VALIDATING SIMULATED AGENTS WITH PHARMACEUTICAL SUPPLY CHAIN GAME DATA

Souri Sasanfar[1], Omid Mohaddesi[1], Noah Chicoine [1], Jacqueline Griffin[1], and Casper Harteveld[2]

[1]Dept. of Mechanical and Industrial Eng., Northeastern University, Boston, MA, USA
[2]Dept. of Arts+Design, Northeastern University, Boston, MA, USA

ABSTRACT

Human decision-making under uncertainty poses a significant challenge in modeling disrupted pharmaceutical supply chains. This study introduces a behaviorally grounded simulation framework systematically integrating empirical data from participatory game experiments using an advanced methodological pipeline. Specifically, we combine Principal Component Analysis (PCA) for dimensionality reduction, Longest Common Prefix (LCP) clustering for behavioral sequence segmentation, and Hidden Markov Models (HMMs) for dynamic behavioral archetype extraction. These integrated methods identify three archetypes: Hoarders, Reactors, and Followers, which we embed into agent-based simulations to assess impacts under varying disruption and information-sharing scenarios. Agent realism is rigorously validated using Kolmogorov-Smirnov tests that compare simulated and empirical behavioral data, demonstrating strong alignment in general but revealing notable discrepancies among Hoarders under disruption-duration conditions. The findings highlight the methodological necessity of capturing dynamic behavioral heterogeneity, guiding future research on alternative clustering methods, multiphase clustering, and dynamic behavioral transition modeling for enhanced supply chain resilience.

1 INTRODUCTION

Pharmaceutical supply chains operate with high stakes and are increasingly vulnerable to disruptions, ranging from production halts to global health emergencies (Snyder et al. 2014). Traditional simulation models typically assume rational and consistent behaviors from supply chain actors. However, real-world disruptions reveal decision-makers often engage in panic ordering, hoarding, and strategic shifts due to incomplete or evolving information (Sheffi 2005). Accurately capturing this behavioral variability is crucial for improving supply chain resilience and informing robust policy interventions.

This paper extends recent work integrating behavioral data from human-in-the-loop game experiments into supply chain simulations (Mohaddesi et al. 2022). Specifically, we demonstrate how behavioral archetypes—Hoarders, Reactors, and Followers—can be extracted using Principal Component Analysis (PCA), Longest Common Prefix (LCP) clustering, and Hidden Markov Models (HMMs). These archetypes are embedded in agent-based simulations to evaluate their impact in varying information-sharing and disruption scenarios.

Although focused on pharmaceutical supply chains, our methods have a broader logistic applicability, facilitating resilience testing against disruptions such as pandemics or natural disasters, and informing policies around information exchange and behavioral nudging (Ergun et al. 2022; Croson et al. 2014; Sterman and Dogan 2015; Doroudi et al. 2018).

Our contributions are threefold. First, we propose a methodological pipeline that integrates PCA, LCP clustering, and HMM to identify behavioral response modes across supply chain phases (stable, disrupted, recovery). Second, we calibrate agent-based decision rules using empirically derived deviations from base-stock policies. Third, we rigorously validate agent realism through statistical comparisons of simulation

output and empirical data, highlighting strengths and limitations of clustering methods in the capture of complex behavioral dynamics.

Results demonstrate behaviorally-informed agents reproduce aggregate outcomes and critical patterns, like disruption-induced over-ordering spikes. However, detailed validation reveals discrepancies, especially among Hoarders under certain informational conditions, motivating future methodological improvements in multi-phase clustering and dynamic behavioral transitions.

The paper proceeds as follows: Section 2 reviews related work; Section 3 presents our application to a pharmaceutical supply chain; Section 4 details our proposed modeling methodology; Section 5 presents results of our simulation application; Section 6 discusses methodological insights, challenges, and limitations; Section 7 outlines future methodological extensions. We conclude with broader implications for simulation practice.

2 RELATED WORK

2.1 Agent-Based Models (ABMs) in Supply Chain Applications

Agent-Based Modeling (ABM) is widely used to simulate complex interactions in multi-echelon supply chains, where each echelon (e.g. manufacturer, distributor, wholesaler) acts independently with localized objectives and partial information (Snyder and Shen 2019; Ergun et al. 2022). This decentralized perspective allows ABMs to capture phenomena such as the bullwhip effect, where minor demand fluctuations amplify upstream and lead to rising costs (Lee, Padmanabhan, and Whang 1997). By incorporating heterogeneous agent behaviors, ABMs effectively model irrational or disruptive actions like over-ordering or hoarding, enhancing realism and predictive power (Ergun et al. 2022; Sterman and Dogan 2015).

