizational maturity of those adopting the technology, and the presence of interoperability standards that help minimize fragmentation. Ultimately, the main contribution of this thesis is the introduction of an extended TCE–Web3 framework. This framework not only revisits transaction cost theory for the decentralized economy but also serves as a practical tool for assessing technological make-or-buy decisions and crafting hybrid business models. In doing so, it provides valuable insights for both academics looking to refine economic theory and practitioners making strategic decisions in an increasingly digital and decentralized business environment.

Keywords:

Web3
Transaction Cost Economics (TCE)
Blockchain
Smart contracts
DAO
Emerging costs

INDEX

Introduction

Over the last twenty years, we've seen a remarkable shift in the way the digital economy is structured.

The emergence of large-scale platforms in the Web2 paradigm has generated new business models based on network effects, data-driven strategies, and platform-mediated transactions, at the same time, the concentration of power in a few hands has fueled intense debate on market concentration, privacy, surveillance capitalism, and user attention extraction.

Seen in this way, massive data collection and algorithmic manipulation become central to value creation processes (Zuboff, 2019). Despite generating global connectivity and profits, this model has revealed structural limitations such as poor interoperability, lock-in, and dependence on intermediaries, which have also been highlighted by data management scandals and the information asymmetry induced by recommendation algorithms. These critical issues have undermined user confidence and reveal the fragility of a highly centralized ecosystem.

Against this context, the emergence of Web3 technologies has been described as a potential paradigm shift.

Based on blockchain infrastructure, smart contracts, and decentralized autonomous organizations (DAOs) , Web3 proposes a decentralized architecture for the digital ecosystem so instead of central platforms, the goal is to redistribute decision-making and control power, allowing individuals and organizations to coordinate, transact, and govern themselves in a more transparent and trustless manner. As Davidson, De Filippi, and Potts have observed, blockchain "replaces traditional trust with algorithmic transparency," reconfiguring coordination mechanisms and mitigating the dependencies typical of Web2 (Buterin, 2014; Davidson, De Filippi, & Potts, 2018).

The relevance of this research lies at the intersection of the evolution of digital business models and transaction cost theory. Transaction Cost Economics (TCE), pioneered by Coase (1937) and systematized by Williamson (1985), offers a powerful lens for understanding the boundaries of the firm and the efficiency of different forms of governance. By focusing on the costs of search, negotiation, monitoring, enforcement, and adaptation, TCE explains why some activities are internalized while others are entrusted to the market or hybrid forms.

The rise of Web3 technologies thus reopens classic questions but with new implications: can decentralized architectures reduce transaction costs compared to Web2 systems? , and do they introduce cost categories that "classical" TCE does not fully capture?

The objective of the thesis is therefore with two main goals. Firstly, to analyze the impact of Web3 on transaction costs through a structured application of the TCE framework and on the other hand, to develop a conceptual extension of TCE that includes emerging categories : implementation and training, distributed governance, code security and auditability, scalability, usability, and protocol resilience that arise when transactions are enabled by open and programmable infrastructures. In this extension, both the relational dimension (stakeholder relationships, trust, cooperation) and that of institutional resilience are important, as discussed

in the most recent literature as well as in our mapping of "emerging costs" developed in Chapter 3.

From a methodological point of view, this thesis adopts a comparative case study design in the field of supply chain management. The comparison is between two solutions: Oracle SCM (a widely used Web2 benchmark in large companies) and VeChain (a Web3 platform focused on supply chain traceability and transparency). The analysis is conducted along three key dimensions of TCE(ex-ante, ex post, and adaptation costs) integrating the assessment of the emerging cost categories mentioned above. This approach allows for a detailed assessment of the trade-offs highlighting the conditions under which decentralized infrastructures outperform or underperform centralized and hybrid alternatives.

The expected contribution is threefold. First of all, the thesis shows that Web3 does not uniformly reduce transaction costs but tends to redistribute them across different stages of the exchange (marked reductions in monitoring and enforcement when the object is natively digital but also potential increases in adaptation and governance costs when uncertainty and distributed coordination increase). Second, it proposes an extended TCE-Web3 framework, which incorporates relational and resilience dimensions, offering a more comprehensive lens for the study of decentralized technologies. Third, it provides empirical evidence through comparative cases, useful to scholars and practitioners for assessing the conditions of efficiency and sustainability of Web3 adoption.

The structure of the thesis is as follows: Part I reconstructs the evolution from Web2 to Web3, highlighting the limitations of centralized models and the opportunities of decentralized infrastructures. Part II presents Transaction Cost Economics (origins, key concepts, applications to digital) and the last part constitutes the analytical core: Chapter 3 examines the enabling technologies of Web3 (blockchain, smart contracts, DAOs) and their effects on transaction costs, introducing emerging cost categories. Chapter 4 develops the comparative cases of Oracle SCM vs. VeChain in light of TCE, discussing the trade-offs between Web2 and Web3 architectures. The Conclusions summarize the results, indicate strategic implications for businesses, and propose avenues for future research. In this sense, the thesis does not intend to make a definitive judgment on the superiority of Web3 over Web2, but to advance a conditional perspective that is to say that the efficiency of Web3 depends on contextual factors (sector, investment specifics, technological and institutional uncertainty) and on the balance between reduced intermediation costs and new organizational frictions related to security, governance, and scalability. In TCE terms, the 'organizational optimum' can shift when technological innovation introduces costs not foreseen by the original framework, as a matter of facts the most recent literature shows, that blockchain reduces verification and networking costs however can increase adaptation and coordination costs in turbulent environments.

Chapter 1 – From Web1 to Web3: Evolution, Limits and Opportunities

1.1 The Origins and Evolution of the Web

The World Wide Web represents one of the most significant technological and cultural transformations of the late 20th century. Conceived in 1989 by Tim Berners-Lee at CERN, the Web was created as a system for sharing documents and research data through hypertext links between distributed computers (Berners-Lee, 1999).

A crucial moment in its success came on April 30, 1993, when CERN decided to make the Web software public, free of licenses, thus ensuring open and universal accessibility.

This institutional choice of openness, contrary to proprietary logic, was decisive: it enabled the rapid global spread of the Web and accelerated its large-scale adoption, first in scientific contexts, then in industry, and finally in society as a whole.

The first phase of this technological evolution was later defined by the term Web 1.0. This phase, which can be roughly placed between 1989 and the early 2000s, was characterized by static sites, poor interactivity, and a highly unidirectional information structure (Cormode & Krishnamurthy, 2008). The architecture was based mainly on static languages such as HTML and client-server logic, with pages linked by simple hyperlinks. Users could only read and consume information created by a small group of content producers, generally institutions, companies, or individual webmasters.

In Web 1.0, "content creators were few, while the vast majority of users were consumers of content" (Cormode & Krishnamurthy, 2008). For this reason, Web 1.0 is often described as a "read-only web," a medium in which participation was limited to consulting information, without the possibility of modifying or enriching it.

The nature of Web 1.0 reflected the socio-technical conditions of the time. Born in an academic and scientific context, it was not designed as a tool for commercialization or mass interaction, but rather as a neutral infrastructure for the dissemination of knowledge.

Gillies and Cailliau (2000) emphasize how the project embodied ideals of openness and universality, distinguishing itself from other proprietary systems in circulation at the time, such as America Online (AOL) or CompuServe, which offered closed, subscription-based ecosystems. While these systems ended up stagnating or disappearing the Web thrived precisely because of its universal accessibility and interoperability, which transformed it into a global communication infrastructure.

Technically, Web 1.0 sites were rudimentary and static, built mostly with simple text, hyperlinks, and basic images, with no dynamic content.

Possible interactions were minimal and often limited to experimental tools as guestbooks, rudimentary forums, or simple email contact forms. In this sense, the usage model was closer to a digital library than a participatory space, as reiterated by Berners-Lee himself (1999), who envisioned the Web primarily as a tool for universal information sharing rather than a collaborative platform. From an economic point of view, this early phase was not yet dominated

by complex digital platforms: revenue models were essentially based on online advertising and e-commerce catalogues very similar to traditional commercial logic.

Towards the end of the 1990s, however, signs of change began to emerge. The growing spread of personal computers and the expansion of Internet connections stimulated new expectations: users no longer wanted to limit themselves to passively consuming content, but sought more active forms of participation (Manovich, 2001), at the same time, technological innovations introduced greater possibilities for interactivity.

The development of scripting languages such as JavaScript and the introduction of technologies such as Flash made it possible to enrich websites with multimedia elements, interactive menus, and more complex applications. These advances paved the way for the emergence of personal blogs, thematic forums, and collaborative platforms such as Wikipedia, founded in 2001, which encouraged users to contribute directly to the creation of knowledge (Lih, 2009). These were still partial and unsystematic innovations, but they signaled the gradual erosion of the boundary between content producers and consumers.

In retrospect, Web 1.0 can be interpreted as a phase of "digital infancy" (Fuchs, 2017): an open and global infrastructure that made access to information possible but proved to be limited in terms of participation.

As Castells (2001) points out, the late 1990s saw the convergence of technological innovations and new socio-cultural demands, which would soon favor the transformation of the Web from a simple information archive to a participatory ecosystem.

It was on this basis that the next phase, the so-called Web 2.0, would establish itself. However, before addressing its characteristics and critical issues, it is important to recognize how the original vision and open architecture of Web 1.0 represented the fundamental pillar on which all subsequent iterations of the Web were built.

1.2 The Features and Structural Limits of Web 2.0

The concept of Web 2.0 was introduced in the early 2000s to describe a new phase in the evolution of the Web characterized by greater interactivity, user participation, and platform centrality. The definition was anticipated by Darcy DiNucci in 1999 and later popularized by Tim O'Reilly in 2004 at the "Web 2.0" conference (O'Reilly, 2005). Although the label suggests a new formal version of the network, the term actually refers to a set of socio-technological transformations that have profoundly changed the way we conceive and use the Web, transforming it from a simple information archive into a participatory and dynamic space (Fuchs, 2010).

One of the central features of Web 2.0 is the emergence of users as active content producers. Blogs, forums, social networks, and video sharing platforms such as YouTube are paradigmatic examples of an ecosystem in which individuals do not merely consume information, but participate directly in its creation, modification, and dissemination.

This phenomenon has led to the emergence of the so-called prosumer, i.e., the user who is both a producer and consumer of content, embodying the participatory and collaborative logic of Web 2.0 (Flew, 2008). From a technical point of view, this phase was based on the adoption of new dynamic technologies, such as AJAX and HTML5, capable of offering a more fluid, interactive, and real-time browsing experience (Cormode & Krishnamurthy, 2008). At the same time, the spread of cloud computing has enabled scalable and flexible access to digital resources, reducing technological and infrastructural barriers.

Alongside the benefits in terms of interactivity, collaboration, and democratization of content production, Web 2.0 has nevertheless shown significant structural limitations. The first concerns the growing centralization of power in the hands of a few large platforms that operate as gatekeepers of the digital economy. Giants such as Google, Apple, Facebook, and Amazon (GAFA) now control most online traffic and data, imposing lock-in conditions that make it costly for users to migrate to alternative solutions (Farrell & Klemperer, 2007).

In this scenario, the contents and data generated by users, Zuboff (2019) observes, do not remain under their direct control, but become the property of the platforms that exploit them to fuel extractive economic models based on targeted advertising.

This model, described by Zuboff (2019) as surveillance capitalism, is based on the systematic transformation of personal data and user attention into economic resources. Platforms collect, analyze, and monetize digital behavior, generating value through profiling and the intensive use of predictive algorithms in this way, users unwittingly become suppliers of 'free raw material' without any real economic return, with a clear asymmetry in the distribution of value (Momtaz, 2022).

The absence of real digital ownership by individuals translates into fragility: accounts, data, and relationships are in fact subject to the terms of service of platforms, which can unilaterally change the rules, suspend profiles, or remove content (Helmond, 2015).

A further limitation concerns the quality and reliability of information in the Web2 ecosystem. Recommendation algorithms, optimized to maximize engagement, have favored the formation of echo chambers and the rapid spread of polarizing content or misinformation.

Echo chambers, as Pariser (2011) points out, are closed information environments in which users are exposed almost exclusively to content that confirms their pre-existing opinions, thus reducing the diversity of sources and strengthening social and political polarization.

At the same time, scandals such as Cambridge Analytica have highlighted the risks of concentrating data in the hands of a few players, in 2018, indeed it emerged that the company had improperly collected data from approximately 87 million Facebook users, using it to build psychometric profiles and influence, through targeted advertising, crucial political choices such as the Brexit referendum and the 2016 US presidential election. This episode highlighted not only shortcomings in privacy protection, but also the potential for manipulative use of algorithms, undermining user trust in digital platforms.

Finally, the centralized model of Web 2 has also proven vulnerable to disruptions, cyberattacks, and inefficiencies related to the absence of truly interoperable protocols. While Web 2 has democratized content creation and enabled the emergence of new digital business models, it has also generated new forms of inequality, surveillance, and dependence on intermediaries.

Digital inequalities do not only concern access to the network, but also the ability of users to have the tools and skills to protect their security and data, thus accentuating the structural fragilities of the Web 2 paradigm. (Madden, 2017) It is precisely these limitations that have paved the way for the development of new alternatives, such as Web 3, conceived as a decentralized and transparent response to the intrinsic criticalities of Web 2.

1.3 The Emergence of Web3: Foundational Principles and Infrastructures

The Web3 paradigm was born as a direct response to the structural limitations of Web2, marked by the centralization of power, the opaque extraction of personal data, and infrastructural vulnerability, proposing a project of institutional and economic reform of the digital world based on "decentralization," "transparency," "distributed ownership," and "programmable trust" (Buterin, 2014).

Decentralization is better understood as an institutional principle rather than a technical feature, since blockchains function as shared and immutable ledgers that validate transactions collectively and verifiably instead of through a central authority. Consequently, as Davidson, De Filippi, and Potts observe, "blockchain replaces traditional trust with algorithmic transparency" (Davidson, De Filippi, and Potts, 2018), reconfiguring trust mechanisms and enabling forms of peer-to-peer coordination that mitigate the dependencies and intermediation powers typical of Web2.

Another pillar concerns the redefinition of property rights and information control. In Web2, users' personal data has become the raw material of the "attention economy" (Zuboff, 2019).

Web3 aims to overturn this logic by proposing the idea of self-sovereign identities and cryptographic primitives (wallets, attestations, verifiable credentials) that allow individuals to regain control over the management and monetization of their data.

In parallel with tokenization and the introduction of NFTs (non-fungible tokens), the very notion of digital ownership is undergoing a radical transformation: users can own unique assets, transfer them, monetize them, and above all, maintain full control over them (Catalini & Gans, 2016). This reduces forms of technological lock-in and opens up previously unthinkable possibilities for economic participation.

Governance is also being rethought as a native component of the infrastructure. Decentralized Autonomous Organizations (DAOs) encode decision-making rules in smart contracts and orchestrate participation through voting mechanisms and economic rights expressed in tokens: "DAOs encode governance rules in open protocols, enabling transparent and collaborative decision-making" (Hassan & De Filippi, 2021).

This reconceptualization shifts the axis from vertical models, centered on proprietary platforms, to forms of programmable community self-governance, where economic incentives (tokens) and coordination tools (smart contracts) are intertwined.

Talking about infrastructure, Web3 operates on a layered architecture. This means that blockchain and consensus mechanisms ensure integrity and finality. Smart contracts allow for conditional logic to be executed without needing to trust one another. Tokens define rights, incentives, and the flow of money, while DAOs offer flexible organizational structures. From an economic and organizational standpoint, this setup can be viewed through the lens of transaction cost economics. The shared transparency and automated execution can help cut down on costs related to searching, negotiating, monitoring, and enforcing agreements between organizations. However, it also brings about new challenges and constraints, like technical complexity, the risks associated with governance by code, and reliance on oracles.

Web3 opens space for new native markets and business models, too. Decentralized finance (DeFi) demonstrates the on-chain replicability of financial functions: exchange, lending, market-making, and insurance without traditional intermediaries, coordinated by token-based incentives and community governance (Schär, 2021).

In general the tokenization of real and digital assets and community-owned models aim to realign the distribution of value from proprietary platforms to the communities that generate it, hypothesizing a competitive rebalancing in digital markets through widespread ownership and protocol interoperability moreover, through distributed ownership and programmable trust mechanisms, Web3 opens up the possibility of giving users back not only autonomy, but also a renewed ability to actively influence economic and organizational processes.

1.4 Critical Issues and Challenges of Web3

Despite the promises of decentralization and transparency, the Web3 paradigm is not without structural challenges that limit its diffusion and its effective ability to replace the Web2 model. Three aspects emerge as particularly critical: scalability, complexity of use, and fragmentation of the ecosystem.

The first major challenge concerns scalability. Public blockchains, particularly those based on consensus mechanisms such as Proof of Work or Proof of Stake, suffer from inherent limitations in transaction processing speed as a matter of facts while centralized payment systems can process tens of thousands of transactions per second, networks such as Bitcoin or Ethereum can process significantly fewer (Gervais, 2016). As Xu et al. (2019) point out, "scalability remains one of the most pressing limitations for blockchain networks, preventing their widespread adoption." This gap reduces the competitiveness of Web3 applications in contexts that require speed and high processing capacity, such as mass digital payments or real-time supply chain management.

The second critical aspect is the complexity of use and cognitive barriers. While Web3 offers distributed ownership and participatory governance models, the user experience often remains difficult. The use of cryptographic wallets, private keys, and complex security procedures is a significant obstacle for the average user.

Huang et al. (2022) argue that, "the usability gap between Web2 and Web3 applications is a crucial barrier for mainstream adoption", this means that, although more experienced users are able to interact with dApps and DAOs, the mass of digital consumers continues to prefer the simplicity and familiarity of Web2 platforms.

The last critical issue is the fragmentation of the Web3 ecosystem: unlike Web2, which is dominated by a few large, centralized players, Web3 is characterized by a plurality of blockchains, protocols, and standards that are not always interoperable with each other. This creates technological silos that hinder the user experience and reduce the benefits of decentralization: "the lack of interoperability across decentralized networks leads to fragmentation and limits the formation of cohesive digital ecosystems" Schär (2021). This fragmentation affects not only the technical aspect, but also the regulatory and institutional aspects: different countries are adopting divergent regulations on cryptocurrencies, smart contracts, and DAOs, creating a complex and uncertain legal mosaic.

