

AI-Driven Performance Feedback and Worker Persistence Under Income Uncertainty

Cesar Cohen

Department of Computer Science, University of Central Florida, Orlando, FL, USA.
hellocesar@ucf.edu

Jesse C. Norris

Department of Computer Science, University of Alabama at Birmingham, Birmingham, AL, USA.
norris1975@uab.edu

Milos A. Wells

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

Abstract

The proliferation of platform-mediated contingent work has made income uncertainty a defining feature of contemporary labor markets, while artificial intelligence increasingly mediates performance feedback in these environments. This paper presents a system-level examination of how AI-driven performance feedback interacts with worker persistence under conditions of income volatility. Drawing on interdisciplinary scholarship spanning computer science, organizational behavior, behavioral economics, and law, we analyze the sociotechnical architectures that algorithmic feedback systems inhabit, and the structural trade-offs they impose on governance, fairness, privacy, and sustainability. We argue that AI feedback functions as a coordinating mechanism that can either amplify or attenuate the motivational consequences of income variability, depending on its design transparency, temporal granularity, adaptivity, and embeddedness in broader platform governance regimes. The analysis reveals fundamental tensions between precision personalization and worker autonomy, between predictive optimization and interpretive agency, and between feedback-induced persistence and the erosion of long-term skill development. Using a multi-paradigm lens, the paper explores how different socio-infrastructure models yield divergent outcomes for persistence, inequality, and platform resilience. Policy implications related to openness mandates, algorithmic auditing, and cooperative data trusts are evaluated, offering a forward-looking synthesis for scholars, engineers, and regulators grappling with the transformation of effort regulation by AI in uncertain income settings.

Keywords

AI feedback, gig economy, worker persistence, income uncertainty, algorithmic management, platform governance, sociotechnical systems.

1. Introduction

The contemporary platform economy has reconfigured the nature of work by decoupling labor from the institutional buffers of traditional employment, substituting stable salaries with piece-rate, task-based compensation subject to continuous variation in demand, pricing, and

algorithmic coordination. In these settings, workers face persistent income uncertainty, which has been shown to shape effort allocation, emotional wellbeing, and long-term career trajectories. Simultaneously, platform operators have deployed ever more sophisticated artificial intelligence systems that deliver real-time, individualized performance feedback, often in the form of nudges, rankings, badges, and predictive guidance. This confluence of algorithmic feedback and income volatility creates a unique sociotechnical landscape for understanding worker persistence, a construct that captures the sustained engagement in productive tasks despite adverse or ambiguous conditions.

Existing research on gig work has thoroughly documented the mechanisms of algorithmic control and information asymmetries that characterize platforms [1]. Empirical evidence on the size and composition of the contingent workforce further underscores the magnitude of these structural shifts [2]. Meanwhile, classic theories of self-determination posit that feedback influences persistence by affecting the satisfaction of basic psychological needs for competence, autonomy, and relatedness [3]. Bringing these streams together, we posit that the interaction between AI-driven feedback and income uncertainty constitutes a complex adaptive system whose behavioral and organizational outcomes cannot be reduced to linear cause-effect relationships. Rather, persistence emerges from the interplay between algorithmic design choices, institutional arrangements, data infrastructures, and the interpretive frames workers deploy when making sense of feedback under financial precarity.

The present analysis adopts a system-level perspective to examine the architectural configurations, governance models, and structural trade-offs that shape the impact of AI feedback on persistence. We forgo granular mathematical modeling in favor of a conceptual synthesis that highlights key tensions: the tension between precision and opacity, between personalization and surveillance, and between short-term task persistence and long-term capability erosion. The paper proceeds by first unpacking the microfoundational mechanisms linking feedback, income volatility, and persistence (Section 2), then characterizing the architecture of AI-driven feedback systems on digital labor platforms (Section 3), followed by analyses of governance and agency (Section 4), structural trade-offs (Section 5), fairness and distributional consequences (Section 6), infrastructure scalability and robustness (Section 7), sustainability (Section 8), and policy implications (Section 9). A concluding synthesis draws together the threads and proposes a research agenda centered on value-sensitive design and pluralistic governance of algorithmic feedback regimes.

