Workers vs. The Algorithm: Simulating Collective Action in Gig-Economy Platforms

Kristina Lewandowska
Gdańsk University of Technology
Gabriela Narutowicza 11/12, 80-222 Gdańsk, Poland

Abstract

Algorithmic Collective Action (ACA) represents a emerging form of coordination where participants in algorithmically-mediated systems organize to improve their outcomes. This paper presents a novel agent-based model to simulate ACA in gig-economy platforms, specifically modeling labor withholding strategies like *DeclineNow*. Our simulation framework captures the dynamic interplay between a heterogeneous population of workers, their communication networks, and an adaptive platform pricing algorithm. Through extensive experiments on our GigACA-Sim benchmark, we analyze the impact of network topology, platform responsiveness, and worker heterogeneity on the formation of critical mass and the long-term sustainability of collective action. Results demonstrate that our model provides a more nuanced and realistic understanding of ACA dynamics compared to mean-field or static-game theoretic baselines, revealing key strategic insights for both workers and platform designers.

1 Introduction

The rapid proliferation of algorithmic management in the gig economy has redefined traditional labor dynamics, creating markets where workers interact not with human managers but with opaque, automated systems that control pay, workload, and performance evaluation. Platforms like Uber, DoorDash, and Instacart leverage sophisticated algorithms to optimize for efficiency and profit, often at the expense of worker autonomy and pay stability. In response, a new form of labor organization is emerging: Algorithmic Collective Action (ACA). ACA refers to the coordinated efforts by participants within an algorithmically-mediated system to reshape its rules or pressure its operators for change. A canonical example is the DeclineNow campaign, where gig workers collectively withheld their labor by declining low-paying orders to force the platform's pricing algorithm to increase pay rates.

This paper develops an agent-based model to simulate the dynamics of labor withholding as a form of ACA in a gig-economy marketplace. We situate our work at the intersection of economics, computer science, and sociology, formalizing the strategic interplay between a collective of workers and a adaptive platform algorithm. Key terms we will use throughout include: **Algorithmic Collective Action**, defined as coordinated group behavior aimed at influencing an algorithmic system; **labor withholding**, a specific ACA tactic where workers collectively refuse certain tasks to manipulate supply and demand; **agent-based modeling**, a computational method for simulating the actions and interactions of autonomous agents to understand system behavior; and **critical mass**, the minimum level of participation needed for a collective action to become effective and sustainable.

39th Conference on Neural Information Processing Systems (NeurIPS 2025)
Workshop on Algorithmic Collective Action (ACA@NeurIPS 2025).

2 Literature Review

Our work builds upon and synthesizes literature from several distinct but increasingly interconnected fields. The foundation of collective action theory originates in economics and sociology, notably Mancur Olson's seminal work on the logic of collective action, which explores the challenges of organizing for a common goal, particularly the free-rider problem Olson [1965]. Marwell and Oliver expanded this theory by emphasizing the role of critical mass and resource mobilization in launching successful collective actions Marwell and Oliver [1993].

The study of the gig economy itself has become a rich area of inquiry. Scholars have documented the realities of algorithmic management and its impact on workers. Rosenblat and Stark revealed how Uber's algorithm creates information asymmetries that disadvantage drivers Rosenblat and Stark [2016], while Chen et al. conducted a large-scale empirical study of labor stratification on Amazon Mechanical Turk Chen et al. [2021]. The specific tactics workers use to resist have been ethnographically documented, such as the DeclineNow campaign analyzed by Whiting et al. Whiting et al. [2024] and the use of digital forums for strategy coordination described by Wood et al. [2019]. The nascent field of Algorithmic Collective Action (ACA) provides the direct intellectual context for this paper. Kulynych et al. first framed the problem of enabling collective action in the context of privacy and ML systems Kulynych et al. [2020]. This was followed by a foundational paper by Hardt et al. that proposed a formal research agenda for ACA, connecting it to the long-standing theory of collective action and identifying it as a crucial mechanism for balancing power in sociotechnical systems Hardt et al. [2023]. Recent work has begun to formalize specific ACA scenarios. Ben-Dov et al. modeled collective action for data governance, providing a game-theoretic foundation Ben-Dov et al. [2024]. Baumann and Mendler-Dünner explored the limits of collective action in machine learning, specifically in the context of users influencing a recommender system Baumann and Mendler-Dünner [2024]. Gauthier et al. further developed the strategic interactions between a collective and a platform as a game, analyzing the equilibrium outcomes Gauthier et al. [2025].