These critical issues are compounded by other emerging problems. Web3 networks, based on economic incentives and tokenized models, can generate new forms of speculation and financial volatility (Momtaz, 2022). Furthermore, the energy consumption of some blockchains, although declining with the introduction of more efficient algorithms, remains an open question in terms of environmental sustainability (Mora, 2018).

Web3 still has significant hurdles to overcome before it can truly establish itself as an alternative to Web2. The promise of a more equitable and decentralized ecosystem risks remaining unfulfilled unless issues related to scalability, interface complexity, and technological and regulatory fragmentation are addressed. It is precisely in this context that Transaction Cost Economics becomes relevant, offering a tool for assessing whether the benefits introduced by Web3 truly outweigh the new costs and inefficiencies it generates.

To better frame this transition, the following comparative matrix (Tab.1) illustrates the distinctive features of Web 1.0, Web 2.0, and Web 3.0 across their core dimensions, highlighting the progressive shift from static information consumption to participatory interaction, and finally to decentralized ownership and governance.

Tab 1- Comparative Table: Web 1.0 vs Web 2.0 vs Web3.0

Aspect	**Web 1.0**	**Web 2.0**	**Web 3.0**
Core Philosophy	"Read-only": users are passive consumers	"Read-write": user participation and sharing	"Read-write-own": user ownership and decentralized coordination
Content Creation	Institutions, companies,	Users via blogs, social media, platforms	Users & communities; tokenized/on-chain assets
Data Ownership	Website owners	Controlled by platforms User data monetized via surveillance capitalism	Users own and manage their own data via self-sovereign identities and wallets
Governance	Webmaster/admin control	Platforms impose Terms of Service (centralized)	Transparent and participatory governance via DAOs and open-source smart contracts (decentralized)
Value Distribution	Almost none to users	Concentrated in platforms	Shared among users, creators, communities
Examples	Personal pages, Yahoo!	META, YouTube,	Ethereum, DAOs
Risks & Limitations	Low interactivity, poor scalability	Privacy issues, lock-in, censorship, misinformation	Technical risks, complexity, regulatory uncertainty

Chapter 2 – Theoretical Framework: Foundations of Transaction Cost Economics

2.1 Origins and Key Concepts of TCE (Coase, Williamson)

Transaction Cost Economics (TCE) stems from the reflections of Ronald H. Coase, who in his famous article The Nature of the Firm (1937) asked a radical question: why do firms exist in a theoretical context where the market should ensure the optimal coordination of economic activities through the price mechanism? If the market were perfectly efficient, all transactions would take place within it, without the need for hierarchical organizations. Coase's answer was that the use of the market is not free but involves costs: "the main reason why it is profitable to establish a firm would seem to be that there is a cost of using the price mechanism" (Coase, 1937). Among these, he cited the difficulty of discovering relevant prices, the costs of negotiating and entering contracts, and the need for supervision to ensure compliance with the terms.

This insight led to an explanation for the existence of firms: they arise, therefore, not to replace competition, but to internalize transactions and reduce these costs, ensuring greater efficiency where the market is too costly as a coordination mechanism. Coase thus introduced a concept that was destined to revolutionize the economics of organizations: that of "transaction costs," for example, the set of costs associated with searching for information, negotiating, contracting, and monitoring economic relationships.

Coase's idea remained underdeveloped for decades, until Oliver Williamson, starting in the 1970s, transformed it into a systematic theoretical framework, giving rise to a veritable line of research that would profoundly influence the economics of organizations.

In seminal works such as the Economic Institutions of Capitalism (1985), Williamson developed a comparative theory of forms of economic governance based on the analysis of transaction costs. The unit of analysis becomes the transaction, which Williamson defines as the act that occurs when a good or service is transferred across a technologically separable interface: "a transaction occurs when a good or service is transferred across a technologically separable interface"(Williamson,1985).

To fully understand the origin of these costs, it is essential to consider the behavioural assumptions of TCE. Williamson and subsequent literature identify four constitutive conditions: bounded rationality, opportunism, asset specificity, and environmental uncertainty.

Bounded rationality draws on the teachings of Herbert Simon, according to whom individuals are intentionally rational but limited in their cognitive abilities (Simon, 1957). In contracts, this translates into the impossibility of predicting all future conditions, making contracts inevitably incomplete: "all complex contracts are unavoidably incomplete" (Williamson, 2002). Williamson emphasizes that this incompleteness is one of the main sources of vulnerability in exchange relationships (Williamson, 1985).

The second assumption, opportunism, refers to the tendency of economic actors to pursue their own interests through cunning, misinformation, or opaque behaviour. As Williamson states,

"opportunism is self-interest seeking with guile". The consequence is that, without surveillance and sanction mechanisms, transactions risk degenerating as a result the importance of building contracts and governance systems capable of reducing incentives for opportunistic behaviour (Poppo & Zenger, 2002).

The third hypothesis concerns asset specificity, i.e., the degree to which investments and resources are dedicated to a transaction and are difficult to reuse in other contexts without loss of value. Klein, Crawford, and Alchian (1978) showed how high specificity exposes the parties to the so-called hold-up problem, i.e., the risk that, once the investment has been made, the other party will attempt to renegotiate the terms to its own advantage. Williamson (1985) identifies different forms of specificity: physical, human, temporal, and geographical. All of these increase switching costs and reinforce mutual dependence, making it more likely that hierarchical or relational forms of governance will be used (Nickerson & Silverman, 2003).

The fourth condition is environmental uncertainty. This term refers to the degree of unpredictability of external conditions that affect the performance of a transaction. It does not depend on the behavior of the actors involved, but on external and dynamic factors beyond their control.

Uncertainty can manifest itself in various forms as market uncertainty, when supply and demand vary unpredictably, exposing companies to the risk of sudden fluctuations in raw material prices or changes in consumer tastes; technological, when rapid innovation renders previous investments or contracts obsolete, as is the case in digital sectors where a technology signed today may be outdated in a few years; or institutional and regulatory, when new laws, tax regulations, or trade policies radically change the conditions of reference.

These four factors limited rationality, opportunism, asset specificity, and environmental uncertainty form the basis of ECT and determine the contractual risk associated with a transaction; it follows that the central problem of economic organization becomes the alignment between transaction attributes and the most efficient forms of governance to minimize overall costs: the market is efficient in standardized and non-specific transactions; hierarchy (integrated companies) is preferable when opportunism and specificity are high; and hybrid solutions, such as joint ventures or relational contracts, are suitable in intermediate cases (Williamson, 1991; Ménard, 2004).

The idea of transaction costs was later broadened by Douglass North, who described it as "the costs of operating the economic system" (North, 1990). This definition highlights how these costs are everywhere and influence the overall operation of markets and institutions. Over the years, Transaction Cost Economics (TCE) has found its way into various areas, including industrial economics, business law, corporate strategy, and even digital markets and online platforms (Rindfleisch & Heide, 1997).

In short, ECT assumes that the economic system does not automatically function efficiently thanks to the free play of the market, but that there are invisible and often underestimated costs associated with the management of economic relations. (De Filippi & Potts, 2018). For this very reason, the theory has proved to be an extraordinarily influential interpretative framework, capable of explaining phenomena such as vertical integration, outsourcing, franchising and, more recently, new forms of coordination linked to the platform economy and Web3.

Fig. 1. As shown below, the canonical Transaction Cost Economics framework depicts how firms aim to minimize transaction costs shaped by bounded rationality, opportunism, environmental uncertainty, and asset specificity.

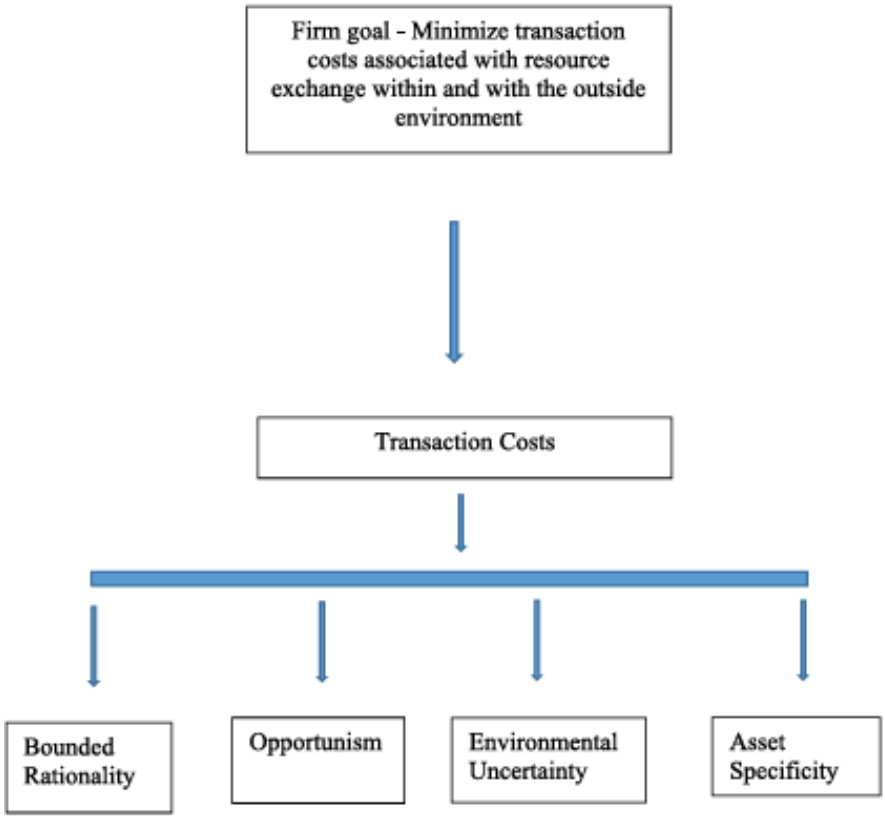

Fig. 1 Transaction Cost Economics framework (Source: Ahluwalia, S., Mahto, R. V., & Guerrero, M., 2020)

2.2 The Three Categories of Transaction Costs

Within the framework of Transaction Cost Economics, transaction costs do not constitute a uniform block, but are divided into three analytical moments that mark the entire cycle of exchange: an ex ante phase, linked to the preparation of the agreement (information search, partner selection, negotiation, and contract drafting); an ex post phase, connected to the implementation and monitoring of the contract (monitoring, compliance, enforcement, and dispute management); and adaptation costs, which arise when, following economic, technological, or institutional shocks, the agreement must be modified or renegotiated.

Looking at this from a different angle, research has indicated that having more precise upfront investment information and safeguards can help cut down on some of the costs that come later, like monitoring and litigation. However, this approach might lead to higher adaptation costs if the environment shifts. On the flip side, simpler clauses can save money in the short term but push those adaptation costs into the future as a matter of facts practical studies reveal that the ideal balance is not static, it varies based on how specific the investments are, the level of uncertainty involved, and how well the parties can renegotiate without escalating tensions.

Evidence from long-term contracts (energy, procurement, supplies) shows how regulatory or price shocks can shift the center of gravity of costs from ex post control activities to renegotiation, while more recent research, including in digital contexts, indicates that automated enforcement mechanisms can reduce certain ex post costs but make adaptation more rigid. In summary, Williamson's tripartite division not only organizes the concept of transaction costs but also guides a comparative assessment of governance choices capable of bringing together prevention, oversight, and contractual resilience.

2.2.1 Ex ante costs (search, negotiation, contracting)

In the Transaction Cost Economics world, ex ante costs refer to the expenses that pile up before an agreement is finalized. These costs are all about getting ready for the exchange relationship: gathering and checking information, choosing the right partners, setting up safeguards, and drafting the necessary clauses. According to Williamson's classic framework, this category specifically covers "the costs of drafting, negotiating, and safeguarding an agreement," which helps define the boundaries of how we think about the preventive side of contractual governance (Williamson, 1985). Essentially, this preparatory phase helps manage some of the expected uncertainty and potential opportunism by creating protective measures and coordination strategies before the actual exchange takes place.

It is important to evidence that the intensity of these costs increases with the complexity of the object being exchanged, the asymmetry of information, and the specificity of investments, requiring more due diligence, technical expertise, and legal advice. Reviews by Shelanski and Klein to Macher and Richman place ex ante costs at the center of the link between transaction attributes and governance choices, emphasizing that initial (ex ante) design can reduce some of the monitoring and litigation costs (ex post), but often at the price of less future flexibility (Macher & Richman, 2008).

As Macher and Richman, observe about this preparatory phase, "the ex-ante phase is crucial because it is the time when the parties attempt to reduce contractual uncertainty through forecasts, clauses, and safeguards."

Sectoral evidence confirms this logic.In the agri-food supply chain, Hobbs shows that mapping transactional risks guides ex ante contract writing, allocation of responsibilities, quality standards, audits, and prevents opportunistic drift and coordination problems (Hobbs, 1996). Complementarily, the summary by Rindfleisch and Heide in the Journal of Marketing highlights that ex ante costs include not only research and negotiation, but also the provision of specific safeguards (hostages, penalties, guarantees) as instruments of credible commitment in market relations (Rindfleisch & Heide, 1997).

A great example to consider is public procurement, especially when it comes to handling complex projects. Here, the initial design of the contractual framework really needs to factor in the chances of future renegotiations. The model created by Bajari and Tadelis makes it clear that there's a balancing act between offering incentives upfront and preventing costly transaction issues later on due to renegotiation. As they put it, "the intuition for our central result stems from a tension between providing ex ante incentives and avoiding ex post transaction costs due to costly renegotiation." (Bajari & Tadelis, 2001). This leads to the well-supported idea that cost-plus contracts are often better than fixed-price contracts when dealing with complex projects. Why? Because they allow for more flexibility after the fact, shifting resources from the initial stages to adapt as needed later on.

In a nutshell, ex ante costs are all about the initial investment that parties make to turn a naturally uncertain and imperfectly informed situation into a clear contractual agreement. According to transaction economics, this upfront investment doesn't completely wipe out uncertainty; instead, it shifts when that uncertainty hits. By spending more at the beginning on things like information, contract clauses, and protective measures you might actually save money down the line on monitoring, enforcement, and legal battles, as long as the agreement isn't too rigid for future adjustments. This idea is at the heart of the balance between prevention and operational oversight in transaction cost economics, which we'll dive deeper into in the upcoming sections.

Table 2 synthesizes the ex-ante component of transaction costs, search, negotiation, and contracting, by translating Williamson's taxonomy into concrete managerial activities and examples.

Table 2 — Ex-ante transaction costs in Transaction Cost Economics

Category	Description	Concrete examples
Search costs	Resources spent to identify exchange partners and opportunities	Market analysis; supplier screening; collecting price/quality information; preliminary due diligence
Negotiation costs	Time and expenses to define the terms of exchange	Meetings and iterations; legal/technical consultations; drafting term sheets; coordination time
Contracting costs	Costs to formalize and safeguard the agreement	Contract drafting; compliance and regulatory checks; review and final approval

2.2.2 Ex post costs (monitoring, enforcement, conflict resolution)

Ex post costs refer to all costs incurred after the contract is signed to ensure that the exchange takes place "as agreed": measuring and verifying performance, monitoring quality, correcting deviations, handling complaints and disputes, applying penalties, or renegotiating terms when necessary.

The root of these costs lies in the inevitable incompleteness of complex contracts: as Williamson points out, "all complex contracts are unavoidably incomplete," so during execution "the parties will be called upon to adapt" to gaps, errors, omissions, and unforeseen shocks (Williamson, 2002). It follows that ex post costs "arise when contract execution is misaligned" and require continuous attention: monitoring, adjustments, enforcement, renegotiations, up to mediation/arbitration or, in extreme cases, litigation.

A first family of ex post costs relates to performance measurement and verification. In many transactions, this means inspections, audits, testing, sampling, reporting, and benchmarking on contractual KPIs. The literature on measurement costs has shown that the more difficult it is to observe or standardize the qualities of a good or service, the greater the expense of measuring and certifying them, with direct effects on ex post costs (Barzel, 1982).

In TCE terminology applied to marketing, these costs are explicitly included in the "monitoring and enforcing agreements" that follow the contract (Rindfleisch & Heide, 1997). Concrete examples include quality controls along supply chains (e.g., agri-food or components) and compliance checks in outsourced services, where the buyer incurs costs to track SLAs, resolve discrepancies, and prepare corrective measures.

There’s another aspect to consider: the costs associated with keeping an eye on and preventing opportunistic behaviour that can pop up after an agreement is made. This includes things like shirking, strategic delays, manipulating information, or making unreasonable requests for renegotiation. When it comes to relationships involving highly specific assets, these costs often rise because the balance of bargaining power shifts once investments are on the table. This shift

can lead to hold-ups and necessitate more intense protective measures (Crocker & Masten, 1991). In these situations, just the act of "checking comparable external prices, ensuring they're relevant, and managing adjustment clauses" can create extra information and administrative headaches. If negotiations happen on the fly, the costs of haggling like managerial time, consulting fees, and potential alternative dispute resolution processes can really add up as everyone tries to get back on track (Crocker & Masten, 1991).

A third ex post component is the cost of managing non-compliance and disputes: from the application of penalties to conflict resolution, through mediation or arbitration. TCE reviews explicitly document these items as part of post-signing costs, distinguishing them from ex ante research/negotiation costs (Rindfleisch & Heide, 1997). In accounting terms, many of these expenses are included in "overhead" items (legal, audit, compliance), but analytically they represent ex post transaction costs related to contract execution and control.

Numerous empirical studies have measured these costs in specific contexts. In IT and business process outsourcing, for example, contractual complexity (control clauses, incentives, penalties, flexibility) is correlated with the levels of ex post costs incurred to measure quality, correct output, enforce terms, and resolve conflicts (Barthélemy & Quélin, 2006).

In analyses of samples of European contracts, the authors show how the operationalization of ex post transaction costs includes managerial time, resources dedicated to monitoring, renegotiation costs, and losses from relational malfunctions.

Even in the world of digital markets and platforms, there are certain types of post-transaction costs that come into play. These include expenses for fraud prevention, tackling fake reviews, maintaining trust and safety systems, and ensuring compliance with policies. Typically, these costs arise after a transaction agreement is made, such as moderation fees, chargebacks, and the costs associated with anti-fraud systems. While automation through smart contracts can help lower some of these enforcement costs, as discussed in Transaction Cost Economics (TCE), it can also lead to new post-transaction expenses. These might include costs for security audits, fixing bugs, managing oracles, and addressing situations that the code doesn't cover, which can make corrections during implementation quite pricey (Vatiero,2022). In short, we see a shift: less spending on manual compliance but more on ensuring technical compliance.

