

Sustainable Manufacturing through AI-Governed Capacity Sharing and Cooperative Resource Utilization

Ningyang Luo

Department of Computer Science, University of Houston, Houston, TX, USA.

ningyangluo91@uh.edu

Abstract

The transition toward sustainable manufacturing demands a systemic departure from isolated, overbuilt production capacities toward orchestrated, cooperative resource networks. This paper presents a comprehensive analysis of how artificial intelligence (AI) can govern capacity sharing across heterogeneous manufacturing firms, enabling collaborative utilization of machinery, logistics assets, and human skill pools while advancing environmental and economic sustainability goals. We examine the structural trade-offs between centralized coordination and decentralized autonomy, and argue that AI-driven governance architectures must balance efficiency with fairness, transparency, and long-term relational trust. The discussion spans multi-agent system design, trust and reciprocity dynamics, incentive alignment through smart contracts, algorithmic fairness, and the integration of digital twins with industrial data spaces. Particular attention is given to how AI governance can incorporate behavioral insights, ensuring that capacity-sharing ecosystems remain robust against strategic manipulation, supply disruptions, and evolving regulatory landscapes. The paper further explores deployment challenges at scale, the role of interoperability standards, and the policy frameworks necessary for fostering cross-sector circular manufacturing networks without undermining competition or data sovereignty. By synthesizing insights from operations management, artificial intelligence, industrial ecology, and socio-technical systems research, we outline a forward-looking research agenda that positions AI-governed capacity sharing as a foundational pillar of regenerative industrial systems.

Keywords

sustainable manufacturing, capacity sharing, cooperative resource utilization, artificial intelligence, governance, circular economy, industrial symbiosis.

1. Introduction

Contemporary manufacturing systems face intensifying pressure to decouple value creation from resource consumption while maintaining supply chain viability in an era of heightened volatility. The traditional model of building dedicated, firm-owned production capacity to meet peak demand leads to chronic underutilization, redundant capital investments, and excessive environmental footprints. As regulatory mandates tighten and stakeholder expectations shift toward net-zero operations, the search for structural alternatives has directed scholarly and industrial attention toward capacity sharing and cooperative resource utilization paradigms. Industrial symbiosis and the sharing economy have long demonstrated the potential of inter-firm collaboration to reduce material and energy intensity, but scaling such cooperation beyond bilateral, trust-intensive relationships has proven exceedingly difficult. The emergence of advanced artificial intelligence (AI) technologies, coupled with

ubiquitous sensing and secure data exchange layers, now opens the possibility of governing dynamic, multi-party capacity-sharing networks in real time. This paper investigates how AI can serve not merely as an optimization engine but as the governance backbone for sustainable manufacturing ecosystems that coordinate production assets across organizational boundaries.

The problem is fundamentally socio-technical. On the one hand, technical developments in cloud manufacturing, digital twins, and industrial internet-of-things platforms have made it feasible to expose and trade idle capacity as a service. On the other hand, the viability of such ecosystems depends on behavioral and institutional factors: firms must trust that their proprietary data will not be exploited, that the reciprocal obligations will be honored under uncertainty, and that the allocation mechanisms will not systematically disadvantage certain participants. Recent research in sustainable manufacturing has emphasized the need to transcend efficiency-centric views and incorporate lifecycle impacts, circular strategies, and social dimensions into system design [1]. The industrial symbiosis literature similarly highlights that the emergence of synergistic exchanges is path-dependent and contingent on relational capital, geographic proximity, and information transparency [2]. These findings suggest that AI governance models must be architected to nurture cooperative relationships rather than simply extract short-term transactional gains.

The contribution of this paper is a multilevel analysis of AI-governed capacity sharing that connects architectural choices, governance mechanisms, infrastructure requirements, resilience considerations, and policy implications. We integrate perspectives from several disciplines, showing how cloud-based manufacturing-as-a-service platforms [3] can evolve into self-governing ecosystems when augmented with decentralized trust layers, multi-agent negotiation protocols, and adaptive incentive structures. While the sharing economy has demonstrated in consumer domains that peer-to-peer access can dramatically improve asset utilization [4], the industrial context introduces distinct constraints related to safety, quality assurance, intellectual property, and long planning horizons. The application of AI across supply chains is rapidly maturing, with predictive analytics, digital control towers, and autonomous planning agents showing tangible operational benefits [5], yet the governance of resource pooling across competing entities remains an under-explored frontier. Accordingly, we structure the analysis to provide a system-level perspective on how AI-driven capacity sharing can become an integral enabler of the circular economy in manufacturing [6].