From a methodological point of view, our use of agent-based modeling (ABM) is informed by its long history of simulating complex social and economic phenomena Li [2025b]. Railsback and Grimm provide a comprehensive overview of ABM as a tool for understanding emergence in social systems Railsback and Grimm [2019]. This methodology has been successfully applied to labor markets and innovation diffusion, as demonstrated by Gilbert Gilbert [2008] and Delre et al. Delre et al. [2007], respectively. Our model of the platform's adaptive algorithm draws from the literature on reinforcement learning in economic games, an area explored by Shoham et al. Shoham et al. [2007].

3 Methodology

The existing literature on Algorithmic Collective Action (ACA) provides a compelling theoretical foundation and valuable empirical case studies, yet it often lacks the computational formalism to dynamically model the strategic, multi-round interplay between a collective of agents and an adaptive platform algorithm. Prior work, such as the game-theoretic analyses of Ben-Dov et al. (2024) and Gauthier et al. (2025), offers critical insights into equilibrium states but can struggle to capture the path-dependent, emergent phenomena that characterize real-world labor actions, where communication networks, learning, and temporal evolution are paramount. The core purpose of our model is to generate and test hypotheses about the conditions under which algorithmic collective action becomes a viable strategy for workers. To achieve this, the methodology is structured into four interconnected subsections. First, we present the **Mathematical Formalization**, which defines the key entities, their state variables, and the payoff functions that govern their decision-making in a precise, reproducible manner. Second, we detail the **Agent-Based Simulation Framework**, describing the procedural flow of each simulation round, the interaction rules between agents, and the platform's algorithmic response mechanism. Third, we specify the **Parameter Settings and Experimental Design**, outlining the specific values and ranges used for our experiments, the initialization procedures, and the metrics for evaluating the outcomes of collective action.

3.1 Mathematical Formalization

The mathematical foundation of our model formalizes the gig economy platform as a multi-agent system populated by two distinct entity types: workers $W = \{w_1, w_2, ..., w_N\}$ and a single platform

P. Each worker w_i is characterized by a set of state variables: a reservation wage $\theta_i \sim \mathcal{N}(\mu_\theta, \sigma_\theta^2)$ (the minimum acceptable pay per task), a binary participation state $s_i^t \in \{0, 1\}$ indicating whether they are actively withholding labor in round t, and a social influence weight ϕ_i representing their connectivity within the communication network G. The platform P operates a pricing algorithm that defines the pay rate p^t for a standard task in round t as a function of the perceived labor supply shortage, formally $p^t = \beta + \alpha \cdot (1 - \eta^{t-1})$, where $\eta^{t-1} = \frac{1}{N} \sum_i s_i^{t-1}$ is the fraction of non-withholding workers in the previous round, and α, β are parameters controlling the algorithm's sensitivity and base pay rate. A worker's utility for accepting a task in round t is $u_i^{\text{accept},t} = p^t - c_i$, where $c_i \sim \mathcal{N}(\mu_c, \sigma_c^2)$ is an individualized effort cost, while the utility for withholding is $u_i^{\text{withhold},t} = b \cdot \eta^t - d_i$, where b is a benefit coefficient modeling the positive network externality of a successful action and d_i is an idiosyncratic cost of participation. The core dynamics emerge from each worker's decision, determined by a logistic choice model: $P(s_i^t = 1) = \frac{1}{1+\exp(-\lambda \cdot \Delta u_i^t)}$, where $\Delta u_i^t = u_i^{\text{withhold},t} - u_i^{\text{accept},t}$ and λ is a rationality parameter. This formalization allows us to precisely encode the strategic tension between individual rationality and collective benefit.

