

# AI-Driven Productivity Nudges and Task Completion Behavior on Digital Labor Platforms

Chetan A. Agarwal

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

Aarav D. Iyer

Department of Electrical Engineering and Computer Science, University of Kansas, Lawrence, KS, USA.  
aaravmail@ku.edu

Timothy Watson

Department of Computer Science, University of Houston, Houston, TX, USA.  
timothywatson968@uh.edu

Yash M. Srinivasan

Department of Computer Science, University of North Texas, Denton, TX, USA.  
yash1997@unt.edu

## Abstract

Digital labor platforms have fundamentally restructured the relationship between workers, tasks, and algorithmic management systems. Central to this transformation is the deployment of artificial intelligence-driven productivity nudges, which are subtle, often personalized interventions designed to influence worker behavior without imposing explicit mandates. This paper presents a comprehensive analysis of how such nudges affect task completion behavior on digital labor platforms, examining the interplay between system architecture, worker autonomy, and platform governance. We develop a conceptual framework that situates productivity nudges within the broader context of socio-technical systems, emphasizing the structural trade-offs between algorithmic efficiency and human agency. Drawing on empirical evidence from field experiments and observational studies, we analyze how nudge design parameters, including timing, frequency, framing, and personalization, interact with individual worker characteristics to produce heterogeneous behavioral responses. The paper further investigates the infrastructural dependencies underlying nudge deployment, including real-time data pipelines, feedback loops, and recommendation engines, and examines how these technical components introduce latent biases and fairness concerns. We discuss the sustainability of nudge-based interventions over time, considering habituation effects, worker resistance, and the potential for algorithmic fatigue. Governance implications are explored through the lens of platform accountability, transparency obligations, and the ethical boundaries of behavioral modification. The analysis reveals that while AI-driven nudges can enhance short-term productivity metrics, their long-term efficacy is contingent upon careful calibration of system architecture, respect for worker autonomy, and robust institutional oversight. We conclude by proposing design principles for responsible nudge systems that balance platform objectives with worker well-being, and we identify critical directions for future interdisciplinary research at the intersection of artificial intelligence, labor economics, and human-computer interaction.

## Keywords

artificial intelligence, productivity nudges, digital labor platforms, algorithmic management, task completion behavior, socio-technical systems, platform governance.

## 1. Introduction

The proliferation of digital labor platforms has introduced a new paradigm of work organization in which algorithmic systems mediate nearly every aspect of task allocation, performance monitoring, and compensation [1]. These platforms, spanning domains from ride-hailing and food delivery to micro-task markets and freelance project management, rely on sophisticated AI-driven infrastructures to coordinate large, geographically distributed workforces. Within this ecosystem, productivity nudges have emerged as a prominent mechanism for influencing worker behavior without resorting to direct commands or coercive measures [2]. These nudges, which may take the form of goal reminders, progress notifications, comparative performance feedback, or personalized task recommendations, are designed to steer workers toward behaviors that align with platform objectives, such as increased task acceptance rates, faster completion times, or higher engagement levels [3].

The theoretical underpinnings of productivity nudges draw from behavioral economics, particularly the concept of choice architecture articulated by Thaler and Sunstein, which posits that subtle changes in the presentation of options can significantly alter decision outcomes while preserving freedom of choice [4]. In the context of digital labor platforms, AI systems operationalize this principle by dynamically tailoring nudges to individual workers based on real-time behavioral data, historical performance patterns, and contextual variables [5]. This personalization capability represents a significant departure from static, one-size-fits-all interventions, enabling platforms to optimize nudge effectiveness at scale. However, the deployment of AI-driven nudges also raises fundamental questions about the boundaries of acceptable influence, the distribution of power between platform operators and workers, and the long-term consequences for worker autonomy and well-being [6].

This paper addresses a critical gap in the existing literature by providing a systems-level analysis of AI-driven productivity nudges and their implications for task completion behavior on digital labor platforms. While prior research has examined nudge effectiveness in controlled experimental settings [7] or focused on specific platform contexts [8], less attention has been devoted to the architectural, infrastructural, and governance dimensions that shape nudge design and deployment in real-world systems. Our analysis is motivated by the recognition that nudges do not operate in isolation but are embedded within complex socio-technical systems comprising data collection pipelines, machine learning models, feedback mechanisms, and organizational policies [9]. Understanding these interdependencies is essential for evaluating both the benefits and risks of nudge-based interventions.

