Game-Theoretic Capacity Pooling under AI-Predicted Supply Chain Disruptions

Enzo Andersson
School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR, USA.
enzo1979@oregonstate.edu

Abstract

The growing frequency and severity of supply chain disruptions driven by geopolitical turbulence, climate change, and pandemics compel industrial systems to move beyond firm-level redundancy toward collaborative capacity pooling. This paper presents a system-level analysis of how artificial intelligence-based disruption forecasting can be fused with game-theoretic mechanisms to design stable, efficient, and fair capacity sharing arrangements across independent enterprises. We examine the structural trade-offs inherent in architectures that translate probabilistic AI predictions into cooperative capacity allocation rules, highlighting the interplay between predictive performance, incentive compatibility, strategic behavior, and operational robustness. The discussion develops a layered view of infrastructure, encompassing data integration fabrics, AI prediction services, game-theoretic coordination engines, and smart contract-based settlement layers. Key challenges are unpacked, including the governance of asymmetric information, the calibration of fairness criteria under uncertainty, adversarial manipulation of prediction models, and regulatory constraints on cross-border data sharing. We contrast formal mechanism design approaches with trust-based relational contracts, identifying complementarities that are essential for sustaining capacity pools over extended time horizons. The paper further considers the implications of such systems for sustainability, industrial policy, and the reconfiguration of global supply networks toward collective resilience. By integrating perspectives from operations management, machine learning, economic mechanism design, and digital infrastructure, the analysis provides a road map for building AI-mediated capacity pooling platforms that do not merely optimize short-term cost but also strengthen long-term systemic robustness and equity.

Keywords

capacity pooling, supply chain disruptions, game theory, artificial intelligence, mechanism design, resilience, governance.

1. Introduction

Global supply chains have entered an era of polycrisis in which overlapping shocks threaten the smooth flow of materials, components, and finished goods. The COVID-19 pandemic, the blockage of the Suez Canal, extreme weather events, and armed conflicts have revealed deep fragilities in extended production and logistics networks, exposing the limits of lean, single-source strategies and prompting a fundamental rethink of resilience paradigms. In response, both industry and academia have increasingly turned their attention to cooperative capacity pooling, wherein independent firms share warehousing space, transportation assets, manufacturing lines, or inventory buffers to absorb shocks more efficiently than any single organization could achieve alone. Such pooling arrangements promise substantial efficiency gains by leveraging statistical economies of scale across a larger risk pool, yet they introduce intricate challenges of trust, strategic misrepresentation, and equitable distribution of benefits. Simultaneously, the rapid maturation of artificial intelligence (AI) offers new capabilities for anticipating disruptions before they propagate through the network. Machine learning models trained on satellite imagery, news feeds, port congestion data, and historical incident databases can generate probabilistic forecasts of disruption timing, location, and severity, enabling proactive capacity reallocation decisions that exceed human reactive planning. The conjunction of these two developments, cooperative capacity pooling and AI-driven predictive intelligence, creates a novel space for system design that is both promising and poorly understood. This article critically examines the structural, behavioral, and policy dimensions of such systems through the lens of game-theoretic design, without resorting to mathematical formalism, focusing instead on the architectural trade-offs, governance mechanisms, and deployment challenges that will determine whether AI-mediated capacity pooling can evolve from theoretical curiosity to operational reality.

2. The Imperative of Capacity Pooling in an Era of Disruption

Modern supply chains are characterized by deep interdependencies among geographically dispersed nodes, which transform localized disruptions into cascading failures across multiple tiers. The academic literature on supply chain resilience has extensively documented how a disturbance at a single supplier can propagate via the ripple effect, generating inventory shortages, production stoppages, and demand amplification further downstream [1] [2]. Traditional risk mitigation strategies emphasize inventory buffers, dual sourcing, and flexible manufacturing, yet these approaches are individually costly and often fail to provide adequate coverage against low-probability, high-impact events. Capacity pooling represents a fundamentally different philosophy: rather than each firm independently maintaining surge capacity that remains idle under normal conditions, a collective pool is established in which spare capacity can be dynamically redirected to whichever network participant faces an imminent shortfall. This transformation turns fixed, dedicated redundancies into shared, variable resources, thus reducing the total idle capacity required to achieve a given level of system-wide service reliability.