2. The Sustainability Imperative and the Structural Logic of Capacity Sharing

The sustainability pressures acting on manufacturing are multidimensional, encompassing climate mitigation, resource scarcity, circular material flows, and social equity. Conventional responses have largely been incremental, focusing on energy efficiency, waste reduction, and renewable sourcing within individual facilities. While important, these measures fail to address the systemic overcapacity that characterizes many sectors. Firms routinely invest in production assets sized for peak demand plus safety margins, leading to average utilization rates that seldom exceed 60% in discrete manufacturing and can fall below 30% in process industries with pronounced seasonality. This idle capacity represents embedded carbon, locked-in capital, and lost opportunities for smaller enterprises that lack access to high-end equipment. Sharing capacity across firms can raise aggregate utilization, defer new capital expenditures, and enable smaller players to compete on quality without bearing the full fixed-cost burden. The sustainability logic thus extends beyond resource efficiency to encompass industrial inclusivity and resilience through diversification.

Digital twin technologies have reached a maturity that enables real-time, high-fidelity representations of production assets, processes, and constraints [7]. When coupled with AI-driven orchestration, these digital twins can act as programmable interfaces through which firms advertise available capacity, specify technical requirements, and negotiate terms automatically. Yet the mere availability of technology does not spontaneously generate cooperative behavior. The history of industrial symbiosis illustrates that even economically compelling exchanges often fail due to misaligned incentives, information asymmetry, and fears of dependency. A manufacturing SME may hesitate to rely on a larger competitor's milling capacity even when it is demonstrably underused, fearing that the partner could prioritize its own orders in times of demand spikes or extract proprietary process knowledge. Such concerns underscore the need for an AI governance layer that actively manages relational dynamics and not only instantaneous resource matching.

From a systems perspective, capacity sharing can be conceptualized as a cooperative resource pooling problem where the collective welfare gains must be distributed in a manner perceived as fair to sustain participation. The operational management of shared capacity involves decisions about scheduling, prioritization, quality assurance, and conflict resolution that are far more complex than in single-firm settings. AI can handle this complexity through learning-based optimization and negotiation, but the governance framework must define the rules under which these algorithms operate. The structural choice between a centralized clearinghouse and a decentralized peer-to-peer network has profound implications for scalability, trust distribution, and regulatory compliance. A fully centralized platform, while technically simpler to optimize, becomes a single point of control and potential failure that may deter firms concerned about data concentration. Decentralized architectures built on multi-agent systems and distributed ledger technologies can reduce the need for a central authority, but they introduce coordination overhead and may struggle to guarantee globally efficient outcomes. The governance design space therefore involves navigating trade-offs between efficiency, autonomy, and accountability.

3. Architectural Foundations for AI-Governed Capacity-Sharing Ecosystems

An AI-governed capacity-sharing ecosystem can be understood as a layered architecture comprising physical assets, their digital representations, a coordination intelligence layer, and a governance infrastructure that encodes rules, sanctions, and dispute resolution mechanisms. The physical layer includes production machinery, logistics fleets, testing laboratories, and skilled human operators whose availability can be shared under predefined service-level agreements. The digital representation layer relies on digital twins and industrial internet-of-things streams to capture asset status, maintenance schedules, process capabilities, and real-time deviations. The coordination intelligence layer hosts the AI agents responsible for matchmaking, scheduling, dynamic pricing, and anomaly detection. The governance layer oversees identity management, access rights, contract enforcement, and performance auditing, often leveraging programmable smart contracts and decentralized identity frameworks. This layered model separates concerns and allows independent evolution of each stratum, yet it demands careful integration to ensure that the high-level governance rules can effectively constrain the learning and decision-making processes in the coordination layer.