The paper proceeds as follows. Section 2 develops a conceptual framework for analyzing productivity nudges within digital labor platforms, drawing on theories of algorithmic management and socio-technical systems. Section 3 examines the mechanisms through which nudges influence task completion behavior, including psychological, cognitive, and social pathways. Section 4 investigates the system architecture and infrastructural dependencies that enable nudge deployment, with attention to data flows, model training, and real-time inference. Section 5 addresses issues of fairness, bias, and equity in nudge design and implementation. Section 6 explores the sustainability of nudge interventions over time, considering habituation, resistance, and algorithmic fatigue. Section 7 discusses governance

and policy implications, including transparency, accountability, and regulatory frameworks. Section 8 concludes with design principles and future research directions.

## **2. Conceptual Framework: Nudges in Algorithmic Management Systems**

The concept of algorithmic management refers to the use of software algorithms and data-driven decision-making to coordinate and control labor processes that were traditionally managed by human supervisors [10]. Digital labor platforms exemplify this approach, as they rely on algorithms to assign tasks, evaluate performance, determine compensation, and shape worker behavior through various incentive mechanisms. Within this framework, AI-driven productivity nudges represent a distinct class of algorithmic interventions that operate through informational and motivational channels rather than through direct commands or financial incentives [11].

To understand how nudges function within algorithmic management systems, it is useful to adopt a socio-technical systems perspective that emphasizes the interdependence of technical components and social actors [12]. From this viewpoint, a nudge is not merely a message delivered to a worker but rather an intervention that emerges from a complex chain of data collection, algorithmic processing, and human interpretation. The technical infrastructure includes sensors and logging systems that capture worker actions, databases that store historical behavioral data, machine learning models that predict worker responses, and delivery systems that transmit nudges through platform interfaces [13]. The social dimension encompasses worker perceptions of nudge legitimacy, their cognitive and emotional responses, and their strategies for resisting or accommodating algorithmic influence.

A key structural trade-off in nudge design concerns the balance between personalization and standardization. Highly personalized nudges, which leverage detailed individual data to tailor content, timing, and framing, may achieve greater behavioral impact but also raise privacy concerns and increase the risk of manipulative targeting [14]. Standardized nudges, while less intrusive, may fail to account for individual differences in motivation, skill level, and contextual constraints, thereby reducing overall effectiveness. Platform operators must navigate this trade-off while also considering the computational costs of personalization, the interpretability of nudge algorithms, and the potential for feedback loops that amplify initial behavioral differences across workers [15].

Another critical dimension of the conceptual framework involves the temporal dynamics of nudge effects. Nudges are typically designed to produce immediate behavioral changes, such as accepting a task or completing it within a specified timeframe. However, the cumulative effects of repeated nudging may differ substantially from short-term responses. Workers may habituate to certain types of nudges, requiring platforms to continuously vary intervention strategies to maintain effectiveness [16]. Alternatively, workers may develop resistance or reactance, particularly if they perceive nudges as undermining their autonomy or as being motivated solely by platform profit maximization [17]. These dynamic considerations necessitate a longitudinal perspective on nudge evaluation, one that accounts for adaptation, learning, and strategic behavior on the part of workers.

## **3. Mechanisms of Nudge Influence on Task Completion Behavior**

The influence of AI-driven productivity nudges on task completion behavior operates through multiple psychological and cognitive mechanisms. One primary mechanism is goal priming, wherein nudges that remind workers of their stated or implied goals activate motivational processes that increase effort and persistence [18]. For example, a nudge that displays a

worker's progress toward a daily earnings target may reinforce commitment to that goal, leading to higher task acceptance rates and faster completion times. The effectiveness of goal priming depends on the specificity and attainability of the goal, as well as the worker's identification with the goal as personally meaningful rather than externally imposed [19].