It is also crucial to note that measurement, surveillance, enforcement, and dispute handling are interdependent: better investments in metrics and reporting reduce disputes and penalties; conversely, high specificity and behavioural uncertainty amplify the need for inspections, audits, and renegotiations. It is precisely this interdependence throughout the exchange lifecycle that explains why, in TCE, ex-post cost analysis cannot be isolated from what was (or was not) predicted ex ante.

Table 3 summarizes the ex-post component of transaction costs, showing how contractual incompleteness reallocates resources to ongoing measurement, deterrence, and renegotiation during execution.

Table 3 — Ex-post transaction costs in Transaction Cost Economics

Category	Description	Concrete examples
Measurement & verification	Costs to monitor and certify contractual performance	Quality inspections; audits and testing; KPI reporting; supplier site visits; supply-chain checks
Surveillance and deterrence of opportunism	Efforts to prevent/contain post-contract shirking or hold-up	Tracking delays; validating external reference prices; applying adjustment clauses; managerial haggling
Non-conformity & dispute handling	Expenses to correct deviations and resolve conflicts	Claims management; penalties; mediation/arbitration; legal fees; compliance overhead
Emerging digital enforcement	New ex-post costs in digital exchanges and automated execution	Content moderation; fraud prevention/chargebacks; security audits; bug fixing; oracle management

2.2.3 Adaptation costs (technological shocks, institutional changes)

By adaptation costs, we mean the expenses that arise after signing when the world deviates from what was imagined ex ante and the agreement must be modified, recalibrated, or renegotiated. The theoretical basis for this lies in the inevitable incompleteness of complex contracts ("all complex contracts are unavoidably incomplete" (Williamson, 2002), which makes it natural, during execution, to identify contingencies and realign operations and documentation

Adaptation costs typically include: diagnosis of the deviation (measuring how prices, standards, demand, or rules have changed compared to contractual forecasts); impact analysis (assessing which clauses and specifications are inconsistent); renegotiation (prices, quantities, technical specifications, timelines); legal reformulation (addenda, updating of SLAs and guarantees); and operational realignment (testing, qualifications, reconfiguration of processes and data, staff training), to which are added opportunity costs from delays and downtime. In classical taxonomy, these burdens manifest themselves as maladaptation costs (when execution remains "out of alignment") and haggling costs (the time/resources spent in corrective negotiations), items that TCE literature identifies as among the main ex post requirements once deviations from the ex-ante agreed scenario have emerged.

Measuring contingency is the crucial first step in the costing process. When price fluctuations or regulatory changes hit, we need to figure out just how significant they are and convert that into numbers we can act on like indicators, indexation bases, and thresholds. Research on long-term contracts highlights that having shared metrics (or not) plays a big role in how often and how costly revisions are, which directly impacts the 'cost' of adapting (Joskow, 1987; Goldberg

& Erickson, 1987). These studies reveal that when specific investments are involved, parties tend to commit to longer contract durations and avoid frequent negotiations at the outset. This strategy is all about saving on measurement and renegotiation costs later. However, when shocks do happen, it's unavoidable and often expensive to rework the numbers and contracts. The next step is the actual renegotiation process: this involves gathering for technical and legal meetings, swapping proposals, and ensuring everything aligns with existing rules and standards. This is where the costs of back-and-forth discussions (like managerial time, consulting fees, and ongoing interactions) and drafting new terms (such as adjustment clauses, penalties, or guarantees) come into play. Plus, there are the costs associated with finding reliable external benchmarks to back up the new terms (Crocker & Masten, 1991).

Even with adjustment clauses, filters, and thresholds in place, real-world experience shows that reaching a negotiation "landing" often requires verification and testing sessions, which can lead to significant administrative and organizational expenses.

The third step is operational realignment: integrating changes into processes, systems, and data. In complex management systems (ERP/SCM), post-implementation literature reminds us that a significant portion of the TCE in the life cycle stems precisely from adaptive activities: updates, reconfigurations, regression testing, data migrations, training, and support. in some estimates, post-implementation costs (maintenance and adaptations) can reach very significant proportions of the total cost in the long term (Law & Chen, 2010).

When we break it down technically, these activities are categorized as adaptive and perfective maintenance, alongside corrective and preventive maintenance, as highlighted in studies on ERP maintenance (Nah, Faja & Cata, 2001). They essentially represent the "operational arm" of the adaptation costs mentioned in Transaction Cost Economics (TCE). This means that turning a renegotiation into real configurations, data, and procedures takes time, resources, and a fair bit of coordination.

Table 4 organizes adaptation-related transaction costs incurred when reality diverges from contractual assumptions into diagnostic, renegotiation, legal, operational, and opportunity-loss components, highlighting how execution requires responsive change.

Table 4 — Adaptation costs in Transaction Cost Economics

Phase/Item	Description	Concrete examples
Contingency assessment	Detecting the gap between forecast and reality (price, regulatory, tech shocks)	Analyzing price movements; tracking new regulatory standards; applying indexation indicators
Impact analysis	Evaluating clauses/specifications now misaligned with the new scenario	Identifying breached SLAs; mapping newly applicable legal/compliance obligations
Contract renegotiation	Negotiation and legal work to realign exchange terms	Technical/legal meetings; expert consultancies; haggling; revising prices/quantities
Legal reformulation	Drafting addenda and updating contracts and guarantees	Contract addenda; adjustment/price-escalation clauses; revised penalties
Operational realignment	Translating changes into processes, systems, and procedures	ERP/SCM updates; regression testing; data migration; staff training
Opportunity costs	Losses due to delays, downtime, and inefficiencies during adaptation	Service interruptions; foregone production; managerial time diverted

The diagram below (Fig.2) clearly lays out the key subcomponents of transaction costs, breaking them down into ex ante, ex post, and adaptation costs. This structured approach offers a thorough framework for understanding how these different elements function in economic exchanges.

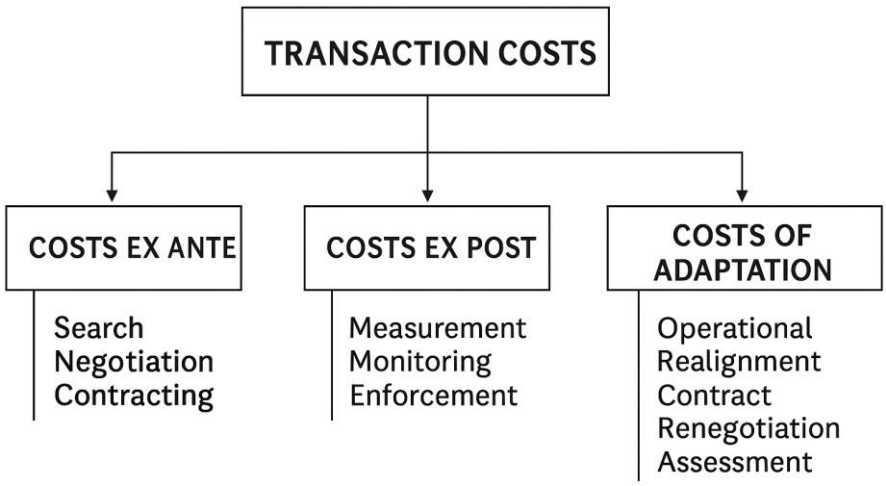

Fig. 2 Transaction Cost Economics framework

Chapter 3 – Web3 technologies and their impact on transaction costs

3.1 Introduction to the chapter and TCE–Web3 logic

Transaction Cost Economics (TCE) was conceived to explain the boundaries of the firm and governance choices in a context where transactions are inherently costly because they are influenced by factors such as limited rationality, opportunism, and asset specificity, highlighting in particular how the difficulty of predicting all contingencies makes problems of adaptation, monitoring, and enforcement inevitable.

In the original context of TCE, companies are faced with a choice between three modes of governance: market, hybrid, and hierarchy, based on each one's ability to minimize transaction costs. The underlying assumption is that coordination mechanisms are not neutral but produce different costs and benefits, which must be evaluated in light of the characteristics of the transaction.

With the rise of the digital economy, Transaction Cost Economics (TCE) has gained fresh traction as a valuable tool for analyzing how information technology influences transaction costs. Back in the 1980s, Malone, Yates, and Benjamin (1987) demonstrated that the emergence of electronic communication systems significantly reduced coordination costs. In a similar vein, Clemons, Reddi, and Row (1993) pointed out that information technologies enabled more efficient management of complex contracts, which in turn lowered uncertainty and minimized opportunistic risks. Since then, applying TCE to digital technology has become well-established, helping to explain trends like online platforms, IT outsourcing, and digital supply chains (Grover & Malhotra, 2003).

Web3 is a part of this movement, characterized by a collection of decentralized technologies such as blockchain, smart contracts, decentralized autonomous organizations (DAOs), and identity and reputation systems. These innovations are designed to create new ways of coordinating efforts, ultimately aiming to lower the costs that TCE has highlighted as crucial. As Catalini and Gans (2020) note, while blockchain doesn’t eliminate transaction costs, it does change how they are structured, particularly affecting verification, networking, and enforcement costs. In this way, Web3 serves as a "natural laboratory" for exploring the explanatory power of TCE in a landscape of radical innovation.

The TCE–Web3 logic, which this chapter aims to develop, starts from a fundamental question: to what extent do decentralized technologies reduce the transaction costs identified by Williamson and in which cases, on the contrary, do they generate new ones? The perspective is not simply to compare old and new forms of coordination, but to develop an extension of the TCE framework that takes into account the specificities introduced by decentralized infrastructures. Smart contracts, for example, can reduce enforcement costs but increase the costs of adapting to unforeseen events, precisely because their rigidity reduces the scope for legal interpretation and renegotiation (Vatiero, 2022). DAOs, on the other hand, represent an attempt to reduce agency costs, but pose new challenges in terms of decision-making efficiency and internal conflict management (DuPont, 2020).

TCE is not superseded, but reinterpreted so the traditional categories of ex ante, ex post, and adaptation remain valid, but must be reinterpreted in light of new forms of decentralized governance.

This chapter sets out to explore how Web3 technologies influence different aspects of transaction costs. It will identify the areas where we see a notable decrease in these costs and delve into the new expenses that don’t quite fit into Williamson's traditional framework. Ultimately, the aim is to suggest an extension of Transaction Cost Economics (TCE) that incorporates the dynamics of Web3, paving the way for a refreshed theoretical framework to better understand today’s digital economy.

3.2 Enabling Technologies and Transaction Cost Reduction

At the heart of the Web3 paradigm lies the introduction of a set of enabling technologies: blockchain, smart contracts, and decentralized autonomous organizations (DAOs) that aim to transform the structure of transaction costs. Transaction Cost Economics (TCE), as shown by Williamson, highlights how search, bargaining, monitoring, and enforcement constitute physiological barriers to exchange. Web3, through its technical innovations, offers an institutional and technological response capable of redesigning these mechanisms.

Paragraphs 3.2.1–3.2.3 will analyse three technological vectors ordered "from the bottom up," from the infrastructure layer to the organizational layer:

(i) blockchain, as a distributed ledger and consensus system.
(ii) smart contracts, as deterministic executive logic that automates clauses and payments.
(iii) DAOs, as on-chain governance and decision-making architectures. This sequence reflects a hierarchical dependency: smart contracts presuppose a shared ledger and a consensus mechanism, while DAOs combine contracts and voting/treasury primitives to implement collective decision-making and coordination rules (Weking 2020; Schmidt & Wagner, 2019).

For each technology, we'll use a consistent analytical framework to assess the technical-economic relationship based on Transaction Cost Economics (TCE) principles. First, we'll define the specific technology and explain how it operates, highlighting the essential components needed for it to function. Next, we'll connect these characteristics to their effects on transaction costs, categorizing them into the three types identified by TCE: ex ante, ex post, and adaptation.

In a nutshell, blockchain, smart contracts, and DAOs contribute, with varying intensity and in different ways, to redesigning the cost categories identified by Transaction Cost Economics.

Blockchain mainly affects ex ante costs, simplifying research and negotiation thanks to the transparency and reliability of the distributed ledger.

Smart contracts significantly reduce ex post costs by automating measurement and enforcement, but important trade-offs emerge during the adaptation phase: the deterministic 'if-this-then-that' logic tends to stiffen efficient realignments in the face of unexpected shocks, reviving the classic issue of incomplete contracts.

DAOs mainly affect ex post governance and coordination costs, reducing monitoring and verification expenses, but at the same time they can amplify adaptation costs: the need for collective on-chain decisions entails risks of slowness, forks, coordination problems, and possible capture phenomena.

3.2.1 Impact of Blockchain within the TCE Framework

In technical and managerial language, blockchain is described as a distributed digital ledger which, in most implementations, is immutable and stores a verified sequence of transactions over time without resorting to a trusted central authority. In other words, it is a database shared among multiple participants in which anyone can propose the addition of data, but no individual can retroactively modify what has been validated and written to the chain.

Validation takes place through consensus rules and cryptographic signatures that guarantee integrity and synchronization between the nodes of the network (Christidis & Devetsikiotis, 2016). This vision of a "distributed, consensus-based and (mostly) immutable ledger of transaction records" is now standard in operations and supply chain literature.

Blockchain is part of a larger group known as Distributed Ledger Technologies (DLT). This term refers to a type of database architecture that enables records to be stored and shared in a decentralized way, ensuring their integrity through consensus protocols and cryptographic signatures. What sets blockchain apart from other DLTs, like Hashgraph or Directed Acyclic Graphs (DAG), is its unique process for handling new transactions. These transactions are spread across the network, bundled into blocks, validated collectively, and then linked together in chronological order using hash functions, forming a genuine “chain of blocks” (Catalini & Gans, 2016).

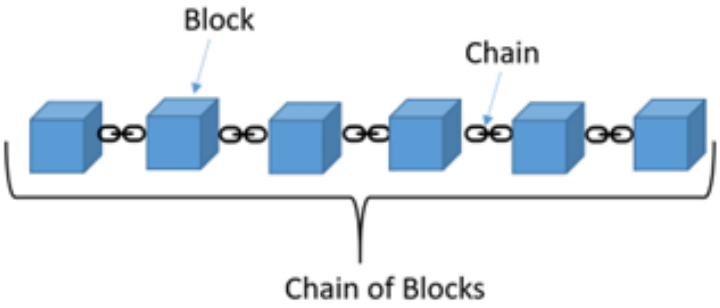

Fig. 3: Blockchain architecture (Source: Medium)

From a practical standpoint, each block in the chain includes (i) a collection of transactions, (ii) a timestamp, (iii) the hash of the previous block (often referred to as the parent), and (iv) metadata that helps with hash verification (like the nonce in proof-of-work systems, which is essentially a random number used for verification). The hash acts as the key for the block, while the transaction details are kept secure through encryption. Because of how hashes are combined, even the slightest change in a block will alter its hash value, creating a ripple effect on all the following blocks (Nofer et al., 2017). This "hash-linked" design ensures that the ledger remains practically immutable and allows us to trace the entire history back to the very first block, known as the genesis block.

The entire ledger is replicated among participants: each node maintains a synchronized copy and can locally verify the validity of new information before it is added to the chain. This peer-

to-peer replication, combined with distributed consensus, replaces the "personal trust" typical of bilateral relationships with a form of "system trust" anchored to the rules of the protocol.

The nature of a "shared ledger" clarifies the difference from traditional networks, in which each party maintains its own records or entrusts trust to an intermediary: a blockchain provides a constantly updated common ledger (Gupta, 2017).

In this configuration, all counterparties work, directly or through validators, on the same source of truth (shared ledger). This conceptual shift is well illustrated in Figure 3 by the comparison between "traditional networks" and "networks using blockchain technology."

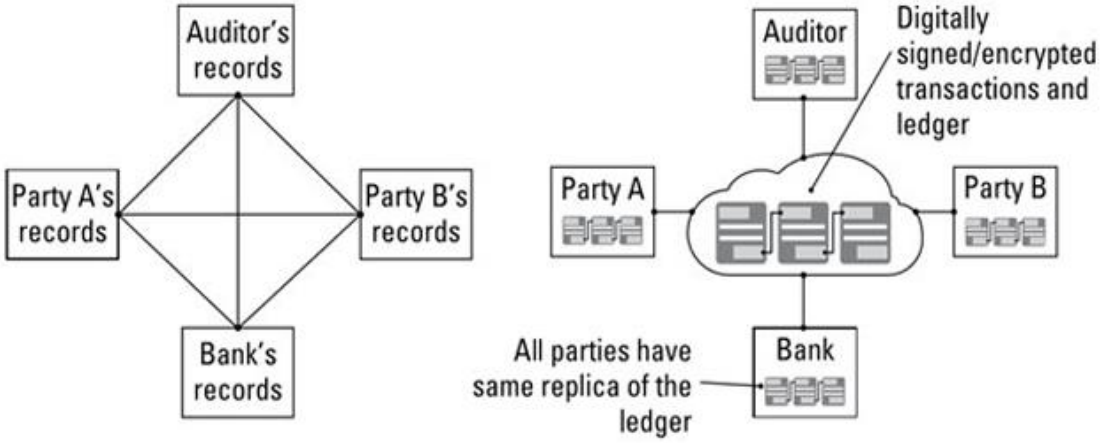

Fig. 4: Traditional networks vs. networks using blockchain technology. Source: Gupta (2017)

Blockchain configurations vary in two key institutional dimensions. The first concerns who can participate in the network: permissionless (no prior authorization, as in Nakamoto's original design) versus permissioned (access and roles managed by an entity).

The latter concerns who can see what: public ledger (data visible to anyone) versus private ledger (controlled access and visibility). Architectural choices along these dimensions affect performance, security, governance, and economic impacts.

Consensus, i.e., the mechanism by which the network agrees on the "true" state of the ledger, is the other fundamental element. In permissionless chains, the most studied schemes are Proof-of-Work (PoW) and Proof-of-Stake (PoS): the former links validation to the solution of energy-intensive computational puzzles, the latter to the staking of economic stakes by validators, reducing computational expenditure in favor of on-chain incentives/penalties. In both cases, consensus manages synchronization between nodes and transaction ordering, ensuring that every copy of the ledger converges on the same state.