The theoretical foundations of pooling in logistics and operations management build on the classic risk-pooling principle, which demonstrates that aggregating demands across locations reduces relative variability. Extended to capacity under disruption, the concept implies that a coalition of firms facing partially correlated disruption exposure can achieve a higher collective service level at lower aggregate cost than the sum of their independent efforts [3] [4]. However, moving from principle to practice requires solving a constellation of coordination problems. Capacity is heterogeneous, with different manufacturing technologies, geographic footprints, and lead-time characteristics; disruptions are partially private information, often detected earlier by one firm than by others; and the benefits of pooling are unevenly distributed across participants, creating persistent incentives for free-riding or strategic withdrawal. Early attempts at industrial symbiosis, such as shared logistics platforms and joint inventory consortia, illustrate both the potential gains and the fragility of collaborative models when not supported by robust institutional mechanisms. This fragility motivates the search for theoretically grounded allocation rules that can align self-interest with collective welfare, a search that naturally leads into the domain of game theory.

3. Artificial Intelligence for Predictive Disruption Intelligence

AI technologies are profoundly altering how supply chain actors perceive and respond to emerging disruptions. Whereas classical risk management relied on historical loss data and expert judgment, contemporary AI pipelines ingest heterogeneous, high-velocity data streams including natural language processing of news media, sentiment analysis of social platforms, satellite imagery for monitoring port congestion and factory activity, weather model outputs, and Internet of Things sensor data from containers and vehicles. These inputs feed deep learning architectures, graph neural networks, and hybrid models that detect weak signals and generate probabilistic forecasts of disruption occurrence, estimated duration, and spatial propagation paths [5] [6]. The output is not a deterministic alarm but a continuously updated distribution over future states, which can be communicated to decision-makers as likelihood scores, risk heatmaps, or scenario trees.

The integration of such predictive intelligence into capacity pooling systems substantially alters the cooperative game. When participants possess heterogeneous AI capabilities, information asymmetry becomes a first-order design concern. A firm with superior predictive models may detect an impending shortage of a critical component days before its partners, creating an opportunity to preemptively hoard shared capacity or to manipulate the pooling mechanism by misreporting private forecasts. Conversely, if all participants contribute data to a centralized prediction engine, they face the dual challenge of protecting commercially sensitive information and ensuring that model outputs are not deliberately corrupted by adversarial data injections. These tensions suggest that the technical quality of AI forecasts cannot be evaluated in isolation; it must be assessed within the broader socio-technical context of the strategic environment in which those forecasts are generated, shared, and acted upon. The coupling of AI with game-theoretic coordination mechanisms must therefore be designed to be robust not merely against statistical noise but against rational manipulation by agents who may benefit from distorting the information base on which collective decisions rest.

4. Game-Theoretic Mechanisms for Cooperative Capacity Sharing

Game theory provides a rigorous language for analyzing situations in which multiple self-interested agents must agree on joint actions and the distribution of resulting payoffs. In the context of capacity pooling, the cooperative game consists of a set of firms that can form coalitions to share resources, where each coalition achieves a certain value defined as the minimum expected cost of meeting demand in the face of uncertain disruptions, given the pooled capacity of its members. A core question is whether there exists an allocation of these coalitional gains that no subgroup would have an incentive to reject, a condition known as stability in the core. Early surveys of game-theoretic supply chain analysis established conceptual frameworks for thinking about coalition formation and profit allocation in joint production and inventory systems, laying groundwork that can be extended to disruption-driven capacity sharing [7] [8].