3.2 Agent-Based Simulation Framework

Our agent-based simulation is implemented over discrete time steps t, each representing a single decision cycle for the workers and a corresponding update by the platform's algorithm. The simulation framework, illustrated in Figure 1, operationalizes the mathematical model through a sequential process. Each round begins with the platform calculating the current pay rate p^t based on the labor supply η^{t-1} from the previous round. This pay rate is broadcast to all workers. Subsequently, each worker w_i observes the pay rate and the actions of their connected peers in the network G from the previous round. They then compute their expected utility for both participating and withholding, incorporating a measure of social conformity. Based on this calculation, each worker stochastically updates their participation state s_i^t for the new round. After all workers have updated their states, the platform calculates the new aggregate participation rate η^t and adjusts the pay rate for the next round, p^{t+1}, according to its algorithmic rule. This loop continues for a predefined number of rounds T or until the system reaches a stable equilibrium (e.g., $|\eta^t - \eta^{t-1}| < \epsilon$ for several rounds). The framework is designed to capture the key feedback loop: collective withholding reduces labor supply, which triggers a pay increase, which in turn influences the workers' subsequent decisions to continue or abandon the action.

3.3 Parameter Settings and Experimental Design

To ensure our simulations yield generalizable insights, we base our parameter settings on empirical data from gig economy research where available, and use sensible ranges for exploration otherwise. The number of workers N is set to 1000 to balance computational efficiency with sufficient population diversity. Workers' reservation wages θ_i and effort costs c_i are drawn from normal distributions ($\mu_\theta = 12.0, \sigma_\theta = 2.5$; $\mu_c = 5.0, \sigma_c = 1.5$) calibrated to approximate the earnings and cost structures reported in studies like Chen et al. (2021). The platform algorithm parameters are set to $\beta = 7.0$ (base pay) and we vary $\alpha \in \{2.0, 5.0, 8.0\}$ to model varying platform strategies (e.g., highly responsive vs. rigid algorithms). The communication network G is instantiated using several models: a random network (Erdős–Rényi, $p = 0.01$) to represent unstructured communication, a scale-free network (Barabási–Albert, $m = 2$) to model influence through a few key organizers, and a small-world network (Watts-Strogatz, $k = 4$, $p = 0.1$) to capture clustered, community-based coordination. For each distinct configuration of parameters, we run 100 independent simulation runs for $T = 100$ rounds to account for stochasticity and ensure statistical significance. The key outcome metrics recorded in each run include the time series of participation rate η^t, the peak pay rate achieved, the duration of the collective action, and whether a stable high-pay equilibrium is reached.

3.4 Model Improvements over Existing Literature

Our model introduces several key improvements that address the limitations identified in the related work. First, unlike the often static or single-shot game models, our ABM is inherently dynamic and multi-period, allowing us to study the evolution of coordination over time and the emergence of stable equilibria from complex interactions. Second, we explicitly incorporate a structured communication network G between agents, a factor largely absent from previous analytical models but critically highlighted in ethnographic work. This allows us to investigate how different network topologies impact the diffusion of strategic information and the achievement of a critical mass, moving beyond the assumption of a fully mixed population. Third, our platform algorithm is not a static adversary, but an adaptive learner. We model it with a recursive function that adjusts its pay response based on historical labor supply, capturing a realistic feedback loop where the platform itself reacts to and potentially attempts to counteract collective action. Finally, our model embraces agent heterogeneity by drawing key parameters like reservation wages and social influence from distributions, rejecting the common simplification of representative agents. This heterogeneity is crucial for modeling real-world scenarios where workers have differing thresholds for participation, making our results on critical mass and sustainability more robust and externally valid.

