20th IMEKO TC10 Conference
"Measurement for Diagnostics, Optimization and Control to Support Competitiveness and Innovation"
Lisbon, Portugal, September 28–29, 2026

Simulation-Based Measurement of Information Value in Circular Economy Networks

László Zahorán[1,2,3], József Váncza[1], Gábor Nick[1,2], Ádám Szaller[1,2]

[1]*HUN-REN Institute for Computer Science and Control, Budapest, Hungary*
[2]*EPIC InnoLabs, Budapest, Hungary*
[3]*Corvinus University of Budapest, Budapest, Hungary*

***Abstract* – Circular economy networks require not only physical recovery capabilities, but also timely information about inventories, repair status, remanufacturing progress, and expected availability. This paper presents a simulation-based framework for measuring the operational value of information sharing in cooperative circular production networks and reports full-scale results from a demonstrative European spare-part fulfilment network implemented in AnyLogic. Three information-sharing scenarios, each evaluated over an identical set of 722,910 demand instances, are compared: local decision-making without shared information, network-wide inventory sharing, and combined inventory and due-date sharing. Order service level rose from 3.25% under fragmented decision-making to 57.57% once inventory became visible network-wide, and reached 82.70% once expected repair and remanufacturing output was also made visible ahead of its completion date, provided it could still reach the customer in time. New production correspondingly fell from 96.75% to 42.43% and to 17.30%, while total transport cost dropped by roughly 30% relative to the no-sharing baseline. The remaining 17.30% of demand still requiring new production marks the point at which the network's physical and temporal capacity, rather than a lack of visibility, becomes the binding constraint. Separating this information gap from the capacity gap is the core diagnostic contribution of the proposed framework.**

Keywords* – Circular economy, value of information, information sharing, simulation, AnyLogic, remanufacturing, repair, closed-loop supply chain, diagnostics, Responsible Consumption and Production

I. INTRODUCTION

Circular economy implementation in manufacturing extends well beyond recycling or end-of-life recovery. Once products return from the field, they must be assessed, routed, and reintegrated through repair, remanufacturing, or redistribution into active fulfilment processes. This creates a class of operational dependencies that has no direct equivalent in linear supply chains. Recovered items may be geographically dispersed, in varying states of condition, and only become usable after processing delays of uncertain length. The practical consequence is that circularity is not purely a materials problem; it is equally a problem of operational visibility. What makes this particularly challenging is that physical availability does not imply decision-making availability. A repaired component may exist in the network and become ready before a customer's due date, yet trigger unnecessary new production simply because no actor in the chain is aware of it. A regional service point may place an emergency order while adequate stock sits idle at a nearby distributor. These are not failures of physical capacity but failures of information. The problem is especially acute in systems where repair and remanufacturing introduce a temporal gap between product return and renewed availability. If that expected future output is not visible to fulfilment logic, circular resources remain operationally inaccessible even when they are physically present. This paper proposes a simulation-based approach for measuring the operational value of information sharing in a cooperative circular economy network. Three information-sharing scenarios are compared within a fixed physical topology and demand structure, so that observed performance differences can be attributed directly to differences in information availability. The evaluation covers service reliability, circularity indicators such as new production avoidance, and economic performance.

This paper reports full simulation results for all three scenarios, evaluated over an identical set of 722,910 demand instances. The results show that due-date visibility is a major lever on top of inventory sharing, while the residual demand that no amount of visibility can serve in time separates the information gap from the capacity gap of the network. The work remains a proof-of-concept, but this diagnostic separation is directly transferable to calibration with industrial data.

II. RELATED RESULTS IN THE LITERATURE

The circular economy literature has converged on a shared understanding that resource loops must be closed, slowed, and narrowed by extending product life and reducing dependence on virgin material inputs [1]. Translating this into production systems requires more than strategic intent; it demands operational mechanisms for reuse, repair, refurbishment, remanufacturing, and recycling, each of which involves a distinct set of actors, timelines, and coordination requirements. The diversity of definitions and conceptual framings surrounding the circular economy reflects how broad this transition is in practice [2]. As Farooque et al. [3] note, circular supply chain management cannot be reduced to a material-flow problem. It requires information exchange and decision support structures that actively enable recirculation across network boundaries. In spare part networks specifically, this uncertainty is compounded by the fact that demand for remanufactured components is often driven by installed base dynamics rather than by predictable market signals.