On various platforms, smart contracts introduce a layer of general-purpose computation to the transaction ledger. Traditionally, they're defined as "a computerized transaction protocol that executes the terms of a contract" (Szabo, 1997). In simpler terms, these are on-chain programs that automatically run when certain conditions are met. In account-based platforms like Ethereum, contracts function as bytecode-controlled accounts: users send transactions with

specific execution parameters, and the Ethereum Virtual Machine follows deterministic rules that influence the state of the shared ledger (Wood, 2014). This programmability opens the door to automation and interaction between applications think conditional payments, escrow, oracles, and tokenization, greatly broadening the range of use cases beyond just "simple" value transfers (Christidis & Devetsikiotis, 2016). While we won't dive deep into smart contracts here, as they'll be explored in the next chapter, it's important to highlight their role in transforming the blockchain from merely a "ledger" into a fully programmable transactional platform.The intrinsic functioning of blockchain offers a number of benefits to network participants, which have been repeatedly highlighted in the literature. Abeyratne and Monfared (2016) identify four main advantages. First, durability: the decentralized structure eliminates the risk of a "single point of failure," distributing risks across all nodes and making blockchain more resilient to external attacks and data loss. Second, transparency: each participant has a synchronized copy of the ledger, with real-time visibility of transactions. Third, immutability: validated transactions become permanent and traceable thanks to the cryptographic link between blocks (Nofer et al., 2017). Fourth, process integrity: the rules defined in the protocols are automatically executed by the code, reducing the possibility of manipulation.In addition to these aspects, there is also a reduction in dependence on intermediaries. Instead of relying on separate registers and a third party to certify their validity, the parties share a single common ledger. This eliminates duplication and reduces the overall time and cost of transactions (Gupta, 2017; Nofer et al., 2017).

In the context of Transaction Cost Economics, transaction costs are broken down in operational and measurable terms across four key procurement activities: supplier selection (search), negotiation and contract definition (contracting), performance monitoring (monitoring), and dispute resolution (enforcement). The taxonomy we will use is derived from Dyer's empirical analysis, which links these activities to corresponding "types" of costs and proposes proxy indicators for comparative measurement between companies (Dyer, 1997).

The theoretical question is therefore whether and how blockchain, as a shared and programmable recording infrastructure, reduces ex ante information and negotiation costs and, ex post, control and enforcement costs. Catalini and Gans emphasize that the technology mainly affects "two key costs... the cost of verification of state, and the cost of networking," i.e., the cost of verifying the state (authenticity, ownership, outcome) and the cost of network coordination. This perspective perfectly complements Dyer's breakdown and allows us to see where cost reductions are generated along the supply cycle.

Search/selection costs: when relevant supplier data (certifications, quality, delivery history) is recorded on-chain, the reliability of the information increases. This reduces the need for scouting and qualification activities, as trust shifts from unilateral declaration to shared and validated proof. Procurement studies show that on-chain storage of relevant data "makes the data more trustworthy," with a positive effect on selection and, by inference, a reduction in search costs.

Fig. 5: Blockchain effect on search costs

Monitoring costs: on-chain rules and integration with external data sources (IoT, ERP, quality systems) enable the automation of compliance checks, tracking, and exception management. In TCE terms, this lowers performance measurement costs in the presence of behavioral uncertainty. However, efficiency depends on data reliability: if the information pipeline is not properly governed, the "oracle problem" emerges, i.e., the risk that unreliable data will be permanently recorded. (This part will be explored in more detail in the chapter on smart contracts).

Enforcement costs: ensuring data integrity, traceability, and clear execution rules can significantly lower the chances of disputes. And when disagreements do pop up, having a shared and indisputable log helps resolve them more swiftly. Research shows that maintaining "process and information integrity... can prevent many conflicts or help resolve them quickly," which ultimately leads to lower policing and enforcement costs. In the context of Transaction Cost Economics (TCE), blockchain technology minimizes the risk of opportunistic behavior after the fact by clarifying uncertainties around deliveries, quality, timing, and payments.

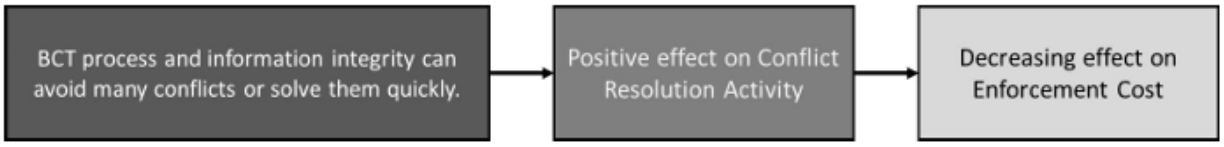

Fig. 6: Blockchain effect on enforcement costs.

The sum of these effects on the three main buckets (search, monitoring, enforcement) leads to the conclusion that blockchain reduces overall transaction costs. This result is consistent with economic theory, which attributes cost reduction to the simultaneous decrease in verification and networking costs in digital markets (Catalini & Gans, 2016).

Catalini and Gans clearly explain that the heart of innovation lies in the possibility of "cheaply verifying the state, including information about past transactions and their attributes, and current ownership in a native digital asset."

This has a direct impact on how digital markets are designed and how competition plays out between different platforms. In terms of Transaction Cost Economics (TCE), the drop in verification costs influences both the pre-transaction phase (like information screening, due

diligence, and adverse selection) and the post-transaction phase (including reconciliation, auditing, performance measurement, and enforcement). The authors point out that the initial effects are seen "on the intensive margin of production," as businesses and organizations start moving their processes to shared ledgers to "cut down on settlement and reconciliation costs." Essentially, a lot of the costs that are currently tied up in redundant controls, database alignments, and accounting checks can be streamlined into a single source of shared truth. The most significant improvement is seen with natively digital assets think tokens, usage rights, and computational resources on the network where "the cost of verifying transaction attributes and enforcing simple contracts for self-contained tokens can be extremely low" (Catalini & Gans, 2016).

This creates a twofold effect: on one side, the costs associated with measuring and enforcing incomplete contracts are lowered; on the other, the fragmentation of property rights (like micropayments and conditional payments) becomes economically viable, thanks to dependable and affordable verification.

However, the authors point out that when the exchange involves the digital representation of off-chain goods/offers, "costless verification" is difficult to achieve. In such cases, blockchain reduces costs only if investments are made in organizational and technological complements (data entry standards, IoT devices, KYC/AML procedures). They clearly emphasize this: "when entries on a shared ledger are digital representations of offline identities, products, services... costless verification is difficult to achieve" (Catalini & Gans, 2016).

When these frictions are contained, the dynamic changes radically: "decentralized verification goes from being costly, scarce, and prone to abuse, to being cheap and reliable" (Catalini & Gans, 2016). In concrete terms, "Expensive audits and due diligence can be progressively substituted with more frequent and fine-grained verification."

This combination of native auditability and continuous low-cost verification realigns incentives, reduces information asymmetries, and compresses both measurement and enforcement costs.

Viewed through the lens of Transaction Cost Economics, blockchain operates primarily as a "general-purpose verification technology": it streamlines the verification of shared status and the attribution of rights, reducing information asymmetries and redundant reconciliations throughout the exchange cycle ("we identify two key costs affected by the technology: the cost of verification and the cost of networking") (Catalini & Gans, 2016).

At execution time, the auditable ledger shifts compliance verification from discretionary controls to traceable and (partially) automated checks: in this sense, the technology "offers a distinct way of enforcing agreements and achieving cooperation and coordination compared to traditional contractualism and relational governance" (Lumineau, Wang, & Schilke, 2021), and on-chain audit trail mechanisms demonstrate how continuous monitoring becomes less costly and more reliable.

On the adaptation front, timely data sharing mitigates some of the environmental uncertainty and facilitates reallocations and contractual adjustments without eliminating the need for renegotiation, because efficiency remains conditioned by the institutional context and the specificity of the assets (Schmidt & Wagner, 2019; see also TCE's analysis of uncertainty as a driver of unexpected adaptations in Williamson).

To visually clarify the different impacts, below is a summary matrix showing which TCE cost families are most ad least affected by blockchain technology.

Table 5: Impact of Blockchain within the TCE Framework

Cost category	**Impact of blockchain**	**Why**
Ex ante (research, information, negotiation)	**High**	Shared ledger Drastic reduction in verification and networking costs; fewer reconciliations, streamlined due diligence, symmetrical information prior to signing.
Ex post (monitoring, enforcement)	Moderate (High when combined with smart contracts)	Immutable trace Low-cost monitoring and auditing; enforcement only becomes truly automatic when the rules are codified in smart contracts.
Adaptation (realignment to changes and shocks)	Moderate	Blockchain provides a common, up-to-date information base, which reduces coordination costs in renegotiations and reorganizations and lowers uncertainty in exchange relationships. However, the effect on adaptation remains moderate, because the technology simplifies information alignment but does not replace contract renegotiation.

3.2.2 Impact of Smart Contracts within the TCE Framework

The term smart contract originated in the 1990s, before the advent of blockchain. The reason why the two concepts are now closely associated is that blockchain has enabled a qualitative leap forward, increasing trust, reliability, and security through the ability to enter codes into the network without the risk of counterfeiting or modification (Avarello, 2021).

In general terms, smart contracts are, in Nick Szabo's classic formulation, "a computerized transaction protocol that executes the terms of a contract" (Szabo, 1997). In other words, they are programs stored on-chain that automatically execute agreed clauses when predetermined conditions are met, reducing reliance on intermediaries and increasing the predictability and enforcement of the agreement. The idea, developed between 1994 and 1997, was subsequently transposed into technical and informational terms by the Ethereum white paper, which introduced a Turing-complete computation layer (EVM) to distribute and execute contract code on a network of nodes, with deterministic execution and tracking on the ledger (Buterin, 2014).

Operationally, a smart contract is an on-chain application that exposes functions that can be invoked via transactions: given the same state and input, all nodes obtain the same output and update the ledger consistently. The logic is typically of the if/when... then... type: when the event or condition is verified, the contract transfers digital assets, records states, sends signals, or enables rights without the need for ex post authorizations.

The immutability of the blockchain ensures that completed transactions cannot be altered, thus ensuring reliable and transparent execution; at the same time, the persistence of the distributed ledger ensures the immutability of the code already implemented and the non-repudiability of executions. Together, these features enable a significant reduction in time, automation of settlement processes, and a reduction in administrative costs compared to traditional contracts (Fig.7).

Fig. 7: Comparison between smart contracts and standard contracts (Source: smart-contracts)

The life cycle of a smart contract develops through distinct but closely interconnected phases. First, there is creation, i.e., the translation of the agreement into computer code; this is followed by deployment, which consists of publishing the contract on the blockchain, thus making it accessible and executable. Next, in the execution phase, the functions are automatically activated when the predefined conditions are met, until completion, which involves updating the statuses and assets recorded in the distributed ledger. However, it is the design phase that is the most complex: the parties must precisely define the representation of on-chain data, conditional rules, applicable exceptions, and any dispute resolution mechanisms to ensure the correct implementation and reliability of the contract.

Ethereum has made this paradigm mainstream thanks to the Ethereum Virtual Machine: a "world machine" that executes compiled bytecode (e.g., from Solidity). Execution is "virtually unstoppable" unless the contract incorporates pause or interrupt functions ex ante; this makes self-execution as resilient to censorship as the underlying blockchain. In this setting, smart contracts are not limited to the transfer of cryptocurrencies: they can define rules, escrow, marketplaces, and governance modules, applying them automatically and verifiably.

Since blockchain natively "sees" only on-chain data, a crucial element is access to information from the outside world (off-chain), through oracles to import prices, event outcomes, logistics statuses, etc. into the chain.

Oracles are software or hardware gateways that provide the contract with inbound or outbound data. The former collect information from digital sources, the latter from physical sensors; incoming oracles insert data into the blockchain, while outgoing oracles transfer information from the contract to the outside world. In supply chains, for example, RFID tags and sensors can act as oracles to certify the departure/arrival of goods at designated locations and trigger parametric obligations.

Technological architectures reveal certain fundamental properties that clearly distinguish smart contracts from traditional contracts. The first characteristic is determinism, i.e., the need for distributed execution to produce the same result on all nodes with the same status and input. To ensure this consistency, smart contracts avoid using non-reproducible sources of randomness or external dependencies that are not anchored on-chain; when randomization is necessary, it is constructed using verifiable cryptographic techniques or delegated to oracles with appropriate precautions.

A second feature concerns immutability and traceability. The bytecode distributed on-chain and all contract invocations remain recorded in the distributed ledger, ensuring ex post auditability and resistance to unilateral manipulation. Although forms of updatability are possible through proxy patterns, they require explicit governance mechanisms, such as multiple administrators, in order to preserve user trust and avoid malicious updates.

This is linked to the properties of verifiability and transparency. Code and execution status can be inspected directly on the blockchain, except in cases of permissioned networks or privacy-preserving solutions. This openness reduces information asymmetries and allows any participant to check the exact fulfillment of the contract, with a level of predictability and non-repudiation that is difficult to achieve in traditional contracts (Avarello, 2021).

As noted in the literature, "anyone can verify on-chain the correct execution and ownership of transferred assets, within a predictable and irrevocable system of conditions" (Hoffmeister & Stossberger, 2018).

However, smart contracts operate in a limited context: they can only access their own internal state, transaction parameters, and block metadata. Anything external requires the intervention of oracles, with the consequent risk that incorrect or manipulated data will produce irreversible effects, such as unrecoverable money transfers (Clack, Bakshi, & Braine, 2016). Although this constraint guarantees determinism, it also represents a design limitation that requires the definition of "design for verifiability" architectures.

Smart contracts introduce a crucial property: the principle of code is law: the behavior of the parties is directly bound by the computer code. In other words, "if the condition 'if' is true, then the action 'then' follows automatically". This approach promises to reduce opportunism and interpretative ambiguity, but it also leads to rigidity in the face of unforeseen contingencies. As Vatiero has shown, "because of the need for adaptation to mutable and unpredictable occurrences, smart contracts may incur higher transaction costs than traditional contracts" (Vatiero, 2022).

The scope of application of smart contracts is also limited by the requirement of tokenizability : the assets being exchanged must be natively digital or representable by tokens. Where the contractual object is off-chain, it is necessary to introduce a reliable anchoring mechanism (oracles, IoT sensors, custodians, certifications) that reports the relevant information on-chain. In the absence of such a link, automatic execution cannot fully cover the agreement.

Finally, an essential aspect concerns security and technical costs. The reliability of a smart contract depends on the virtual machine that executes it, the language used, and the quality of the code. Known vulnerabilities such as re-entrancy, overflow, weak access control, or front-running have caused significant losses in the sector (Brina, 2022). In addition, each operation is subject to computational costs (gas/fees), both during deployment (bytecode size and storage initialization) and during execution (opcodes, accesses, logs). For this reason, tools such as formal verification, advanced testing, and peer review are now an integral part of development best practices.

All these properties can be summarized along three fundamental axes: transparency, because the code is publicly verifiable; automatism, because execution is self-enforcing and does not require human intervention; and computational binding, i.e., the idea that it is the code itself that conclusively determines the outcome of transactions (Arruñada, 2020). These features make smart contracts unique tools in the digital contracting landscape, with clear benefits in terms of speed, accuracy, reduced administrative costs, and greater trust between counterparties, especially in transnational contexts where traditional intermediaries would slow down processes (Schmidt & Wagner, 2019).

The impact of smart contracts on transaction cost structure, in line with TCE, can be understood by distinguishing between three distinct (but interdependent) stages: ex ante, ex post, and adaptation. On the ex ante side, programmability and self-execution reduce the costs of research, negotiation, and writing: some of the clauses become machine-interpretable (or even machine-executable), with less dependence on intermediaries for document verification and

procedural closing. As Werbach and Cornell observe, smart contracts "are self-executing digital transactions using decentralized cryptographic mechanisms for enforcement" and promise a "digital bypass" of traditional contract law mediation (Werbach & Cornell, 2017).

In this sense, the pooling of templates and modules (conditional payments, escrow, oracles) can reduce drafting costs and settlement times.

On the ex post front, the gain is even clearer when performance is verifiable on-chain: monitoring, compliance, and enforcement are incorporated into the code ("if/when... then...") and replicated by the network. Legal literature has highlighted how conclusive technological enforcement can drastically reduce disputes in standardizable areas (Arruñada, 2018).

In TCE terms, the risk of ex post opportunism is mitigated because the credible threat is no longer (only) the sanction of the third party, but the impossibility of deviating from what is codified and validated by the blockchain.

The picture changes when we move on to adaptation costs. TCE emphasizes that institutional efficiency depends on the ability to realign the relationship when the state of the world changes (price shocks, logistical disruptions, force majeure events).

This is precisely where the structural friction of smart contracts lies: "because of the need for adaptation to mutable and unpredictable occurrences... smart contracts may incur higher transaction costs than traditional contracts" (Vatiero, 2022). This result stems from two channels. First, the absence of legal gap-filling: the code avoids (or minimizes) ex post judicial intervention, including corrective interpretation and gap-filling, which, in "semantic" contracts, allows for the management of unforeseen contingencies at lower costs than "writing everything" (Shavell, 2009).

Second, the technical governance of adaptation: to regain flexibility, ex ante (proxy/diamond upgrade, emergency stop, redundant oracles) and ex post (multisig quorum, DAO voting) mechanisms are needed, which transfer costs from the external forum to on-chain design/management (Bank of Italy, 2023). Furthermore, the extreme "system" adaptation valve of network consensus up to the hard fork can generate majority maladaptation, split the expectations of the parties, and multiply coordination and migration costs.

To clarify, Vatiero (2022) distinguishes between four enforcement mechanisms: (i) external legal enforcement typical of "semantic contracts"; (ii) hierarchical enforcement within the company (Coase, 1937); (iii) relational enforcement based on reputation; and (iv) blockchain enforcement that characterizes smart contracts, where "code is law" and the code itself represents the only conclusive form of constraint. In this scheme, smart contracts reduce the risk of opportunism ex ante, but the very "conclusiveness" makes ex post adaptation extremely costly.

