

Multi-Agent Learning and Strategic Goal Competition Among Platform Workers

Wenhaojie Du

School of Information Technology, University of Cincinnati, Cincinnati, OH, USA.
du1991@uc.edu

Isawaan Mathur

Department of Computer Science, Binghamton University, Binghamton, NY, USA.
ishaan.mathur@binghamton.edu

Ananya Jase

Department of Computer Science, Colorado State University, Fort Collins, CO, USA.
ananyajoshi78@colostate.edu

Prakash Sean

Department of Computer Science, University of Houston, Houston, TX, USA.
prakash.work@uh.edu

Abstract

Digital labor platforms have reconfigured the nature of work by embedding large populations of independent workers within algorithmically managed ecosystems. Workers act as autonomous decision-makers who continuously adapt their strategies in response to dynamic incentives, peer behavior, and opaque platform governance. This paper adopts a systems-oriented perspective to examine platform labor as a large-scale multi-agent learning environment in which strategic goal competition fundamentally shapes market outcomes, individual welfare, and systemic stability. Rather than treating workers as isolated optimizers, we conceptualize the platform as a coupled socio-technical infrastructure where agents learn from sparse feedback, compete for spatially and temporally distributed tasks, and form individual earning or activity goals that interact through shared market mechanisms. We analyze the structural trade-offs inherent in centralized algorithmic architectures, the role of emerging multi-agent reinforcement learning paradigms in modeling these dynamics, and the consequences of goal-driven competition for fairness, robustness, and long-term ecosystem sustainability. Through a synthesis of interdisciplinary research spanning computer science, labor studies, and regulatory scholarship, we examine how platform design choices act as meta-level governance interventions that shape the learning landscape, often producing unintended emergent equilibria such as oversupply clustering, social dilemma deadlocks, and distributive inequities. The discussion foregrounds infrastructural resilience, the ethical dimensions of automated management, and policy frameworks that can realign platform objectives with social welfare. By reframing platform work as a multi-agent co-adaptation problem, the paper advances a rigorous conceptual apparatus for evaluating the systemic consequences of algorithmic labor coordination and informs the design of more sustainable and equitable platform architectures.

Keywords

platform work, multi-agent learning, goal competition, algorithmic governance, socio-technical systems, fairness, strategic adaptation.

1. Introduction

The proliferation of on-demand digital platforms has created a new class of labor markets in which millions of workers operate as independent contractors, matched with tasks through real-time algorithmic dispatch and dynamic pricing mechanisms [1]. These platforms represent far more than simple two-sided marketplaces; they function as tightly orchestrated socio-technical systems in which worker behavior, platform algorithms, and market-level outcomes co-evolve continuously [2]. Workers are not passive recipients of algorithmic instructions but active strategic agents who set their own goals, learn from experience, observe their peers, and adjust their decisions in an effort to maximize earnings, maintain satisfaction, and manage the precarity inherent in gig arrangements [3]. The central premise of this paper is that platform labor can be productively understood as a multi-agent learning system in which strategic goal competition is a defining structural force rather than an incidental behavioral nuance. This perspective shifts analytical attention from individual optimization or single-agent models toward the macroscopic properties that arise from the interplay of many adaptive workers coupled through shared market infrastructure. The systems view exposes critical trade-offs related to architectural design, governance, fairness, robustness, and long-term sustainability that are invisible in traditional microeconomic accounts.

Contemporary platform architectures embed workers in an information environment characterized by selective visibility, algorithmic pricing, reputation scoring, and frequently opaque task allocation rules [4]. Within this environment, each worker develops a learning strategy to interpret signals, adjust effort, select spatial or temporal zones of operation, and set personal earning targets. These learning processes are inherently multi-agent: a driver's choice to relocate to a surge zone changes the supply density for all other drivers, a freelancer's bid on a project influences the reference price that others will encounter, and a rider's decision to log off in the afternoon alters the competitive pressure during the evening peak. The resulting dynamics exhibit the hallmarks of complex adaptive systems, including feedback loops, emergent coordination, and strategic externalities that are difficult to manage through centralized algorithmic control alone [5]. Multi-agent reinforcement learning (MARL) provides a formal vocabulary for describing such environments, but the phenomena observed in deployed platforms also demand a broader interdisciplinary lens that integrates organizational psychology, science and technology studies, and labor economics. This paper develops such an integration, advancing a system-level analysis of learning, goal competition, and governance in platform work.