The foundational literature on reverse logistics and closed-loop supply chains established many of the core challenges relevant to this context. Fleischmann et al. [4] demonstrated that return flows introduce planning complexity well beyond what forward supply chain models address, primarily because recovered products are heterogeneous in timing, quantity, and reuse potential. Govindan et al. [5] extended this perspective through a comprehensive review showing that recovery systems require specific coordination mechanisms to manage the uncertainty inherent in collection, sorting, and reprocessing decisions. These insights remain relevant for circular production networks, where the coordination challenge is compounded by the need to connect demand, inventory, ongoing repair processes, and expected future availability within a single decision framework.

The value of information sharing has a long tradition in forward supply chain research. Gavirneni et al. [16] quantified the benefit of shared order and inventory information in capacitated supplier–retailer settings, Cachon and Fisher [17] showed that shared demand and inventory data reduce supply chain costs beyond what faster and cheaper order processing alone achieves, and Lee et al. [18] demonstrated that the value of shared demand information grows with demand correlation over time. The circular setting studied here differs from this forward-chain tradition in one essential respect: the shared information concerns not downstream demand but upstream recovery output and its expected timing, which is precisely the information that repair and remanufacturing lead times make uncertain.

The role of information in this context has received growing attention. Jäger-Roschko and Petersen [11] synthesised 84 peer-reviewed articles on information sharing in circular economy settings and found that interorganisational visibility is a prerequisite for keeping materials and products in the loop. Without reliable information on what exists, where it is, and when it will be ready, circular resources cannot be effectively assigned to demand. Zhang et al. [12] examined information sharing in a closed-loop supply chain with dual recycling channels and showed that the value of downstream demand information to upstream remanufacturing decisions depends on market conditions and channel power structure. Khan and Abonyi [13] approached the same problem from a digital infrastructure perspective, arguing that supply chain interoperability is a structural condition for circular economy goals, particularly under Industry 5.0 principles.

Simulation has emerged as a natural methodological choice for studying these systems, precisely because circular production networks are dynamic and distributed in ways that static optimisation models cannot capture [6], [7]. Such approaches have been applied to circular business model adaptation [8], decision support in circular manufacturing [9], and operational evaluation of circular supply chains [10]. Massari et al. [14] showed that the interaction between information flows and network structure determines how effectively circular resources are used.

A specific gap remains in the literature concerning the operational value of due-date information in spare part fulfilment networks. Gupta et al. [15] showed that accurate demand forecasting for remanufactured parts is critical for inventory optimisation across distribution networks, and that information deficits lead directly to delayed fulfilment or costly substitution with new parts. The present paper addresses a related but distinct question: not how to forecast demand for remanufactured parts, but how visibility into expected repair and remanufacturing output changes the fulfilment decisions that are made before that output becomes available.

III. DESCRIPTION OF THE METHOD

The simulation model was implemented in AnyLogic as a demonstrative circular spare-part fulfilment network covering a set of European regions. The network consists of customer regions, distributors, service centres, collection points, repair centres, and remanufacturing facilities, connected by directed transport links. Each link carries attributes for transport mode, distance, lead time, cost, and operational status. The topology follows a regional mesh structure rather than a hub-and-spoke configuration, a deliberate choice to avoid concentrating all routing decisions through a single central node. This matters because the operational value of a circular resource is not determined solely by its existence; it also depends on whether it can be reached within acceptable time and cost constraints from the point of demand. The model captures both the forward fulfilment flow and the circular recovery logic running alongside it. Customer regions generate stochastic demand for spare parts. Distributors and service centres hold inventory available for immediate fulfilment. Collection points and service centres receive returned products and recoverable cores, which are then routed to repair centres or remanufacturing facilities depending on their condition. Once a repair or remanufacturing order is completed, its output re-enters the network as circular availability. The distinction between stock that is currently on hand and output that is expected at a future date is central to the design of the third scenario, since it is precisely this expected future availability that due-date information makes visible. The network topology is illustrated in Fig. 1. Proximity to demand is not the only criterion for source selection; the model evaluates candidate sources on quantity, reachability, lead time, and transport cost jointly, rather than ranking them on availability alone.