A foundational design decision concerns the degree of autonomy granted to AI agents that act on behalf of participating firms. In lightly governed systems, each firm deploys its own agent that negotiates bilaterally, relying on the governance layer only for dispute resolution *ex post*. This agent-based manufacturing paradigm has been studied extensively, with multi-agent

systems demonstrating emergent coordination and adaptive scheduling in shop floors and supply chains [8]. However, scaling agent negotiations to large ecosystems without a coordinating intelligence can lead to unstable bidding cycles, strategic misrepresentation of capacities, and deadlocks when agents pursue locally optimal strategies. Hybrid architectures that embed a lightweight supervisory AI, perhaps operated by a neutral industry consortium, can guide the market toward efficient equilibria while preserving firm autonomy for tactical decisions. Such supervisory agents can compute reference prices, suggest coalitional structures, and detect anticompetitive behaviors without taking over execution authority. The governance implications are significant: the supervisory agent must be transparent, auditable, and governed by a clear mandate that prevents it from favoring any coalition of participants.

The integration of blockchain technology introduces a credible transparency layer that can underpin trust in decentralized settings. A permissioned distributed ledger can record capacity commitments, delivery confirmations, quality inspections, and financial settlements in a tamper-evident manner, significantly lowering the verification costs for all parties [9]. When combined with tokenized incentives, the ledger can also facilitate novel reward mechanisms that compensate firms not only for providing capacity but also for contributing to the ecosystem's environmental performance, such as reducing the collective carbon footprint through load shifting. Nevertheless, the computational overhead, latency, and governance of the blockchain consensus itself must be carefully aligned with the real-time requirements of manufacturing coordination. Layered approaches that use blockchain sparingly for critical auditing and leave high-frequency matching to off-chain AI engines are emerging as a practical compromise.

4. Governance Mechanisms: Trust, Reciprocity, and Algorithmic Fairness

The long-term viability of capacity-sharing ecosystems depends fundamentally on the establishment and maintenance of trust among participants. While technology can reduce the need for interpersonal trust by providing cryptographic guarantees and transparent records, it cannot eliminate the behavioral dimensions of repeated interactions. Empirical insights from operations management demonstrate that trust and reciprocity are powerful drivers of firms' willingness to share capacity, often outweighing purely economic considerations in longitudinal relationships [10]. This finding suggests that AI governance mechanisms should be designed not as static rule engines but as adaptive systems that recognize and nurture relational norms. For instance, an AI orchestrator can track the history of reciprocal favors and adjust priority weights or pricing tiers to reward firms that have reliably honored their commitments during past disruptions. Such a reputation system, grounded in behavioral economic principles, can sustain cooperation even in environments where formal contracts are incomplete.

Incentive alignment remains a central challenge. Cooperative game-theoretic models have shown that while capacity pooling can generate significant total surplus, the distribution of gains must belong to the core of the cooperative game to prevent sub-coalitions from seceding [11]. AI governance can operationalize these theoretical insights by embedding cost-and-surplus allocation algorithms that compute Shapley-value-based or nucleolus-based distributions and enforce them through smart contracts. However, algorithmic allocation raises profound fairness questions. If the AI inadvertently produces allocations that systematically disadvantage certain firms—for example, smaller enterprises with less bargaining power or those located in higher-cost regions—the ecosystem can unravel despite aggregate efficiency. Algorithmic fairness in capacity sharing must consider multiple metrics,

including individual fairness (treating similar entities similarly), group fairness (avoiding disparate impact on protected categories), and procedural fairness (explainability of decisions). The ethics of algorithmic decision-making in sociotechnical systems has been extensively debated, and the lessons about opacity, bias amplification, and accountability apply directly to manufacturing governance [12]. An AI system that denies a small manufacturer access to a critical resource without a comprehensible justification can destroy the trust that the ecosystem relies upon.

Reciprocity mechanisms can be formalized through tokenized credit systems that allow firms to earn priority tokens by releasing capacity during peak community demand or by participating in environmental improvement projects. These tokens can then be redeemed for preferential access or fee discounts during their own demand surges. The AI governance layer manages the issuance, valuation, and redemption of such tokens while preventing inflationary manipulation and gaming. This token economy adds a temporal smoothing dimension that complements instantaneous market-clearing mechanisms. Firms that might otherwise exit the ecosystem during periods of sustained self-sufficiency are incentivized to remain engaged because the tokens hold future option value, thereby preserving the network's collective capacity buffer and enhancing its sustainability by avoiding the need for each firm to maintain its own costly reserve. The design of these token economies must comply with regulatory frameworks concerning digital assets and taxation, adding another layer of complexity that the governance infrastructure must address.