2. The Microfoundations of Persistence: Feedback, Income Volatility, and Motivation

Worker persistence in the face of income uncertainty is not a monolithic phenomenon but arises from a mesh of cognitive, affective, and social processes. Uncertainty about earnings can elevate stress and reduce intrinsic motivation, yet it may also spur heightened effort when feedback signals provide clear, proximate targets. Ethnomethodological accounts of on-demand workers' lived experience reveal that income insecurity is a constant background condition that shapes how platform prompts are interpreted; when a budget shortfall looms, a personalized nudge to accept more rides or tasks may be perceived as a helpful boost toward a daily target, or alternatively as a coercive manipulation [5]. Hence, the same feedback content can produce diametrically opposite effects on persistence depending on the financial mental accounting frame workers adopt.

The Janus-faced nature of AI-generated feedback is well-established: while such feedback can enhance short-run performance, it may simultaneously undermine deeper learning and self-regulatory capacity [6]. In gig platforms, this duality is amplified because workers rarely

receive developmental feedback that would support skill accumulation; they are instead inundated with evaluative metrics that foreground output quantity and customer ratings. This narrow feedback ecology can erode the sense of mattering that is critical for sustained engagement under precarious conditions [7]. Without relational acknowledgment of worker context, AI feedback systems risk reducing the motivational calculus to a mechanical stimulus-response pattern that collapses under prolonged income shocks.

Behavioral economic research on reference points and effort provision offers a complementary lens. When workers form reference incomes based on recent earnings or peer comparisons, the gap between actual and reference income determines the marginal utility of additional effort [8]. AI feedback that adjusts targets dynamically can shift these reference points, sometimes in ways that exploit loss aversion to compel persistence beyond what a rational agent model would predict. Yet such exploitation can backfire if workers perceive the platform as deliberately manipulating reference points to depress effective wages. The literature on accountability further underscores the importance of justifiability in algorithmic interventions; without transparent rationale, workers may attribute negative income outcomes to unfair algorithmic decisions, corroding trust and undermining persistence [9].

Large-scale empirical studies of ride-hailing labor supply have demonstrated that drivers exhibit a remarkable degree of flexibility in responding to time-varying wages, yet the same studies reveal persistent earnings targeting behavior that is inconsistent with standard neoclassical models but consistent with income reference dependence [10]. In an environment where an algorithm provides real-time surge pricing information alongside ride acceptance prompts, the feedback system essentially serves as a simultaneous wage signal and effort directive. The ambiguity regarding whether the algorithm is optimizing for platform throughput or driver welfare introduces an interpretive layer that shapes the motivational potency of the feedback. Older taxonomies of formative feedback emphasize that effective feedback must close the gap between current and desired performance by clarifying goals, reducing cognitive load, and promoting self-assessment [11]. These principles are often violated in the lean, notification-driven feedback architectures of gig platforms.

Another critical insight emerges from studies of information asymmetries in platform-mediated work. When algorithmic rating systems determine task allocation and future earning opportunities, feedback is not a neutral signal but an instrument of power that can lock workers into cycles of persistent overwork to maintain a high rating [12]. This structural coupling of feedback to eligibility directly links persistence to survival under income uncertainty, transforming a potentially supportive feature into a disciplinary apparatus. The interplay of these microfoundational mechanisms sets the stage for considering the architectural choices that platform designers face.

Recent field experiments bring empirical weight to the idea that worker persistence can be systematically influenced by the structure of goal-setting mechanisms embedded in platform interfaces. In an environment where income fluctuation is the norm, being prompted to set a personal performance goal can increase both effort and earnings, yet the durability of such effects depends on the richness and timeliness of the feedback that follows [13]. When AI systems merely track goal progress without offering learning-oriented guidance, the motivational boost may be transient, dissipating once financial pressures intensify. Thus, the microfoundational landscape reveals a deep connectivity among feedback design, income reference frames, interpretive justice, and the institutional envelope that governs platform-worker interactions.