System dynamics (SD) offers similar robust modeling of feedback loops, time delays, and complex interactions in supply chains (Sterman 1989; Minegishi and Thiel 2000; Georgiadis et al. 2005; Sterman and Dogan 2015). SD studies have examined the impacts of demand variability or production stoppages on sectors such as food and pharmaceuticals, including stock-out risks following severe disruptions (Georgiadis et al. 2005; Strohhecker and Größler 2015). Early research (Sterman 1989) demonstrates that even seemingly rational actors can inadvertently generate system-wide bullwhip effects, while (Lee et al. 1997) shows how minor demand perturbations trigger volatile orders. Therefore, capturing realistic decision-making heuristics is essential in both SD and ABM frameworks, especially in pharmaceutical supply chains, where product recalls or capacity breakdowns critically impact supply (Azghandi et al. 2018).

Despite extensive research on disruption management, few studies specifically address pharmaceutical product recalls or multi-node disruptions simultaneously (Doroudi et al. 2018). Moreover, many analytical models assume rational, homogeneous decision-makers, overlooking behavioral biases like panic ordering or hoarding. We address this limitation by empirically incorporating behavioral differences into our agent-based models, enhancing realism in pharmaceutical supply chain disruption simulations. Furthermore, systematic data-driven approaches to calibrating behavioral archetypes under dynamic disruptions remain scarce. Our study fills this gap through an integrated PCA, clustering, and sequential modeling pipeline for empirically grounded agent calibration.

2.2 Participatory Simulations and Game-Based Research

To overcome the limitations of purely computational models, researchers increasingly integrate human decision makers directly into simulations to capture behavioral nuances accurately. Early methods, termed companion modeling, combined Role-Playing Games (RPGs) with Multi-Agent Systems (MAS) to actively engage stakeholders in iterative refinement (Guyot and Honiden 2006; Bousquet and Trébuil 2005). More recently, agent-based participatory simulations allow participants to control agents directly, validating or tuning decision-making models (Meijer 2009; Anand et al. 2016).

In supply chains, human biases such as panic ordering, hoarding, and strategic shifts critically influence system outcomes (Croson et al. 2014; Sterman and Dogan 2015). Short interactive scenarios known

as "gamettes" effectively capture real-time decision-making, enabling realistic calibration of simulation models (Mohaddesi et al. 2020). For instance, (Doroudi et al. 2018) integrated gamette-derived data with a supply chain flow simulator, continuously tracking orders, shipments, and inventory. Participants, acting as wholesalers, demonstrated behaviors such as over-ordering or rationing under uncertainty (Azghandi et al. 2018).

These participatory frameworks reduce reliance on rational assumptions by capturing human decisions under various disruption or information scenarios. However, robust analysis and validation of the empirical data remain critical limitations. Our study addresses this by developing a validated analytical pipeline that ensures that simulated agent behaviors accurately reflect empirically observed decisions.

2.3 Modeling Methodology: Clustering, PCA, and HMM

Real-world supply chain simulations produce extensive high-dimensional data (e.g. inventory, backlog, ordering histories), which complicates pattern detection, particularly during disruptions (Sterman and Dogan 2015; Mohaddesi et al. 2022). Principal Component Analysis (PCA) commonly reduces dimensionality, identifying key variance drivers, and clearly labeling overarching states like *stable*, *disrupted*, and *recovery* phases (Sterman and Dogan 2015; Mohaddesi et al. 2022). Once these broad states are identified, clustering methods (e.g., hierarchical, k-means, Longest Common Prefix (LCP)) group similar decision sequences, facilitating finer segmentation (Marsella et al. 2004).

For sequential decision-making contexts, Hidden Markov Models (HMMs) effectively identify latent response modes from longitudinal ordering data, estimating hidden-state transition probabilities (Rabiner 1989; Mohaddesi et al. 2022).Othersemployed HMMs to classify participant deviations from recommended orders, revealing distinct archetypal responses (e.g., panic-induced over-ordering) triggered by disruptions (Croson et al. 2014; Sterman and Dogan 2015).

Empirical validation methods critically bridge computational analyses with real-world behavior, comparing simulated outputs against empirical distributions using statistical tests like Kolmogorov-Smirnov (K–S) (Azghandi et al. 2018). Our approach leverages PCA for identifying behavioral phases, LCP clustering for segmenting decision sequences, and HMMs for uncovering hidden behavioral patterns. Carefully validating these empirically derived behaviors embeds realistic, human-informed decision-making into pharmaceutical supply chain simulations. While PCA, clustering, and HMM methods are individually established, our methodological innovation integrates them into a unified behavioral modeling pipeline, enhancing predictive accuracy and realism under dynamic disruptions.

3   APPLICATION TO PHARMACEUTICAL SUPPLY CHAIN

In this section, we demonstrate the practical applicability of our integrated PCA, LCP clustering, and HMM-based methodology through a pharmaceutical supply chain case study. We clearly describe the supply chain setup, the gamette interaction, and the conditions for information sharing.