In traditional contracts, courts can fill gaps, interpret ambiguous clauses, or modify terms to avoid inefficiencies (Shavell, 2009), allowing parties to intentionally leave remote contingencies open and save writing costs. In the deterministic world of smart contracts, on the other hand, either the improbable is also coded ex ante, paying very high writing costs, or the risk of inefficient outcomes is accepted if that contingency occurs. The example of Vatiero is emblematic: a sudden exponential increase in production costs (as happened during the

COVID-19 pandemic) can render a fixed price codified in a smart contract inefficient; while a traditional contract could be adapted via legal gap-filling, the smart contract remains rigid and inefficient.

The second axis concerns the costs of "mal-adaptation." Pro-blockchain rhetoric emphasizes that consensus and forks make the system flexible and adaptable, but this is precisely where risks arise. A fork consists of a split of the blockchain into two parallel versions, due to a change in the consensus rules or disagreements between validators: in the case of a soft fork, the network remains compatible with the previous version, while a hard fork creates two divergent and incompatible chains (Nakamoto, 2008; Werbach, 2018). Historical examples include the Ethereum/Ethereum Classic fork in 2016 (following the hack of The DAO) and the Bitcoin forks in 2017 (Bitcoin Gold). These events demonstrate how the network adaptation mechanism can degenerate into institutional uncertainty: majority coalitions can change the rules to their advantage, penalizing minorities and increasing coordination and migration costs. In comparative terms, in a state governed by the rule of law, rule changes are subject to burdensome procedures (qualified majorities, constitutional guarantees), which sacrifice flexibility but guarantee stability. In blockchain, on the contrary, the threshold is lower: this increases adaptability but reduces legal certainty.

This results in a well-defined trade-off in terms of TCE. Where performance is standardizable, data is verifiable, and outcomes are tokenizable, smart contracts reduce ex ante and ex post costs. Where, on the other hand, the environment is uncertain and highly variable, adaptation costs can erode and sometimes exceed the savings achieved.

The comparative efficiency condition outlined by Vatiero is clear: self-executing code is advantageous when writing costs are low, legal costs are high, and the consensus infrastructure remains impartial; otherwise, the balance tips toward semantic contracts. A possible compromise, as Werbach observes, is hybridization: "blockchains may supplement, complement, or substitute for legal enforcement," building mixed structures where code ensures repetitive and verifiable execution, while the judge remains a fallback for unprogrammed contingencies.

With this in mind, Vatiero (2022) suggests measures to reduce the costs of maladaptation, such as raising the thresholds for forks or providing compensation to minorities penalized by majority decisions, similar to constitutional guarantees. The message is clear: blockchain is not a "free lunch."

Smart contracts are more efficient when they operate as a complement to existing institutional arrangements rather than as complete substitutes, and their efficiency depends on the level of environmental uncertainty and the quality of network governance.

To visually clarify the different impacts, below is a summary matrix showing which TCE cost families are most and least affected by smart contracts.

Table 6 – Effects of Smart Contracts on Transaction Costs (TCE)

Cost category	Impact of smart contracts	Why
Ex ante (research, information, negotiation)	Moderate to high	Codified clauses and public interfaces reduce ambiguity and reconciliations; regulation by technology sets exchange rules; code and states are verifiable before signing.
Ex post (monitoring, auditing, enforcement)	Strong reduction	Automatic execution of conditional payments and escrow; immutable logs and continuous on-chain auditing; transparent and deterministic "if-this-then-that" enforcement.
Adaptation (realignment to changes and shocks)	Risk of increase	Code rigidity raises writing costs for covering contingencies; adaptation requires proxy upgrades and on-chain governance, generating organizational and migration costs; risk of malicious/bad adaptation.

In summary, smart contracts are consistent with the logic of Transaction Cost Economics: on the one hand, they significantly reduce negotiation and enforcement costs thanks to the automation of performance and the transparency of distributed ledgers; on the other hand, they can increase adaptation costs in contexts characterized by uncertainty and changeability, where reality exceeds the scope of the code.

From this perspective, the operational recommendation is not to "codify everything," a choice that would lead to disproportionate writing costs and inefficient rigidity, but rather to allocate adaptation where it is least costly and most sustainable. This can be done ex ante, during the design phase, through parametric clauses or pause and upgrade mechanisms with adequate guarantees; on-chain, through formalized governance rules (deliberative quorums, change-management systems, formal verification procedures); or off-chain, leveraging traditional institutions such as arbitration or courts, where discretionary interpretation and gap-filling are more efficient (Arruñada, 2020).

This perspective highlights how the efficiency of smart contracts does not derive from a total replacement of existing legal and organizational forms, but rather from a calibrated combination of them. Digital tools guarantee speed and deterministic execution in standardized and highly verifiable relationships, while legal institutions remain crucial for managing unprogrammable contingencies and safeguarding the certainty of contractual expectations.

It follows that the real challenge is not technological, but one of institutional design: identifying the optimal mix of enforcement for each class of transactions, so as to maximize the benefits of automatic codification without sacrificing the adaptive flexibility required in an inherently uncertain economic and social context.

In this sense, the question of adaptation leads directly to the topic of new organizational forms born in Web3. If smart contracts are the technical tool that governs individual transactions, Decentralized Autonomous Organizations (DAOs) extend this logic to entire collective governance structures.

DAOs are in fact the most radical experiment in transposing "rule by code" into the organizational sphere, with the aim of replacing or integrating traditional decision-making and coordination mechanisms. Analyzing their properties and their impact on transaction costs means taking a further step in the application of TCE to Web3, shifting the focus from individual contractual relationships to the institutional architectures that support them.

This is precisely where the next chapter comes in, dedicated specifically to DAOs and their role in reducing (and transforming) transaction costs.

3.2. Impact of DAOs within the TCE Framework

In Web3 terminology, Decentralized Autonomous Organizations (DAOs) are natively digital entities in which coordination between participants, operating rules, and a significant portion of decision-making processes are implemented and executed via smart contracts on public blockchains.

Put simply, a DAO is a new kind of organization that runs on a blockchain. It allows people to come together and make decisions collectively, following rules that are written directly into code, these are the so called smart contracts. Unlike traditional organizations, there’s no central authority calling the shots; instead, power is distributed among all participants, who can have a real say in how things are run. This means that the rules are transparent, actions are automatica

lly enforced by the code, and the entire system is open to anyone with a stake, making DAOs both collaborative and self-governing by design.

In this sense, the DAO does not coincide with the underlying blockchain network, but is a digitally native organization that lives "above" the blockchain through smart contracts that define membership, rights, decision-making procedures, and resource allocation.

The historical evolution of the concept is rooted in discussions within the early crypto community about the Decentralized Autonomous Corporation (DAC), which emerged after the introduction of Bitcoin and popularized by Vitalik Buterin, who in 2014 shifted the focus from a "dividend" idea of a crypto-enabled corporation to an organization that deliberates through digitalized voting open to members and whose operational relationships are governed and executed in a decentralized manner via software (Buterin, 2014). This transition marks the evolution from a restricted financial use to a new code-enabled corporate governance approach; the point is then empirically tested with The DAO (2016), an on-chain venture capital

experiment formalized in Jentzsch's white paper and quickly becoming a case study for the limits and possibilities of algorithmic governance (Jentzsch, 2016).

Functionally, a DAO 'lives' on-chain: it aggregates resources in a programmable treasury, defines membership and voting mechanisms, and entrusts the execution of decisions to deterministic code paths that activate expenses, parameter changes, or contract module updates. The literature notes that, compared to a traditional organization, this transfers portions of internal regulation from law and administrative procedures to executable code, producing ex ante transparency and ex post auditability of processes.

From this perspective, DAOs do not correspond to a single business model, but to a general organizational model applicable to heterogeneous functions (from the management of DeFi protocols to collective financing, to cultural and public interest initiatives).

The distinctive properties of DAOs emerge from the intertwining of blockchain architecture and governance design: like smart contracts, they inherit the transparency, verifiability, and resilience of the distributed ledger and, at the same time, introduce key organizational elements. To better understand how DAOs change the rules of the game, let’s break down their main features in plain language:

Institutional automation (rule by code): once approved, decisions become self-enforcing through calls to on-chain functions, reducing executive discretion and increasing the predictability of enforcement, while leaving open the issue of secure code updatability (De Filippi & Wright, 2018).

Native accountability: balance sheets, voting histories, quorums, and executive logs can be inspected on the ledger, reducing information asymmetries between insiders and outsiders and agency costs.

Tokenization of participation: political and economic rights are often conveyed by fungible or non-fungible tokens that enable token-weighted voting schemes or alternatives such as delegated/liquid democracy, quadratic voting, and reputation-based mechanisms; the infrastructure also allows for "exit" tools such as rage quit, which reduce lock-in and internal conflicts in more cooperative models.

Finally, composability: "DAO-as-a-service" frameworks provide reusable modules for membership, proposal-voting-execution, treasury multisig, and governance parameters, reducing set-up costs and promoting standardization.

A decentralized autonomous organization (DAO) functions through a continuous series of interconnected actions instead of distinct phases. Membership is determined by the code itself , in simple words, your participation and voting rights are tied to tokens whether they’re fungible, non-fungible, or based on reputation. This setup strikes a balance between openness, pseudonymity, and protection against attacks, all governed by specific admission rules like token gating, vesting schedules, or whitelists. Within this structure, proposals are put forward, discussed, and voted on during set timeframes using various counting methods (like majority, quorum, or delegation). Once approved, these decisions trigger on-chain transactions that happen automatically. The organization’s treasury is also programmable, enabling collective discussions to directly allocate funds or modify protocol parameters. All these processes depend

on careful parameter settings and the ability to securely upgrade or pause operations, with thresholds, time-locks, and spending limits established in advance to ensure both flexibility and security.

If we look at DAOs through the lens of Transaction Cost Economics (TCE), several effects on transaction costs emerge.

From an ex ante perspective, DAOs tend to reduce information and negotiation costs because they make rules, rights, and constraints observable before they are agreed upon: by-laws are codified in smart contracts, admission criteria (token gating, thresholds, delegability), treasury parameters, and voting procedures are public and verifiable on-chain, so that due diligence can shift from document collection to technical verification of contractual status and functions (De Filippi & Wright, 2018). This is equivalent to lowering the cost of verification and, with it, the cost of networking necessary to coordinate databases and counterparties (Catalini & Gans, 2020). The modularity of DAO-as-a-service frameworks (Aragon, DAOstack, Colony) allows for the reuse of verified templates for the proposal-vote-execution pipeline and for treasury, reducing drafting and set-up costs and shifting the focus to parametric choices rather than ad hoc writing (Hassan & De Filippi, 2021). However, these savings are offset by new ex ante costs:

i) a governance design cost (thresholds, quorums, timelocks, proxies) and a cognitive cost for members, who are required to understand the code-law interface and its implications.
ii) a regulatory/legal cost linked to the classification of tokens and membership, as shown by the SEC's DAO Report on the possible qualification of tokens as securities (SEC, 2017) and, more generally, uncertainty about legal personality and the allocation of responsibilities.

That means that DAOs replace part of the traditional negotiation costs with institutional design and compliance costs, where standardization is high, the balance tends to be favourable, instead in less mature contexts, it can be neutral or negative.

On the ex post front, the gain is generally more pronounced when outcomes are verifiable on-chain: the shared ledger generates continuous auditability of proposals, quorums, proxies, and treasury movements; Above all, the conditional logic that guides the entire process from proposal to vote, through to any waiting period and subsequent execution, means that the application of decisions no longer depends on the discretion of an external authority, but takes place automatically and irrevocably: a shift from rule of law to rule by code (De Filippi & Wright, 2018).

In repetitive or parametric contexts (updating risk parameters, unlocking conditional payments), this drastically reduces policing, reconciliations between registers, and disputes over "who did what, when" (Lumineau, 2021).

Furthermore, in DAOs, the adjustment of fees, debt ceilings, and treasury operations is traceable and justifiable on a shared database, reducing inspection and agent-monitoring costs. The native accountability of public logs can thus reduce agency costs and moral hazard, as the behavior of delegates and proponents is observable ex post and linked to accounting outcomes.

(Fritsch, Müller, & Wattenhofer, 2024).However, the reduction is not automatic: evid ly documents concentration of voting power, abstentionism, and dependence on a few pivotal delegates, which reintroduce a "de facto hierarchy" and shift part of the monitoring savings toward internal political control costs. In summary, when the object is tokenizable and performance is measurable on-chain, ex post costs (monitoring/enforcement) decrease; when political power is concentrated or on-chain voting is costly, a governance cost emerges that mitigates the TCE gain

The adaptive dimension is the most delicate area in the analysis of DAOs through the lens of Transaction Cost Economics (TCE). Williamson had already highlighted how institutional efficiency depends on the ability to realign a relationship when the state of the world changes, for example following price shocks, the emergence of new information, or technical defects (bugs). In this context, DAOs are all about adapting through structured collective methods. Just think about things like code upgrades (which are often done through proxies), tweaks to how things operate, emergency stop tools, and in some serious situations, even chain forks.

These mechanisms allow shocks to be managed without resorting to courts or external authorities, but they introduce coordination costs and risks of maladaptation: decision thresholds that are too low can accelerate corrections but open the door to opportunistic coalitions, while thresholds that are too high protect minorities but slow down action in time-sensitive situations (Wang 2025).

The literature on blockchain governance has proposed viewing distributed ledgers as a tertium genus of coordination with respect to contract law and fiduciary relationships, capable of reducing ex post discretion by shifting enforcement to the code.

On the other hand, his shift does not eliminate the need for corrective institutions that balance security and speed. In practice, decisions are never made entirely on-chain: many choices are discussed off-chain (forums, social media, calls) and only subsequently ratified on-chain, preserving that dimension of social judgment based on deliberation and reputation that code alone cannot replace (Reijers, 2021). When adaptation fails, forking is the last resort.

The case of The DAO in 2016 showed that chain splitting can restore security, but at the cost of institutional uncertainty and migration burdens, generating path dependence and loss of coherence between communities (Hassan & De Filippi, 2021). The most recent empirical evidence also shows that adaptation through collective voting can become costly: reduced effective quorums and a high concentration of proxies lead to polarized outcomes, long delays, and the risk of inefficiency (Fritsch et al., 2024).

Looking at it from this angle, solutions aimed at shifting adaptation to what’s considered the 'least costly point ‘like reputational proxies, predictive mechanisms, or options for users to exit if they’re unhappy don’t really eliminate costs. Instead, they just shift them around: moving from the external court system to the internal governance design. This means that the traditional costs of enforcement transform into costs associated with collective decision-making. It’s at this juncture that Transaction Cost Economics (TCE) can evolve to incorporate Decentralized Autonomous Organizations (DAOs). Now, efficiency isn’t just about cutting down on bargaining or enforcement costs; it’s also about creating institutions that ensure sustainable adaptation while keeping the risks of delays, manipulation, and ongoing conflict to a minimum.

The comparative picture that emerges from the analysis of DAOs through the lens of Transaction Cost Economics confirms the ambivalent nature of this organizational form.

Ex ante, DAOs tend to reduce search and qualification costs thanks to codified by-laws and public track records, reducing costs in standardizable and highly verifiable contexts, while introducing new design and compliance burdens.

Ex post, they significantly reduce monitoring and enforcement costs when the object of the transaction is native-digital and the pipeline is deterministic, exploiting the traceability and self-execution of smart contracts, however, this same feature can reintroduce political control costs, especially in the presence of concentrated voting power or low levels of active participation.

In terms of adaptation, DAOs replace the flexibility "by interpretation" of semantic contracts with flexibility "by collective decision," based on voting and deliberation processes. This mechanism is powerful, as it allows for endogenous realignments without recourse to courts or external authorities, but it is not without costs: it depends heavily on governance parameters, the distribution of rights, and the participation costs incurred by members.

As noted by Williamson (1991) and echoed in contemporary literature, the ability to react to external shocks is a decisive criterion for institutional efficiency.

In this sense, DAOs broaden the spectrum of choices available in institutional make-or-buy, creating hybrid spaces between market and hierarchy, in which part of the authority is replaced by verifiable automatisms and part of the trust by computable transparency (Wang & Schilke, 2021), even thought, as pointed out by De Filippi and Wright (2018), DAOs can increase adaptation costs in environments characterized by high uncertainty and heterogeneity of interests, situations in which classical TCE envisaged the use of hierarchy or legal gap-filling.

The most recent literature highlights that "blockchain governance may, under certain conditions, incur lower transaction costs than contractual and relational governance" (Lumineau, 2021), but the advantage tends to diminish when participation is reduced and decision-making power is concentrated, making adaptation slow, costly, and potentially inefficient.

Considering this , it is possible to summarize the impact of DAOs on transaction costs in the three dimensions of TCE (ex ante, ex post, adaptation). The following table provides a structured comparison highlighting both areas of actual cost reduction and critical points where the balance can become negative due to technical complexity, participation issues, or risks of maladaptation.

Table 7 – Effects of DAOs on Transaction Costs (TCE)

Cost category	Impact of DAOs	Why
Ex ante (search, information, negotiation)	Medium–high	Public governance code, treasury records, and decision history enable faster due diligence and comparison; standardization of participation modules reduces initial set-up. However, costs may increase due to the need for technical–legal expertise and uncertainty over legal personality and liability.
Ex post (monitoring, auditing, enforcement)	Strong reduction	Proposal–vote–execution pipeline and immutable logs allow continuous auditing and technological enforcement over treasury and parameters; lower monitoring and agency costs.
Adaptation (realignment to shocks)	Risk of increase	Collective decisions can be slow or polarized; risks of power concentration and low participation; potential forks generate high coordination and migration costs.

DAOs, or Decentralized Autonomous Organizations, can be thought of as organizations that run on programmable code, allowing them to manage coordination and enforcement directly on public blockchains.

Their properties make them a laboratory for experimenting with new forms of distributed governance. However, empirical evidence shows that these properties coexist with technical and organizational constraints: the distinction between on-chain and off-chain governance confirms that many processes remain anchored to social and negotiation practices, while the concentration of voting power and low participation pose concrete risks of plutocracy in purely token-weighted systems.

These limitations highlight that the promise of decentralization is not free, but conditioned by participation costs, cognitive complexity, and unequal distribution of wealth among governance token holders. Furthermore, questions remain open regarding the meaning of "decentralization" and the degree of "autonomy" required, as well as the legal qualification of DAOs in terms of subjectivity and responsibility. To summarize DAOs not only redesign traditional transaction costs (ex ante, ex post, adaptation) but also introduce a new category of emerging costs related to governance design, collective coordination, and sustainability of participation. It is precisely these costs, and the conditions under which they can exceed the benefits of classic cost reduction, that will be the focus of the next chapter.