A second mechanism involves social comparison, which leverages the human tendency to evaluate one's own performance relative to others. Platforms frequently deploy nudges that inform workers of their ranking within a peer group, their percentile performance, or the average completion times of similar workers [20]. Such comparative feedback can induce competitive motivation, particularly among workers who value status or who perceive their performance as lagging behind peers. However, social comparison nudges also carry risks, including demotivation for low-performing workers, increased anxiety, and the potential for unhealthy competition that undermines collaboration and knowledge sharing [21]. The design of comparison frames, such as whether to highlight top performers or average performance, significantly moderates these effects.

A third mechanism is temporal anchoring, wherein nudges influence behavior by altering the perceived salience of time-related factors. For instance, countdown timers that display remaining time until a task must be accepted or completed can create a sense of urgency that accelerates decision-making [22]. Similarly, nudges that highlight the opportunity cost of delaying task acceptance, such as projected earnings loss, can shift workers' time preferences toward immediate action. Temporal anchoring is particularly effective in contexts where tasks have short deadlines or where platform algorithms dynamically adjust task availability based on worker responsiveness [23]. However, excessive reliance on temporal pressure may lead to rushed decisions, reduced task quality, and increased worker stress.

The personalization of nudge mechanisms through AI introduces additional complexity. Machine learning models can predict which mechanism is most likely to influence a given worker at a given moment, based on historical response patterns, demographic characteristics, and contextual variables such as time of day or recent task history [24]. This predictive capability enables platforms to dynamically select and sequence nudges, potentially increasing overall effectiveness. However, personalization also raises concerns about algorithmic profiling and the potential for discriminatory outcomes if models encode biases present in training data [25]. Furthermore, workers may develop strategies to game personalized nudge systems, such as deliberately altering their behavior to receive more favorable nudges, thereby undermining the validity of the underlying models.

#### **4. System Architecture and Infrastructural Dependencies**

The deployment of AI-driven productivity nudges relies on a sophisticated technical infrastructure that integrates data collection, storage, processing, and delivery components. At the foundation of this infrastructure lies the data pipeline, which continuously captures worker actions, including task acceptance and rejection events, completion times, idle periods, location data, and interaction patterns with platform interfaces [26]. These data streams are typically ingested into cloud-based data warehouses or data lakes, where they are cleaned, transformed, and aggregated for analysis. The scale of data generation on major digital labor platforms is substantial, with millions of workers generating billions of events daily, necessitating distributed computing frameworks and real-time processing capabilities.

Machine learning models form the core intelligence of nudge systems. These models are trained on historical data to predict worker responses to different nudge types, timings, and

framings. Common modeling approaches include reinforcement learning, which optimizes nudge policies through trial-and-error interaction with workers, and supervised learning, which predicts outcomes such as task acceptance probability or completion time based on worker features and nudge parameters [27]. Model training requires careful consideration of data quality, feature engineering, and validation strategies to avoid overfitting and ensure generalizability across diverse worker populations. Additionally, models must be updated periodically to account for changes in worker behavior, platform policies, and external conditions such as seasonal demand fluctuations.

Real-time inference engines are responsible for delivering nudges at the appropriate moments. These systems evaluate the current state of each worker, including their recent activity, current task availability, and predicted responsiveness, and select the optimal nudge from a predefined set of interventions. The latency requirements for nudge delivery are often stringent, particularly in contexts where workers are actively browsing task listings or engaged in ongoing tasks. Delays of even a few seconds can reduce nudge effectiveness or cause workers to miss time-sensitive opportunities [28]. Consequently, inference engines are typically deployed on edge computing nodes or through content delivery networks to minimize latency.

The infrastructural dependencies of nudge systems introduce several vulnerabilities and trade-offs. Reliability is a primary concern, as system failures or data pipeline disruptions can lead to missed nudges, inconsistent worker experiences, or unintended behavioral consequences. Redundancy mechanisms, failover protocols, and monitoring systems are essential for maintaining operational continuity. Scalability is another critical consideration, as nudge systems must accommodate rapid growth in worker populations without degradation in performance or personalization quality. Cloud-based architectures with auto-scaling capabilities can address scalability requirements, but they also introduce dependencies on third-party infrastructure providers and potential cost implications.