The design space of allocation mechanisms is rich. Approaches based on the Shapley value offer a normative principle in which each participant's share of the surplus is proportional to its average marginal contribution across all possible orderings of coalition formation. While the Shapley value possesses appealing axiomatic properties of fairness and efficiency, its computation in large, dynamically changing coalitions with stochastic disruptions is demanding, and its direct application may ignore strategic manipulation of the coalition formation process itself. Alternative market-based mechanisms, such as iterative auctions for capacity slots, allow participants to express their willingness to pay for access to pooled resources, thereby revealing private disruption risks indirectly through bids [9]. These auction designs can be augmented with AI-generated state predictions that condition the auction parameters, such as reserve prices and allocation limits, on the current estimated probability of disruption. Nonetheless, classical mechanism design confronts significant obstacles when transferred from the frictionless theoretical realm to the messy operational environment of supply chains: contracts are incomplete, monitoring is costly, and renegotiation is pervasive.

Because information sharing is both a prerequisite for effective pooling and a source of competitive vulnerability, confidentiality-preserving frameworks become indispensable. Techniques such as secure multiparty computation and differential privacy can, in principle, allow firms to compute optimal pooled capacity allocations without revealing their individual forecasts or cost structures, but they impose computational and latency overheads that may be incompatible with the real-time demands of crisis response [10]. A comprehensive architectural vision must therefore navigate the tension between the theoretical elegance of full-information mechanism design and the practical necessity of operating under partial, noisy, and strategically distorted information.

5. Architectural Design of an AI-Integrated Pooling Platform

Translating the conceptual frameworks of cooperative game theory and AI prediction into a functioning socio-technical infrastructure demands a layered platform architecture that spans data ingestion, model serving, mechanism execution, and settlement. At the lowest layer, a federated data fabric collects and normalizes operational data from participants' enterprise resource planning systems, transportation management systems, and external data providers while enforcing access controls, data provenance tracking, and usage auditing. This layer must reconcile heterogeneous schemas and time granularities, a non-trivial engineering challenge that has historically hindered interoperability across supply chain ecosystems.

Above the data layer resides the AI prediction service layer, which hosts a suite of models that compute disruption probabilities and impact forecasts. These models can be instantiated as a model hub, allowing participants to bring their own models while also providing a shared, consortium-governed baseline model that serves as a common reference. The tension between centralized and federated learning architectures is particularly acute here: centralized model training on pooled data yields the highest predictive accuracy but raises intellectual property and antitrust concerns, whereas federated learning preserves data locality at the cost of degraded model performance and increased vulnerability to model poisoning attacks. The architectural choice has far-reaching implications for the strategic behavior of participants, because a shared model that systematically underestimates certain types of disruption may induce underinvestment in capacity by the very firms that would suffer most.

The coordination layer implements the game-theoretic mechanism, translating AI-generated state information and participant-reported valuations or claims into capacity allocation decisions. This layer must be designed with strict modularity to allow the mechanism to be swapped or upgraded without destabilizing the entire platform. It must also incorporate robust optimization logic that accounts for forecast uncertainty by adopting minimax or distributionally robust decision rules, ensuring that allocations are not excessively sensitive to small perturbations in the AI predictions. Implementation via smart contracts on a distributed ledger offers automatic execution, tamper-proof audit trails, and the possibility of embedding dispute resolution routines directly into the transaction logic [11]. However, the energy consumption and scalability limitations of public blockchains motivate hybrid architectures in which only critical settlement events are recorded on-chain while off-chain channels handle high-frequency operational messages.

The topmost settlement and governance layer manages financial flows, performance guarantees, and membership rules. Ex-post settlement may involve not only monetary transfers but also capacity credits that are banked and drawn upon in future periods, creating a dynamic game that discourages short-term opportunism. The platform architecture as a whole must be viewed as an evolving institutional technology rather than a one-time engineering deployment; its long-term viability hinges on governance processes that can adapt rules in response to changing disruption patterns, coalition composition, and technological capabilities.