3.1 Supply Chain Network Simulation

We modeled a simplified pharmaceutical supply chain consisting of two manufacturers (MN1, MN2), two wholesalers, and two health centers over 35 weeks. MN1 experienced a 95% capacity disruption starting in Week 28 and lasting several weeks, creating downstream ordering fluctuations and backlogs. This structure captures realistic complexity, facilitating meaningful behavioral responses without overwhelming participants. Figure 1 illustrates the simulated pharmaceutical supply chain network, clearly showing the flow of materials and information.
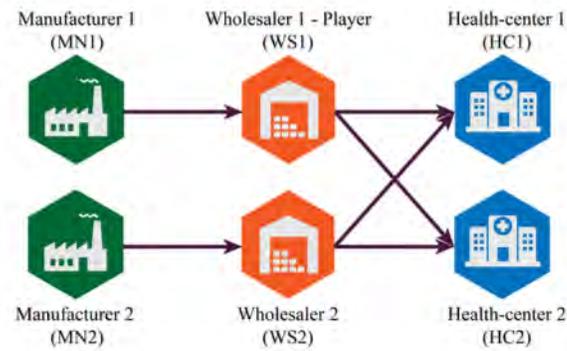
Figure 1: Simplified pharmaceutical supply chain structure used in the case study.

The integrated simulation framework combined a Python-based multi-agent Flow Simulator with a gamette environment, enabling participants to directly interact with the supply chain scenario through weekly decisions on orders and shipments. The Flow Simulator advanced weekly, updating inventory levels, backlogs, and shipments based on participant decisions. Communication between the Flow Simulator and the gamette was facilitated via a RESTful API, ensuring real-time state updates and participant decisions. Figure 2 illustrates this integration, highlighting the real-time interaction.
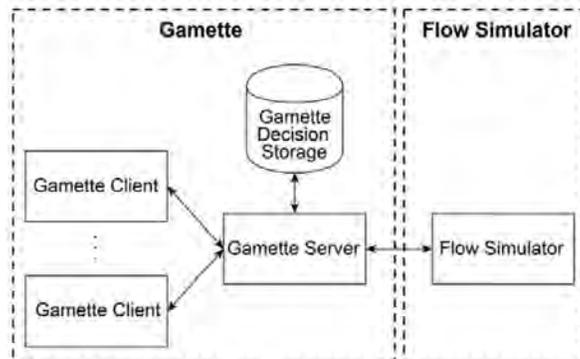

Figure 2: Integration of Gamette and Flow Simulator (adapted from Mohaddesi et al., 2020). Communication is facilitated by a RESTful API, enabling real-time decision-making and immediate state updates.

The disruption scenario allowed us to clearly distinguish and analyze three distinct behavioral phases: Stable, Disruption (Shortage), and Recovery. These phases were identified through PCA analysis of ordering and inventory patterns, categorizing weeks into distinct system states reflecting the operational status of the supply chain (Figure 3).
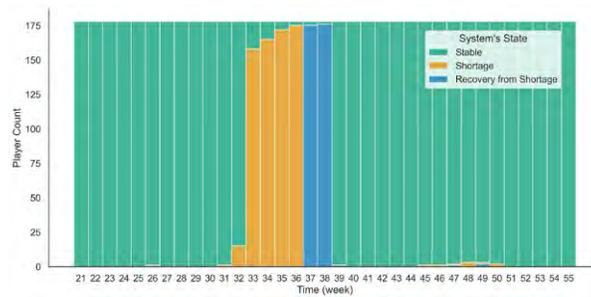
Figure 3: System states and disruption phases identified through PCA, categorized as Stable (green), Shortage (orange), and Recovery (blue) states over the simulation timeline.

3.2 Gamettes

Participants individually interacted with the gamette via an online interface, reviewing weekly inventory levels, backlog, and demand before making ordering decisions. Immediate feedback and cumulative performance summaries were provided to participants.

Participants were recruited via the Prolific online platform, initially totaling 192 individuals. After data cleaning—removing 14 outliers based on predefined statistical thresholds—178 participants remained. These participants were randomly assigned to one of four independent experimental conditions detailed in Section 3.3. Each participant individually controlled the wholesaler node (ds_2), while other nodes (manufacturers and health centers) were automated. Importantly, participants were not grouped; each participant independently interacted with their individual simulation session. The distinct system states (Stable, Shortage, Recovery) identified via PCA and illustrated in Figure 3 were determined from aggregated data and were not directly related to participant grouping. Division into smaller archetype-condition subgroups, necessary for detailed analysis, contributed to the small-sample-size limitations discussed further in Section 6.

3.3 Information Sharing Conditions

Participants were randomly assigned to one of four information-sharing scenarios with the intent to assess how varying levels of transparency influenced decision-making under supply disruptions:

- **Condition 320 (No Information)**: Participants received no additional information beyond basic ordering data.
- **Condition 321 (Disruption Duration Only)**: Participants were informed of the disruption duration explicitly.
- **Condition 322 (Supplier Capacity Only)**: Participants had real-time information about the supplier's current capacity.
- **Condition 323 (Combined Information)**: Participants received both disruption duration and real-time supplier capacity information.