The remainder of the paper is organized as follows. Section 2 examines the platform as a multi-agent architecture and characterizes how workers emerge as strategic learners within that architecture. Section 3 reviews learning mechanisms and the adaptive formation of goals, linking algorithmic feedback structures to established behavioral theories. Section 4 analyzes strategic goal competition as a dynamic process that shapes market equilibria, collective outcomes, and inequality. Section 5 discusses governance infrastructures and algorithmic management as meta-level interventions that constrain and channel agent learning. Section 6 addresses fairness, robustness, and sustainability, highlighting the normative and structural vulnerabilities of competitive multi-agent platform systems. Section 7 explores policy implications, regulatory frameworks, and the prospects for institutional design that respects

both innovation and worker welfare, and Section 8 concludes with a forward-looking synthesis.

2. Platform Architecture and the Emergence of Worker Multi-Agent Systems

A digital labor platform is simultaneously a technical artifact, a governance institution, and an interactive field in which autonomous workers pursue heterogeneous objectives. At the architectural level, the platform maintains a centralized control plane that sets prices, dispatches tasks, manages reputational data, and enforces behavioral policies, while workers retain operational autonomy over when, where, and how to work [3]. This structure creates a distributed decision-making environment that is neither fully centralized nor purely decentralized. Workers function as independent agents whose local actions generate system-wide consequences, and the platform's central algorithms act as an environmental regulator that modulates incentives, information flows, and competitive intensity. The resultant multi-agent system exhibits emergent properties that cannot be adequately captured by decomposing it into isolated worker-platform dyads.

The informational topology of the platform critically shapes how agents learn and compete. Workers typically observe a subset of the state space: their own earnings history, limited aggregate signals such as surge multipliers or heat maps, and occasionally peer benchmarks that vary across platforms in granularity and frequency [6]. Information asymmetries are not mere market imperfections but are deliberate design choices that constrain agent reasoning, discourage certain forms of strategic gaming, and preserve the platform's capacity to steer aggregate behavior. Because workers cannot observe the full distribution of competitor intentions, their learning processes are necessarily approximate and model-based, relying on mental representations of how the market behaves. These representations are updated through reinforcement, imitation, or simple heuristics, and they generate adaptive strategies that differ substantially across individuals, demographic groups, and local submarkets.

The platform's reputation infrastructure amplifies multi-agent coupling. Ratings, acceptance rates, and cancellation penalties create a persistent public record that links past behavior to future opportunity, transforming transient task-level decisions into durable career trajectories [7]. A worker who systematically declines low-value tasks may protect short-term income but eventually faces reduced access to high-quality assignments if algorithms penalize selectivity. This feedback loop introduces a temporal credit assignment problem that is characteristic of reinforcement learning in partially observable environments: workers must infer the long-term consequences of their actions from noisy and delayed signals. Moreover, because reputation thresholds affect all workers simultaneously, adaptive changes in aggregate behavior can shift the distribution of opportunities, producing phase transitions where minor algorithmic adjustments trigger cascades of strategic recalibration across the entire worker population.

The multi-agent nature of platform labor also manifests in the spatial and temporal dimensions of task competition. Ride-hailing drivers position themselves in anticipation of demand surges, delivery couriers cluster near popular restaurant aggregations, and online freelancers time their bids to coincide with employer activity windows [8]. These positioning decisions constitute a continuous spatial game in which each agent's expected utility depends on the evolving distribution of other agents. The architectural features that govern how workers perceive spatiotemporal demand, such as heat maps, estimated waiting times, and surge multipliers, function as control knobs that can either stabilize or destabilize the collective allocation of labor. When signals are too coarse or too lagged, oscillations and herding behavior can degrade system efficiency and increase worker earning volatility.