## **5. Fairness, Bias, and Equity in Nudge Design**

The algorithmic nature of AI-driven productivity nudges raises significant concerns about fairness, bias, and equity. These concerns stem from multiple sources, including biased training data, unequal access to platform features, and differential impacts of nudge interventions across worker subgroups [29]. For instance, if historical data reflects systematic differences in task completion behavior across demographic groups, models trained on such data may perpetuate or amplify these disparities by delivering less effective or more intrusive nudges to certain groups. Similarly, workers with limited digital literacy or language barriers may be disproportionately affected by text-heavy nudges or those that rely on culturally specific references.

Algorithmic bias in nudge systems can manifest in several ways. Selection bias occurs when the data used to train nudge models is not representative of the full worker population, leading to models that perform poorly for underrepresented groups. Measurement bias arises when the metrics used to evaluate nudge effectiveness, such as task acceptance rate or completion time, are themselves influenced by factors unrelated to worker effort, such as task difficulty or platform algorithm quirks. Interaction bias occurs when nudges are designed based on assumptions about worker preferences or motivations that do not hold across diverse cultural or socioeconomic contexts [30]. Addressing these biases requires careful attention to data collection practices, model evaluation methodologies, and iterative testing with diverse worker samples.

Equity considerations extend beyond bias to encompass the distribution of nudge benefits and burdens across workers. Nudges that are highly effective for some workers may be ineffective or counterproductive for others, potentially widening productivity gaps within the workforce. Workers who are already highly motivated may receive fewer nudges or less intensive interventions, while those who are struggling may be subjected to more frequent or more coercive nudges, creating a feedback loop that exacerbates inequality [31]. Platform operators must therefore consider not only average nudge effectiveness but also distributional outcomes when designing and evaluating nudge policies.

Fairness in nudge design also involves transparency and consent. Workers may not be aware that they are being subjected to AI-driven nudges, or they may not understand how nudge algorithms work or what data is being used to personalize interventions. Lack of transparency undermines informed consent and reduces workers' ability to make autonomous decisions about their behavior [32]. Some platforms have begun to provide workers with information about nudge practices, such as disclosure statements or opt-out options, but these measures remain inconsistent across platforms and jurisdictions. Developing standardized transparency frameworks for nudge systems is an important step toward ensuring that workers can exercise meaningful agency in algorithmic management environments.

## **6. Sustainability and Long-Term Dynamics**

The sustainability of AI-driven productivity nudges over extended periods is a critical concern that has received limited empirical attention. Initial evidence suggests that nudge effectiveness tends to decay over time as workers habituate to repeated interventions [33]. Habituation occurs when workers become desensitized to nudges, requiring increasingly frequent or intense interventions to achieve the same behavioral response. Platforms may respond by escalating nudge intensity, such as increasing message frequency, using more salient visual cues, or incorporating stronger emotional appeals. However, this escalation can lead to algorithmic fatigue, wherein workers experience cognitive overload, reduced motivation, or active resistance to nudge influence.

Worker resistance to nudges can take various forms, including ignoring or dismissing nudge messages, deliberately altering behavior to avoid predicted nudge responses, or engaging in collective action to protest algorithmic management practices [34]. Resistance is more likely when workers perceive nudges as manipulative, unfair, or misaligned with their own interests. Platforms that fail to account for worker resistance may find that nudge interventions become counterproductive, reducing trust and engagement rather than enhancing productivity. Building sustainable nudge systems requires ongoing calibration of intervention strategies, incorporating worker feedback and adapting to changing behavioral patterns.

The long-term effects of nudges on worker well-being represent another dimension of sustainability. While short-term productivity gains may benefit platform operators and, in some cases, workers who earn more, the cumulative psychological costs of constant behavioral steering are less well understood. Chronic exposure to productivity nudges may contribute to stress, anxiety, and burnout, particularly for workers who already face precarious working conditions and limited job security [35]. Platforms have a responsibility to monitor worker well-being outcomes and to design nudge systems that respect workers' need for rest, autonomy, and psychological safety.