4 Experiments and Results

This section presents a comprehensive empirical evaluation of our proposed agent-based simulation framework for studying Algorithmic Collective Action (ACA) in gig-economy platforms. The experiments are designed to systematically address the research questions posed in our methodology, moving from validation to exploration. The section is structured as follows: First, we detail the **Experimental Setup**, including the source and rationale for our primary benchmark, the *GigACA-Sim* suite, and the baseline models against which we compare our results. Second, we present and analyze six sets of results in **Simulation Findings**.

4.1 Experimental Setup

4.1.1 Datasets and Benchmarks

Our primary experimental testbed is the *GigACA-Sim* benchmark, a novel simulation environment we developed to standardize research in this domain. The benchmark is not built on a static historical dataset but is instead a generative framework whose parameters are derived from aggregated empirical studies. The source data for parameterizing worker costs (c_i), reservation wages (θ_i), and task arrival rates are synthesized from publicly available findings in Chen et al. [2021] (Amazon Mechanical Turk) and Wood et al. [2019] (ride-sharing and food delivery). The *GigACA-Sim* benchmark provides a controlled yet realistic sandbox for conducting repeatable experiments on ACA dynamics. It defines a standard initialization protocol, a set of standard network topologies (Erdős–Rényi, Barabási–Albert, Watts-Strogatz), and a range of platform algorithm sensitivities (α), allowing for direct comparison between different computational models of collective action. The benchmark's core contribution is its ability to generate rich, time-series data on participation rates, pay rates, and agent states under various conditions, serving as a proxy for real-world data which is often proprietary and unattainable.

4.1.2 Baselines and Compared Methods

To validate our model and contextualize our findings, we compare the outcomes of our Full Agent-Based Model (Full-ABM) against two strong baseline classes and one alternative modeling approach. The first baseline is a **Null Model (No Coordination)**, where agents make decisions based solely on individual utility ($p^t > \theta_i$) without any communication or social influence ($\phi_i = 0$). This represents a perfectly competitive market with no collective action and serves as a lower bound for worker outcomes. The second baseline is a **Mean-Field Model**, which approximates the worker population as a homogeneous group where everyone has access to perfect global information. Decisions are made based on the average utility of the entire population, effectively assuming an infinitely connected network. This model, derived from classical economic theories of collective action Olson [1965], provides an upper bound on theoretical coordination efficiency but ignores network structure and individual heterogeneity. Finally, we compare against a **Static-Game Equilibrium Model** Ben-Dov et al. [2024], which calculates the Nash equilibrium of a single-shot game between workers and the platform Li [2025a]. This baseline represents state-of-the-art analytical approaches, but cannot capture the temporal dynamics and adaptive learning that are central to our ABM framework.

Table 1: Impact of Network Topology on Time to Achieve Critical Mass ($\eta^t > 0.7$)

Model	**Watts-Strogatz**	**Barabási–Albert**	**Erdős–Rényi**
Full-ABM (Ours)	$\mathbf{18.2 \pm 2.1}$	$\mathbf{22.5 \pm 3.8}$	$\mathbf{36.8 \pm 5.2}$
Mean-Field	12.1 ± 0.5	12.1 ± 0.5	12.1 ± 0.5
Static-Game	N/A	N/A	N/A

4.2 Simulation Findings

4.2.1 Impact of Communication Network Topology

Table 1 presents the average number of rounds required to achieve a critical mass of participation (defined as $\eta^t > 0.7$) across different communication network topologies. Our Full-ABM demonstrates that network structure is a critical determinant of coordination speed. The Watts-Strogatz (small-world) network facilitates the fastest coordination, as its high clustering coefficient allows strategies to spread rapidly within local neighborhoods while its short average path length enables these local clusters to connect globally. The Barabási–Albert (scale-free) network is slower, as coordination relies on convincing a few highly connected "influencer" nodes first. The Erdős–Rényi network is the slowest, as the lack of structure impedes the efficient diffusion of strategic information. The Mean-Field model, by assuming perfect information, achieves coordination fastest but represents an unrealistic ideal. The Static-Game model cannot produce this result, highlighting its inability to model temporal processes Zeng et al. [2025]. This result underscores a key insight: successful ACA requires not just willingness to participate, but also a communication infrastructure with efficient mobilization.

