

Algorithmic Goal Recommendations and Labor Supply Responses among Gig Economy Workers

Jack Koskinen

Department of Computer Science, University of North Texas, Denton, TX, USA.
hellojack@unt.edu

Kabir Bhatia

Department of Computer Science, Binghamton University, Binghamton, NY, USA.
kabirbhatia743@binghamton.edu

Huayuan Yang

Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY, USA.
huayuan.yang423@buffalo.edu

Abstract

The proliferation of digital labor platforms has fundamentally restructured the relationship between work, worker autonomy, and algorithmic management. Among the most pervasive yet underexamined features of these platforms is the algorithmic goal recommendation system, which suggests personalized productivity targets to gig economy workers. This paper presents a comprehensive systems-level analysis of how such algorithmic goal recommendations shape labor supply decisions, focusing on the structural, behavioral, and governance implications of these socio-technical interventions. We argue that goal recommendation algorithms function as a form of soft paternalistic control, leveraging behavioral nudges to influence worker effort without explicit coercion. Drawing on a synthesis of empirical evidence from field experiments, platform data analyses, and behavioral labor economics, we examine the mechanisms through which goal recommendations affect worker output, scheduling, and retention. The analysis reveals a fundamental tension: while goal recommendations can increase short-term productivity and help workers self-regulate, they also introduce new vulnerabilities including goal manipulation, effort distortion, and the erosion of intrinsic motivation. We further explore the architectural design choices that determine the fairness and robustness of these systems, including the role of reference group selection, dynamic adjustment rates, and feedback loops between platform governance and worker behavior. The paper contextualizes these findings within broader debates on algorithmic management, digital Taylorism, and the future of work. We conclude by proposing a set of design principles for transparent, equitable, and sustainable goal recommendation systems that balance platform efficiency with worker welfare. This research contributes to the growing interdisciplinary literature on algorithmic governance in labor markets and offers actionable insights for platform designers, policymakers, and labor advocates.

Keywords

algorithmic management, gig economy, labor supply, goal setting, behavioral nudges, platform governance, worker autonomy.

1. Introduction

The digital transformation of labor markets has given rise to a new class of socio-technical systems in which algorithms mediate nearly every aspect of work, from task allocation and performance evaluation to compensation and scheduling. Gig economy platforms such as ride-hailing services, food delivery networks, and freelance marketplaces now rely on algorithmic management to coordinate millions of workers operating in decentralized, on-demand environments [1], [2]. Within this ecosystem, one particularly influential algorithmic intervention is the goal recommendation system, which uses historical worker data, peer benchmarks, and platform objectives to suggest personalized productivity targets. These recommendations, often displayed as daily or weekly earnings goals, trip counts, or completion rates, are intended to guide worker behavior without the overt directives characteristic of traditional employment [3]. However, the systematic effects of these algorithmic nudges on labor supply remain poorly understood, particularly with respect to their long-term consequences for worker welfare, platform sustainability, and market efficiency.

The central question motivating this paper is how algorithmic goal recommendations influence the labor supply decisions of gig workers, and what structural trade-offs emerge from the deployment of such systems at scale. Labor supply in the gig economy is characterized by high elasticity, with workers exercising substantial discretion over when, how much, and for how long they work [4]. This flexibility is both a defining feature and a source of instability, as workers must constantly navigate the tension between income maximization, leisure, and fatigue. Algorithmic goal recommendations intervene in this decision process by providing a salient reference point that anchors worker expectations and effort levels. While goal setting has long been studied in organizational psychology and behavioral economics as a powerful motivator [5], its algorithmic instantiation introduces novel dynamics, including real-time personalization, opacity of design, and the integration of platform-level incentives that may diverge from worker interests.

This paper adopts a systems-level perspective, treating the goal recommendation algorithm not as an isolated feature but as a component of a larger socio-technical infrastructure that includes data collection pipelines, feedback mechanisms, governance rules, and market structures. We argue that the effectiveness and fairness of goal recommendations depend critically on architectural choices: how reference groups are defined, how goals are updated over time, how goal attainment is rewarded or penalized, and how workers can contest or customize their targets. These design parameters shape not only individual labor supply but also aggregate market dynamics, including surge pricing, worker churn, and the distribution of earnings across demographic groups. By examining these interdependencies, we aim to provide a holistic account of algorithmic goal recommendation systems that bridges the gap between technical design and social outcomes.