Conversely, overly precise real-time revelation of competitor locations may induce rapid convergence to degenerate equilibria where all workers flock to a small set of hypercompetitive zones, leaving large areas underserved.

3. Learning Mechanisms and Adaptive Goal Formation

Workers in platform ecosystems learn from a heterogeneous mixture of personal experience, social observation, and platform-mediated feedback. This learning can be characterized along multiple dimensions: the degree of model-based reasoning, the reliance on social versus individual signals, and the temporal horizon over which outcomes are evaluated. The simplest form is reactive reinforcement, where workers repeat actions that led to satisfactory earnings and avoid those associated with poor outcomes. More sophisticated model-based learning involves constructing causal narratives about market dynamics, such as inferring that working on Tuesday evenings in a particular district yields higher average hourly earnings because of a concert venue schedule [9]. Both forms coexist in large worker populations, and the relative prevalence of each is shaped by the complexity of the platform environment and the availability of interpretable feedback.

Goal setting plays a central role in structuring this learning process. Workers do not merely learn which actions are rewarding; they set objectives that define what counts as a satisfactory outcome in the first place. Daily income targets, minimum hourly rates, and aspired utilization levels serve as reference points that moderate the interpretation of experience and guide exploration behavior. Foundational psychological research has demonstrated that specific, challenging goals enhance performance through mechanisms of directed attention, effort mobilization, and persistence [10]. In platform settings, however, goal formation occurs under conditions of considerable uncertainty, and the goals that workers adopt are themselves subject to adaptive revision. A driver who consistently fails to reach a self-imposed daily earnings target may lower aspirations over time, shift strategies, or exit the platform entirely. Conversely, workers who observe peers achieving high earnings may raise their own benchmarks, leading to an escalation of effort and risk-taking.

Empirical investigations have begun to document these dynamics. Studies of ride-hailing drivers reveal that earnings targets function as flexible anchors that are influenced by recent income, peer comparisons, and even the framing of informational nudges embedded in the application interface [11]. Crowdworkers on microtask platforms similarly adjust their reservation wages and effort intensity based on historical task availability and observed pay distributions. In the food delivery sector, couriers develop finely calibrated spatial and temporal heuristics that balance expected order volumes against the physical toll of cycling or driving in dense urban traffic [12]. These diverse empirical patterns underscore the need for a unified framework in which goal competition is understood not as a static motivational input but as a dynamic, co-evolving component of the multi-agent learning ecosystem. Recent field experimental evidence further confirms that explicitly encouraging workers to set personal goals can significantly modulate effort and market participation patterns, although the systemic consequences of widespread goal-setting interventions remain an open and underexplored question [17].

The informational infrastructure provided by platforms heavily mediates which learning strategies are feasible. When platforms display detailed historical earnings dashboards, workers can engage in more accurate credit assignment and develop finer-grained expectations. When feedback is limited to aggregate weekly summaries, learning becomes less directed and more susceptible to superstition and noise. The design choice between these

two poles represents a structural trade-off: richer feedback may improve individual decision quality but can also intensify goal competition by making social comparisons more salient, thereby fueling an arms race of effort escalation that can harm worker well-being without corresponding productivity gains. Multi-agent reinforcement learning research offers analogous results, showing that agents with access to richer state information may converge to more efficient equilibria or, under certain competitive conditions, to more aggressively exploitative strategies that unravel cooperative arrangements [13]. Translating these insights to human platform work requires careful attention to the perceptual and cognitive constraints that distinguish human agents from algorithmic ones.

4. Strategic Goal Competition: Dynamics, Conflicts, and Emergent Equilibria

When many workers pursue self-defined goals within a shared platform environment, their strategies inevitably interact through congestion, price competition, and reputation externalities. This strategic goal competition can be analyzed as a non-cooperative multi-agent game in which each worker selects a goal level and an effort strategy to achieve it, while the realized payoffs depend on the aggregate actions of all participants. The resulting dynamics frequently display characteristics that are well known in computational multi-agent systems: resource races, tragedy-of-the-commons patterns, and cyclical instability. In ride-hailing, for example, drivers collectively converging on a surge zone can oversaturate the area so rapidly that the surge multiplier dissipates before most drivers secure a fare, leaving many with lower earnings than if they had spread across the city more evenly [14]. This is structurally analogous to the overgrazing dynamic in reinforcement learning agents competing for a common resource, where individually rational actions produce collectively suboptimal outcomes.