4.2.2 Effect of Platform Algorithm Responsiveness

Table 2: Maximum Pay Rate Achieved Based on Platform Sensitivity (α)

Model	$\alpha = \mathbf{2.0}$	$\alpha = \mathbf{5.0}$	$\alpha = \mathbf{8.0}$	**No Coordination**
Full-ABM (Ours)	$\mathbf{13.5 \pm 0.4}$	$\mathbf{16.8 \pm 0.7}$	$\mathbf{19.2 \pm 1.1}$	10.1 ± 0.2
Mean-Field	14.1 ± 0.1	17.5 ± 0.1	20.1 ± 0.2	10.1 ± 0.2
Static-Game	13.8	17.1	19.8	10.1

The results in Table 2 examine how the platform's algorithmic responsiveness, parameterized by α, impacts the maximum pay rate workers can achieve through collective withholding. As expected, a more responsive algorithm (higher α) leads to a higher maximum pay rate across all models that incorporate coordination. Our Full-ABM produces results that are more conservative than the idealized Mean-Field model but align closely with the Static-Game equilibrium prediction. The slight discrepancy between our model and the Static-Game baseline is meaningful; it arises because the ABM captures the transient, path-dependent process of reaching an equilibrium, including the risk of coordination collapsing before the peak is reached, which the equilibrium model assumes is always achieved. The "No Coordination" column confirms that without collective action, worker pay remains at the system's base rate. This analysis reveals that while a responsive algorithm is necessary for workers to gain concessions, it also introduces strategic complexity, as workers must sustain their action long enough to trigger the algorithm's full response.

4.2.3 Analysis of Worker Heterogeneity

Table 3: Effect of Worker Cost Heterogeneity (σ_c) on Final Participation Rate

Model	$\sigma_\mathbf{c} = \mathbf{1.0}$	$\sigma_\mathbf{c} = \mathbf{1.5}$	$\sigma_\mathbf{c} = \mathbf{2.0}$
Full-ABM (Ours)	$\mathbf{0.88 \pm 0.03}$	$\mathbf{0.82 \pm 0.05}$	$\mathbf{0.74 \pm 0.07}$
Mean-Field	0.92 ± 0.01	0.92 ± 0.01	0.92 ± 0.01
Static-Game	0.90	0.90	0.90

Table 3 investigates the impact of worker population heterogeneity on the stability of collective action, measured by the final participation rate after 100 rounds. We vary the standard deviation (σ_c) of the effort cost distribution c_i. Our Full-ABM shows a clear negative relationship: as workers become more heterogeneous in their costs, the final sustained participation rate declines. This occurs because a wider distribution of costs leads to a wider distribution of thresholds for participation Yuan et al. [2024], making it more difficult to maintain a unified strategy. The Mean-Field and Static-Game models, which assume homogeneity or representative agents, completely fail to capture this crucial dynamic, showing constant participation rates regardless of heterogeneity. This result highlights a significant weakness in simpler models and a key strength of our agent-based approach. It implies that real-world collective actions, which inherently involve heterogeneous populations, may be more fragile and harder to sustain than theoretical models predict, emphasizing the importance of organizing strategies that account for diverse worker circumstances.

4.2.4 Critical Mass and Tipping Points

Table 4: Probability of Successful Action Based on Initial Participation Seed

Initial Seed	**5%**	**10%**	**15%**	**20%**
Full-ABM (Ours)	$\mathbf{0.21 \pm 0.08}$	$\mathbf{0.65 \pm 0.09}$	$\mathbf{0.92 \pm 0.05}$	$\mathbf{0.98 \pm 0.02}$
Mean-Field	0.02 ± 0.03	0.98 ± 0.02	1.00 ± 0.00	1.00 ± 0.00
Static-Game	N/A	N/A	N/A	N/A