The remainder of this paper is organized as follows. Section 2 reviews the relevant literature on algorithmic management, goal setting theory, and labor supply in the gig economy. Section 3 develops a conceptual framework for analyzing goal recommendation systems as socio-technical interventions. Section 4 presents empirical evidence from field experiments and platform data, highlighting key findings on productivity, effort distortion, and worker welfare. Section 5 examines the architectural and governance dimensions of these systems, including fairness, transparency, and robustness. Section 6 discusses policy implications and design recommendations. Section 7 concludes with a summary of contributions and directions for future research.

2. Literature Review

The study of algorithmic management has expanded rapidly over the past decade, driven by the proliferation of digital platforms that replace human supervisors with automated decision systems. Early work in this area focused on the mechanisms of control, including rating systems, task assignment algorithms, and dynamic pricing [1], [6]. These studies established that algorithmic management differs fundamentally from traditional management in its scale, speed, and opacity, creating new forms of power asymmetry between platforms and workers. More recent scholarship has turned to the behavioral effects of algorithmic interventions, examining how workers adapt their strategies in response to algorithmic cues [7]. Goal recommendation systems represent a particularly salient case because they directly target worker motivation and self-regulation, domains that have traditionally been shaped by personal agency and social norms.

Goal setting theory, originating from organizational psychology, provides a foundational lens for understanding the effects of goal recommendations. The theory posits that specific, challenging goals lead to higher performance than vague or easy goals, provided that individuals are committed to the goals and receive feedback on their progress [5]. This framework has been validated across numerous contexts, from manufacturing to education. However, its application to algorithmic goal setting raises important questions. In traditional settings, goals are typically self-set or negotiated with a supervisor who has contextual knowledge of the worker's capabilities and constraints. In contrast, algorithmic goals are generated by models that may not account for individual circumstances, such as health status, family obligations, or local market conditions [8]. Moreover, the platform's objective function, typically maximizing throughput or revenue, may conflict with the worker's preference for stable income or reduced fatigue.

Empirical research on goal setting in the gig economy is still nascent but growing. Field experiments on ride-hailing platforms have shown that personalized earnings goals can increase the number of trips completed per hour, particularly among newer workers who lack experience-based benchmarks [8]. These effects are mediated by the salience of the goal display and the frequency of feedback. However, the same studies reveal heterogeneous effects: workers who consistently fall short of algorithmic goals report lower satisfaction and higher turnover intentions, suggesting that goal recommendations can backfire when they are perceived as unattainable or unfair [9]. Furthermore, goal chasing behavior may lead to effort distortion, where workers prioritize quantity over quality, engage in unsafe practices, or work during periods of low demand simply to meet a target [10].

The labor supply literature in economics offers additional insights into the behavioral mechanisms at play. Standard models of labor supply assume that workers optimize a utility function that balances income and leisure, with marginal decisions driven by wage rates. In the gig economy, however, wage rates are not fixed but vary dynamically with demand, surge multipliers, and algorithmic bonuses [4]. Goal recommendations introduce a non-linear reference point into this decision space, shifting worker focus from marginal wage comparisons to cumulative target attainment. This framing is consistent with prospect theory, which predicts that individuals are loss-averse and will exert extra effort to avoid falling short of a reference point [11]. The algorithmic goal thus functions as an anchor that can increase labor supply at the intensive margin, but may also lead to suboptimal scheduling decisions if the goal is poorly calibrated.

A critical gap in the existing literature is the lack of attention to the system-level consequences of widespread goal recommendation deployment. Most studies focus on individual-level outcomes, such as hours worked or earnings, without examining how aggregate labor supply patterns affect market stability, platform pricing, or the distribution of work across time and space. Additionally, the interaction between goal recommendations and other algorithmic management features, such as surge pricing or acceptance rate penalties, remains underexplored. This paper addresses these gaps by adopting a systems perspective that integrates insights from computer science, economics, and organizational behavior.