3.3 Emerging Costs in Web3: New Cost Categories Beyond Classical TCE

The analysis conducted in the previous paragraphs has shown how blockchain, smart contracts, and DAOs reduce, with varying intensity, the transaction costs in the three families identified by Transaction Cost Economics (ex-ante, ex-post, and adaptation). In the transition from Web2 to Web3, however, TCE does not lose its validity,it remains a "fully applicable" theoretical framework, but it needs to be extended to capture the specific frictions of open, programmable, and constantly evolving infrastructures.

Certain characteristics of Web3 make it possible to reduce ex-ante (search and negotiation) and ex-post (monitoring and enforcement) costs thanks to traceability, auditability, and automation. at the same time, studies emphasize that technological uncertainty, network effects, and the quality of input data can introduce new frictions that shift the efficient governance structure (Schmidt & Wagner, 2019), that means that, TCE retains its explanatory power, but the organizational "optimum" may shift when technological innovation gives rise to cost categories not anticipated by the original framework.

The notion of 'emerging costs' of Web3 arises investments in implementation and training, distributed governance and participation costs, security and technical insurance costs, as well as expenses related to scalability, usability, and protocol resilience.

The point is not that TCE is disproved, but rather that its domain is expanding alongside the costs of individual transactions, it becomes necessary to map the infrastructure costs and collective choice costs that make the transactions themselves possible. While classic TCE focused primarily on friction between contractual parties, in Web3, "upstream" friction (setup, standards, hardening) and "system" friction (distributed network coordination) emerge, which can absorb, shift, or even exceed the savings generated "downstream" by automation.

An initial group of emerging costs is associated with implementation, integration, and training. Empirical studies on adoption in supply chains show converging barriers such as high setup costs, the need for specialist skills, staff training, and inter-organizational orchestration, all of which weigh heavily on the ex-ante and transition phases (Queiroz & Wamba, 2019).

Systematic reviews report "high implementation costs," technological immaturity, and managerial knowledge gaps as recurring drivers of failure or delay (Kafeel, 2023). These outcomes are consistent with the observation that, at present, Web3 poses distinct challenges for systems, developers, and users: from scalability to interoperability, from code security to data recoverability, which require significant investments in upskilling and process redesign (Huang, 2024). Furthermore, the scarcity of Web3/smart contract talent, often multi-project freelancers, raises the cost of skilled labor and dependence on external suppliers, affecting "make-or-buy" decisions.

A second obstacle is the cost of distributed governance. The promise of reducing agency costs through on-chain rules coexists with coordination frictions that can re-emerge in new forms: voting apathy, concentration of voting power in large token holders, procedural complexity, and

vulnerability of proposals. Empirical evidence on DAOs documents very low average participation and strong effective centralization of control, with impacts on decision-making legitimacy and allocative efficiency.

From a TCE perspective, the reduction of certain agency costs may be offset by recurring coordination costs (deliberation, quorum, proposal auditing, proxies) and new risks of "governance attacks" (e.g., leverage via flash loans), which require anti-manipulation mechanisms and thus additional design and control expenses. Furthermore, in the absence of flexible renegotiation channels, the 'if-this-then-that' rigidity of smart contracts can generate 'bad adaptation' and therefore higher transaction costs than semantic contracts when unexpected events require rapid deviations from the codified structure (Vatiero, 2022). Recent evidence and mapping on DAOs confirm that process design (voting phases, voting systems, proxies, contingency plans) is an organizational cost item that is independent of the traditional hierarchy.

the final block, dedicated to security, implies a recurring "security budget": code audits (including formal ones), remediation and patching activities, bug bounty programs, runtime monitoring, and insurance coverage/risk transfer instruments are needed. The literature on smart contract vulnerabilities and development practices shows that defects are recurrent and that analysis tools are heterogeneous, as a result, verification costs are not marginal and should be considered not only as an initial investment (CAPEX) but also as an ongoing operating expense (OPEX). The challenges of secure development and the various technical taxonomies of key attack surfaces, such as runtime environment/VM, languages, and code, all center on one fundamental aspect, namely the need for multiple planned tests, independent validation, and controlled upgrade procedures.

When we talk about infrastructure, it's important to understand that ensuring consensus security comes with its own set of costs. To prevent issues like double-spending and transaction reorganizations, we need to implement proper fees, staking, and incentives. There really aren’t any shortcuts if we want to maintain the finality of transactions. Additionally, how we distribute validation power through methods like pooling, slashing, and setting barriers to entry plays a significant role in managing risks and oversight expenses. So, in a nutshell, the strength of our consensus isn’t just a technical matter; it’s also an economic decision.

Furthermore, value extraction practices related to transaction ordering (MEV) can reduce the fairness and predictability of settlement, forcing actors to introduce mitigation mechanisms (protected routing, anti-reorg defenses, dedicated auctions), with additional burdens for developers and users.

When the application depends on oracles or bridges, risks and reliability costs arise for off-chain and cross-chain components.

Input validation, source redundancy, and contractual risk management become necessary and ongoing activities, because these elements are known attack vectors and reintroduce points of trust into the system. Again, the expense is not a one-off but an ongoing commitment.

In brief security in Web3 is not a "closed chapter" at go-live, but a structural item to be planned on a recurring basis between CAPEX and OPEX, which includes code and infrastructure

hardening, mitigations related to consensus and fairness, as well as the management of risks introduced by oracles and bridges.

Scalability, usability, and "protocol resilience" constitute the fourth block and, in operational terms, mean that in order to grow, more technical complexity and more organizational processes must be budgeted for. Scalability reviews show structural trade-offs: increasing throughput and latency requires either upgrading the base layer (Layer-1, i.e., the "parent" blockchain that performs consensus and finality) or shifting part of the load to higher layers (Layer-2/Layer-3, i.e., networks or rollups that execute transactions "outside" L1 and publish proof of them on L1).

In both cases, architectural complexity increases network sharding, rollup protocols, data availability mechanisms, and a share of the burden shifts to application integration (more complex toolchains, new end-to-end latencies, retry and synchronization logic between domains), as documented by technical surveys on scalability and, more generally, analyses of the Web3 landscape. In this sense, studies note that many off-chain or second-level solutions perform well precisely because they forego some of the "native" properties of the base layer, which reopens issues of trust, transparency, and security that must be managed at the project and governance levels.

These same L2/L3 solutions introduce new risk surfaces and therefore new recurring operating costs. In rollups, for example, the 'sequencer' (the actor that orders and forwards transactions to the rollup) can become a point of censorship or congestion, while 'bridges' (connectors of assets and messages between chains/layers) expose interoperability and security risks: both of which require mitigation procedures (sequencer redundancy, more robust cryptographic tests, bridging controls) and the management of "back-pressure fees," i.e., pricing mechanisms that increase under load to slow down the flow of transactions and stabilize the queue. Recent literature on Web3 clearly categorizes these trade-offs (scalability-interoperability-privacy at the system level; code security and incentives at the development level; key recovery and interfaces at the user level), confirming that scale comes at the cost of coordination, monitoring, and ongoing maintenance (Huang., 2024).

On the user side, HCI (Human–Computer Interaction) research, "a discipline concerned with the design, evaluation, and implementation of interactive computing systems for humans" converges in pointing out persistent friction in wallets and key management: understanding and secure storage of the seed phrase, recovery in case of loss, and readability of signatures converges in pointing out persistent friction in wallets and key management.

Understanding and secure storage of seed phrases, recovery in case of loss, and readability of signatures/permissions remain critical issues that generate often underestimated support, training, and incident management costs.

Empirical studies on Web3 projects show that complex interfaces and inadequate recovery mechanisms increase the risk of error and phishing, transferring training and support burdens to developers and platform operators.

Finally, the "institutional resilience" of the protocol, i.e., its ability to absorb shocks such as forks, rule changes, or phases of non-finality, does not come free of charge: maintaining it

requires well-designed upgrade processes, network audits, client interoperability testing, stakeholder training, and governance capable of coordinating heterogeneous actors. Management literature that has followed real cases of sectoral transformation on DLT (Distributed Ledger Technology) shows that, as the scale grows, organizational capital (training plans, first-level support, budgets for agile iterations) and public-private cooperation mechanisms are needed; and how, in some contexts, the decision is even made to migrate away from DLT if the cost of maintaining resilience and usability exceeds the expected benefits. These organizational and reputational costs are often externalities that "classical" Transaction Cost Economics, calibrated to centralized markets, tends to underweight; in fact, the most recent TCE studies call for incorporating technological uncertainty and continuous adaptation costs when evaluating blockchain governance structures (Schmidt & Wagner, 2019).

Emerging costs show that Transaction Cost Economics remains valid, but must be extended to include new items typical of Web3. In addition to traditional costs of negotiation and research (ex-ante), monitoring and enforcement (ex-post), and adaptation, the following must also be considered: (i) implementation, integration, and training costs, exacerbated by a shortage of specialist skills; (ii) distributed governance costs, which reflect the difficulties of coordinating decision-making and the risks of vote manipulation; (iii) security costs, which require audits, bug bounties, and constant monitoring at the code, protocol, and external component levels, such as oracles and bridges; (iv) costs related to the scalability, usability, and resilience of protocols, which manifest themselves in multi-level architectures, complex interfaces, and continuous upgrade processes.

The literature confirms that many Web3 "economies" do not result from the elimination of costs, but from their shift: relational verification and intermediation costs are reduced, but those for skills, tools, governance, and infrastructure hardening increase.

In the absence of adequate technical and institutional adaptation mechanisms, these items may even exceed the expected savings, making adoption less cost-effective than Web2 or hybrid solutions.

The following table compares the impact of emerging Web3 costs on the three classic categories of transaction costs identified by Transaction Cost Economics (ex-ante, ex-post, adaptation), showing how the theoretical framework extends to include infrastructure and collective choice costs.

Table 8 – Emerging Cost Categories in Web3 and Their Impact on TCE

Emerging cost category	Ex-ante (search, contracting)	Ex-post (monitoring, enforcement)	Adaptation (renegotiation, resilience)
Implementation & Training	High (setup, training, skill gaps)	Medium (system integration, technical support)	Medium (continuous updates, re-skilling)
Distributed Governance (DAO, voting)	Medium (design of rules and processes)	Medium (proposal auditing, anti-manipulation)	High (voter apathy, forks, contractual rigidity)
Security (audits, bug bounties, consensus)	High (initial audits, insurance)	High (patching, remediation, MEV mitigation)	Medium (protocol upgrades, slashing)
Scalability & Usability	Medium (architectural setup, L2/L3)	High (network monitoring, user support)	High (upgrades, forks, institutional resilience)

The mapping of emerging costs shows that the adoption of Web3 does not depend solely on the ability to reduce traditional transaction costs, but also on the sustainability of the institutional and infrastructural costs that arise from it. The real convenience lies in balancing savings with new implementation, governance, security, and resilience costs.

The next chapter will take a closer look at the conditions that make Web3 adoption possible or impossible, distinguishing between favourable contexts and structural limitations related to technological and economic trade-offs.

3.4 Conditions for Web3 Adoption

The previous paragraphs have highlighted how Web3 not only helps to reduce certain transaction costs but also introduces new categories of expenditure compared to the classic framework of Transaction Cost Economics.

It follows that the adoption of decentralized solutions cannot be evaluated solely on the basis of the promise of disintermediation or automation: the economic and institutional conditions in which these technologies are implemented must also be considered.

In other words, Web3 is not a "universal optimum": it works better in some contexts, while in others it may be less efficient than hybrid or centralized structures (Schmidt & Wagner, 2019).

Recent literature emphasizes this conditional dimension of adoption. On the one hand, there are sectors and organizational models in which the benefits of transparency, auditability, and automation clearly outweigh the costs of setup and distributed coordination. This is the case, for example, with traceable supply chains or decentralized finance (DeFi), where the marginal value of "automated" trust is particularly high (Queiroz & Wagner, 2019). Wamba, 2019). On the other hand, there are scenarios in which technological rigidity and emerging costs make Web3 less convenient, especially when the environment requires frequent adaptations or high usability for inexperienced users.

In this perspective, the central question is no longer whether Web3 reduces transaction costs, but under what conditions such a reduction is effective and sustainable. It is therefore necessary to examine the contexts favorable to adoption and identify the factors that amplify the comparative advantages of decentralized solutions.

Web3 tends to be cost-effective when the expected benefits of transparency, auditability, and automation are combined with organizational and institutional conditions capable of sustaining the emerging costs described in the previous paragraph over time. The literature shows that the most favorable environments are characterized by high interdependence among actors, persistent information asymmetries, and the need for reliable shared logs in the absence of a central authority, provided that the ecosystem develops robust forms of collaboration, clear governance rules, and a flexible (even hybrid) technological path capable of adapting to the scaling phases.

Projects that manage to overcome the stalemate associated with the pilot phase show that success depends on the convergence of several factors: the construction of a shared vision of the sector, the activation of public-private partnerships, the structuring of an economically sound governance model, the articulation of complementary roles among partners, investment in training and support programs for stakeholders, and the adoption of technology management strategies based on agility, progressive learning, and scalability.

Even when, after market entry, there is a shift towards more centralized solutions, it is often the organizational capital built around the blockchain that makes scalability possible (Leshinsky & Junnila, 2025).

This enabling condition is particularly visible in complex financial and production supply chains, where trust is fragmented and access to data is unequal.

In supply chain finance, for example, adoption is more sustainable when actors address ex ante the barriers that typically block projects (market uncertainty, skills gaps, integration burdens, regulatory risks), equipping themselves with shared methodological tools to identify and weigh them (prioritization frameworks, objective/subjective weights, decision support tools). Analyses show that if the supply chain ecosystem establishes coordination and joint learning mechanisms, the benefits of data traceability and reliability tend to outweigh the recurring infrastructure costs (tooling, auditing, integration), with positive effects on the cost of capital and time-to-market for financial innovation (Wang, 2025).

Favorable contexts are those where we can meet both throughput and latency requirements by scaling effectively. This means using a mix of on-chain and off-chain solutions like layer-2 rollups, state channels, or sharding/sidechains that boost throughput and cut down on latency, all while keeping security and verifiability intact.

Systematic reviews show that sharding, state channels, and second-level solutions significantly improve performance, at the cost of greater architectural complexity and new coordination tasks (data availability, cross-domain synchronization); adoption is therefore more efficient when the organization is prepared to internalize these trade-offs as predictable operating costs over the solution's lifecycle.

There are domains where data sensitivity and regulatory constraints make information centralization very costly or impractical. In these multi-party data collaboration scenarios, the combination of decentralization, immutable traceability, and cryptographic mechanisms preserves privacy and integrity, improving accountability and audit trails without concentrating trust and decision-making power in a single intermediary. Coordination between multiple parties is more reliable when there are recurring investments in security (CAPEX+OPEX for logging, verification, attack mitigation, and updates) and when technical and organizational governance remains aligned with the risk and compliance objectives of the ecosystem (Orabi, Emam, & Fahmy, 2025).

In summary, Web3 finds favorable conditions for adoption when:

(i) gains in transparency, auditability, and automation have high marginal value for multi-actor networks with fragmented trust;
(ii) organizational capabilities exist to co-design governance, incentives, and training to sustain institutional and infrastructure costs over time;
(iii) performance requirements can be met with composite scaling paths, accepting and governing the related architectural trade-offs;
(iv) data protection and accountability are strategic complements to the business model, not mere ancillary requirements.

In contexts where these conditions are present, and only in these contexts, the reduction in traditional transaction costs is more likely to exceed the new implementation, governance, security, and resilience costs, making the adoption of Web3 economically and institutionally sustainable.

To bring the convenience criteria mentioned in this paragraph to life, we suggest a scorecard that outlines the conditions for adoption (see Table 9). This scorecard captures the common decision-making factors found in Web3 projects. It’s not meant to serve as a strict test, but rather as a framework that connects the ecosystem's goals to the effects on transaction costs both before and after transactions, as well as during adaptations and introduces new cost categories like implementation and training, distributed governance, security, scalability and usability, and resilience. The indicators are presented as guiding questions and qualitative "signals" that can help identify prerequisites, bottlenecks, and areas for improvement before investing significant resources.

Table 9 — Conditions for Web3 Adoption

Key condition	**Guiding question**	**Favorable signals (Go)**	**Caution signals (Delay / No-go)**	**Predominant TCE impact**
Regulation & compliance	Is the legal perimeter clear?	Applicable rules defined; identity checks in place and legal responsibilities clear	Regulatory uncertainty; unmitigated legal risks	Ex ante
Data interoperability	Are common standards adopted across partners?	Adoption of open standards; shared formats and vocabularies	Proprietary formats; high lock-in	Ex ante
Distributed governance	Are roles and decision processes clear?	RACI defined; stable quorum/delegations;	Role ambiguity; low participation	Adaptation
Security	Is the security posture credible?	Periodic audits; patching plans; bug bounty	No audits; long remediation times	Ex post
Performance	Are throughput and latency adequate?	Requirements met in realistic tests	Predictable bottlenecks	Ex post
Scalability (combined scaling)	Is there a clear path to scale?	L2/channels/sharding planned; data availability & bridging managed	No scaling plan; unmanaged sync risks	Adaptation
Economic sustainability	Is the economic model viable mid-term?	Recurring costs compatible with expected volumes	Disproportionate fixed/variable costs	Ex ante

Usability & adoption	Can users operate without undue friction?	High task success; lightweight onboarding; sustainable support	Frequent errors; heavy training burden	Ex post
Ecosystem onboarding	Are critical partners on board?	Qualified majority of key partners	Limited or uncertain partner commitment	Ex ante
Sensitive data & privacy	Is privacy built into the architecture?	Sensitive data off-chain with proofs/attestations;	Exposed data; weak controls	Ex post
Oracles & input integrity	Are external data sources reliable?	Redundant oracles with SLAs and fallback	Single point of failure	Ex post
Resilience & upgrades	Can the system handle incidents and upgrades?	Runbooks, rollback plans, defined release windows	No continuity plans	Adaptation

The scorecard should be read as a strategic triage tool: if compliance and security are adequate and the majority of areas show favorable signs, adoption appears justifiable; in the presence of red flags (regulatory, security, or governance), it is preferable to opt for postponement with corrective actions or non-adoption.