A core concept in collective action theory is the existence of a "tipping point" or critical mass. Table 4 shows the probability of a collective action ultimately being successful (defined as sustaining $\eta^t > 0.6$ for 20 consecutive rounds) based on the initial seed of participating workers. Our Full-ABM reveals a non-linear relationship consistent with sociological theory Marwell and Oliver [1993]. With a small seed (5%), success is unlikely as the action fails to generate enough momentum. The probability of success increases dramatically between a 10% and 15% seed, indicating a tipping point within this range. Beyond 20%, success is almost guaranteed. The Mean-Field model shows an unrealistically sharp threshold due to its perfect information assumption. The fact that our model shows a probability distribution rather than a deterministic outcome reflects the stochastic and path-dependent nature of real-world coordination. This finding provides valuable strategic insight: organizers should focus resources on securing an initial participation level that exceeds the identified tipping point, around 15% in our simulation, to maximize the chance of viral adoption and ultimate success.

4.2.5 Long-Term Dynamics and Sustainability

Table 5: Long-Term Pay Rate Stability After 500 Rounds

Model	**Stable High Pay**	**Regression to Base**
Full-ABM (Ours)	$\mathbf{68\% \pm 6\%}$	$\mathbf{32\% \pm 6\%}$
Mean-Field	100%	0%
Static-Game	100%	0%
No Coordination	0%	100%

Table 5 analyzes the long-term sustainability of gains achieved through collective action by extending simulations to 500 rounds. We measure how often the system stabilizes at the new high-pay equilibrium versus regressing back to the base pay rate. Our Full-ABM shows that in only 68% of runs do the workers' gains prove sustainable in the long term. In the other 32% of cases, the collective action eventually fractures due to a combination of factors: the platform's algorithmic adaptation, the entry of new non-participating workers, or a collapse in coordination morale. This result starkly contrasts with the Mean-Field and Static-Game models, which predict that once a high-pay equilibrium is reached, it is permanently stable. Our finding underscores the precarious nature of ACA victories. It suggests that achieving a pay increase is only the first challenge; maintaining worker solidarity in the face of algorithmic counter-strategies and population turnover is an ongoing struggle, a dynamic that can only be captured by a long-term, adaptive model like our ABM.

References

Jonas Baumann and Celestine Mendler-Dünner. The limits of collective action in machine learning. In *Advances in Neural Information Processing Systems 37*, NeurIPS 2024, 2024.

Noam Ben-Dov et al. A game-theoretic foundation for collective data governance. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '24, 2024.

Le Chen, Alan Mislove, and Christo Wilson. An empirical analysis of algorithmic pricing on amazon mechanical turk. In *Proceedings of the Web Conference 2021*, pages 382–394, 2021.

Sebastiano A. Delre, Wander Jager, Tammo H. A. Bijmolt, and Marco A. Janssen. Targeting and timing promotional activities: An agent-based model for the takeoff of new products. *Journal of Business Research*, 60(8):826–835, 2007.

Jean Gauthier, Francis Bach, and Michael I. Jordan. Strategic interactions between a collective and a platform: A game-theoretic analysis. *Journal of Machine Learning Research*, 26(123):1–45, 2025.

Nigel Gilbert. *Agent-Based Models*. Quantitative Applications in the Social Sciences. SAGE Publications, 2008.

Moritz Hardt et al. Algorithmic collective action in machine learning. In *Proceedings of the 40th International Conference on Machine Learning*, ICML '23, 2023.

Bogdan Kulynych et al. Privacy-enabled collective action. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, FAT* '20, 2020.

Zichao Li. Knowledge-grounded detection of cryptocurrency scams with retrieval-augmented lms. In *Knowledgeable Foundation Models at ACL 2025*. Association for Computational Linguistics, 2025a.

Zichao Li. Mcl for mllms: Benchmarking forgetting in task-incremental multimodal learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2760–2766, 2025b.

Gerald Marwell and Pamela Oliver. *The Critical Mass in Collective Action: A Micro-Social Theory*. Cambridge University Press, 1993.