3. Conceptual Framework: Goal Recommendations as Socio-Technical Interventions

We conceptualize algorithmic goal recommendations as socio-technical interventions that operate at the intersection of behavioral design, data infrastructure, and platform governance. Unlike traditional goal setting, which is embedded in social relationships and organizational routines, algorithmic goal setting is mediated by software systems that are opaque, adaptive, and optimized for platform-level objectives. This shift has profound implications for how goals are generated, communicated, and enforced.

At the core of any goal recommendation system is a predictive model that estimates a worker's future performance based on historical data, peer comparisons, and contextual features such as time of day, location, and demand forecasts. The model outputs a target value, typically expressed as a number of trips, deliveries, or earnings, that is displayed to the worker through the platform interface. The recommendation may be accompanied by progress indicators, time remaining, and potential rewards for goal attainment, such as bonuses or badges. The system may also adjust goals dynamically in response to worker performance, creating a feedback loop in which past behavior influences future targets [12].

The design of the reference group against which goals are calibrated is a critical architectural choice. Platforms may use absolute benchmarks, such as the median earnings of all workers in a city, or relative benchmarks, such as the performance of workers with similar tenure or location. Each approach carries trade-offs. Absolute benchmarks are simpler to compute but may fail to account for individual differences in capacity or market conditions. Relative benchmarks are more personalized but can create ratchet effects, where high performers are penalized with increasingly difficult goals, while low performers may be disengaged by targets that seem unattainable [13]. The choice of reference group also has fairness implications, as workers in disadvantaged areas or with less favorable schedules may be systematically assigned goals that do not reflect their actual opportunities.

Another key dimension is the frequency and granularity of goal updates. Some platforms set daily goals, while others use weekly or per-session targets. Frequent updates can increase salience and provide real-time feedback, but they may also induce anxiety and encourage short-term thinking at the expense of long-term planning. The timing of goal display is equally important: goals presented at the beginning of a shift may influence the decision to start working, while goals updated mid-shift may affect the decision to continue or stop [14]. These temporal dynamics interact with the worker's internal state, including fatigue, motivation, and competing demands, to produce complex labor supply patterns.

We also emphasize the importance of feedback loops between goal recommendations and platform governance. For example, if goal attainment is linked to access to higher-paying tasks or priority in task assignment, workers may engage in gaming behaviors, such as working during periods of artificially low demand to accumulate trips, or accepting tasks they

would otherwise reject [15]. These behaviors can distort the intended effects of the goal system and create inefficiencies in the broader market. Conversely, if goals are purely informational and carry no consequences, their motivational power may be limited. Finding the right balance between influence and coercion is a central governance challenge.

4. Empirical Evidence and Behavioral Mechanisms

The empirical literature on algorithmic goal recommendations provides a nuanced picture of their effects on labor supply. Field experiments conducted on large ride-hailing platforms have demonstrated that personalized earnings goals can increase the number of trips completed by five to ten percent on average [8]. This effect is most pronounced among workers with low tenure, who lack internal benchmarks for what constitutes a good day of work. For these workers, the algorithm serves as a substitute for experience, providing a credible signal of achievable performance. However, the same studies show that the effect diminishes over time, as workers develop their own heuristics and may become skeptical of algorithmic recommendations that do not align with their subjective experience.

A critical finding from the literature is the heterogeneity of treatment effects. Workers who are consistently below the recommended goal exhibit higher dropout rates and lower platform engagement in subsequent weeks [9]. This suggests that goal recommendations can have a demotivating effect when they are perceived as unattainable, consistent with the predictions of goal setting theory regarding goal difficulty and commitment. Moreover, workers who are close to achieving a goal at the end of a shift are more likely to extend their working time, even when the marginal returns are low due to declining demand or fatigue [10]. This behavior, known as goal gradient effect, can lead to suboptimal labor supply decisions from both the worker's and the platform's perspective.

Effort distortion is another documented consequence. Workers who are focused on meeting a quantitative goal, such as a trip count, may sacrifice service quality, take shortcuts, or ignore safety protocols. In the context of food delivery, for example, workers may rush deliveries or accept orders that require long travel distances, increasing the risk of accidents [16]. In ride-hailing, goal chasing may lead to more aggressive driving or acceptance of rides in unsafe neighborhoods. These externalities are often invisible to the platform because they are not captured in the metrics used to evaluate goal attainment, creating a misalignment between the algorithm's objective and broader social welfare.