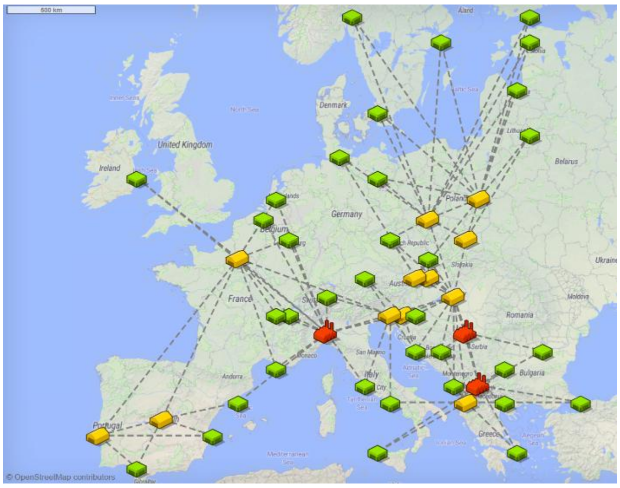

Fig. 1. Demonstrative European circular economy network used in the simulation model. Green markers denote customer regions, yellow markers distributors, service centres, and collection points, and red markers repair centres and the remanufacturing facility; the centralised production facility is co-located with the Budapest distribution hub.

Three information-sharing scenarios structure the experimental comparison. In the first, no shared information exists; each customer region can draw only on its local inventory, and any unmet demand is immediately escalated to new production. This represents a baseline of fragmented decision-making, in which circular resources elsewhere in the network are physically present but operationally invisible. This baseline deliberately excludes local replenishment policies in order to isolate the pure effect of information availability; as a consequence, the absolute service level of the no-sharing scenario should be read as a lower bound rather than as a realistic estimate of an uncoordinated network. In the second scenario, network-wide inventory visibility is introduced. The fulfilment logic can now search across all nodes for currently available stock and select a source based on the joint criteria above. The third scenario adds due-date sharing to inventory sharing: when no suitable current stock is found, the model can also consider expected future output from active repair and remanufacturing orders. The three scenarios therefore represent increasing levels of digital coordination, from isolated local decisions to network-wide visibility of both current stock and expected future output.

Performance is assessed through a multidimensional indicator framework designed to avoid redundancy across dimensions: order service level, capturing whether demand is fulfilled from a circular source at all; quantity and circularity performance, captured jointly here by the new production avoidance ratio; time performance, measured by average fulfilment lateness relative to the requested due date; and economic performance, measured by total transport cost. Two further indicator groups defined at the design stage, namely future-availability utilisation as a metric distinct from the circular fulfilment ratio, and inventory, emissions, and recovery/recycling performance, are not yet computed in the present iteration of results and are discussed as outlook in Section V.

The model is implemented as a data-driven simulation in which the topology, inventory positions, and demand records are not hard-coded but loaded at run time from a relational schema through parametrised SQL queries. Nodes are read from a node table (identifier, code, name, type, country, coordinates, daily capacity, and an active flag); transport links are read from an edge table specifying source, target, transport mode, distance, lead time, and cost; on-hand inventory is aggregated per node and product code as the quantity on hand net of reserved quantity, so that stock of one part number can no longer be used to satisfy demand for another; and expected future availability is read separately from repair-order and remanufacturing-order tables, the latter discounted by a recovery-yield coefficient representing imperfect output from the remanufacturing process. This separation between current inventory and future circular availability is what allows the three scenarios to be implemented as three variants of the same fulfilment search rather than as three different models. The demonstrative network instance used in this study comprises 51 nodes: 35 customer regions, four distributors, four service centres, four collection points, two repair centres, one remanufacturing facility, and the centralised production facility in Budapest, spanning 30 European countries, with daily node capacities ranging from 60 to 220 units.

