

Too Many Specialists: Emergent Inefficiencies and Bottlenecks for Multi-agent Ad-hoc Collaboration

Extended Abstract

Benjamin Panny
University of Pittsburgh
Pittsburgh, PA, United States
BMP83@pitt.edu

Shashank Mehrotra
Honda Research Institute USA, Inc.
San Jose, CA, United States
shashank_mehrotra@honda-ri.com

Zahra Zahedi
Honda Research Institute USA, Inc.
San Jose, CA, United States
zahra_zahedi@honda-ri.com

Teruhisa Misu
Honda Research Institute USA, Inc.
San Jose, CA, United States
tmisu@honda-ri.com

Kumar Akash
Honda Research Institute USA, Inc.
San Jose, CA, United States
kakash@honda-ri.com

ABSTRACT

Computational models of collaboration without prior coordination often overlook how heterogeneous agent traits and complex task structures jointly produce systemic bottlenecks, inefficiencies, and contribution inequalities. We address this by using an agent-based model of ad-hoc teamwork in a kitchen environment. Our model integrates diverse agent personas with tasks that combine serial and parallel dependencies. We identify a *specialist's dilemma*, where rigid role assertion generates system-level bottlenecks, amplifies workload inequality, and fosters fragmented, homophilous networks. We also find that team size and communication overhead interact with problem structure to generate diminishing returns and redundant collaboration. Linking micro-level behavior to macro-level outcomes provides insights into emergent collaboration and design principles for effective multi-agent teamwork.

KEYWORDS

Multi-agent systems, Agent-based models, Social simulation, Collaborative Networks

ACM Reference Format:

Benjamin Panny, Shashank Mehrotra, Zahra Zahedi, Teruhisa Misu, and Kumar Akash. 2026. Too Many Specialists: Emergent Inefficiencies and Bottlenecks for Multi-agent Ad-hoc Collaboration: Extended Abstract. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 3 pages. <https://doi.org/10.65109/CYXP1261>

1 INTRODUCTION

Collaboration requires agents to balance initiative, adaptability, and mutual awareness under uncertainty [13], particularly in ad-hoc settings where agents lack prior coordination. Computational models of teamwork can alleviate these challenges by predicting the conditions under which collaboration will succeed or fail [11, 12].

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Proceedings of Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)* (AAMAS '26).

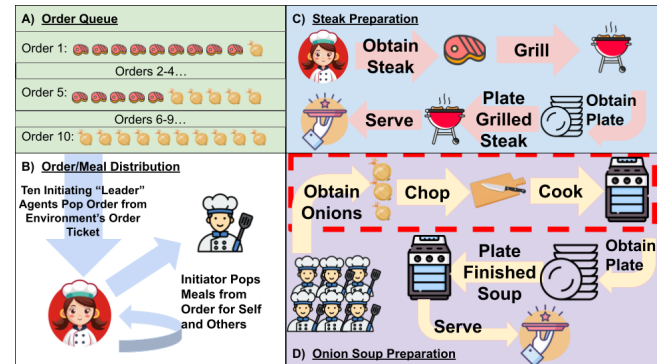


Figure 1: Kitchen Environment and Problem Structure. (A) Orders vary the ratios of parallel (Soup) vs. serial (Steak) meals. (B) Agents claiming tasks become “leaders.” (C) Serial Steak sequence. (D) Parallelizable Soup cycle.

Frameworks such as the Transactive Systems framework [4, 10, 14, 15] and Collective Adaptation [2] highlight how socially networked reasoning and knowledge enable team reconfiguration when environments change. However, current models rarely scale beyond small teams or account for serial-parallel task structures [6, 7]. As a result, few models formally link macro-level performance, bottlenecks, and inequalities observed in real teams to micro-level agent behaviors observed in the individuals that constitute these teams.

In this paper, we address this gap through an agent-based model of ad-hoc teamwork in a physically grounded kitchen environment. By introducing heterogeneous agent personas endowed with social traits—such as collaboration initiative, agreeableness, and skill assertion—we show how bottlenecks, inequalities, and subgroup formations can emerge from simple local interaction rules. Our findings offer mechanistic levers for human-AI team design: seeding proactive leadership, mitigating rigid specialization, and bridging silos through adaptive coordination.

2 COLLABORATION MODEL

We developed an agent-based model (ABM) where agents collaborate to complete cooking tasks in a 2D grid environment (Cf.