Weekly participant data—ordering decisions, inventory levels, and backlogs—were systematically recorded, providing a rich empirical dataset for analyzing behavioral responses to varying disruption conditions.

4 PROPOSED MODELING METHODOLOGY

In this section, we introduce a novel methodological pipeline that systematically integrates Principal Component Analysis (PCA), Longest Common Prefix (LCP) clustering, and Hidden Markov Models (HMMs)

to capture and embed realistic human behavioral archetypes into agent-based supply chain simulations. This integrated methodological approach addresses critical gaps in existing behavioral simulation methodologies by explicitly modeling dynamic and heterogeneous decision-making patterns, especially during disruption scenarios.

4.1 Overview of Methodological Pipeline

The core contribution of this study is our systematic pipeline that combines PCA-based dimensionality reduction, LCP clustering for behavioral sequence segmentation, and HMMs for extracting dynamic behavioral archetypes. The pipeline involves the following sequential steps:

1. **Dimensionality Reduction via PCA**: We standardized six wholesaler-level variables—Inventory (Inv), downstream demand to HC_1 and HC_2 (Dem_1, Dem_2), Backlog (Blg), Received-Shipment (Shp), and On-Order stock (Oor)—and performed PCA. The first two principal components explain 77% of the variance, clearly distinguishing *stable*, *disrupted*, and *recovery* behavioral phases (Figure 3; detailed loadings in Table 2 of Mohaddesi et al., 2022).

2. **Sequence Segmentation via LCP Clustering**: Building on the phase labels, we encode each participant's order-deviation series as a sign string (e.g., + + 0 − −) and compute the longest-common-prefix distance (Marsella et al. 2004). Cutting the dendrogram at $d_{LCP} = 0.5$ yields three clusters whose modal prefixes ++, 0, and + correspond to Hoarder, Follower, and Reactor archetypes, respectively (full dendrogram in Fig. 8 of Mohaddesi et al., 2022), enabling finer-grained segmentation of adaptive decision-making.

3. **Dynamic Behavioural Archetype Extraction via HMMs**: We retain an eight-state Hidden Markov Model (states N1, N2, C, and P1–P5) chosen by the Bayesian Information Criterion to capture latent response modes in the sign-encoded order-deviation sequences. Trained with the Baum–Welch algorithm, the model estimates transition probabilities between hidden states and thus reveals archetypal reactions, such as panic-induced overordering triggered by disruptions reported by (Croson et al. 2014; Sterman and Dogan 2015). State-specific Student-*t* emission parameters and the full 8 × 8 transition matrix are provided in Mohaddesi *et al.* (2022, App. B).

This methodological integration significantly improves our ability to embed empirically derived behavioral realism in agent-based simulations, thus improving predictive precision and robustness, as evidenced by the aggregate K–S statistics reported in Section 5 (0.14–0.28 with $p = 0.23 - 0.97$) and by the model maintaining this fit in all four information sharing conditions and three behavioral archetypes.

4.2 Behavioral Archetypes: Definition and Formulation

Using the pipeline described above, we empirically identified and defined three distinct behavioral archetypes—*Hoarders (H)*, *Reactors (R)*, and *Followers (F)*—across the supply-chain states (Stable, Disruption, Recovery). Each archetype captures distinctive patterns in ordering decisions:

- **Hoarders (H)**: Agents characterized by frequent over-ordering in anticipation or response to disruption signals, reflecting panic-induced stockpiling behaviors.
- **Reactors (R)**: Agents who adaptively adjust their ordering behavior upon receiving disruption signals, initially deviating significantly but gradually reverting toward recommended orders as new information emerges or disruptions subside.
- **Followers (F)**: Agents that adhere closely to suggested policies, showing minimal deviations, but leaving supply chains potentially vulnerable to sudden disruptions.

Behavioral decision-making at each time step *t* is mathematically formulated by decomposing an agent's order $O(t)$ into deterministic and stochastic components:

$$O(t) = D(t) + \Delta(t) \tag{1}$$

The deterministic base-stock ordering component $D(t)$ is defined as:

$$D(t) = \max\{S - I(t) - O_{\text{pending}}(t) + B(t),\ 0\} \tag{2}$$

Here, S denotes the base-stock level, $I(t)$ is inventory on-hand, $O_{\text{pending}}(t)$ are pending orders, and $B(t)$ represents backlogs at time t.

The stochastic deviation $\Delta_{k,p}(t)$, specific to behavioral archetype $k \in \{H, R, F\}$ in supply-chain phase $p \in \{S, D, R\}$, is drawn from a Student's t-distribution:

$$\Delta_{k,p}(t) \sim t_{\nu_{k,p}}(\mu_{k,p},\ \sigma_{k,p}) \tag{3}$$

with parameters $(\mu_{k,p}, \sigma_{k,p}, \nu_{k,p})$ denoting the location, scale, and degrees-of-freedom, empirically estimated from experimental data.