## **7. Governance and Policy Implications**

The governance of AI-driven productivity nudges on digital labor platforms involves a complex interplay of platform self-regulation, industry standards, and public policy. Currently, most platforms operate with limited external oversight of nudge practices, relying on internal guidelines and ethical review processes that vary widely in rigor and transparency [36]. This governance gap is concerning given the potential for nudges to influence significant aspects of workers' livelihoods, including earnings, scheduling flexibility, and career progression. Developing robust governance frameworks requires input from multiple stakeholders, including workers, platform operators, researchers, policymakers, and civil society organizations.

Transparency is a foundational principle for responsible nudge governance. Workers should have access to clear information about what data is collected, how nudge algorithms work, and what behavioral changes are being targeted. Transparency enables workers to make informed decisions about their participation in nudge systems and to hold platforms accountable for unfair or manipulative practices. However, transparency alone is insufficient if workers lack the technical literacy to interpret algorithmic processes or if platforms provide only superficial disclosures. Meaningful transparency requires accessible explanations, opportunities for worker input, and mechanisms for contesting nudge decisions.

Accountability mechanisms are equally important. Platforms should be responsible for the outcomes of their nudge systems, including any unintended negative effects on worker well-being, fairness, or autonomy. This accountability can be operationalized through regular audits of nudge algorithms, impact assessments that evaluate distributional outcomes, and grievance procedures that allow workers to seek redress for harmful nudge practices [37]. Regulatory frameworks, such as the European Union's General Data Protection Regulation and proposed Artificial Intelligence Act, provide some baseline requirements for algorithmic transparency and accountability, but their application to nudge systems on digital labor platforms remains an area of active development.

Policy interventions may also address the structural power imbalances that underlie nudge deployment. For example, requirements for worker representation in platform governance, including participation in the design and evaluation of nudge systems, could help ensure that worker perspectives are incorporated into decision-making processes. Similarly, policies that mandate minimum standards for nudge transparency, consent, and opt-out options could empower workers to exercise greater control over their interactions with algorithmic management systems. The challenge for policymakers is to craft regulations that are specific enough to address the unique features of nudge systems while remaining flexible enough to accommodate technological evolution and platform diversity.

## **8. Conclusion**

AI-driven productivity nudges represent a powerful but contested tool for shaping task completion behavior on digital labor platforms. This paper has provided a comprehensive analysis of nudge systems from a socio-technical perspective, examining the mechanisms of influence, infrastructural dependencies, fairness concerns, sustainability challenges, and governance implications. Our analysis reveals that while nudges can enhance short-term productivity metrics, their long-term efficacy and ethical acceptability depend on careful system design, robust oversight, and genuine respect for worker autonomy. The structural trade-offs inherent in nudge design, including those between personalization and standardization, effectiveness and fairness, and short-term gains and long-term sustainability, require ongoing deliberation and empirical investigation.

Several directions for future research emerge from this analysis. First, longitudinal studies that track nudge effects over months or years are needed to understand habituation, resistance, and well-being outcomes. Second, comparative research across different platform types, worker populations, and cultural contexts would illuminate the contextual factors that moderate nudge effectiveness and fairness. Third, interdisciplinary collaborations between computer scientists, labor economists, organizational psychologists, and legal scholars are essential for developing holistic frameworks that address the technical, social, and regulatory dimensions of nudge systems. Finally, participatory design approaches that involve workers as co-creators of nudge interventions could lead to systems that better balance platform objectives with worker values and preferences.

In conclusion, the responsible deployment of AI-driven productivity nudges requires a shift from a purely optimization-oriented mindset to one that embraces ethical reflection, stakeholder engagement, and adaptive governance. Platforms that invest in transparent, fair, and sustainable nudge systems are likely to build greater trust with workers, reduce resistance, and achieve more durable productivity improvements. As digital labor platforms continue to evolve, the principles and practices developed for nudge governance will inform broader debates about the role of artificial intelligence in shaping the future of 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. Ye, T., & Kankanhalli, A. (2020). Nudges in digital platforms: A review and research agenda. *Journal of the Association for Information Systems*, 21(5), 1189-1224.
3. Caraway, B. R. (2024). Algorithmic management and worker resistance on digital labor platforms. *New Media and Society*, 26(2), 789-809.
4. Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press.
5. 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.
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. Dai, W., & Li, Y. (2022). The effects of goal-setting nudges on worker productivity: Evidence from a field experiment. *Management Science*, 68(4), 2456-2474.
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. Jarrahi, M. H., & Sutherland, W. (2019). Algorithmic management and algorithmic competencies: Understanding and managing the relationship between algorithms and workers. *Journal of the Association for Information Science and Technology*, 70(12), 1337-1352.
10. Rosenblat, A., & Stark, L. (2016). Algorithmic labor and information asymmetries: A case study of Uber's drivers. *International Journal of Communication*, 10, 3758-3784.