The temporal structure of goal competition further complicates the picture. Many platform workers set daily or weekly targets, which means that their strategy selection is governed by a finite horizon with a strong reference point. Workers who are behind their targets late in the day may accept riskier or lower-paying tasks to close the gap, a behavior that intensifies competition during the very periods when demand is waning. This loss-aversion dynamic, coupled with deadline effects, creates predictable oscillations in service quality, driver availability, and market-clearing prices that platform algorithms must buffer. From a system design perspective, these oscillations are not mere noise but emergent features of a multi-agent learning process in which goal gradient asymmetries shape the temporal distribution of labor supply. Failure to account for these dynamics when tuning surge multipliers or dispatch rules can degrade both reliability and equity.

Inequality is an almost inevitable byproduct of unmoderated strategic goal competition in platform environments. Workers who are faster learners, who possess better prior knowledge of local demand geography, or who have access to superior hardware and connectivity will systematically outperform others, accumulating reputational advantages and gaining privileged access to high-value tasks. Over time, small initial advantages can compound into durable stratification, as reinforced learning algorithms on the part of the platform amplify performance differentials [15]. The resulting ecological structure resembles a competitive niche formation process in which some agents occupy high-return specialist niches while others are relegated to residual, low-return segments. This outcome raises profound questions about fairness and the distributional consequences of algorithmically managed markets, questions that extend beyond the technical efficiency concerns that dominate much of the MARL and operations management literature.

Despite the competitive framing, platform workers are not entirely devoid of cooperative strategies. In several countries, informal coordination through messaging groups enables couriers to share information about restaurant waiting times, surge patterns, and algorithmic changes, effectively creating a distributed collective intelligence that partially offsets the platform's informational advantage [16]. Such cooperative behaviors can be understood as emergent forms of multi-agent learning where agents share policies or value functions through social channels. However, the sustainability of informal cooperation is fragile; it relies on trust, shared identity, and a perceived common threat from the platform's algorithmic authority. When platform design deliberately atomizes workers, anonymizes performance comparisons, or discourages direct communication, it erodes the social substrate on which cooperative learning depends, nudging the system toward a more purely competitive regime.

5. Governance Infrastructures and Algorithmic Management

The platform operates as a meta-agent that shapes the multi-agent learning environment through a suite of algorithmic governance mechanisms. Surge pricing, task dispatching, batch assignment, cancellation penalties, and acceptance rate thresholds are not neutral market-clearing devices but active policy instruments that encode platform objectives such as throughput maximization, service level maintenance, and demand-side satisfaction [18]. These instruments co-determine the reward structure that workers experience, effectively engineering the reinforcement learning curriculum through which workers adapt their strategies. The platform thus occupies a dual role: it is both a market facilitator and a behavioral designer, constantly tuning the parameters of a vast distributed learning experiment without fully controlling or even observing the internal states of the human agents who populate it.

One of the most consequential governance choices concerns the transparency of the scoring and allocation functions. Opaque algorithms create an adversarial learning setting in which workers must infer hidden rules from incomplete outcome data, a process that is cognitively intensive and prone to systematic errors [19]. Workers who incorrectly model the platform's objective function may pursue strategies that appear locally rational but are globally counterproductive, leading to frustration, churn, and welfare losses. From a systems perspective, opacity can be defended as a mechanism to prevent gaming, but it simultaneously undermines the quality of agent learning and erodes trust in the institutional architecture. Transparency, on the other hand, may invite sophisticated exploitation, but in multi-agent systems with large populations, perfect transparency often enables more efficient equilibration by allowing agents to coordinate without wasteful exploration. The optimal degree of transparency is therefore a non-trivial design parameter that must balance anti-gaming resilience against learning efficiency and procedural fairness.