Empirical deviations from suggested orders were symmetric about zero yet exhibited markedly heavier tails than a Gaussian. Distributions restricted to positive support (e.g. lognormal, gamma, beta) are therefore unsuitable, and a Cauchy distribution—while heavy-tailed—has infinite variance and over-predicts extreme values, complicating cost estimates. A Student-t distribution with $\nu > 2$ preserves finite variance while capturing the observed tail mass, making it the most appropriate choice for modelling $\Delta_{k,p}(t)$.

The final scaled deviation integrated into agent orders is defined as:

$$\Delta(t) = s\,\Delta_{k,p}(t) \tag{4}$$

Here s is the empirical standard deviation of the raw (un-standardised) deviations for the corresponding archetype–phase pair; multiplying by s converts the unit-scale draw $\Delta_{k,p}(t)$ back into order-quantity units.

4.3 Agent Integration and Formulation

The final methodological step involves embedding these behavioral archetypes into an agent-based simulation framework to rigorously examine how heterogeneous decision-making influences system-level outcomes under varying disruption scenarios.

Agents were systematically integrated into the Flow Simulator environment, as illustrated in Figure 2 in Section 3.1, Supply Chain Network Simulation. Each agent, calibrated to empirically-derived archetype parameters $(\mu_{k,p}, \sigma_{k,p}, \nu_{k,p})$, follows the above-defined ordering rules at each simulation step. Calibration parameters were empirically estimated via PCA, LCP clustering, and HMM analysis of the original human-in-the-loop experimental data. Each agent thus realistically represents empirically validated deviations from classic base-stock policies.

To validate agent realism, we conducted a series of 50 simulation replications per behavioral archetype under each experimental condition (320–323). Simulated ordering behaviors were rigorously compared to empirical distributions using Kolmogorov–Smirnov (K–S) tests. Successful calibration was concluded when K–S tests indicated no significant distributional differences between simulated and empirical data. In scenarios where significant deviations occurred, particularly for Hoarders under specific information-sharing conditions, iterative recalibration was performed to improve alignment.

This methodological integration and rigorous calibration provide a robust and transparent foundation for analyzing complex supply chain behaviors and responses, especially under disruption scenarios. The Gamette exchanges data with the Flow Simulator once per simulated week, i.e. 35 request–response cycles per session. Server logs during the study show a median round-trip latency of $\sim$90 ms (95[th]<130 ms), which is negligible relative to the player's think-time budget and does not interrupt play. A detailed comparison of calibration accuracy before versus after each on-line update requires more space than the current paper allows; we therefore leave a complete ablation study of update frequency, latency tolerance, and error reduction to future work.

4.4 Empirical Validation Methodology

To comprehensively evaluate whether our agents accurately replicated observed human behavior, we adopted a thorough empirical validation approach grounded in scenario replication and statistical distribution testing. Specifically, we conducted multiple simulation replications to systematically assess the fidelity of our calibrated agent-based models against actual gameplay-derived behavioral data. We form two empirical samples for each archetype–condition pair: $\{o_t^{\text{game}}\}_{t=1}^n$ from gameplay and $\{o_t^{\text{sim}}\}_{t=1}^m$ from the 50 simulation replications. The two-sample K–S statistic is then $D = \sup_x |F_{\text{game}}(x) - F_{\text{sim}}(x)|$, and behavioural validity is accepted whenever the associated *p*-value exceeds the $\alpha = 0.05$ threshold.

We ran 50 independent simulation replications for each identified behavioral archetype (Hoarders, Reactors, Followers) under each of the four experimental conditions (320–323). Each simulation replication captured the ordering behavior of agents across defined phases (Stable, Disruption, Recovery). The distributions of ordering behaviors generated by our simulations were then methodically compared to empirical data obtained from the human-in-the-loop experiments. Validation was conducted at two distinct levels: aggregate validation across all participants and detailed behavioral validation separately for each behavioral archetype. The primary statistical method for empirical validation was the Kolmogorov–Smirnov (K–S) test, chosen specifically because it allows for detailed comparison of distributional similarity rather than testing differences in central tendencies. Formally, the K–S test evaluates the null hypothesis that two samples (simulation-generated and empirically observed) originate from the same underlying distribution. Failure to reject the null hypothesis (typically at the conventional significance threshold of $\alpha = 0.05$) indicates that our behavioral agents successfully capture the distributional characteristics of real-world ordering decisions.

In scenarios where the K–S test indicated significant discrepancies—particularly notable among hoarders under the condition providing disruption-duration information (condition 321)—we performed further diagnostic analyses. These discrepancies prompted iterative recalibration of the behavioral parameters to better align the simulated distributions with empirical observations, ensuring methodological robustness and realism.

This systematic empirical validation framework is a core methodological strength of our approach, carefully ensuring that our integrated PCA, LCP, and HMM-based modeling pipeline produces agent behaviors that realistically reflect complex human decision-making observed in practice.