For any candidate source–customer pair, the shortest path is computed twice over the directed transport graph using a Dijkstra-type search: once with edge weight equal to lead time and once with edge weight equal to transport cost. This yields, for each pair, both the fastest feasible lead time and the cheapest feasible cost, which are used as the ranking criteria in source selection: candidates are ranked first by lead time, then by cost, and, in case of a tie, by the available quantity at the source. In the no-sharing scenario this search space is restricted to the customer's own node; in the inventory-sharing scenario it is extended across all nodes carrying on-hand stock of the requested product; in the combined scenario, if no on-hand stock satisfies the request, the search is extended a second time over the set of active repair and remanufacturing orders, matched on product code and quantity. A future-availability record is only accepted as a source if the sum of its expected completion delay and the transport lead time to the customer does not exceed the customer's due-in-days window; candidates that would arrive late are discarded rather than accepted and penalised afterwards. Among the remaining candidates, the one with the shortest total lead time is chosen, cost being the tie-breaker. Once a future-availability record is used to satisfy a demand, its remaining quantity is decremented immediately, so that subsequent demands cannot draw twice on output that has already been committed.

Demand is represented at the level of individual order lines: each of the 722,910 demand instances specifies a customer node, a product code, a requested quantity, a priority level, and a request/due-date pair from which the allowable due-in-days window is derived, or, where explicit dates are unavailable, a stochastic three-to-eight-day window. Demands are processed in the order they would realistically need attention, by request date, then due date, then priority level, so the greedy allocation described above at least follows a sensible dispatching rule rather than an arbitrary one. Whenever none of the above sources can satisfy a demand, the model escalates it to new production at a single, centralised manufacturing node in Budapest, applying a fixed four-day production lead time in addition to the transport lead time and cost to the customer; if no transport route can be resolved, a conservative fallback lead time and cost are applied instead. Because this fallback route is centralised, the new-production ratio reported in Section IV also reflects, indirectly, the cost and lead-time penalty of relying on a single production hub rather than on the distributed circular capacity already present in the network. The total demand volume of 722,910 instances, and the underlying physical network, were held identical across all three scenario runs, satisfying the controlled-comparison design on which the framework depends.

IV. RESULTS AND DISCUSSIONS

All three scenarios were executed over the identical set of 722,910 demand instances described in Section III, so that the differences reported below can be attributed to information availability rather than to differences in demand or network capacity. Table 1 summarises the resulting indicators, while Fig. 2 shows how service level evolves over the course of the simulation for each scenario. Fig. 2 shows that this end-state comparison conceals a distinctly non-uniform trajectory. All three scenarios begin near full service level, simply because early in the run only a handful of demand instances have been processed and the network's initial inventory and repair pipeline are still intact. The no-sharing scenario collapses almost immediately, while inventory sharing and inventory-plus-due-date sharing remain close to each other and near 100% up to around 400,000 processed instances, before visibly diverging. This delayed divergence indicates that the operational value of due-date visibility only becomes decisive once the network-wide inventory buffer is depleted and expected repair and remanufacturing output becomes the main remaining source of circular fulfilment.

Metric	No sharing	Inventory sharing	Inventory + due-date
Fulfilled demands	23,502	416,188	597,841
Service level	3.25%	57.57%	82.70%
New production ratio	96.75%	42.43%	17.30%
Shortage events	699,408	306,722	125,069
Average lateness (days)	1.44	0.64	0.53
Total transport cost index (no sharing = 100)	100.0	75.9	69.8

Table 1. Summary of simulation results across the three information-sharing scenarios (722,910 demand instances each).