3. Architecture of AI-Driven Feedback Systems in Contingent Work Platforms

The technical infrastructure that delivers AI-generated feedback in contingent work contexts comprises multiple interconnected layers: data acquisition and fusion, predictive modeling, feedback generation and personalization, delivery interface, and monitoring for model drift. Data acquisition draws from a heterogeneous sensor and log stream that includes GPS traces, task completion times, acceptance rates, customer ratings, cancellation events, and, increasingly, biometric and behavioral telemetry. These data are fused to construct rich worker profiles that feed into machine learning pipelines designed to predict effort elasticity, churn risk, and responsiveness to incentives.

The predictive engine often employs gradient-boosted trees, deep neural networks, or reinforcement learning agents to estimate the counterfactual impact of a feedback message on a worker's next-period productivity. Such models are trained on historical A/B test data, creating a closed loop where the system continuously refines its understanding of worker sensitivity. However, the architecture's reliance on observational data introduces endogeneity that complicates causal inference; workers who receive certain feedback are systematically different from those who do not, and the platform's own allocation algorithms confound the treatment effect. Without careful debiasing, the feedback model may overfit to transient patterns and produce recommendations that are not robust to shifts in the macroeconomic environment that drive income uncertainty.

The feedback delivery layer often takes the form of mobile push notifications, in-app dashboards, and summary reports. Designers face a crucial architectural choice between synchronous, just-in-time feedback nudges delivered during task execution and asynchronous, post-hoc performance summaries. The former can exploit momentary cognitive states to maximize persistence in the current work session, while the latter may support reflective learning but lose temporal relevance. In practice, many platforms adopt a hybrid approach where reinforcement learning policies govern the timing and channel of feedback, balancing real-time optimization against longer-term user experience metrics. Architectural decisions about data latency, edge versus cloud processing, and personalization granularity directly shape the system's ability to adapt feedback to the instantaneous income context of the worker.

An underappreciated architectural dimension is the feedback system's interface with other platform subsystems: dynamic pricing, task matching, and dispute resolution. The interdependencies create emergent behaviors. For example, if the feedback engine encourages a driver to stay online during a demand lull through gamified milestones, but the pricing algorithm simultaneously lowers surge multipliers due to an increased supply of drivers responding to those nudges, the net effect on worker income may be negative. Such multi-objective entanglement is a classic large-scale systems challenge requiring coordinated optimization that respects Pareto boundaries between platform profit, consumer welfare, and worker wellbeing. Failure to architect these interfaces thoughtfully can produce feedback loops that systematically bias effort under uncertainty in ways that undermine platform legitimacy.

Furthermore, the push toward explainable AI in feedback architectures has introduced tension between model performance and interpretability. Complex ensemble models may more accurately predict which feedback variant will maximize persistence, but their opacity hinders the auditing that labor advocates and regulators increasingly demand [20]. Architectural compromises, such as training a simpler surrogate model to approximate the black-box

predictor's output, are technically feasible but introduce additional layers of approximation error. These design tensions reverberate throughout the governance structure, shaping the accountability relationships that will be examined next.

4. Governance, Agency, and the Algorithmic Management of Effort

The deployment of AI-driven feedback systems within platform-mediated work necessitates a fundamental rethinking of governance, because these systems do not merely inform workers but actively configure the incentive landscape in which they exercise agency. Traditional employment relationships embed feedback within hierarchical reporting structures and managerial discretion that, for all their flaws, afford some procedural recourse through labor law and organizational grievance mechanisms. In the platform context, the algorithm assumes the role of supervisor, performance evaluator, and allocator of future work opportunities, yet it operates through a technical architecture that is largely shielded from democratic oversight.

This algorithmic management creates a governance gap. The feedback system's rules for determining which worker receives what nudge, ranking, or penalty are encoded in software and data pipelines that evolve rapidly through automated model retraining. The speed and opacity of this evolution mean that workers cannot negotiate the terms of their evaluation the way they might bargain over a performance review rubric in a traditional firm. As a result, worker agency becomes constrained to a set of behavioral adaptations within the feedback gamut, a phenomenon that scholars have characterized as the emergence of algorithmic cages that limit genuine autonomy [4]. When income uncertainty already curtails the worker's bargaining power, the addition of non-negotiable algorithmic feedback further tilts the power asymmetry, potentially inducing what can be termed compelled persistence: sustained effort that is less a reflection of volitional commitment than of an inability to decode or contest the system's directives.