6. Governance, Trust, and Fairness in Capacity Alliances

While game-theoretic mechanisms provide a formal apparatus for aligning incentives, sustaining capacity pooling over time requires cultivating trust and perceptions of procedural fairness that formal contracts alone cannot guarantee. In real-world supply chain relationships, firms often observe each other's behavior over repeated interactions, and reputational concerns can substitute for or complement explicit contractual enforcement. Experimental studies have documented that capacity sharing is significantly influenced by trust and reciprocity, even when monetary incentives are held constant [14]. This finding suggests that the institutional design of AI-mediated pooling must move beyond the narrow assumption of purely selfish rational agents. Governance structures should create transparency regarding the rules of the game and the computations that drive allocations, while simultaneously preserving the confidentiality of individual firm data. Explainable AI techniques that can articulate, at a high level, why a particular disruption probability has been estimated and how it influenced the allocation of capacity, can enhance participants' confidence that the system is neither rigged nor arbitrary.

Fairness is a multidimensional construct in capacity pooling. Allocations that satisfy coalitional stability in the core may still be perceived as inequitable if one firm consistently benefits more than another due to structural asymmetries unrelated to effort or investment. The tension between marginal contribution-based allocation principles and need-based or egalitarian norms mirrors broader debates in distributive justice and informs the design of weighting schemes that can be adjusted through democratic governance processes. Moreover, the dynamic nature of disruption risk implies that fairness must be evaluated intertemporally: a firm that subsidizes another during a crisis may reasonably expect reciprocity in a future, inverse scenario. Mechanisms that encode such reciprocal obligations, such as capacity credit markets or rolling baseline adjustments, can embed long-term fairness without requiring explicit, state-contingent contracting for every possible future. The integration of trust-based relational contracts with formal game-theoretic mechanisms thus provides a more robust social foundation for capacity alliances than either approach could achieve alone.

7. Robustness, Security, and Adversarial Threats

Platforms that tightly couple AI predictions with automated coordination introduce novel attack surfaces that have no counterpart in traditional supply chain risk management. Adversarial machine learning research has demonstrated that small, carefully crafted perturbations to input data can cause deep neural networks to produce catastrophically incorrect outputs [15]. In a capacity pooling context, an attacker with access to the data pipeline could inject fraudulent sensor readings or fabricate news signals to manipulate the AI's disruption forecast, thereby inducing the mechanism to misallocate capacity in their favor. Such attacks are particularly dangerous because they can be executed remotely and may remain undetected if verification mechanisms are weak. Beyond external adversaries, participating firms themselves have incentives to game the system by misreporting their internal state, inflating their anticipated needs, or strategically downgrading the capacity they make available to the pool. Mechanism design theory offers a class of solutions under the rubric of incentive-compatible mechanisms in which truth-telling is a dominant strategy, but the conditions required for such mechanisms, such as transferable utility and no limits on side payments, are often violated in practice.

Robust mechanism design relaxes the strong common knowledge assumptions of classical Bayesian mechanism design and seeks allocation rules that perform adequately across a broad set of possible information structures and behavioral models. Applying this philosophy to AI-mediated capacity pooling implies designing the system to be resilient not only to statistical noise in predictions but also to strategic manipulation of those predictions. Technical countermeasures include the use of ensemble models that combine orthogonal data sources, anomaly detection on input streams, and deterministic fallback rules that activate when AI confidence drops below a threshold. Equally important are institutional safeguards such as post-event audits, penalty schemes for detected manipulation, and the ability of coalition members to collectively suspend or recalibrate the mechanism in response to suspected integrity breaches. These safeguards transform the platform from a purely algorithmic allocator into a sociotechnical immune system that can learn from adversarial encounters and adapt its defenses.