Platform governance also extends to the temporal granularity of control. Some platforms enforce real-time behavioral constraints through mandatory acceptance windows, automatic log-out penalties, and immediate deactivation for rule violations, creating a tight feedback loop that resembles a continuous control system. Others adopt a softer approach with periodic reviews, cumulative quality scores, and delayed consequences, which yields a looser coupling between actions and outcomes. Tight coupling can stabilize system behavior in the short term but risks over-constraining worker autonomy, inducing learned helplessness, and suppressing the exploratory variation that is essential for adaptive learning and innovation in dynamic environments [20]. Loose coupling preserves agent diversity and adaptability but may allow

undesirable strategies to persist and spread through social imitation. Designing an appropriate coupling regime requires system-level modeling that integrates agent heterogeneity, environmental volatility, and platform objective functions.

The advent of sophisticated multi-agent reinforcement learning algorithms in simulation environments offers new possibilities for ex ante evaluation of governance policies before deployment on human workers. Platforms can construct digital twins of their markets, populate them with heterogeneous learning agents calibrated on behavioral data, and stress-test alternative incentive schemes, communication architectures, and fairness constraints [21]. Such simulation-based governance design, while ethically complex, could reduce the frequency of disruptive policy pivots that currently force millions of workers to repeatedly relearn the rules of the game. Nevertheless, simulation fidelity remains a central challenge; human workers bring emotional states, fairness preferences, and social motivations that are poorly captured by current economic or reinforcement learning agent models, necessitating continued investment in interdisciplinary behavioral research.

6. Fairness, Robustness, and Sustainability of Platform Ecosystems

The convergence of multi-agent learning and strategic goal competition raises fundamental concerns about fairness that are not reducible to aggregate efficiency metrics. Individual fairness demands that similar workers be treated similarly by the platform's allocation and pricing algorithms, yet learning-induced divergences in strategy and reputation can cause initially identical workers to experience vastly different outcomes over time. Group-level fairness is challenged when algorithmic management perpetuates or amplifies historical disparities rooted in geography, mobility constraints, digital literacy, or socioeconomic background [22]. Procedural fairness requires that workers be able to understand and contest the rules that govern their livelihoods, a condition that is frequently violated by the black-box character of production-oriented machine learning models. Addressing these multiple dimensions of fairness requires moving beyond single-agent optimization and embracing a multi-agent justice framework that evaluates platform architectures against criteria of relational equality.

Robustness refers to the capacity of the platform-work ecosystem to maintain functional integrity under shocks such as sudden demand spikes, policy interventions, or adversarial manipulation by coordinated worker groups. Multi-agent learning systems can exhibit remarkable resilience through rapid adaptation, as seen during the COVID-19 pandemic when ride-hailing drivers swiftly reallocated to food delivery as demand patterns shifted [23]. However, the same adaptive capacity can generate brittleness when learning algorithms converge to fragile strategies that perform well only under narrow environmental assumptions. A platform whose driver labor supply model has been implicitly tuned to a particular spatial demand distribution may discover that workers' learned positioning heuristics break down catastrophically when that distribution changes, producing severe service gaps and cascading dissatisfaction among all stakeholders. Building robustness thus demands that platform architects anticipate the co-adaptation of a population of strategic learners and design mechanisms that preserve sufficient strategic diversity to buffer against unknown future contingencies.

Sustainability encompasses both the economic viability of platform work as a livelihood and the broader ecological health of the socio-technical system. In purely competitive multi-agent regimes, workers can become trapped in a rat race of escalating effort, leading to burnout, turnover, and long-run labor supply degradation that eventually harms platform profitability

and service quality. The structural incentives of many platforms, which benefit from surplus labor supply in the short term, can create a self-undermining dynamic when high churn generates a permanent class of novice learners who provide lower quality service and require constant retraining investment [24]. A sustainable equilibrium requires that platforms internalize the externality that their algorithmic design choices impose on worker welfare and aggregate human capital. This insight shifts the design problem from one of simple revenue maximization to one of co-optimizing platform performance and worker thriving over extended time horizons, a challenge that aligns with growing calls for stakeholder capitalism in digital markets.