Governance solutions must therefore address not only the transparency of the feedback logic but also the procedural mechanisms through which workers can influence its design. Co-governance models, such as platform-worker councils with binding input on algorithmic parameters, represent one institutional innovation being piloted in some jurisdictions. In parallel, the labor protection frameworks being debated in Europe and elsewhere seek to extend collective bargaining rights to algorithmic decision-making processes, a move that would fundamentally alter the governance environment [14]. These legal developments are not merely external constraints on system design; they are constitutive elements of the sociotechnical architecture that can either enable or undermine the feedback-persistence relationship.

The ethical dimension of agency also intersects with the question of informational self-determination. AI feedback systems often rely on granular personal data, including location, communications, and even psychometric inferences, to personalize nudges. Without robust data stewardship frameworks, workers may feel that they are under constant surveillance, a perception that can activate psychological reactance and reduce persistence over the long term. Governance architectures that incorporate data trusts or worker-owned data cooperatives can restore a measure of control by allowing workers to set boundaries on what data feeds into the feedback algorithm, thus rebalancing agency in the human-machine relationship [17]. The trade-off, however, is that restricting data flows may degrade the personalization accuracy that some studies have shown to be critical for motivating persistence among heterogeneous worker populations. Navigating this trade-off requires a deliberative process that integrates

technical feasibility studies with normative stakeholder dialogues, a hallmark of responsible socio-technical systems design.

5. Structural Trade-Offs: Precision, Privacy, and Worker Autonomy

The design space for AI-driven feedback is riddled with structural trade-offs that cannot be optimized away but must be negotiated through intentional architectural choices. The first and most conspicuous trade-off lies between precision and privacy. To deliver feedback that accurately predicts the type, timing, and framing of a message most likely to sustain persistence during an income shortfall, the system requires high-resolution behavioral and contextual data. Yet the accumulation of such data inevitably intrudes upon the worker's informational privacy, creating a surveillance infrastructure that can be repurposed for disciplinary ends beyond performance support. Differential privacy techniques and federated learning architectures offer partial mitigations, but they introduce noise that diminishes the feedback model's ability to capture rare but critical situations, such as a worker on the verge of quitting the platform due to cumulative financial stress. In such edge cases, the loss of precision may translate directly into failed interventions that push a vulnerable worker out of the labor market, raising equity concerns that go beyond aggregate metrics of platform throughput.

A second trade-off involves the tension between adaptive personalization and worker autonomy. Highly adaptive feedback systems use real-time intent prediction to tailor goals and incentives, effectively guiding the worker along a platform-optimized effort trajectory. While this can boost persistence in the short run, it simultaneously reduces the worker's opportunity to exercise metacognitive control, set self-determined work rhythms, and develop the adaptive expertise needed to manage income uncertainty independently. The educational literature on formative feedback consistently warns that overly directive feedback undermines self-regulated learning [11]. Translated to the gig context, this insight suggests that a feedback regime that maximizes immediate persistence might atrophy the worker's capacity to forecast market conditions, budget effectively, and diversify income sources, capacities that are essential for long-term resilience.

A third trade-off concerns the granularity of feedback timing. High-frequency, session-level feedback can counteract the motivational dips caused by income uncertainty by providing micro-rewards for continued effort, exploiting the psychological principle of small wins. Yet such granular feedback can also contribute to cognitive overload and decision fatigue, particularly when income uncertainty already imposes a high cognitive tax. Platform designers must therefore balance the benefits of feedback density against the costs of attentional fragmentation. This trade-off is reminiscent of classic discussions in human-computer interaction regarding interruption management, with the additional gravity that, in the gig economy, income depends on how workers navigate these interruptions.