5   RESULTS OF SIMULATION APPLICATION

To evaluate our methodological integration, we compared the aggregate outcomes of the simulated agents with the empirical results of the participants, explicitly focusing on the inventory, backlog, and cost metrics under conditions (320-323). Validation was performed both at the aggregate level and separately for each behavioral archetype.

Figure 4 shows that our agent-based model accurately replicates trends in Conditions 320 (No Information), 322 (Supplier Capacity Only), and 323 (Combined Information). However, Condition 321 (Disruption Duration Only), despite similar means, exhibited significant distributional differences (aggregate Kolmogorov–Smirnov statistic = 0.47, $p = 0.002$), prompting further behavioral archetype-level analysis. The left panel therefore contains 50 replications of agents whose decision rules were calibrated from the gameplay data via the PCA–LCP–HMM pipeline described in Section 4, while the right panel shows the original participant outcomes.
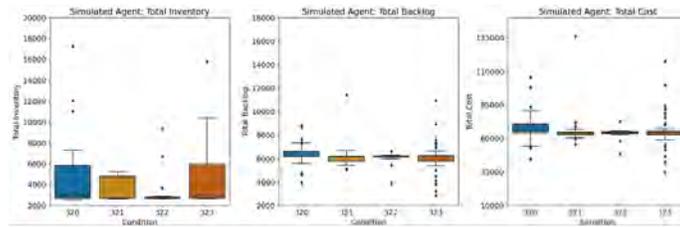
(a) Simulated agents' aggregate results.
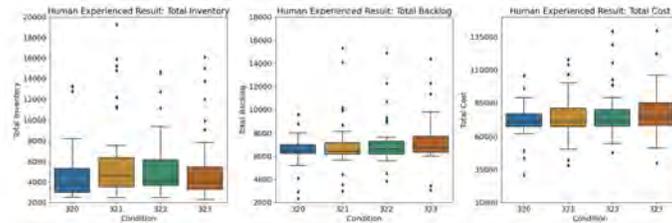
(b) Human participants' aggregate results.

Figure 4: Box plots comparing total inventory, backlog, and cost across conditions.

Further behavioral breakdown (Figure 5) revealed discrepancies driven primarily by Hoarders in Condition 321. This underscores the importance of archetype-level validation, demonstrating our approach's capability to isolate and address nuanced behavioral discrepancies.
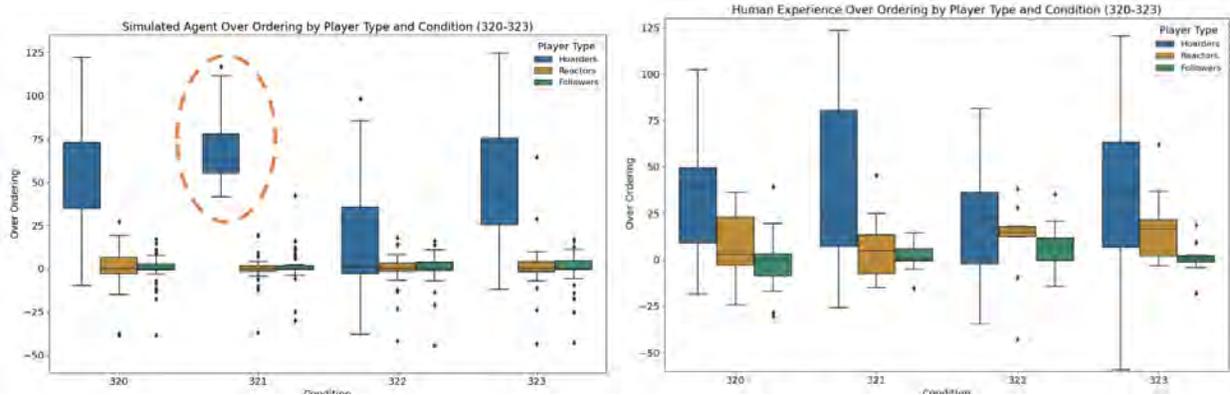
(a) Simulated agents' over-ordering by type.

(b) Human participants' over-ordering by type.

Figure 5: Box plots of over-ordering for Hoarders, Reactors, and Followers under each information-sharing condition. *(a)* 50 simulated replications generated with the PCA–HMM–LCP–calibrated agents. *(b)* empirical outcomes of the 178 human participants.

This clearly demonstrates the effectiveness of our methodological integration, highlighting its ability to detect, interpret, and respond to behavioral discrepancies within supply chain simulations.

6 DISCUSSION AND METHODOLOGICAL INSIGHTS

6.1 Interpretation of Methodological Findings

Our results show that behavioral heterogeneity—captured by the HMM-based classification of *Hoarders, Reactors, and Followers*—strongly shapes how closely agent-based simulations mirror actual human decision-making. Aggregate comparisons between simulations and empirical gamette data align closely under most conditions, with distinct archetype responses:

- **Hoarders** significantly over-order when informed about disruption duration, highlighting panic-induced stockpiling.
- **Reactors** adjust ordering upon receiving disruption signals but revert closer to baseline if disruptions are brief or additional information arises.
- **Followers** closely follow suggested policies, potentially leaving systems vulnerable to unexpected shocks.