The analysis must also be contextualized (sector, scale, risk) and updated cyclically throughout the project lifecycle, as interoperability, onboarding, performance, and usability typically improve through agreements, architectural tuning, and training courses.

Chapter 4 – Comparative cases in Supply Chain Management: Oracle (Web2) vs. VeChain (Web3)

This chapter applies the lens of Transaction Cost Economics (TCE) to a "controlled" comparison between two technological archetypes currently available for supply chain management: on the one hand, a cloud-centric and widely used ERP/SCM system such as Oracle SCM (Web2 paradigm), and on the other, a Web3 infrastructure geared towards traceability such as VeChain, based on public blockchain, smart contracts, and IoT integration. The goal is not to establish an "absolute" winner, but to measure how the three families of TCE costs: ex ante, ex post, and adaptation are redistributed as the architecture changes, taking into account the new cost categories that have emerged with Web3 (implementation/training, distributed governance, security/technical auditability, scalability/usability, protocol resilience). The analysis dialogues with the literature that, from Coase and Williamson onwards, has linked the boundaries of the firm to the frictions of search, negotiation, monitoring, enforcement, and contractual realignment (Coase, 1937; Williamson, 1985), showing how digital technologies can reduce some frictions and create new ones, especially when verification depends on off-chain data and oracles (Catalini & Gans, 2016/2018; Schmidt & Wagner, 2019).

4.1 Brief description of the two systems

4.1.1 Oracle Corporation and the Web2 Paradigm

Founded in 1977 in the United States, Oracle Corporation is among the leading global players in enterprise software and cloud services, with a portfolio that integrates management applications, infrastructure services, and data technologies. Its offering is structured across three complementary layers: the data layer (with Oracle Database and, in the open-source domain, MySQL), the application layer (Oracle Fusion Cloud Applications: ERP, SCM, HCM, CX), and the infrastructure layer (Oracle Cloud Infrastructure – OCI). Within this ecosystem, the Oracle Fusion Cloud SCM platform is one of the most widely adopted cloud-centric systems, integrating modules for planning, procurement, manufacturing, order management, logistics, and product lifecycle management, supported by continuous updates and real-time analytics.

This model reflects the distinctive traits of the Web2 paradigm, based on centralized platforms, managed identities and roles, interactions through standardized APIs, and provider-governed updates. Oracle's emphasis on delivering "one cloud for the supply chain" aims to reduce the typical fragmentation of legacy stacks, namely outdated and often highly customized proprietary systems accumulated over time, which generate data duplication, compatibility issues, and high maintenance costs. The goal is to replace such fragmentation with a centralized backbone in which planning, execution, and monitoring processes occur in an integrated manner.

At the architectural level, Oracle adopts a multi-tenant model, which in simple terms means that multiple companies can use the same shared cloud platform, but each has its own separate "space" (tenant) where data and configurations remain isolated and secure. This architecture relies on the use of Application Programming Interfaces (APIs), which allow different software systems to "communicate" with each other. In practice, an API functions like a translator enabling two programs to exchange information: for example, if a warehouse management

system (WMS) and a transportation management system (TMS) need to share shipping data, APIs allow real-time data exchange without manual intervention. Alongside APIs, Oracle also supports Electronic Data Interchange (EDI), established protocols for the electronic exchange of business documents (orders, invoices, delivery notes) between trading partners. Thanks to these solutions, the platform can interoperate with external systems such as MES (Manufacturing Execution Systems), WMS, TMS, PLM (Product Lifecycle Management), or B2B portals, while maintaining centralized security and logging (Oracle, 2025a).

The literature on cloud-based SCM solutions highlights the benefits of this approach in terms of reduced capital expenditure (CAPEX), lower maintenance costs, and faster implementation compared to traditional on-premise systems (Saari et al., 2025). However, the multi-tenant nature also brings challenges: integration with databases and external systems remains complex, resulting in reconciliation costs and the risk of data silos when multiple supply chain actors work on different platforms (Schmidt & Wagner, 2019). In multi-enterprise contexts, interactions with partners and suppliers typically occur through API/EDI and iPaaS connectors, with logs and audits centralized at the tenant level (for example, ingestion of Oracle Integration Cloud logs into Logging Analytics). This enables operational traceability and troubleshooting but leaves the "truth" of transactions fragmented across separate databases, making periodic reconciliations and contractual assurance mechanisms necessary (Oracle, 2024).

These dynamics are consistent with the literature on ERP–SCM integration, which emphasizes that technology, while reducing certain barriers, does not fully eliminate the need for organizational alignment and inter-firm governance (Rajapakse, 2023; Mhaskey, 2024).

On the technological side, Oracle adopts a Software-as-a-Service (SaaS) architecture based on Oracle Cloud Infrastructure (OCI), which provides advanced services such as data warehousing, machine learning, and the Internet of Things (IoT). With the SaaS model, the company no longer purchases software as a license to be installed on its own servers but instead uses it online as a subscription-based service, paying a periodic fee. In this scheme, the software is entirely hosted and managed by the provider (in this case Oracle), which handles updates, maintenance, security, and availability, while the client company accesses the applications via browser or dedicated interfaces. This approach reduces initial investments in hardware infrastructure (CAPEX) and shifts costs to a more flexible subscription model (OPEX).

Such tools ensure high scalability and centrally managed security levels, but remain under the provider's control, which governs access, updates, and the roadmap. The result is a closed ecosystem, typical of the Web2 paradigm, in which interoperability processes are enabled through established standards such as EDI and proprietary APIs, and where information flows are governed by a single entity that guarantees service continuity and reliability (Schmidt & Wagner, 2019).

4.1.2 VeChain: A Web3 Infrastructure for Traceability and IoT Integration

Founded in 2015 in Singapore by Sunny Lu, VeChain is one of the first blockchain platforms created with a specifically enterprise-oriented mission. Managed by the VeChain Foundation, the company positions itself as a global provider of blockchain- and IoT-based solutions for product traceability, supply chain certification, and digital assurance. From its inception, VeChain has targeted industrial sectors with high requirements for transparency and reliability, such as food, fashion, luxury, automotive, logistics, and healthcare, building strategic partnerships with major corporations and international certification bodies (VeChain Foundation).

From a technological perspective, VeChain presents itself as an enterprise-oriented Web3 infrastructure, designed to enable traceability, certification, and attestations across the entire product lifecycle through the combination of public blockchain, smart contracts, and IoT integration. Its main network, VeChainThor, is compatible with the Ethereum Virtual Machine (EVM), which allows developers to easily adapt and reuse smart contracts originally designed for Ethereum.

In terms of consensus, VeChain employs a Proof-of-Authority (PoA) mechanism, where validators are authorized and publicly known, under the governance of the VeChain Foundation. With the evolution to PoA 2.0, also known as SURFACE, the network introduced transaction finality, ensuring that once transactions are confirmed, they cannot be reversed. This feature strengthens operational reliability, which is crucial in business-critical contexts that require the certification of physical processes and supply chains without uncertainty. Compared with probabilistic models such as Proof-of-Work, finality reduces the risk of double-spending and increases security for certification and traceability purposes (VeChain Foundation, 2020; VeChain Foundation, 2022).

The application layer is centered on VeChain ToolChain®, a “ready-to-use” platform that provides standardized templates for multiple industries (food, fashion, luxury, manufacturing, healthcare). ToolChain also offers web SDKs to customize interfaces and processes, as well as connectors for physical devices such as QR codes and NFC chips, which function as “oracles” bringing data and events from the physical world onto the blockchain. A key strength lies in VeChain’s collaboration with DNV, a leading international certification body, which has developed reusable traceability and assurance models on VeChain, thus ensuring credibility and standardization.

The partner ecosystem is indeed one of VeChain’s defining characteristics. DNV leverages the blockchain as a public ledger for digital assurance services, such as the My Story™ solution, which integrates user management, IoT devices, and independent audits to certify product compliance and performance (DNV, 2018). In the retail sector, the Walmart China case, developed with the support of PwC China, illustrates the use of VeChainThor to enhance food safety through a multi-level traceability system and tamper-proof data sharing among all supply chain participants (PR Newswire, 2019).

These examples demonstrate how the network was designed to combine on-chain controls with off-chain audit and certification processes conducted by qualified third parties, transforming data provenance into immutable, end-to-end verifiable evidence.

From a technical and operational standpoint, VeChainThor integrates four main elements:

(i) EVM compatibility, which facilitates the development and adoption of smart contracts.
(ii) A dual-token economic model that provides greater predictability of total cost of ownership (TCO) for enterprises;
(iii) A PoA consensus with transaction finality, balancing performance, integrity, and transparent governance;
(iv) An ecosystem of industrialized tools (ToolChain, SDKs, traceability templates, integration guidelines) that shorten the time-to-value of enterprise projects. In this sense, VeChain positions itself as a purpose-built infrastructure for supply chains, capable of combining the immutability of distributed ledgers and the programmability of smart contracts with certified external audits, thereby bridging the gap between the physical and digital worlds (VeChain,).

In summary, VeChain represents a public and open solution focused on traceability, interoperability, and certified attestations, combining predictable cost mechanisms (VET/VTHO), transaction finality, and IoT integration. The availability of certified templates and the collaboration with assurance partners such as DNV and PwC have fostered the diffusion of the platform in real-world scenarios ranging from food to luxury, making it an emblematic case study for understanding how distributed verification can be orchestrated in complex supply chains. These characteristics will serve as the main point of comparison with the Web2 archetype analyzed in section 4.2, which will be further developed in the following chapter through the lens of Transaction Cost Economics (TCE).

4.2 Comparative TCE analysis

The comparison between Oracle SCM and VeChainThor, viewed through the lens of Transaction Cost Economics (TCE), highlights how two radically different technological architectures redistribute transaction costs along the entire contractual cycle. In Web2 systems such as Oracle SCM, preliminary costs concentrate on supplier scouting and qualification, data mapping and reconciliation, contract negotiation, and the configuration of internal processes (workflows, roles, catalogs). The standardization enabled by the cloud reduces part of the complexity compared to on-premise ERP, but it does not eliminate integration costs when the supply chain involves multiple partners and legacy systems: aligning heterogeneous databases requires time and resources, and informational symmetry still depends on certifications, external audits, and reconciliation cycles, with persistent contractual and legal costs as a result (Gupta, 2020). In the Web3 paradigm, the presence of a distributed and shared ledger drastically reduces preliminary verification and networking costs between parties: the logic of "verifying state" on blockchain provides a single, immutable, and synchronized source of truth, with particularly evident benefits when qualifications, certifications, and performance histories are recorded on-chain and can be reused without repeated validations (Catalini & Gans, 2016). However, this advantage is not universal: when the objects of the transaction are off chain (information originating from external environments), efficiency depends on the quality and reliability of the sources. This is where the oracle problem emerges. Consequently, while Web3 reduces scouting and due diligence costs, it simultaneously introduces new project costs linked to data standardization, IoT integration, and the design of reliable oracles.

This difference becomes clearer when focusing on concrete technologies. Oracle adopts a centralized multi-tenant SaaS architecture on Oracle Cloud Infrastructure, which provides access to an integrated application ecosystem (planning, procurement, logistics, manufacturing) with periodic updates and API/EDI interoperability. This model reduces ex ante implementation costs thanks to a mature ecosystem of partners and best practices, but it does not eliminate the need to reconcile data when multiple enterprises interact: the "transactional truth" remains fragmented across separate databases, and reliance on certifications, contracts, and external audits remains essential. VeChain, by contrast, embeds "shared truth" into its infrastructure: the public ledger records every event immutably and reusably, and the use of RFID, QR, and NFC devices links physical data to the blockchain. In this context, ex ante costs decrease in terms of verification and due diligence (also thanks to reusable attestations, such as those provided by DNV along the supply chain), but they reappear in the need to design reliable standards, sensors, and oracles to avoid irreversible errors.

On the ex post side, Oracle offers audit trails and advanced analytics with near real-time monitoring capabilities within the tenant domain; however, inter-enterprise enforcement remains external to the system: in the case of disputes, reconciliation requires legal or relational procedures because truth remains distributed across different databases. With VeChain, the programmability of smart contracts and the finality guaranteed by Proof-of-Authority 2.0 consensus make traceability endogenous to the system and, in many performance domains, allow for the automatic execution of contractual clauses, reducing the risk of post-contractual opportunism.

The Walmart China example represents one of the most significant cases to understand VeChain's potential in complex supply chains. The initiative, developed with the support of PwC China, demonstrated how the use of blockchain to record tamper-proof data throughout the entire supply chain drastically reduced inspection times and, most importantly, the risk of conflicts among partners. The ability to share a unique and immutable truth enabled certification of product origins, enhanced consumer safety, and improved transparency in relationships between supply chain actors. In this way, the platform contributed to reducing uncertainty and opacity, two factors traditionally at the root of reconciliation costs and contractual disputes. However, the case also highlights a critical aspect: system effectiveness crucially depends on the quality of the data entered. Without adequate governance of information, redundancy of sources, and cross-validation mechanisms, blockchain risks transforming errors from the external environment (off-chain) into permanent on-chain records, potentially triggering faulty automatic executions. In other words, even a technically perfect ledger can preserve "imperfect errors" if sources are not properly governed. This shows how VeChain, while reducing ex post monitoring and enforcement costs, introduces new ex ante costs to ensure data reliability, through investments in data standards, sensor infrastructures, redundant validation, and auditing practices that guarantee input accuracy.

As for adaptation costs, Oracle demonstrates greater elasticity through modular updates managed by the vendor and API-based integration, which allow relatively predictable reactions to new regulations or processes, albeit at the price of lock-in and dependence on the provider's roadmap. VeChain instead operates under the logic of "code is law": smart contracts reduce ambiguity and enforcement times but make rapid adjustments more costly in turbulent contexts. As Vatiero (2022) notes, in such contexts smart contracts may generate higher transaction costs than traditional contracts; to mitigate them, mechanisms such as proxy upgrades, pause functions, on-chain governance (quorums, multisig, DAOs), and legal fallback clauses are employed, which nevertheless entail coordination costs largely externalized to the provider in the Oracle model.

Emerging costs further accentuate the distance between the two archetypes. With Oracle SCM, companies benefit from a consolidated ecosystem of integrators, training, and support, with standardized change management processes, official documentation, and continuous assistance tools. The availability of widespread expertise and preconfigured models makes implementation and adaptation costs relatively predictable and shared with a broad ecosystem; here, emerging costs manifest primarily as vendor dependency and the need to integrate heterogeneous systems. With VeChain, onboarding instead requires less common expertise in tokenization, verifiable data standards, IoT integration, and, above all, oracle design; adoption entails building infrastructures for high-quality and resilient data collection, integrating hardware, software, and governance of information flows. Beyond the initial phase, public blockchains require recurring budgets for code audits, bug bounty programs, consensus optimization, and vulnerability management, as well as organizational investments to ensure validator legitimacy and transparency under the PoA model.

Three transversal dimensions complete the picture: scalability, usability, and protocol resilience. In terms of scalability, Oracle expands "vertically" by leveraging the provider's data centers and the economies of scale of an industrialized infrastructure; for the client, scaling simply means purchasing additional capacity or modules. In VeChain, increasing throughput and reducing latency require architectural interventions such as sharding, secondary layers (L2/L3), and consensus optimizations, introducing new coordination nodes (e.g., sequencers

and data availability mechanisms) and new costs of design and governance. From the usability perspective, Oracle offers familiar interfaces, preconfigured KPIs, and an experience consistent with enterprise standards, reducing adoption costs for personnel. VeChain, by contrast, may present frictions linked to cryptographic key management and the readability of digital signatures, requiring additional training. Finally, protocol resilience introduces Web3-specific costs: public blockchains require procedures for forks, upgrades, and client interoperability, with potentially costly collective interventions in the event of bugs or governance disputes; in the Oracle model, resilience is centrally provided by the vendor through versioning, security patches, and updates applied without directly involving client companies.

In terms of governance and security, the difference is equally pronounced. With Oracle, strategic and technical decisions are centralized within the contract with the provider: roadmap, upgrades, patches, and new functionalities are defined by the vendor and regulated by SLAs, reducing managerial uncertainty but increasing dependency. VeChain redistributes part of the decision-making through a distributed authority model based on PoA with known validators selected by the VeChain Foundation, which supervises their rotation and accountability; this reduces the costs of broad consensus but introduces expenses and complexity for validator legitimization and oversight. On security, Oracle relies on certified infrastructures (e.g., ISO, SOC, GDPR) and centralized audits included in the service contract; VeChain, on the other hand, requires continuous budgets for third-party audits, bug bounty programs, and mitigation of threats such as MEV, in addition to rigorous procedures to ensure oracle reliabilitya systemic cost not present in Oracle's centralized flows.

In summary, Oracle SCM represents the efficiency of a centralized system, suited to scenarios requiring planned interoperability, usability, and rapid adaptation, but where enforcement remains anchored to legal instruments and periodic reconciliations. VeChain proves more efficient where the value of traceability, auditability, and shared trust is high, and where actors are willing to invest in data standards, IoT, and distributed governance: under these conditions, ex post monitoring and enforcement costs are significantly reduced, while adaptation and oracle management costs may increase in turbulent contexts or with low-quality data. There is, therefore, no absolute winner: each model redistributes transaction costs differently, and the choice depends on whether priority is given to internal efficiency, inter-firm trust, or adaptive capacity.

Based on the analysis conducted, it is possible to summarize the points of convergence and divergence between Oracle SCM and VeChain in light of Transaction Cost Economics.The following table synthesizes the comparison between Oracle SCM and VeChainThor across the main TCE dimensions, while also identifying the emerging costs associated with the Web2 and Web3 paradigms.