8. Policy, Regulation, and Cross-Border Coordination

The deployment of AI-mediated capacity pooling platforms at industrial scale inevitably intersects with a complex web of regulatory regimes. Competition authorities have long scrutinized horizontal collaboration among competitors, even when such collaboration promises efficiency gains, because of the risk of tacit collusion that could spill over from capacity sharing to price fixing or market allocation. Proponents must therefore design pooling arrangements with clear firewalls between operational coordination and strategic market conduct, a separation that is easier to articulate than to enforce in digitally integrated platforms where data flows are pervasive. Privacy regulations, particularly the European Union's General Data Protection Regulation and emerging AI governance frameworks, impose constraints on the cross-border transfer of data and on automated decision-making that significantly affects commercial interests. Capacity allocation algorithms that rely on firm-level proprietary data must incorporate privacy-by-design principles and be auditable by regulators without exposing sensitive trade secrets. Cross-border pooling consortia, which are especially valuable for global supply chains, must reconcile divergent national standards on data sovereignty, cybersecurity, and algorithmic accountability. International trade law and digital economy agreements are beginning to address these frictions, but the legal landscape remains fragmented [17]. Policy instruments such as regulatory sandboxes, where firms can test collaborative platforms under regulatory supervision, may accelerate learning about the real-world implications of game-theoretic capacity pooling while containing systemic risks. Governments can also act as conveners and trusted intermediaries, providing neutral data infrastructure and certification services that reduce the trust barriers to coalition formation.

9. Long-Term Sustainability and Societal Impact

Beyond the immediate operational benefits, AI-driven capacity pooling carries profound implications for environmental sustainability and social welfare. By reducing the aggregate idle capacity that must be maintained across an industry, pooling lowers the material footprint of supply chain infrastructure, curbing resource extraction, energy consumption, and waste associated with building and maintaining redundant warehouses, fleets, and production lines.

This dematerialization effect aligns with circular economy principles that prioritize utilization intensity over asset ownership. However, the sustainability calculus is not unambiguously positive; the digital infrastructure required for real-time AI coordination itself consumes substantial energy and rare-earth minerals, and the economic incentives created by pooling might, in some scenarios, encourage a race to the bottom in environmental standards as firms seek to minimize their contribution costs. An ethically designed platform must embed sustainability metrics into its allocation criteria, potentially by adjusting capacity credits based on verified carbon footprints or circularity scores.

From a societal perspective, capacity pooling can democratize resilience by allowing small and medium-sized enterprises to access the kind of surge protection that was previously available only to large multinationals with deep pockets. This can reduce supply chain concentration and enhance the robustness of regional economic ecosystems. Yet the benefits will not materialize automatically; they depend on deliberate choices about platform governance, membership criteria, and the pricing of data and computational services. If AI prediction capabilities and participation rights become concentrated among a few dominant firms, pooling could reinforce rather than counteract existing power asymmetries. Policy interventions that promote open data standards, interoperability mandates, and subsidized access for smaller actors can steer the evolution of capacity pooling toward a public-good orientation. The long-term vision is one in which supply chain capacity is managed as a shared digital commons, governed by transparent algorithms, accountable institutions, and a shared commitment to collective resilience, a shift that would represent as profound a transformation in industrial organization as the transition from vertical integration to outsourcing in the late twentieth century.

10. Conclusion

The convergence of AI-based disruption forecasting and game-theoretic capacity pooling opens a transformative pathway for supply chain resilience that moves beyond siloed firm-level hardening toward systemic collective intelligence. This paper has argued that realizing that potential requires a holistic system architecture that harmonizes predictive analytics, incentive-compatible allocation mechanisms, trust-building governance, and robust defense against strategic manipulation. The design space is replete with trade-offs: centralized versus federated model training, formal enforcement versus relational trust, allocative efficiency versus dynamic fairness, and rapid algorithmic response versus deliberate human oversight. Resolving these trade-offs demands not only technical ingenuity but also institutional innovation that spans antitrust, privacy, and industrial policy domains. Future research should deepen the empirical understanding of how human decision-makers interact with AI-generated predictions in cooperative settings, explore the use of reinforcement learning agents to discover novel allocation rules, and prototype cross-industry testbeds that can validate theoretical models under realistic operational constraints. As global supply chains confront an increasingly turbulent world, the ability to pool capacity intelligently, fairly, and securely may well become a defining feature of competitive and sustainable industrial systems.