The PCA-based segmentation effectively captured transitions through stable, disrupted, and recovery states, offering nuanced insights into how information conditions influenced behavioral archetypes. Condition 321 (Disruption Duration Only) notably elicited the highest Hoarder over-ordering, underscoring the critical interaction between archetypes and information availability.

6.2 Methodological Challenges and Limitations

A key limitation involves small subgroup sizes when segmenting participants by archetype and experimental condition, compromising statistical validation reliability (e.g., Kolmogorov–Smirnov tests). Further, the Longest Common Prefix (LCP) clustering predominantly aligns sequences based on initial behaviors, potentially overlooking critical behavioral shifts occurring later, especially post-disruption. Accurate parameter estimation for agent-based models was challenging due to limited subgroup data, emphasizing the need for improved calibration methods to enhance simulation realism.

7 FUTURE METHODOLOGICAL EXTENSIONS

Future research should refine clustering methodologies, leveraging alternatives like Dynamic Time Warping (DTW) to better capture complex behavioral dynamics around disruptions, especially for archetypes inadequately distinguished by LCP.

Multiphase clustering is another promising approach to explicitly capture evolving participant strategies across pre-disruption, disruption, and recovery phases. Expanding participant numbers for underrepresented archetype-condition combinations, particularly Hoarders under disruption-duration scenarios, will enhance subgroup-level statistical power and generalizability.

Developing behavioral transition matrices from empirical data to represent transitions between archetypes (e.g., Followers to Hoarders) can further improve simulation responsiveness. Lastly, creating phase-specific agents specifically tailored to stable, disrupted, and recovery conditions may more accurately reflect observed behavioral shifts, strengthening strategic planning simulations for pharmaceutical and logistics contexts.

8 CONCLUSION

This study introduces a novel methodological integration of PCA-based state segmentation, LCP clustering, and HMM-based behavioral modeling to systematically identify and characterize distinct behavioral archetypes—Hoarders, Reactors, and Followers—from participatory supply chain gameplay data.

Our methodology initially validated well at the aggregate level, accurately reproducing average ordering behaviors and effectively representing behavioral variations during supply chain disruptions. However, deeper validation—performed separately by archetype under varying informational conditions—revealed

important discrepancies. Specifically, the LCP clustering method successfully grouped initial behaviors but did not consistently produce simulation agents that validated accurately under all conditions, as clearly illustrated by the Hoarder archetype in Condition 321 (disruption-duration information only).

This important finding highlights a key methodological insight: clustering methods that effectively group empirical behaviors do not automatically guarantee valid or accurate agent behaviors when implemented in simulations.

Future methodological extensions include exploring alternative sequence-alignment and clustering methods, applying multiphase clustering to better represent behavioral shifts throughout disruptions, and refining behavioral transition matrices to dynamically capture evolving decision-making patterns. These ongoing methodological developments will significantly enhance the rigor and practical relevance of pharmaceutical supply chain simulations, as well as broader applications in other complex supply-chain contexts characterized by uncertainty and human behavioral complexity.

REFERENCES

Anand, N., D. Meijer, J. Van Duin, L. Tavasszy, and S. Meijer. 2016, October. "Validation of an agent based model using a participatory simulation gaming approach: The case of city logistics". *Transportation Research Part C: Emerging Technologies* 71:489–499 https://doi.org/10.1016/j.trc.2016.08.002.

Azghandi, R., J. Griffin, and M. S. Jalali. 2018, January. "Minimization of Drug Shortages in Pharmaceutical Supply Chains: A Simulation-Based Analysis of Drug Recall Patterns and Inventory Policies". *Complexity* 2018(1):6348413 https://doi.org/10.1155/2018/6348413.

Bousquet, F., and G. Trébuil. 2005. "Synergies between muti-agent systems and role-playing games in companion modeling for integrated natural resource management in southeast Asia". Asian Institute of Technology.

Croson, R., K. Donohue, E. Katok, and J. Sterman. 2014, February. "Order Stability in Supply Chains: Coordination Risk and the Role of Coordination Stock". *Production and Operations Management* 23(2):176–196 https://doi.org/10.1111/j.1937-5956.2012.01422.x.

Doroudi, R., R. Azghandi, Z. Feric, O. Mohaddesi, Y. Sun, J. Griffin, *et al*. 2018, December. "An Integrated Simulation Framework for Examining Resiliency in Pharmaceutical Supply Chains Considering Human Behaviors". In *2018 Winter Simulation Conference (WSC)*, 88–99. Gothenburg, Sweden: IEEE https://doi.org/10.1109/WSC.2018.8632387.