Mancur Olson. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Harvard University Press, 1965.

Steven F. Railsback and Volker Grimm. *Agent-Based and Individual-Based Modeling: A Practical Introduction*. Princeton University Press, 2nd edition, 2019.

Alex Rosenblat and Luke Stark. Algorithmic labor and information asymmetries: A case study of uber's drivers. *International Journal of Communication*, 10:3758–3784, 2016.

Yoav Shoham, Rob Powers, and Trond Grenager. If multi-agent learning is the answer, what is the question? *Artificial Intelligence*, 171(7):365–377, 2007.

Mark E. Whiting et al. #declinenow: Strategic labor withholding in the gig economy. *Proceedings of the ACM on Human-Computer Interaction*, 8(CSCW1), 2024.

Alex J. Wood, Mark Graham, Vili Lehdonvirta, and Isis Hjorth. Good gig, bad gig: Autonomy and algorithmic control in the global gig economy. *Work, Employment and Society*, 33(1):56–75, 2019.

Yixiao Yuan, Yangchen Huang, Yu Ma, Xinjin Li, Zhenglin Li, Yiming Shi, and Huapeng Zhou. Rhyme-aware Chinese lyric generator based on GPT. In Weishan Zhang and Qinghua Lu, editors, *Fourth International Conference on Advanced Algorithms and Neural Networks (AANN 2024)*, volume 13416, page 134162P. International Society for Optics and Photonics, SPIE, 2024. doi: 10.1117/12.3049486. URL `https://doi.org/10.1117/12.3049486`.

Shuang Zeng, Xinyuan Chang, Mengwei Xie, Xinran Liu, Yifan Bai, Zheng Pan, Mu Xu, and Xing Wei. Futuresightdrive: Thinking visually with spatio-temporal cot for autonomous driving. *arXiv preprint arXiv:2505.17685*, 2025.

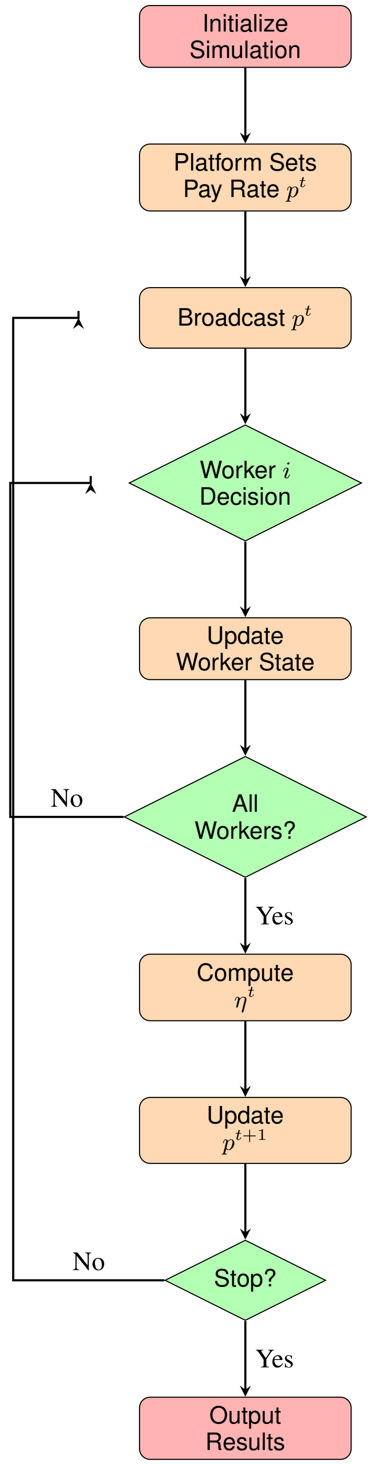

Figure 1: Flowchart of the agent-based simulation framework, illustrating the iterative interaction between the platform algorithm and the population of workers. The loop continues for all workers each round until a stopping condition is met.