References

1. Ivanov, D., & Dolgui, A. (2020). Viability of intertwined supply networks: extending the supply chain resilience angles towards survivability. International Journal of Production Research, 58(10), 2904-2915.

2. Pettit, T. J., Fiksel, J., & Croxton, K. L. (2010). Ensuring supply chain resilience: development of a conceptual framework. Journal of Business Logistics, 31(1), 1-21.

3. Tang, C. S. (2006). Perspectives in supply chain risk management. International Journal of Production Economics, 103(2), 451-488.

4. Kleindorfer, P. R., & Saad, G. H. (2005). Managing disruption risks in supply chains. Production and Operations Management, 14(1), 53-68.

5. Baryannis, G., Validi, S., Dani, S., & Antoniou, G. (2019). Supply chain risk management and artificial intelligence: state of the art and future research directions. International Journal of Production Research, 57(7), 2179-2202.

6. Hosseini, S., Ivanov, D., & Dolgui, A. (2019). Review of quantitative methods for supply chain resilience analysis. Transportation Research Part E: Logistics and Transportation Review, 125, 285-307.

7. Cachon, G. P., & Netessine, S. (2004). Game theory in supply chain analysis. In D. Simchi-Levi, S. D. Wu, & Z. M. Shen (Eds.), Handbook of quantitative supply chain analysis: Modeling in the e-business era (pp. 13-65). Springer.

8. Chen, Y., & Yang, S. (2019). A survey on game-theoretic approaches for supply chain collaboration. Annals of Operations Research, 281, 1-30.

9. Li, L., & Zhang, H. (2008). Confidentiality and information sharing in supply chain coordination. Management Science, 54(8), 1467-1481.

10. Anupindi, R., Bassok, Y., & Zemel, E. (2001). A general framework for the study of decentralized distribution systems. Manufacturing & Service Operations Management, 3(4), 349-368.

11. Kemahlıoğlu-Ziya, E., & Bartholdi III, J. J. (2011). Centralizing inventory in supply chains by using Shapley value to allocate the profits. Manufacturing & Service Operations Management, 13(2), 146-162.

12. Ivanov, D., Dolgui, A., & Sokolov, B. (2019). The impact of digital technology and Industry 4.0 on the ripple effect and supply chain risk analytics. International Journal of Production Research, 57(3), 829-846.

13. 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.

14. Hu, X., & Caldentey, R. (2023). Trust and reciprocity in firms' capacity sharing. Manufacturing & Service Operations Management, 25(4), 1436-1450.

15. Bergemann, D., & Morris, S. (2005). Robust mechanism design. Econometrica, 73(6), 1771-1813.

16. Christidis, K., & Devetsikiotis, M. (2016). Blockchains and smart contracts for the internet of things. IEEE Access, 4, 2292-2303.

17. Goldfarb, A., & Trefler, D. (2018). AI and international trade. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), The economics of artificial intelligence: An agenda (pp. 463-492). University of Chicago Press.

18. Carter, C. R., & Rogers, D. S. (2008). A framework of sustainable supply chain management: moving toward new theory. International Journal of Physical Distribution & Logistics Management, 38(5), 360-387.

19. Seuring, S., & Müller, M. (2008). From a literature review to a conceptual framework for sustainable supply chain management. Journal of Cleaner Production, 16(15), 1699-1710.

20. Ivanov, D., & Das, A. (2020). Coronavirus pandemic and supply chain disruptions: A risk management perspective. Annals of Operations Research, advance online publication.

21. Pournader, M., Kach, A., & Talluri, S. (2020). A review of the existing and emerging topics in supply chain risk management literature. Decision Sciences, 51(4), 867-919.

22. Govindan, K., & Hasanagic, M. (2018). A systematic review on drivers, barriers, and practices towards circular economy: a supply chain perspective. International Journal of Production Research, 56(1-2), 278-311.