Ergun, , J. Griffin, N. Chicoine, M. Gong, O. Mohaddesi, Z. Raziei, *et al*. 2022. "Supply Chain Resilience: Impact of Stakeholder Behavior and Trustworthy Information Sharing with a Case Study on Pharmaceutical Supply Chains". In *Tutorials in Operations Research: Emerging and Impactful Topics in Operations*, edited by M. Chou, H. Gibson, B. Staats, D. Shier, and H. Greenberg, 133–159. Catonsville, MD: INFORMS https://doi.org/10.1287/educ.2022.0246.

Georgiadis, P., D. Vlachos, and E. Iakovou. 2005. "A System Dynamics Modeling Framework for the Strategic Supply Chain Management of Food Chains". *Journal of Food Engineering* 70:351–364 https://doi.org/10.1016/j.jfoodeng.2004.06.030.

Guyot, P., and S. Honiden. 2006. "Agent-Based Participatory Simulations: Merging Multi-Agent Systems and Role-Playing Games". *Journal of Artificial Societies and Social Simulation* 9(4).

Lee, H. L., V. Padmanabhan, and S. Whang. 1997. "Information distortion in a supply chain: The bullwhip effect". *Management science* 43(4):546–558.

Marsella, S. C., D. V. Pynadath, and S. J. Read. 2004. "PsychSim: Agent-based modeling of social interactions and influence". In *Sixth International Conference on Cognitive Modeling*, 243–248. Psychology Press.

Meijer, S. 2009, January. *The organisation of transactions: Studying supply networks using gaming simulation*. Brill | Wageningen Academic https://doi.org/10.3920/978-90-8686-659-5.

Minegishi, S., and D. Thiel. 2000, dec. "System Dynamics Modeling and Simulation of a Particular Food Supply Chain". *Simulation Practice and Theory* 8(5):321–339 https://doi.org/10.1016/S0928-4869(00)00026-4.

Mohaddesi, O., J. Griffin, O. Ergun, D. Kaeli, S. Marsella, and C. Harteveld. 2022, April. "To Trust or to Stockpile: Modeling Human-Simulation Interaction Supply Chain Shortages". In *CHI Conference on Human Factors in Computing Systems*, 1–18 https://doi.org/10.1145/3491102.3502089. arXiv:2201.02694 [cs].

Mohaddesi, O., Y. Sun, R. Azghandi, R. Doroudi, S. Snodgrass, O. Ergun, *et al*. 2020, April. "Introducing Gamettes: A Playful Approach for Capturing Decision-Making for Informing Behavioral Models". In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–13. Honolulu, HI, USA: ACM https://doi.org/10.1145/3313831.3376571.

Rabiner, L. R. 1989. "A tutorial on hidden Markov models and selected applications in speech recognition". *Proceedings of the IEEE* 77(2):257–286.

Sheffi, Y. 2005. *The Resilient Enterprise: Overcoming Vulnerability for Competitive Advantage*. Cambridge, MA: MIT Press.

Snyder, L. V., M. P. Scaparra, M. S. Daskin, and R. L. Church. 2014. "Planning for Disruptions in Supply Chain Networks". In *INFORMS TutORials in Operations Research*, 234–257 https://doi.org/10.1287/educ.1063.0025. Published online: 14 Oct 2014.

Snyder, L. V., and Z.-J. M. Shen. 2019. *Fundamentals of supply chain theory*. John Wiley & Sons.

Sterman, J. D. 1989. "Modeling managerial behavior: Misperceptions of feedback in a dynamic decision making experiment". *Management science* 35(3):321–339.

Sterman, J. D., and G. Dogan. 2015, November. ""I'm not hoarding, I'm just stocking up before the hoarders get here.": Behavioral causes of phantom ordering in supply chains". *Journal of Operations Management* 39-40(1):6–22 https://doi.org/10.1016/j.jom.2015.07.002.

Strohhecker, J., and A. Größler. 2015. "Dynamic Life Cycle Inventory Policies for Consumer Goods with Severe Production Disruptions". In *Proceedings of the 22nd EurOMA Conference*, 78.

ACKNOWLEDGMENTS

AUTHOR BIOGRAPHIES

SOURI SASANFAR is a PhD student of Industrial Engineering at Northeastern University. Her email address is sasanfar.s@northeastern.edu.

OMID MOHADDESI is a Data Scientist and Researcher. His email address is mohaddesi.s@northeastern.edu.

NOAH CHICOINE is a PhD student of Industrial Engineering at Northeastern University. His email address is chicoine.n@northeastern.edu.

JACQUELINE GRIFFIN is Associate Professor of Industrial Engineering at Northeastern University developing OR models for complex decision-making in health and humanitarian systems. Email:ja.griffin@northeastern.edu.

CASPER HARTEVELD is a Professor of Game Design at Northeastern University, focusing on the use of games to study decision making and engage diverse audiences on social issues. Email:c.harteveld@northeastern.edu.