5. Infrastructure, Interoperability, and Deployment Considerations

Deploying AI-governed capacity sharing at scale demands a robust digital infrastructure that extends far beyond individual factory IT systems. The Industry 4.0 vision provides foundational elements such as interoperable communication protocols, edge computing resources, and standardized data models [13], yet full inter-firm sharing requires a higher degree of semantic and syntactic interoperability. Manufacturing firms use heterogeneous machine controllers, enterprise resource planning systems, and process specifications that are often deeply customized. An AI orchestrator cannot optimize shared schedules if the meaning of “available capacity” differs across partners due to divergent definitions of setup times, quality tolerances, or maintenance requirements. Industrial data spaces and reference architecture models have emerged to address this gap by defining common ontologies, data sovereignty rules, and connectors that allow firms to share data without relinquishing control over its subsequent use. These frameworks, initially developed in European research initiatives, provide a governance foundation that complements the AI coordination layer by ensuring that data usage policies are negotiated and enforced automatically.

The integration of digital twins into such data spaces creates a federated simulation environment where the AI can explore capacity-sharing scenarios without moving actual materials. A digital twin of a shared CNC machining cell, for example, can expose its process envelope, current job queue, and predicted wear status, while the AI scheduler can submit simulated job requests to evaluate feasibility and cycle-time impacts. This simulation-based negotiation reduces the risk of disruptive real-world experiments and enables the AI to learn more rapidly from synthetic yet high-fidelity data. However, the fidelity of such simulations hinges on the willingness of asset owners to update their digital twins honestly. The governance framework must therefore incorporate mechanisms to detect and penalize misrepresentation, such as statistical discrepancy checks between simulated and actual performance, supported by the distributed ledger audit trail. In effect, the digital twin becomes

a governed object whose promises are legally and reputationally binding, transforming simulation into a form of algorithmic contracting.

Scalability concerns also arise from the computational complexity of multi-objective optimization over large, dynamic coalitions. Centralized AI planners may become bottlenecks as the number of participating firms and shared asset types grows. Hierarchical and federated learning architectures offer a pathway to scalability by decomposing the decision problem along regional or sectoral boundaries while allowing higher-level coordinators to manage cross-boundary exchanges. For instance, a cluster of automotive suppliers in a given industrial park might self-organize under a local AI coordinator, which then interfaces with a national logistics-sharing network. This nesting of governance scopes aligns with the principle of subsidiarity and can improve responsiveness to local disruptions while maintaining global objectives. The resulting system-of-systems architecture demands clear interface contracts and conflict resolution protocols between coordinators, a topic that remains under-researched in the manufacturing context.

6. Robustness, Resilience, and Risk Management

The robustness of an AI-governed capacity-sharing ecosystem must be evaluated against a spectrum of disturbances, from sudden demand spikes and machine breakdowns to cyberattacks and the withdrawal of key participants. Resilience in supply chain thinking emphasizes the capacity to anticipate, absorb, and adapt to disruptions [14]. When capacity is pooled, the network inherits a structural redundancy that can enhance resilience, as idle capacity from one firm can substitute for a failed asset elsewhere. However, this same interdependence introduces new fragility: if the AI orchestrator or the underlying digital infrastructure fails, all participants may lose access to shared resources simultaneously. Therefore, the governance architecture must incorporate graceful degradation modes, allowing bilateral fallback coordination if the central AI becomes unavailable. This can be implemented through pre-negotiated standby contracts and locally cached decision models that enable agents to operate in a degraded, semi-autonomous fashion.

Cybersecurity in shared manufacturing networks is a paramount concern. The digital interfaces that expose capacity data also expand the attack surface for industrial espionage and sabotage. AI governance can contribute to defense-in-depth strategies by deploying anomaly detection models that learn normal traffic and scheduling patterns and flag deviations that could indicate a compromised node. Importantly, these detection models themselves must be protected against adversarial manipulation, wherein a malicious insider gradually shifts the baseline to disguise unauthorized data exfiltration. Incorporating game-theoretic threat modeling and robust learning techniques into the governance AI layer adds another dimension of complexity, requiring interdisciplinary collaboration between manufacturing engineers, data scientists, and security specialists.