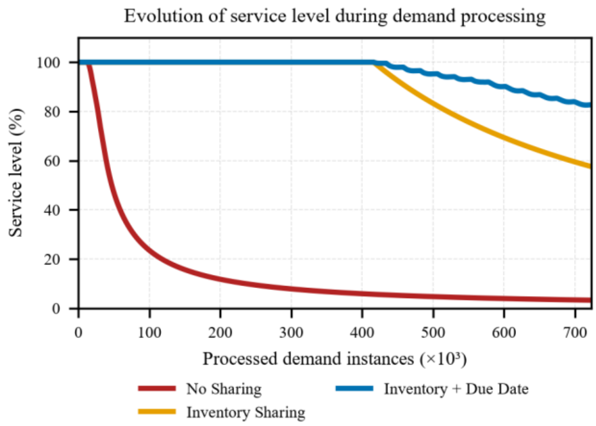

Fig. 2. Evolution of service level over the course of demand processing, shown separately for the three information-sharing scenarios.

Because the underlying indicator logic classifies a demand as fulfilled only when it is satisfied from an existing, repaired, or remanufactured circular source, the service level reported in Table 1 is at the same time the circular fulfilment ratio of the network; the two are therefore not reported separately. The complement of the service level is not lost demand but demand escalated to new production, which is always eventually satisfied, at the cost, lead time, and environmental penalty of manufacturing a new part rather than reusing one already in, or entering, the network. Shortage events in Table 1 accordingly denote the count of such new-production escalations rather than unfulfilled orders in the everyday sense; by construction they equal total demand minus fulfilled demand in every scenario.

Moving from fragmented local decision-making to network-wide inventory sharing raises service level from 3.25% to 57.57% and cuts new production from 96.75% to 42.43%, simply by letting the fulfilment logic see stock that already exists elsewhere in the network. Adding due-date sharing raises service level further, to 82.70%, and brings new production down to 17.30%. It does not remove new production altogether: 125,069 of the 722,910 demand instances, close to a fifth of the total, still end up there even when the fulfilment logic can see both current stock and everything currently being repaired or remanufactured across the network. Total transport cost falls monotonically across the three scenarios, reaching an index of 75.9 and 69.8 relative to the no-sharing baseline (set to 100), a reduction of about 30% overall, even though the volume of demand fulfilled from circular sources grows roughly twenty-five-fold over the same comparison. Average lateness decreases only modestly and monotonically across the three scenarios (1.44, 0.64, and 0.53 days), despite the much larger swings in service level and shortage counts. This is because the lateness indicator is averaged across all demand instances, including those escalated to new production, and new production itself carries a fixed four-day production lead time that is independent of information availability. Lateness therefore mixes a pure information effect with a structural lead-time floor, making it a comparatively insensitive indicator in this model; reporting it separately for circular and new-production fulfilment is left to subsequent iterations.

The residual 17.30% is worth dwelling on, because it is the clearest evidence in these results that due-date sharing does not simply switch off the need for new production. Because of the product-code matching and the due-date feasibility condition described in Section III, the 82.70% service level reflects demand that could actually be matched to a source in time, not demand for which some source, however tardy, existed somewhere in the network. The remaining 17.30% is demand for which, at the moment it needed a decision, no node held the right part and no repair or remanufacturing order due to finish would have reached the customer on time. Whether that gap could be closed by more circular capacity, by different due-date terms, or is simply structural in a network of this size and density is a separate question from the value of information itself, and one the framework is well placed to explore next.

Future-availability reservations are consumed on a first-come, first-served basis, with demands processed by request date, then due date, then priority, which is a sensible dispatching rule but still just one feasible assignment rather than a jointly optimal one; a different processing order, or a scheduler that looked ahead across demands instead of committing greedily, could plausibly close part of the 17.30% gap without adding any new information. Node throughput capacity is also captured in the underlying network data but is not enforced as a hard constraint in the fulfilment search itself, so the results should still be read as an upper bound on what visibility can achieve in this network, not as a claim that the physical system needs no further capacity at all.