The interaction between goal recommendations and platform incentives, such as surge pricing or completion bonuses, further complicates the behavioral picture. When surge pricing is active, the effective wage rate increases, and workers may naturally increase their labor supply without needing a goal nudge. In such cases, goal recommendations may be redundant or even counterproductive, as they divert attention from the most lucrative opportunities [17]. Conversely, during low-demand periods, goal recommendations may be the primary driver of labor supply, leading workers to work during times when earnings are minimal. This raises questions about whether goal systems are designed to benefit workers or to ensure platform capacity during slack times.

5. Architectural Design, Fairness, and Governance

The design of goal recommendation systems involves a series of architectural choices that have profound implications for fairness, robustness, and long-term sustainability. One of the most consequential decisions is the selection of data inputs for the predictive model. Models that rely solely on historical performance may perpetuate existing inequalities, as workers

who have faced structural barriers, such as limited access to high-demand areas or discrimination from customers, will have lower performance histories and thus lower goals [18]. This can create a self-reinforcing cycle of low expectations and low achievement. Models that incorporate exogenous factors, such as local demand forecasts, traffic conditions, or weather, can partially mitigate this issue, but they introduce additional complexity and potential for error.

Transparency is a second critical dimension. Most platforms do not disclose the logic behind goal recommendations, leaving workers to infer the rules through trial and error. This opacity can erode trust and lead to strategic behaviors that undermine the system's effectiveness. For example, workers may deliberately underperform on certain days to avoid being assigned higher goals in the future, a phenomenon known as strategic lowballing [19]. Providing workers with explanations for their goals, such as the reference group used or the factors that influenced the target, could improve acceptance and reduce gaming, but it also risks revealing proprietary algorithms or enabling adversarial manipulation.

The governance of goal adjustment is another area of concern. If goals are updated too aggressively in response to performance improvements, workers may experience a ratchet effect that discourages high effort. Conversely, if goals are too static, they may become irrelevant as market conditions change. Adaptive algorithms that balance responsiveness with stability require careful tuning and ongoing monitoring. Additionally, platforms must decide whether workers can opt out of goal recommendations or customize their targets. Providing such agency can enhance worker satisfaction and autonomy, but it may also reduce the platform's ability to influence labor supply in desired ways [20].

Fairness considerations extend beyond the individual level to encompass group-level disparities. Research has shown that algorithmic systems can produce disparate outcomes across demographic groups, even when they do not explicitly use protected attributes. In the context of goal recommendations, workers in low-income neighborhoods or with non-standard schedules may receive systematically lower goals, which could affect their earnings trajectories and long-term attachment to the platform [21]. Auditing goal recommendation systems for fairness requires access to granular data on worker demographics, performance, and goal assignments, which platforms are often reluctant to share.

6. Policy Implications and Design Recommendations

The findings of this paper have significant implications for platform regulation, worker advocacy, and system design. From a policy perspective, the use of algorithmic goal recommendations raises questions about the boundary between permissible nudging and manipulative control. While goal setting is a well-established motivational tool, its algorithmic deployment in the absence of transparency, consent, and recourse mechanisms may constitute a form of algorithmic paternalism that undermines worker autonomy [22]. Regulatory frameworks developed for algorithmic management, such as the European Union's proposed AI Act, should consider goal recommendation systems as high-risk applications that require impact assessments, transparency obligations, and worker consultation.

From a design perspective, we recommend several principles for building fair and effective goal recommendation systems. First, goals should be based on a combination of individual performance and external market conditions, rather than solely on peer comparisons, to avoid ratchet effects and unfair benchmarks. Second, workers should have the ability to view the rationale behind their goals, adjust their targets within reasonable bounds, and opt out of goal

recommendations without penalty. Third, platforms should conduct regular audits of goal recommendation outcomes to detect and correct disparities across demographic groups. Fourth, goal systems should be integrated with broader worker support mechanisms, such as earnings guarantees or minimum hour commitments, to reduce the pressure to chase unattainable targets [23].

Finally, the long-term sustainability of goal recommendation systems depends on their alignment with worker welfare. Platforms that prioritize short-term productivity gains at the expense of worker burnout or churn will ultimately face higher recruitment and training costs. A more sustainable approach involves using goal recommendations as part of a holistic worker engagement strategy that also includes fair compensation, predictable scheduling, and opportunities for skill development [24]. The goal of algorithmic management should be to enhance, not replace, worker agency.