The risk of cascading failures is particularly acute when shared capacity is concentrated in a few megasuppliers. The AI scheduler's efficiency objective might inadvertently drive consolidation, leading to a high-risk dependency graph that resembles the canonical vulnerabilities of global supply chains. Governing for resilience thus implies encoding diversity constraints and anti-concentration rules into the AI's decision logic, even at the cost of short-term efficiency. These constraints can be formulated as sustainability-oriented guardrails that limit the maximum share of total community demand that any single firm can supply through the platform. Such constraints will likely provoke pushback from dominant players who benefit from scale, highlighting the political economy of cooperative governance.

The AI must therefore balance the technically desirable with the politically feasible, a task that requires not only algorithmic sophistication but also deliberative stakeholder processes that define acceptable risk parameters.

7. Policy Implications and Socio-Technical Integration

The emergence of AI-governed capacity-sharing ecosystems raises novel policy questions that intersect industrial strategy, competition law, data governance, and sustainability regulation. Competition authorities may view information sharing and coordinated capacity scheduling among ostensibly independent firms as potential collusion risks, even when the intent is pro-competitive and environmentally beneficial. Antitrust safe harbors or regulatory sandboxes may be needed to allow experimentation without fear of legal sanction, provided the governance mechanisms include transparency and non-discrimination safeguards. Policy frameworks for the circular economy already promote extended producer responsibility and material loops [16], but they rarely address the operational enablers of resource sharing. Aligning capacity-sharing incentives with circular policy objectives, for instance by recognizing avoided capital expenditures in sustainability reporting standards, could accelerate adoption. Sustainability assessment methodologies that capture lifecycle impacts and social dimensions [15] can be encoded into the AI governance layer to automate triple-bottom-line reporting for shared resources, thereby reducing the administrative burden on participating firms.

The socio-technical integration of AI governance must also address workforce implications. Capacity sharing can shift demand for skilled labor across firms and regions, potentially displacing workers if not managed proactively. On the other hand, it can create new opportunities for specialized service providers and for smaller firms to access advanced manufacturing capabilities without massive capital outlays, potentially revitalizing local manufacturing ecosystems. The governance framework should incorporate mechanisms to monitor labor impacts and, where appropriate, link capacity-sharing privileges to commitments on workforce retraining or community benefit agreements. This extends the notion of sustainability beyond environmental metrics to include social and economic dimensions that are deeply intertwined with industrial resilience.

Interoperability standards and data governance models will likely require public-private coordination to prevent fragmentation. The experience with industrial data spaces demonstrates that multi-stakeholder governance bodies, involving manufacturers, technology providers, trade unions, and regulators, are essential to define the rulebooks that AI systems will enforce [17]. Such rulebooks must address data provenance, usage control, liability allocation when AI decisions cause harm, and the portability of reputation scores across platforms. Without these institutional foundations, the technological sophistication of AI governance will remain confined to pilot projects and will fail to scale across sectors. Extended producer responsibility schemes could be expanded to mandate participation in shared-capacity marketplaces for certain product categories, creating a regulatory pull that complements technology push. This would require careful design to avoid forcing small enterprises into platforms that do not serve their interests, underscoring once more the centrality of fairness and participatory governance [19].

8. Conclusion

The pursuit of sustainable manufacturing compels a fundamental rethinking of how production capacity is owned, utilized, and governed. AI-governed capacity sharing

represents a paradigmatic shift from rigid, firm-centric resource models toward fluid, ecosystem-oriented architectures that can simultaneously improve economic efficiency, environmental performance, and social inclusiveness. Throughout this paper, we have argued that the technical capacity for such sharing is already emerging through digital twins, multi-agent systems, distributed ledgers, and advanced optimization, but that its transformative potential hinges on governance designs that cultivate trust, embed fairness, and align incentives over the long term. The structural trade-offs between centralized coordination and decentralized autonomy, between instantaneous market efficiency and resilient redundancy, and between data transparency and sovereignty must be navigated with careful attention to behavioral realities and institutional contexts. The integration of behavioral insights, as demonstrated by the importance of trust and reciprocity in capacity-sharing relationships, highlights that AI governance cannot be reduced to a purely computational problem but must function as a socio-technical regulatory system in its own right.