7. Policy Implications and Regulatory Frameworks

The systems-level analysis developed in this paper has direct consequences for regulatory and policy debates. Recognising platform work as a multi-agent learning environment underscores the insufficiency of narrow interventions that target isolated practices such as tipping policies or deactivation thresholds. Effective regulation must address the structural conditions that shape the multi-agent learning landscape: algorithmic transparency, data portability, minimum earnings guarantees, and the legal classification of workers as employees or independent contractors with attendant rights. The European Union's Platform Work Directive signals a movement toward algorithmic accountability and rebuttable presumption of employment, principles that can fundamentally alter the incentive structure within which learning occurs [25]. By requiring platforms to disclose key parameters of automated decision-making systems, such policies can reduce information asymmetries and thereby enhance the quality of worker learning, enabling more informed goal setting and strategic adaptation.

Beyond disclosure, regulators can mandate impact assessments that evaluate how algorithmic changes are likely to affect the distribution of outcomes across different worker groups, drawing on simulation and multi-agent modeling techniques. Such assessments could identify potential fairness harms and trigger corrective action before they manifest at scale, shifting the policy stance from reactive enforcement to proactive systemic oversight. Data trust frameworks, in which aggregated and anonymized platform data is made available for independent research, can further empower academic and civil society investigators to study multi-agent dynamics without relying on voluntary corporate disclosure, accelerating the development of an evidence base for governance design. These proposals move regulation beyond static rule compliance toward dynamic, learning-oriented governance that matches the adaptive character of the platform systems it seeks to regulate.

Collective representation also plays a crucial role in rebalancing the multi-agent learning environment. When workers unionize or form cooperatives, they create a meta-agent that can negotiate over the parameters of the platform's algorithmic management regime, in effect introducing a new cooperative learning channel that counterbalances the platform's centralized optimization. Sectoral bargaining, portable benefits schemes, and worker-owned platform alternatives all represent institutional innovations that alter the competitive structure of goal pursuit, shifting the emergent system equilibrium toward greater equity and resilience. The design of these institutions can itself be informed by multi-agent systems thinking, treating regulatory frameworks as meta-level coordination mechanisms that complement and constrain platform-level algorithms in pursuit of broader societal objectives.

8. Conclusion

This paper has argued that platform work is best understood as a large-scale multi-agent learning system in which strategic goal competition is a central generative mechanism of market outcomes. The architectural choices embedded in algorithmic management, the information and feedback infrastructures that shape worker learning, and the governance frameworks that constrain platform behavior collectively determine the fairness, robustness, and sustainability of the resulting ecosystem. By synthesizing insights from multi-agent reinforcement learning, labor sociology, behavioral economics, and regulatory studies, we have shown that the challenges facing platform work cannot be adequately addressed through isolated technical or policy fixes. Instead, they demand a systems-oriented integration that treats workers as adaptive agents, platforms as co-evolving learning environments, and governance as a design problem in multi-agent coordination.

Future research should deepen the connection between computational multi-agent models and empirical platform labor dynamics, developing digital twin simulations that incorporate realistic human cognitive and motivational factors. Interdisciplinary fieldwork that captures the lived experience of learning and goal adjustment under algorithmic governance is equally essential. Ultimately, constructing platform architectures that are both innovative and just requires a collaborative effort across disciplines and stakeholder groups, one that respects the complexity of strategic human behavior while leveraging the analytical power of systems science. The multi-agent learning lens provides a coherent foundation for that endeavor, reframing platform work not as a collection of isolated transactions but as an ongoing process of collective adaptation that can be deliberately shaped toward more equitable and resilient futures.