The consistent decrease in total transport cost alongside the large increase in circularly fulfilled volume is also informative: it indicates that, on average, circular sourcing paths in this mesh topology are both shorter and cheaper than the centralised new-production fallback route, consistent with the topology design rationale discussed in Section III. The results give a first quantitative anchor to the claim of Jäger-Roschko and Petersen [11] that interorganisational visibility is a prerequisite for closing circular loops: due-date visibility alone recovers a further 25 percentage points of service level, while leaving a measurable residual. The pattern also resonates with Zhang et al. [12], where the value of shared information was shown to depend on structural conditions; here the relevant structure is the transport topology rather than channel power, but the underlying logic, that the value of information is contingent on the physical system it acts upon, is the same. Finally, due-date sharing functions in this model similarly to a short-horizon forecast of remanufacturing completion, connecting the present results to the forecasting-deficit argument of Gupta et al. [15].

From a measurement perspective, the quantity being measured is the difference in operational indicators attributable solely to information availability, with the physical network and demand realisation held constant. The dominant sources of uncertainty are the single demand realisation (sampling uncertainty), the stochastic due-date windows assigned where explicit dates were unavailable, and the recovery-yield coefficient. As only one demand draw is evaluated in the present iteration, no variance estimate accompanies the reported figures; replication across demand draws is the first step of the planned follow-up.

V. CONCLUSIONS AND OUTLOOK

This paper presented full simulation results for a diagnostic framework measuring the operational value of information sharing in cooperative circular economy networks. Over an identical set of 722,910 demand instances, network-wide inventory visibility raised order service level from 3.25% to 57.57%, and adding due-date visibility of expected repair and remanufacturing output raised it further to 82.70%, while cutting total transport cost by about 30%. Even with full visibility, however, 17.30% of demand still required new production.

These results show that due-date information is a major lever on top of inventory visibility, not a marginal one, but also that it has a limit. Once a node's product-specific stock and its own repair and remanufacturing pipeline are exhausted, or simply too slow for a given due date, no amount of visibility invents a source that is not there in time. The size of that residual, rather than being a nuisance, is itself a useful number: it separates the part of the fulfilment problem that better information can fix from the part that would need more circular capacity, faster processing, or different due-date terms with customers.

The model remains a proof-of-concept in several respects. Node throughput capacity is present in the network data but not enforced in the fulfilment search, so the reported figures are an upper bound on what visibility alone can deliver rather than a claim that no further capacity is needed. Allocation of future availability is greedy, processed in a fixed and reasonable but still arbitrary order, rather than jointly optimised across demands. The indicator framework is designed to be extensible beyond the dimensions computed here: future-availability utilisation as a metric in its own right, inventory carrying levels, emissions, and recovery/recycling rate are already defined within it and await computation. Subsequent work will address these directly: enforcing node capacity in the fulfilment search, testing look-ahead or optimisation-based allocation against the current greedy rule, replicating scenarios across multiple demand draws, computing the remaining indicators, and calibrating the model against real industrial transaction, repair-yield, and return-quality data.

VI. ACKNOWLEDGMENTS

This work was partially supported by the National Research, Development and Innovation Fund of Hungary under the SNN funding scheme (Project no. SNN 153457). The project identified by EKOP-KDP-CORVINUS-25-132 was realized with the support of the National Research, Development, and Innovation Fund provided by the Ministry of Culture and Innovation, as part of the University Research Scholarship Program – Cooperative Doctoral Program announced for the 2025/2026 academic year.

REFERENCES

[1] M. Geissdoerfer, P. Savaget, N. M. P. Bocken, and E. J. Hultink, "The circular economy: A new sustainability paradigm?", *Journal of Cleaner Production*, Vol. 143, 2017, pp. 757–768.