Looking forward, research must advance on several fronts. The formal specification of governance rulebooks that are simultaneously machine-readable, legally enforceable, and stakeholder-legitimate remains an open challenge. The development of AI algorithms that can provably satisfy fairness constraints in dynamic, multi-objective settings while retaining efficiency and adaptability is a rich area for computational innovation. Empirical studies that examine the longitudinal evolution of trust and participation in real-world capacity-sharing networks will provide essential calibration for simulation models and policy design. Finally, the co-evolution of AI governance with sustainability-oriented industrial policies demands sustained interdisciplinary dialogue that brings together manufacturing science, computer science, law, and political economy. The vision of regenerative industrial systems, in which shared resources are not merely consumed but stewarded, is within reach if the governance architectures match the ambition of the technology.

References

1. Garetti, M., & Taisch, M. (2012). Sustainable manufacturing: trends and research challenges. *Production Planning & Control*, 23(2-3), 83-104.
2. Chertow, M. R. (2000). Industrial symbiosis: literature and taxonomy. *Annual Review of Energy and the Environment*, 25(1), 313-337.
3. Xu, X. (2012). From cloud computing to cloud manufacturing. *Robotics and Computer-Integrated Manufacturing*, 28(1), 75-86.
4. Benjaafar, S., Kong, G., Li, X., & Courcoubetis, C. (2019). Peer-to-peer product sharing: Implications for ownership, usage, and social welfare in the sharing economy. *Management Science*, 65(2), 477-493.
5. Toorajipour, R., Sohrabpour, V., Nazarpour, A., Oghazi, P., & Fischl, M. (2021). Artificial intelligence in supply chain management: A systematic literature review. *Journal of Business Research*, 122, 502-517.
6. Lieder, M., & Rashid, A. (2016). Towards circular economy implementation: a comprehensive review in context of manufacturing industry. *Journal of Cleaner Production*, 115, 36-51.
7. Tao, F., Zhang, H., Liu, A., & Nee, A. Y. C. (2019). Digital twin-driven product design, manufacturing and service with big data. *International Journal of Advanced Manufacturing Technology*, 94, 3563-3576.

8. Monostori, L., Váncza, J., & Kumara, S. R. T. (2006). Agent-based systems for manufacturing. *CIRP Annals*, 55(2), 697-720.
9. Saberi, S., Kouhizadeh, M., Sarkis, J., & Shen, L. (2019). Blockchain technology and its relationships to sustainable supply chain management. *International Journal of Production Research*, 57(7), 2117-2135.
10. Hu, X., & Caldentey, R. (2023). Trust and reciprocity in firms' capacity sharing. *Manufacturing & Service Operations Management*, 25(4), 1436-1450.
11. Körpeoğlu, E., & Cho, S. H. (2018). Incentives for cooperation in supply chains with capacity sharing. *Manufacturing & Service Operations Management*, 20(3), 478-494.
12. Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2), 2053951716679679.
13. Lasi, H., Fettke, P., Kemper, H. G., Feld, T., & Hoffmann, M. (2014). Industry 4.0. *Business & Information Systems Engineering*, 6(4), 239-242.
14. Christopher, M., & Peck, H. (2004). Building the resilient supply chain. *International Journal of Logistics Management*, 15(2), 1-14.
15. Singh, R. K., Murty, H. R., Gupta, S. K., & Dikshit, A. K. (2009). An overview of sustainability assessment methodologies. *Ecological Indicators*, 9(2), 189-212.
16. Kirchherr, J., Reike, D., & Hekkert, M. (2017). Conceptualizing the circular economy: An analysis of 114 definitions. *Resources, Conservation and Recycling*, 127, 221-232.
17. Otto, B., & Jarke, M. (2019). Designing a multi-sided data platform: findings from the International Data Spaces case. *Electronic Markets*, 29(4), 561-580.
18. Hauschild, M. Z., Kara, S., & Röpke, I. (2018). Absolute sustainability: Challenges for life cycle assessment. *CIRP Annals*, 67(2), 681-698.
19. Lindhqvist, T. (2000). Extended producer responsibility in cleaner production: Policy principle to promote environmental improvements of product systems. Lund University.