11. Möhlmann, M., & Zalmanson, L. (2017). Hands on the wheel: Navigating algorithmic management and Uber drivers. *Proceedings of the International Conference on Information Systems*, 1-17.
12. Trist, E. L., & Bamforth, K. W. (1951). Some social and psychological consequences of the longwall method of coal-getting. *Human Relations*, 4(1), 3-38.
13. Cram, W. A., & Wiener, M. (2020). Technology-mediated control: A review and synthesis of the literature. *Journal of Information Technology*, 35(3), 223-247.
14. Acquisti, A., Brandimarte, L., & Loewenstein, G. (2015). Privacy and human behavior in the age of information. *Science*, 347(6221), 509-514.
15. Holman, D., & Rafferty, A. E. (2022). The algorithmic workplace: A multilevel perspective on the effects of algorithmic management on worker well-being. *Journal of Management*, 48(6), 1678-1711.
16. Sunstein, C. R. (2017). Nudges that fail. *Behavioural Public Policy*, 1(1), 4-25.
17. Brehm, J. W. (1966). *A theory of psychological reactance*. Academic Press.
18. 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.
19. Ryan, R. M., & Deci, E. L. (2000). Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *American Psychologist*, 55(1), 68-78.
20. Festinger, L. (1954). A theory of social comparison processes. *Human Relations*, 7(2), 117-140.
21. Buunk, A. P., & Gibbons, F. X. (2007). Social comparison: The end of a theory and the emergence of a field. *Organizational Behavior and Human Decision Processes*, 102(1), 3-21.
22. Zauberman, G., & Lynch, J. G. (2005). Resource slack and propensity to discount delayed investments of time versus money. *Journal of Experimental Psychology: General*, 134(1), 23-37.
23. Milkman, K. L., Rogers, T., & Bazerman, M. H. (2009). Highlights for tomorrow: A field experiment on the timing of goal setting. *Journal of Applied Psychology*, 94(6), 1431-1441.
24. Adomavicius, G., & Tuzhilin, A. (2005). Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, 17(6), 734-749.
25. Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. *California Law Review*, 104(3), 671-732.
26. Chen, M., Mao, S., & Liu, Y. (2014). Big data: A survey. *Mobile Networks and Applications*, 19(2), 171-209.
27. Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (2nd ed.). MIT Press.
28. Dean, J., & Ghemawat, S. (2008). MapReduce: Simplified data processing on large clusters. *Communications of the ACM*, 51(1), 107-113.

29. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 1-35.
30. Danks, D., & London, A. J. (2017). Algorithmic bias in autonomous systems. *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, 4691-4697.
31. Kasy, M., & Abebe, R. (2021). Fairness, equality, and power in algorithmic decision-making. *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 576-586.
32. Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data and Society*, 3(2), 1-21.
33. Halpern, D. (2015). *Inside the nudge unit: How small changes can make a big difference*. WH Allen.
34. Pasquale, F. (2015). *The black box society: The secret algorithms that control money and information*. Harvard University Press.
35. Tarafdar, M., Tu, Q., & Ragu-Nathan, T. S. (2010). Impact of technostress on end-user satisfaction and performance. *Journal of Management Information Systems*, 27(3), 303-334.
36. Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. *Proceedings of the 2019 ACM Conference on Fairness, Accountability, and Transparency*, 59-68.
37. Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., ... & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. *Proceedings of the 2020 ACM Conference on Fairness, Accountability, and Transparency*, 33-44.