7. Conclusion

This paper has presented a comprehensive systems-level analysis of algorithmic goal recommendations and their effects on labor supply among gig economy workers. We have argued that these systems function as socio-technical interventions that shape worker behavior through personalized reference points, feedback mechanisms, and integration with platform governance structures. The empirical evidence suggests that goal recommendations can increase short-term productivity but also introduce risks of effort distortion, demotivation, and inequity. The architectural choices underlying these systems, including reference group selection, goal adjustment rates, and transparency, are critical determinants of their fairness and robustness.

The broader implication of this research is that algorithmic management tools cannot be evaluated solely on their immediate behavioral effects; they must be understood as components of larger systems that interact with market dynamics, worker psychology, and social structures. As digital platforms continue to expand their influence over labor markets, the design of goal recommendation systems will become an increasingly important site of contestation between efficiency and equity. Future research should examine the long-term effects of goal recommendations on worker careers, the interaction between goal systems and other algorithmic management features, and the potential for participatory design approaches that give workers a voice in the algorithms that govern their work.

References

1. 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.
2. Rosenblat, A., & Stark, L. (2016). Algorithmic labor and information asymmetries: A case study of Uber's drivers. *International Journal of Communication*, 10, 3758-3784.
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*, 1603-1612.
4. 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.

5. 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.
6. Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at work: The new contested terrain of control. *Academy of Management Annals*, 14(1), 366-410.
7. Jarrahi, M. H., & Sutherland, W. (2019). Algorithmic management and algorithmic competencies: Understanding and appropriating algorithms in gig work. In *Information in Contemporary Society* (pp. 578-589). Springer.
8. 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.
9. Allon, G., & Babich, V. (2020). Crowdsourcing and crowdfunding in the sharing economy. In *The Sharing Economy* (pp. 75-97). Springer.
10. Horton, J. J., & Chilton, L. B. (2010). The labor economics of paid crowdsourcing. *Proceedings of the 11th ACM Conference on Electronic Commerce*, 209-218.
11. Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263-291.
12. Boudreau, K. J., & Jeppesen, L. B. (2015). Unpaid crowd complements: The platform network effect mirage. *Strategic Management Journal*, 36(12), 1761-1777.
13. Gneezy, U., & Rustichini, A. (2000). Pay enough or don't pay at all. *Quarterly Journal of Economics*, 115(3), 791-810.
14. Kaur, H., Williams, A. C., McDuff, D., Czerwinski, M., Teevan, J., & Iqbal, S. T. (2020). Optimizing for happiness and productivity: Modeling opportune moments for transitions and breaks at work. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1-15.
15. Dai, W., & Liu, D. (2022). The dark side of gamification: How game-like features affect worker behavior on digital platforms. *Management Science*, 68(9), 6821-6840.
16. Ravenelle, A. J. (2019). *Hustle and gig: Struggling and surviving in the sharing economy*. University of California Press.
17. Cachon, G. P., Daniels, K. M., & Lobel, R. (2017). The role of surge pricing on a service platform with self-scheduling capacity. *Manufacturing & Service Operations Management*, 19(3), 368-384.
18. Angwin, J., Larson, J., Mattu, S., & Kirchner, L. (2016). Machine bias. *ProPublica*, May 23.
19. Holtz, D., Carterette, B., Chandar, P., Nazari, Z., Cramer, H., & Aral, S. (2020). The engagement-diversity connection: Evidence from a field experiment on Spotify. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 75-84.
20. Bucher, T. (2018). *If...then: Algorithmic power and politics*. Oxford University Press.
21. Hoffmann, A. L. (2019). Where fairness fails: Data, algorithms, and the limits of antidiscrimination discourse. *Information, Communication & Society*, 22(7), 900-915.
22. Yeung, K. (2017). Hypernudge: Big data as a mode of regulation by design. *Information, Communication & Society*, 20(1), 118-136.

23. Dube, A., & Kaplan, E. (2020). The gig economy and the future of work: A policy agenda. The Hamilton Project, Brookings Institution.
24. Berg, J. (2016). Income security in the on-demand economy: Findings and policy lessons from a survey of crowdworkers. International Labour Organization.