Finally, there is a macro-level trade-off between system-wide optimization and the heterogeneity of worker preferences. A unified feedback policy trained to maximize aggregate platform persistence may perform well for the median worker but can severely disadvantage those whose motivational profiles diverge from the norm, such as workers using platform income to cover unpredictable medical expenses versus those saving for a discrete purchase. Addressing this requires not just algorithmic personalization but an architectural pluralism that allows workers to select among distinct feedback regimes, a design philosophy that treats feedback style as a configurable feature rather than a deterministic output. Such modularity

imposes additional overhead on system design and evaluation infrastructure but may be essential for robust persistence across diverse income uncertainty contexts.

6. Fairness, Bias, and the Distributional Consequences of Feedback Regimes

The distributional ramifications of AI-driven feedback extend beyond individual motivation to influence who persists on a platform and who is silently excluded. Feedback systems trained on historical data embed existing disparities in ratings, task access, and earnings into their predictive models. If the algorithm learns that workers from certain demographic groups or geographic areas are less likely to respond to a particular nudge, it may systematically reduce the feedback quality offered to those groups, creating a self-fulfilling cycle of disengagement that deepens inequality. The concept of opportunity feedback is especially salient here: the informational signal that a worker receives about how to improve their future platform standing can be thought of as a resource, and when that resource is allocated inequitably, it exacerbates income uncertainty for already marginalized populations.

Fairness in algorithmic management is not a single axis but a constellation of concerns including distributive fairness, procedural fairness, and interactional fairness [18]. A feedback regime may be highly effective at increasing aggregate persistence yet fail procedural fairness if workers cannot understand or contest the rationale behind a negative performance alert. Income uncertainty amplifies the perceived injustice of opaque feedback because the financial stakes of misinterpreting a signal are immediate and tangible. Empirical work on the lived experience of algorithmic labor has shown that feelings of injustice in feedback processes can erode trust to the point where workers withdraw effort or exit the platform entirely, even when objective income opportunities remain attractive [12]. Thus, fairness and persistence are causally entangled rather than independent optimization criteria.

Bias can also creep in through the feedback system's goal-setting interface. If the system suggests earnings goals based on typical behavioral patterns, it may inadvertently anchor disadvantaged workers to lower targets, limiting their income aspirations. Recent evidence on self-set goals indicates that the mere framing of a goal prompt can influence persistence outcomes, suggesting that subtle design choices carry significant distributional weight [13]. Designers must therefore audit feedback systems not only for direct discriminatory effects but also for these indirect anchoring channels. The governance implications are profound, requiring periodic fairness audits that go beyond aggregate statistics to include subgroup analyses on persistence, earnings trajectories, and platform exit rates.

The infrastructure required to conduct such auditing is itself a subject of system design. Platforms must invest in logging, annotation, and causal inference frameworks that can detect disparate impact. Yet this creates a paradoxical incentive: the more transparent the auditing infrastructure, the more exposed the platform becomes to legal liability, potentially discouraging investment in the very mechanisms that could ensure fairness. Regulatory mandates for algorithmic impact assessments, such as those emerging in the European Union's platform work directive debates, aim to resolve this collective action problem by establishing a level playing field where all platforms bear the auditing cost [14]. From a system-level perspective, integrating fairness monitoring directly into the feedback pipeline, as a continuous, automated validation layer, may be more sustainable than ex-post adversarial audits.

7. Infrastructure, Scalability, and System Robustness

The infrastructure underpinning AI-driven performance feedback must operate at scale, handling millions of workers across diverse geographic, cultural, and regulatory contexts while maintaining low latency and high availability. Cloud-native microservice architectures are the current industry standard, with feedback logic decoupled into recommendation services, event processors, and notification gateways. Such decomposition facilitates independent scaling of components during demand spikes, such as holiday rushes when platforms intensify feedback nudges to mobilize idle labor. However, this architectural pattern also introduces coupling risks; if the pricing service and the feedback service operate on slightly stale state due to eventual consistency, a worker may receive a nudge to log on to a high-demand zone that no longer exists, undermining the credibility of the feedback system and, over time, worker persistence.