[1]). We vary team size [5], communication costs [9], and agent social/transactive factors [2, 3] (Figure 1).

2.1 Serial and Parallel Kitchen Tasks

Real-world collaboration tasks are rarely perfectly divisible. We designed two recipes modeling this constraint: **(1) Steak**: A serial task with a fixed sequence (get meat, grill, plate, serve). **(2) Onion Soup**: A parallelizable task with repeated substeps (get onion, chop onion, place in stove) that multiple agents can perform concurrently.

2.2 Agent Attributes Determine Agent Rule-sets

Agents' attributes result in actions via a decentralized Affordance-Context-Action (ACA) loop governed by four randomly assigned traits. **Agreeableness** controls the likelihood of accepting help requests, enabling tests of passive cooperation. **Collaboration Initiative** determines whether an agent initiates interactions, allowing us to measure the impact of proactive communicators. **Task Distribution Preference** governs the choice between starting unclaimed tasks or joining existing ones, introducing differences in organizational strategy (i.e., complete started tasks before starting new tasks or not). Finally, **Skill Assertion** determines whether agents refuse tasks outside their expertise; this operationalizes the *specialist's dilemma* [3, 8] by introducing tension between individual competence and team efficiency.

3 RESULTS

3.1 Scaling Effects and Diminishing Returns

Team size and communication affected serial and parallel tasks differently. For the parallelizable onion soup task (Figure 2a), larger teams and lower communication costs produced significantly more meals, though this was subject to diminishing returns. For the serial steak task, larger teams and *higher* communication costs produced steaks faster than their counterparts with low communication costs (Figure 2b). This is because low communication costs promoted larger teams that were effective for parallel tasks, but redundant for serial tasks. Thus, coordination overhead can be beneficial when it restricts redundant team building.




Figure 2: (a) Parallel tasks (Onion Soup) scale with team size and broad communication. (b) Serial tasks (Steak) benefit from higher communication cost, which results in smaller, more efficient teams.

3.2 Bottlenecks from Specialization

We identified a *specialist's dilemma*, where agents asserting their specific skills creates system-level bottlenecks.

A high proportion (100%) of skill-asserting agents resulted in a significant decrease in total meals completed. However, this was partially rectified when specialists were also proactive communicators with task distribution preferences, as this persona enabled specialists to find tasks that fit their specialty earlier in the simulation.

Skill assertion also emerged as the primary driver of workload imbalance. When agents rigidly adhered to roles (e.g., "I only grill"), they remained idle when they could have supported agents working on other tasks.

3.3 Emergent Silos

Agent traits significantly altered the topology of the emergent collaboration network. High levels of skill assertion led to high network assortativity (0.578) and modularity (Figure 3a). Agents naturally formed homophilous clusters, interacting only with those sharing their specialty. This "siloeing" effect fragmented the network into disconnected components, increasing the average shortest path length and reducing global integration. As skill assertion decreased, the network transformed from isolated specialist clusters into a dense, integrated collaborative web.




Figure 3: (a) When 100% of agents assert skills, high assortativity and disconnected components emerge. (b) When 0% of agents assert skills, a dense, integrated collaborative web forms.

4 DISCUSSION & CONCLUSION

Micro-level agent traits and task structures interact to shape the success and failure of ad-hoc teamwork. We show that rigid specialization leads to the specialist's dilemma—generating bottlenecks, inequality, and fragmented social silos. We also show that increasing communication costs can be beneficial if it reduces redundant teamwork. These findings suggest that effective MAS design must actively counteract the tendency for specialists to isolate, perhaps by algorithmically encouraging agents to occasionally accept tasks outside their primary expertise to maintain system-wide resilience.