References

1. Sundararajan, A. (2016). *The sharing economy: The end of employment and the rise of crowd-based capitalism*. MIT Press.
2. Wood, A. J., Graham, M., Lehdonvirta, V., & Hjorth, I. (2019). Good gig, bad gig: Autonomy and algorithmic control in the global gig economy. *Work, Employment and Society*, 33(1), 56-75.
3. Lee, M. K., Kusbit, D., Metsky, E., & Dabbish, L. (2015). Working with machines: The impact of algorithmic and data-driven management on human workers. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI)*, 1603-1612.
4. Rosenblat, A., & Stark, L. (2016). Algorithmic labor and information asymmetries: A case study of Uber's drivers. *International Journal of Communication*, 10, 3758-3784.
5. Lowe, R., Wu, Y., Tamar, A., Harb, J., Abbeel, P., & Mordatch, I. (2017). Multi-agent actor-critic for mixed cooperative-competitive environments. *Advances in Neural Information Processing Systems*, 30, 6379-6390.
6. Leibo, J. Z., Zambaldi, V., Lanctot, M., Marecki, J., & Graepel, T. (2017). Multi-agent reinforcement learning in sequential social dilemmas. *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 464-473.
7. Zhang, K., Yang, Z., & Başar, T. (2021). Multi-agent reinforcement learning: A selective overview of theories and algorithms. In *Handbook of Reinforcement Learning and Control* (pp. 321-384). Springer.

8. Busoniu, L., Babuska, R., & De Schutter, B. (2008). A comprehensive survey of multi-agent reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C*, 38(2), 156-172.
9. Horton, J. J., & Chilton, L. B. (2010). The labor economics of paid crowdsourcing. *Proceedings of the 11th ACM Conference on Electronic Commerce (EC)*, 209-218.
10. Ho, C. J., & Vaughan, J. W. (2012). Online task assignment in crowdsourcing markets. *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, 45-51.
11. Tassinari, A., & Maccarrone, V. (2020). Riders on the storm: Workplace solidarity among gig economy couriers in Italy and the UK. *Work, Employment and Society*, 34(5), 865-883.
12. De Stefano, V. (2016). The rise of the "just-in-time workforce": On-demand work, crowdwork and labour protection in the "gig-economy". *Comparative Labor Law & Policy Journal*, 37(3), 471-504.
13. Lehdonvirta, V. (2018). Flexibility in the gig economy: Managing time on three online piecework platforms. *New Technology, Work and Employment*, 33(1), 13-29.
14. Locke, E. A., & Latham, G. P. (2002). Building a practically useful theory of goal setting and task motivation: A 35-year odyssey. *American Psychologist*, 57(9), 705-717.
15. Cameron, L. K. (2021). "You can't just pick up and move": How algorithmic dispatch shapes worker strategies on ride-hail platforms. *Big Data & Society*, 8(2).
16. Chen, M. K., Chevalier, J. A., Rossi, P. E., & Oehlsen, E. (2019). The value of flexible work: Evidence from Uber drivers. *Journal of Political Economy*, 127(6), 2735-2794.
17. Min, X., Chi, W., Hu, X., & Ye, Q. (2024). Set a goal for yourself? A model and field experiment with gig workers. *Production and Operations Management*, 33(1), 205-224.
18. Calo, R., & Rosenblat, A. (2017). The taking economy: Uber, information, and power. *Columbia Law Review*, 117(6), 1623-1690.
19. Ettlinger, N. (2017). Algorithmic governance and the gig economy: An institutional perspective. *Geoforum*, 81, 188-197.
20. Irani, L. (2015). Difference and dependence among digital workers: The case of Amazon Mechanical Turk. *South Atlantic Quarterly*, 114(1), 225-234.
21. Veen, A., Barratt, T., & Goods, C. (2020). Platform-capital's 'app-etite' for control: A labour process analysis of food-delivery work in Australia. *Work, Employment and Society*, 34(3), 388-406.
22. Rahman, K. S., & Thelen, K. (2019). The rise of the platform business model and the transformation of twenty-first-century capitalism. *Politics & Society*, 47(2), 177-204.
23. Prassl, J. (2018). *Humans as a service: The promise and perils of work in the gig economy*. Oxford University Press.
24. Nowak, M. A., & Sigmund, K. (2004). Evolutionary dynamics of biological games. *Science*, 303(5659), 793-799.
25. European Commission. (2021). Proposal for a Directive on improving working conditions in platform work. COM(2021) 762 final.