Scalability challenges are compounded by the need for cross-platform portability and interoperability. As the gig economy matures, workers increasingly multi-home across platforms, assembling a portfolio of income sources to buffer uncertainty. A feedback system designed in isolation for a single platform may inadvertently encourage short-term persistence at the expense of a worker's overall portfolio optimization, pushing the worker to over-allocate effort to the platform with the most aggressive feedback nudges rather than the one offering the best expected income. An interoperable feedback architecture, while technically complex and commercially contentious, could provide holistic, worker-centric guidance that accounts for the full income stream, potentially enhancing long-term persistence across the ecosystem. Such an architecture would require standardized APIs for effort and income data, federated identity management, and secure multi-party computation to preserve proprietary interests while enabling collaborative feedback optimization.

Robustness pertains to the system's ability to maintain functional persistence effects under distributional shifts in the environment, such as an economic recession that drastically alters both the level and volatility of platform demand. Many reinforcement learning policies used in feedback optimization assume stationary environments and can fail catastrophically when the relationship between feedback features and worker responses changes. Antifragile design principles borrowed from resilient systems engineering suggest that feedback systems should incorporate uncertainty quantification and selective slowdown, that is, the ability to detect regime shifts and revert to simpler, more robust non-AI heuristics when confidence intervals exceed safe thresholds. This graceful degradation property is missing in many current deployments, where the feedback engine continues to issue poorly calibrated nudges even as the underlying economy convulses, potentially worsening the financial distress of already vulnerable workers.

The infrastructure must also contend with adversarial dynamics. Workers, as rational agents, may learn to game the feedback system, for instance by accepting and quickly canceling tasks to meet acceptance rate thresholds while maintaining flexibility. Such gaming degrades the signal used by the platform to match supply with demand, leading to a cat-and-mouse evolution that consumes engineering resources and undermines the trust infrastructure essential for persistence. Designing feedback mechanisms that are strategy-proof without resorting to excessive surveillance is a deep and open systems challenge that sits at the intersection of mechanism design, machine learning, and worker psychology.

8. Sustainability and Long-Term Worker Development

A short-sighted optimization of persistence metrics can erode the human capital base upon which the platform economy ultimately depends. AI feedback systems that prioritize

immediate task throughput tend to crowd out opportunities for skill development, reflection, and occupational identity formation. When a platform's feedback architecture treats a worker as a throughput unit to be nudged, it neglects the worker's own developmental trajectory, leading to a hollowing out of expertise that makes the workforce more fungible and less resilient. Income uncertainty exacerbates this dynamic because workers facing financial pressure are less able to forgo immediate earnings to invest in training or experimentation, a classic poverty trap. Here, AI feedback could play a therapeutic role by nudging workers not just to work more, but to work smarter, for instance by recommending rest periods, upskilling modules, or networking opportunities with peers on the platform [7]. Yet such developmental feedback is rarely observed because the platform's revenue model is tied to transaction volume, not to worker career capital, creating a fundamental misalignment that no amount of algorithmic tweaking can resolve.

The sustainability of the platform workforce is also a matter of public health and social welfare. Persistent engagement under income uncertainty, combined with algorithmic pressure, has been linked to physical exhaustion, mental health deterioration, and social isolation. A sustainable feedback system must incorporate wellbeing metrics as first-class constraints, not as afterthoughts to be addressed by corporate social responsibility initiatives. This implies a re-architecture of the feedback objective function to include worker-sustained persistence, that is, the capacity to maintain engagement over months and years without burnout. Multi-stakeholder optimization, where worker representatives jointly define the cost function with platform owners and regulators, presents a path forward but faces enormous coordination and information challenges. Advanced simulation platforms and digital twins could allow stakeholders to explore the long-term consequences of different feedback policies before deployment, facilitating more democratic deliberation about the trade-offs involved.

From an institutional perspective, sustainability requires moving beyond platform-centric feedback systems toward portable worker profiles that aggregate feedback, ratings, and skill badges across multiple gig engagements. Such portability, championed by proposals for worker data trusts [17], would empower workers to carry their performance history with them, reducing lock-in and enabling platforms to compete on the quality of their developmental feedback rather than on their capacity for behavioral manipulation. This vision encounters significant technical obstacles in data standardization, identity verification, and reputation porting, but its realization could fundamentally alter the structural incentives that currently distort feedback design.