Table 10 — TCE Comparison between Oracle SCM (Web2) and VeChain (Web3)

TCE Dimension & Emerging Cost	**Oracle SCM (Web2)**	**VeChain (Web3)**
Ex ante (research, negotiation, drafting)	Configuration of modules and integration with partners/legacy systems; traditional due diligence; predictable costs thanks to best practices and an established ecosystem of integrators; "truth" distributed across separate databases → reconciliations and contractual/audit mechanisms.	Reduction of preliminary verification and networking costs through a shared ledger and reusable attestations; new costs linked to tokenization, data standards, IoT/oracle design, and "design for verifiability."
Ex post (monitoring and enforcement)	Centralized audit and intra-tenant analytics; inter-firm enforcement remains legal/relational with periodic reconciliations.	Continuous on-chain auditability; smart contracts enable (partial) automated enforcement; dependency on input/oracle quality (oracle problem).
Adaptability	High configurability through vendor-managed modular updates and API integration; trade-off: vendor lock-in and change management dependency.	Rigidity of code ("code is law"); flexibility reintroduced via on-chain governance (quorums, multisig, DAOs), proxy upgrades, and legal fallback clauses; risk of maladaptation and costly migrations.
Implementation & training	Mature ecosystem of partners and training; standardized documentation and change management pathways → reduced project risk.	Scarcity of Web3 skills; need for dedicated toolchains; oracle design and verifiable data pipelines → more costly/slower onboarding.
Governance	Internal policies and contractual agreement with vendor; centralized decision-making (roadmap, SLA, patches).	Hybrid on-chain/off-chain governance; PoA with known validators → lower costs compared to broad consensus, but added expenses for validator legitimization and accountability.
Security & oracles	Traditional IT security: certifications (ISO, SOC, GDPR), centralized audits included in the service.	Recurring code/consensus audits, bug bounty programs, MEV mitigation; oracles/bridges as new risk

		surfaces → redundancy and validation of inputs required.
Scalability / usability	Cloud “vertical scalability” (capacity expansion as a service); familiar enterprise UX, preconfigured KPIs, widespread training and support.	Scalability trade-offs requiring L2/L3, sharding, or consensus optimization; additional operational costs for wallets/key management and usability support.
Resilience	Vendor roadmap, SLAs, versioning, and centrally managed patches (low coordination costs for clients).	Procedures for upgrades/forks, client interoperability, and multi-stakeholder coordination → technical/organizational costs specific to Web3.

Conclusions

This thesis advances a simple claim with complex implications: Web3 does not uniformly lower transaction costs; it reallocates them across the exchange cycle. Blockchain and smart contracts create efficiency gains where measurement, monitoring, and enforcement are the binding frictions by compressing verification costs and aligning ledgers across counterpartiesyet they simultaneously introduce new frictions in adaptation, distributed governance, code security, scalability–usability, and protocol resilience (Catalini & Gans, 2016; Vatiero, 2022). In this sense, Transaction Cost Economics (TCE) remains a powerful lens for understanding digital coordination (Coase, 1937; Williamson, 1985), but it requires an explicit extension to capture Web3-specific cost categories and to model settings in which verifiability depends on off-chain inputs and oracles.

The thesis contributes along two fronts. First, it proposes an extended TCE–Web3 framework that integrates relationality and institutional resilience, recognizing that coordination is sustained by both codified mechanisms and social/institutional scaffolding (Valentinov & Roth, 2024).

Second, it reframes "Web2 vs. Web3" as a contingent design problem rather than a technological contest: there is no absolute winner. The efficient architecture depends on the nature of the traded object (natively digital vs. physical), institutional uncertainty, organizational maturity, and the availability of interoperability standards.

Put differently, Web3 tends to dominate where the social value of reusable attestations and public auditability exceeds the system costs of oracle design, security hardening, and distributed governance (Aldoubaee et al., 2023; Catalini & Gans, 2016).

Theoretically, the mapping offered here aligns with recent work on "code-based" enforcement. Smart contracts reduce measurement and enforcement frictions but can raise adaptation costs when shocks are unanticipated or non-codifiable. As Vatiero argues, "because of the need for adaptation to mutable and unpredictable occurrences, smart contracts may incur higher transaction costs than traditional contracts" (Vatiero, 2022). Efficiency, therefore, hinges on oracle quality, parametrization choices, and governance safeguards. The practical corollary is not "code replaces law," but a code+law hybridization that allocates ex-ante "degrees of freedom" where they are cheapest: parameterized clauses on-chain; legal fallbacks and discretionary remedies off-chain; and on-chain governance with credible safeguards for upgrade and rollback.

Managerially, three implications follow. First, firms should run a full transactive due diligence that augments classical TCE categories (ex-ante, ex-post, adaptation) with Web3-specific "emerging costs": implementation & training, distributed governance, code security & auditability, scalability–usability, and protocol resilience (Saari, 2025). Second, Web3 should be treated as an infrastructural option, efficient when inter-firm information asymmetries are large and when reusable attestations and public logs are especially valuable, less so when context turbulence and unreliable off-chain inputs dominate (Aldoubaee , 2023).

Third, complementarities with AI/IoT should be selective and instrumental: integrate only when they lower verification/networking costs rather than inflating security and interoperability risks;

current evidence on federated learning & blockchain illustrates both benefits (integrity, poisoning resistance) and new coordination overheads (Orabi., 2025).

Institutionally, scalability is not only technical (throughput, L2/L3) but organizational and political: standards, governance, and incentives. The most material barriers to adoption are competitive and organizational uncertainty, especially for SMEs, hence the value of alliance-type platforms with clarified roles, shared runbooks, and early champions (Saari et al., 2025; Wang et al., 2025). In brief, many pilots stall not because hashing is slow, but because coordination is costly.

For research, three avenues are salient:

(a) Causal measurement of TCE–Web3 trade-offs using quasi-experimental designs (before/after; diff-in-diff) and operational indicators for emerging costs (upgrade failure rates, audit effort, quorum costs, UX frictions)
(b) Interoperability and resilience comparisons across architectures (public-only vs. hybrid vs. app-chains) and their effects on adaptation costs and vendor risk.
(c) The political economy of standards: how consortia, foundations, and regulators shape coordination costs and time-to-trust in digital supply chains (Saari., 2025).

In sum, the move from Web2 to Web3 is best read as a re-engineering of trust and cost regimes. Web2 compressed coordination via centralized platforms while accumulating informational power and lock-ins; Web3 promises shared verifiability and programmable rights yet introduces project- and use-frictions that must be anticipated and priced. An updated TCE turns the strategic question from “decentralize, yes or no?” into “how much decentralization, where to anchor trust, and how to keep adaptation options without dissipating the gains from automation and auditability?” Only within this engineering of transaction costs can the Web3 promise mature into sustainable practice, consistent with TCE’s core insights and the evolving needs of platform economies and global value chains.

Bibliography

- Abeyratne, S. A., & Monfared, R. P. (2016). Blockchain-ready manufacturing supply chain using distributed ledger. International Journal of Research in Engineering and Technology, 5(9), 1–10.
- Ahluwalia, S., Mahto, R. V., & Guerrero, M. (2020). Blockchain technology and startup financing: A transaction cost economics perspective. Technological Forecasting and Social Change, 151, 119854.
- Arruñada, B. (2020). Prospects of blockchain in contract and property relations: A transaction cost theory perspective. Journal of Purchasing and Supply Management, 26(2).
- Avarello, G. (2021). The impact of the decentralized economy of the Web3 on business strategies (Master's thesis, Università di Roma Tor Vergata).
- Bajari, P., & Tadelis, S. (2001). Incentives versus transaction costs: A theory of procurement contracts. RAND Journal of Economics, 32(3), 387–407.
- Bank of Italy. (2023). Blockchain and smart contracts: Governance and regulatory insights. Bank of Italy Reports.
- Baron, D. P. (1982). Measurement costs and the organization of markets. Journal of Law and Economics, 25(1), 27–48.
- Barthélemy, J., & Quélin, B. V. (2006). Complexity of outsourcing contracts and ex post transaction costs: An empirical investigation. Journal of Management Studies, 43(8), 1775–1797.
- Barzel, Y. (1982). Measurement cost and the organization of markets. Journal of Law and Economics, 25(1), 27–48.
- Berners-Lee, T. (1999). Weaving the Web: The original design and ultimate destiny of the World Wide Web. HarperCollins.
- Brina, R. (2022). Smart contracts and the vulnerabilities of code: Legal and economic implications. Computer Law & Security Review, 47, 105726.
- Buterin, V. (2014). Ethereum: A next-generation smart contract and decentralized application platform.
- Castells, M. (2001). The Information Age: Economy, Society, and Culture (Vol. I). Blackwell Publishers.
- Catalini, C., & Gans, J. S. (2016). Some simple economics of the blockchain. SSRN Electronic Journal.
- Catalini, C., & Gans, J. S. (2020). The impact of blockchain technology on transaction costs. Review of Economic Studies, 87(6), 3582–3613.
- Christidis, K., & Devetsikiotis, M. (2016). Blockchains and smart contracts for the Internet of Things. IEEE Access, 4, 2292–2303.

- Clack, C. D., Bakshi, V. A., & Braine, L. (2016). Smart contract templates: Foundations, design landscape and research directions. Proceedings of the IEEE International Conference on Software Engineering.
- Coase, R. H. (1937). The nature of the firm. Economica, 4(16), 386–405.
- Cormode, G., & Krishnamurthy, B. (2008). Key differences between Web 1.0 and Web 2.0. First Monday, 13(6).
- Crocker, K. J., & Masten, S. E. (1991). Pretia ex machina? Prices and process in long-term contracts. Journal of Law, Economics, and Organization, 7(1), 1–33.
- Davidson, S., De Filippi, P., & Potts, J. (2018). Blockchains and the economic institutions of capitalism. Journal of Institutional Economics, 14(4), 639–658.
- De Filippi, P., & Wright, A. (2018). Blockchain and the law: The rule of code. Harvard University Press.
- DuPont, Q. (2020). Experiments in algorithmic governance: A history and ethnography of "The DAO." In M. Campbell-Verduyn (Ed.), Bitcoin and beyond: Cryptocurrencies, blockchains, and global governance (pp. 157–177). Routledge.
- Dyer, J. H. (1997). Effective interfirm collaboration: How firms minimize transaction costs and maximize transaction value. Strategic Management Journal, 18(7), 535–556.
- Farrell, J., & Klemperer, P. (2007). Coordination and lock-in: Competition with switching costs and network effects. In Handbook of Industrial Organization (Vol. 3, pp. 1967–2072). Elsevier.
- Fritsch, M., Müller, J., & Wattenhofer, R. (2024). On-chain governance and voting power concentration in DAOs. Journal of Blockchain Research, 2(1), 45–65.
- Fuchs, C. (2017). Social media: A critical introduction (2nd ed.). Sage.
- Gervais, A., Karame, G., Capkun, V., & Capkun, S. (2016). Is Bitcoin a decentralized currency? IEEE Security & Privacy, 12(3), 54–60.
- Gillies, J., & Cailliau, R. (2000). How the Web was born: The story of the World Wide Web. Oxford University Press.
- Grover, V., & Malhotra, M. K. (2003). Transaction cost framework in operations and supply chain management research. Journal of Operations Management, 21(4), 457–473.
- Gupta, V. (2017). The truth about blockchain. Harvard Business Review.
- Hassan, S., & De Filippi, P. (2021). DAOs as a service: The rise of decentralized autonomous organizations in digital governance. Journal of Internet Governance, 10(1), 45–59.
- Helmond, A. (2015). The platformization of the web: Making web data platform ready. Social Media + Society, 1(2).
- Hobbs, J. E. (1996). A transaction cost approach to supply chain management. Supply Chain Management, 1(2), 15–27.

- Hoffmeister, M., & Stossberger, C. (2018). Verifiability of on-chain smart contract execution: Legal and technical challenges. Journal of Information Technology Law, 30(1), 38–52.
- Huang, G. (2024). Scalability and interoperability in Web3 platforms: A survey. IEEE Transactions on Network and Service Management, 21(2), 1035–1046.
- Huang, J., Wu, J., & Su, J. (2022). Barriers to user adoption of decentralized applications: A usability perspective. Information Systems Frontiers, 24(3), 863–879.
- Joo, M. (2021). Walmart China's VeChain blockchain-based traceability system: A case study. Journal of Retailing and Consumer Services, 58, Article 102307.
- Joskow, P. L. (1987). Contract duration and relationship-specific investments: Empirical evidence from coal markets. American Economic Review, 77(1), 168–185.
- Kafeel, S. (2023). Challenges and implementation strategies of Web3 technologies. International Journal of Computer Science and Information Security, 21(3), 112–126.
- Klein, B., Crawford, R. A., & Alchian, A. A. (1978). Vertical integration, appropriable rents, and the competitive contracting process. Journal of Law and Economics, 21(2), 297–326.
- Law, C. C. H., & Chen, C. C. (2010). Post-implementation practices of ERP systems and their impact on performance. International Journal of Production Research, 48(6), 1655–1670.
- Leshinsky, R., & Junnila, A. (2025). Human capital and blockchain organizations. Journal of Organizational Change Management, 38(1), 24–42.
- Lih, A. (2009). The Wikipedia revolution: How a bunch of nobodies created the world's greatest encyclopedia. Hyperion.
- Lumineau, F., Wang, D., & Schilke, O. (2021). Adaptation in decentralized autonomous organizations: Evidence from blockchain governance. Academy of Management Journal, 64(3), 747–777.
- Macher, J. T., & Richman, B. D. (2008). Transaction cost economics: An assessment of empirical research in the social sciences. Business and Politics, 10(1), 1–63.
- Madden, M. (2017). Privacy, security, and digital inequality. Data & Society Research Institute.
- Manovich, L. (2001). The language of new media. MIT Press.
- Malone, T. W., Yates, J., & Benjamin, R. I. (1987). Electronic markets and electronic hierarchies. Communications of the ACM, 30(6), 484–497.
- Ménard, C. (2004). The economics of hybrid organizations. Journal of Institutional and Theoretical Economics, 160(3), 345–376.
- Mhaskey, G. (2024). ERP–SCM integration and organizational alignment. Journal of Enterprise Technology, 17(2), 89–101.

- Momtaz, P. P. (2022). Decentralized finance (DeFi) and the financial system. Journal of Financial Economics, 145(2), 129–146.
- Mora, C., et al. (2018). Bitcoin emissions alone could push global warming above 2°C. Nature Climate Change, 8, 931–933.
- Nah, F. F.-H., Faja, S., & Cata, T. (2001). Characteristics of ERP software maintenance: A multiple case study. Journal of Software Maintenance, 13(6), 399–414.
- Nickerson, J. A., & Silverman, B. S. (2003). Why firms want to organize efficiently: Asset specificity and vertical integration. Administrative Science Quarterly, 48(3), 433–465.
- Nofer, M., Gomber, P., Hinz, O., & Schiereck, D. (2017). Blockchain. Business & Information Systems Engineering, 59(3), 183–187.
- North, D. C. (1990). Institutions, Institutional Change and Economic Performance. Cambridge University Press.
- O'Reilly, T. (2005). What is Web 2.0: Design patterns and business models for the next generation of software. O'Reilly Media.
- Oracle. (2024). Oracle Fusion Cloud SCM documentation and integration guides. Oracle Corporation. https://docs.oracle.com/en/cloud/saas/supply-chain-and-manufacturing/24d/index.html
- Oracle. (2025a). Oracle Cloud Infrastructure and Fusion Applications overview. Oracle Corporation. https://www.oracle.com/applications/cloud-apps-on-oci/
- Orabi, N., Emam, A., & Fahmy, A. (2025). Multi-party data management using blockchain: Challenges and solutions. Data & Knowledge Engineering, 140, Article 102262.
- Pariser, E. (2011). The filter bubble: What the Internet is hiding from you. Penguin Press.***
- Poppo, L., & Zenger, T. (2002). Do formal contracts and relational governance function as substitutes or complements? Strategic Management Journal, 23(8), 707–725.
- Queiroz, M. M., & Wamba, S. F. (2019). Blockchain adoption challenges in supply chain management: An empirical study. International Journal of Production Economics, 208, 87–98.
- Rajapakse, I. (2023). Organizational governance in digital supply chains. Journal of Supply Chain Research, 12(1), 45–67.
- Reijers, W. (2021). Governance of decentralized autonomous organizations. Journal of General Management, 46(3), 179–188.
- Rindfleisch, A., & Heide, J. B. (1997). Transaction cost analysis: Past, present, and future applications. Journal of Marketing, 61(4), 30–54.
- Saari, T., et al. (2025). Cloud ERP benefits and adoption challenges: A systematic review. Information Systems Review, 40(1), 55–84.

- Schär, F. (2021). Decentralized finance: On blockchain- and smart contract-based financial markets. Federal Reserve Bank of St. Louis Review, 103(2), 153–174.
- Schmidt, C. G., & Wagner, S. M. (2019). Blockchain and supply chain relations: A transaction cost theory perspective. International Journal of Production Economics, 219, 230–241.
- Schröder, J. (2021). Decentralized finance and new business models of Web3. Journal of Web Economics, 12(4), 134–149.
- Shelanski, H. A., & Klein, P. G. (1995). Empirical research in transaction cost economics: A review and assessment. Journal of Law, Economics, and Organization, 11(2), 335–361.
- Simon, H. A. (1957). Models of Man: Social and Rational. Wiley.
- Vatiero, M. (2022). Institutional complementarities and adaptation costs in blockchain-based organizations. European Journal of Law and Economics, 54(2), 205–227.
- VeChain Foundation. (2020). VeChainThor consensus and architecture whitepaper. VeChain Foundation. https://www.vechain.org/assets/whitepaper/whitepaper-1-0.pdf
- VeChain Foundation. (2022). PoA 2.0 SURFACE update. VeChain Foundation. https://www.vechain.org/assets/whitepaper/whitepaper-3-0.pdf
- Wang, Y. (2025). Evaluating supply chain finance costs with blockchain technology. Journal of Financial Economics, 142(1), 105–124.
- Weking, J., Hein, A., Böhm, M., & Krcmar, H. (2020). A hierarchical taxonomy of business model patterns. Electronic Markets, 30, 447–468.
- Williamson, O. E. (1985). The economic institutions of capitalism: Firms, markets, relational contracting. Free Press.
- Williamson, O. E. (1991). Comparative economic organization: The analysis of discrete structural alternatives. Administrative Science Quarterly, 36(2), 269–296.
- Williamson, O. E. (2002). The theory of the firm as governance structure: From choice to contract. Journal of Economic Perspectives, 16(3), 171–195.
- Wood, G. (2014). Ethereum: A secure decentralised generalised transaction ledger (Yellow Paper, Version 2).
- Xu, X., Weber, I., & Staples, M. (2019). Architecture for blockchain applications. Springer.
- Zuboff, S. (2019). The age of surveillance capitalism: The fight for a human future at the new frontier of power. PublicAffairs.