REFERENCES

- [1] Micah Carroll, Rohin Shah, Mark Ho, Tom Griffiths, Pieter Abbeel, and Anca Dragan. 2020. Overcooked-ai: A benchmark for multi-agent learning under partial observability. In *Proceedings of the 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*(2020). 2374–2380.
- [2] Mirta Galesic, Daniel Barkoczi, Andrew M. Berdahl, Dora Biro, Giuseppe Carbone, Ilaria Giannoccaro, Robert L. Goldstone, Cleotilde Gonzalez, Anne Kandler, Albert B. Kao, Rachel Kendal, Michelle Kline, Eun Lee, Giovanni Francesco Mas-sari, Alex Mesoudi, Henrik Olsson, Niccolo Pescetelli, Sabina J. Sloman, Paul E. Smaldino, and Daniel L. Stein. 2023. Beyond collective intelligence: Collective adaptation. *Journal of The Royal Society Interface* 20, 200 (March 2023), 20220736. <https://doi.org/10.1098/rsif.2022.0736> Publisher: Royal Society.
- [3] Pranav Gupta. 2022. *Transactive Systems Model of Collective Intelligence: The Emergence and Regulation of Collective Attention, Memory, and Reasoning*. thesis. Carnegie Mellon University. <https://doi.org/10.1184/R1/20039555.v1>
- [4] Pranav Gupta and Anita Williams Woolley. 2021. Articulating the Role of Artificial Intelligence in Collective Intelligence: A Transactive Systems Framework. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 65, 1 (Sept. 2021), 670–674. <https://doi.org/10.1177/1071181321651354c>
- [5] Vincent Larivière, Yves Gingras, Cassidy R. Sugimoto, and Andrew Tsou. 2015. Team size matters: Collaboration and scientific impact since 1900. *Journal of the Association for Information Science and Technology* 66, 7 (2015), 1323–1332. <https://doi.org/10.1002/asi.23266> _eprint: <https://asistdl.onlinelibrary.wiley.com/doi/pdf/10.1002/asi.23266>
- [6] Soo Ling Lim, Peter J Bentley, Randall S Peterson, Xiaoran Hu, and JoEllyn Prouty McLaren. 2023. Kill chaos with kindness: Agreeableness improves team performance under uncertainty. *Collective Intelligence* 2, 1 (Jan. 2023), 26339137231158584. <https://doi.org/10.1177/26339137231158584> Publisher: SAGE Publications.
- [7] Winter Mason and Duncan J. Watts. 2012. Collaborative learning in networks. *Proceedings of the National Academy of Sciences* 109, 3 (Jan. 2012), 764–769. <https://doi.org/10.1073/pnas.1110069108> Publisher: Proceedings of the National Academy of Sciences.
- [8] Mancur Olson. 1965. *The logic of collective action: public goods and the theory of groups*. Harvard Univ. Press, Cambridge, Mass. <https://www.hup.harvard.edu/catalog.php?isbn=9780674537514> Number: 124 Pages: 176 tex.added-at: 2011-02-21T14:53:36.000+0100 tex.interhash: ca423582ffdb8545c0567c16786fa08e tex.intrahash: d52f789a4c4711ba434fc8106806af3b tex.timestamp: 2021-11-04T09:43:47.000+0100.
- [9] Michelle O’Daniel and Alan H. Rosenstein. 2008. Professional Communication and Team Collaboration. In *Patient Safety and Quality: An Evidence-Based Handbook for Nurses*. Agency for Healthcare Research and Quality (US). <https://www.ncbi.nlm.nih.gov/books/NBK2637/>
- [10] Yuqing Ren and Linda Argote. 2011. Transactive memory systems 1985–2010: An integrative framework of key dimensions, antecedents, and consequences. *The Academy of Management Annals* 5, 1 (2011), 189–229. <https://doi.org/10.1080/19416520.2011.590300> Place: United Kingdom Publisher: Taylor & Francis.
- [11] Eduardo Salas, Dana E. Sims, and C. Shawn Burke. 2005. Is there a “Big Five” in Teamwork? *Small Group Research* 36, 5 (Oct. 2005), 555–599. <https://doi.org/10.1177/1046496405277134> Publisher: SAGE Publications Inc.
- [12] Randall Spain, Michael Geden, Wookhee Min, Bradford W Mott, and James C Lester. 2019. Toward Computational Models of Team Effectiveness with Natural Language Processing.. In *TTW@ AIED*. 30–39.
- [13] Milind Tambe. 1997. Towards flexible teamwork. *Journal of artificial intelligence research* 7 (1997), 83–124.
- [14] Daniel M. Wegner. 1987. Transactive Memory: A Contemporary Analysis of the Group Mind. In *Theories of Group Behavior*, Brian Mullen and George R. Goethals (Eds.). Springer, New York, NY, 185–208. https://doi.org/10.1007/978-1-4612-4634-3_9
- [15] Daniel M. Wegner, Toni Giuliano, and Paula T. Hertel. 1985. Cognitive Interdependence in Close Relationships. In *Compatible and Incompatible Relationships*, William Ickes (Ed.). Springer New York, New York, NY, 253–276. https://doi.org/10.1007/978-1-4612-5044-9_12