[2] J. Kirchherr, D. Reike, and M. Hekkert, "Conceptualizing the circular economy: An analysis of 114 definitions", *Resources, Conservation and Recycling*, Vol. 127, 2017, pp. 221–232.

[3] M. Farooque, A. Zhang, M. Thürer, T. Qu, and D. Huisingh, "Circular supply chain management: A definition and structured literature review", *Journal of Cleaner Production*, Vol. 228, 2019, pp. 882–900.

[4] M. Fleischmann, J. M. Bloemhof-Ruwaard, R. Dekker, E. van der Laan, J. A. E. E. van Nunen, and L. N. Van Wassenhove, "Quantitative models for reverse logistics: A review", *European Journal of Operational Research*, Vol. 103, No. 1, 1997, pp. 1–17.

[5] K. Govindan, H. Soleimani, and D. Kannan, "Reverse logistics and closed-loop supply chain: A comprehensive review to explore the future", *European Journal of Operational Research*, Vol. 240, No. 3, 2015, pp. 603–626.

[6] L. Monostori, J. Váncza, and S. R. T. Kumara, "Agent-based systems for manufacturing", *CIRP Annals*, Vol. 55, No. 2, 2006, pp. 697–720.

[7] J. Váncza, L. Monostori, E. Lutters, S. R. Kumara, M. Tseng, P. Valckenaers, and H. Van Brussel, "Cooperative and responsive manufacturing enterprises", *CIRP Annals*, Vol. 60, No. 2, 2011, pp. 797–820.

[8] M. Lieder, F. M. A. Asif, and A. Rashid, "Towards circular economy implementation: An agent-based simulation approach for business model changes", *Autonomous Agents and Multi-Agent Systems*, Vol. 31, 2017, pp. 1377–1402.

[9] M. Roci, N. Salehi, S. Amir, F. M. A. Asif, S. Shoaib-ul-Hasan, and A. Rashid, "Multi-method simulation modelling of circular manufacturing systems for enhanced decision-making", *MethodsX*, Vol. 9, 2022, Article 101709.

[10] N. Salehi, S. Amir, M. Roci, S. Shoaib-ul-Hasan, F. M. A. Asif, A. Mihelič, S. Sweet, and A. Rashid, "Towards circular manufacturing systems implementation: An integrated analysis framework for circular supply chains", *Sustainable Production and Consumption*, Vol. 51, 2024, pp. 169–198.

[11] M. Jäger-Roschko and M. Petersen, "Advancing the circular economy through information sharing: A systematic literature review", *Journal of Cleaner Production*, Vol. 369, 2022, Article 133210.

[12] H. Zhang, X. Hong, and X. Cao, "To share or not to share? The role of retailer's information sharing in a closed-loop supply chain", *Frontiers in Psychology*, Vol. 13, 2022, Article 941952.

[13] A. A. Khan and J. Abonyi, "Information sharing in supply chains – Interoperability in an era of circular economy", *Cleaner Logistics and Supply Chain*, Vol. 5, 2022, Article 100074.

[14] G. F. Massari, R. Nacchiero, and I. Giannoccaro, "Circular supply chains as complex adaptive ecosystems: A simulation-based approach", *Journal of Cleaner Production*, Vol. 475, 2024, Article 143517.

[15] M. Gupta, U. Dayal, S. Horiguchi, and D. Ghosh, "Demand and replacement forecasting for remanufactured parts of industrial products", *Journal of Advanced Manufacturing and Processing*, Vol. 7, 2025, Article e70028.

[16] S. Gavirneni, R. Kapuscinski, S. Tayur, "Value of information in capacitated supply chains", *Management Science*, Vol. 45, No. 1, 1999, pp. 16–24.

[17] G. P. Cachon, M. Fisher, "Supply chain inventory management and the value of shared information", *Management Science*, Vol. 46, No. 8, 2000, pp. 1032–1048.

[18] H. L. Lee, K. C. So, C. S. Tang, "The value of information sharing in a two-level supply chain", *Management Science*, Vol. 46, No. 5, 2000, pp. 626–643.