9. Policy Implications and Regulatory Architectures

The system-level analysis conducted in this paper carries direct implications for the regulation of AI-driven performance feedback in contingent work. Rather than prescribing rigid rules that mandate specific feedback content, which would quickly become obsolete as technology evolves, a more resilient regulatory approach focuses on process mandates and transparency obligations. Requiring platforms to disclose the high-level objectives, data sources, and evaluation criteria of their feedback algorithms would allow workers, researchers, and watchdogs to scrutinize the system for unfair or exploitative patterns. Such transparency is not a panacea; informational overload can dilute its benefits, and platforms may respond with superficial compliance. Hence, regulatory architectures increasingly combine transparency with mandatory algorithmic impact assessments conducted by independent auditors who have access to the system's internals under confidentiality agreements [15].

Another policy lever is the establishment of algorithmic worker bills of rights that grant workers the ability to contest automated feedback decisions, request human review, and receive explanations in plain language. The growing body of work on explainable artificial intelligence provides the technical substrate for such mandates [20]. Yet the translation of technical explainability into meaningful contestability is not straightforward; an explanation that reveals feature importance weights may be incomprehensible to a gig worker without statistical training. Regulatory frameworks must therefore incentivize the development of worker-centered explanation interfaces that situate algorithmic reasoning within the lived experience of income volatility and task context.

The international dimension of platform work complicates governance because feedback algorithms are typically designed in one jurisdiction and deployed globally. Regulatory fragmentation, where the European Union adopts stringent transparency rules while other regions maintain laissez-faire approaches, risks creating a two-tier workforce differentiated by the quality of algorithmic treatment. Harmonization efforts through international labor organizations or trade agreements could establish minimum standards for AI feedback systems, similar to how occupational safety standards have been globalized. The designation of certain forms of manipulative feedback design, such as dark patterns that exploit cognitive biases to coerce overwork, as unfair commercial practices offers another avenue for regulatory convergence [14].

Finally, the tax and social protection infrastructure must adapt. Income uncertainty is not just a psychological variable but a structure of financial risk that feedback systems can exacerbate. Portable benefits schemes that decouple health insurance, retirement savings, and injury protection from any single platform would reduce workers' precarity and, in turn, alter the way they respond to algorithmic feedback, because the downside risk of experimenting with reduced effort becomes less catastrophic. Such social protection architectures indirectly shape the sociotechnical feedback-persistence relationship, underscoring the need for integrated policy design that treats platform regulation, labor law, and social welfare as a coherent system rather than isolated silos.

10. Conclusion

This paper has advanced a system-level analysis of AI-driven performance feedback and worker persistence within the context of income uncertainty, revealing a complex sociotechnical landscape where architectural choices, governance structures, and fairness considerations are deeply intertwined. We have shown that feedback systems are not neutral conveyors of information; they are active sculptors of worker motivation, reference points, and agency, the effects of which are contingent on the broader institutional framework that buffers or amplifies income volatility. The structural trade-offs among precision, privacy, autonomy, and sustainability cannot be eliminated, but they can be managed through deliberate design pluralism, robust governance, and coordinated regulation.

The emerging empirical evidence, including experimental investigations of goal-setting mechanisms among gig workers [13], suggests that persistence can be meaningfully supported through well-designed feedback, yet a narrow operationalization of persistence as short-term task engagement risks undermining the human capital and wellbeing foundations required for a sustainable platform economy. Future research should adopt longitudinal, multi-platform designs that track how feedback architectures interact with evolving financial circumstances, skill trajectories, and mental health outcomes. Engineers and designers must move beyond model accuracy as the gold-standard metric and embrace fairness, explainability,

and developmental impact as equally legitimate performance criteria. Policymakers must recognize that algorithmic management has made feedback governance a central labor issue, and invest in the inspection infrastructure and international coordination needed to ensure that AI feedback serves as an enabler of worker flourishing rather than a vector of digital precarization. The path forward lies in a collaborative, interdisciplinary reimagination of feedback not as a tool of control but as a shared information commons that empowers workers navigating the uncertainties of the platform age.

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. Abraham, K. G., Haltiwanger, J. C., Sandusky, K., & Spletzer, J. R. (2019). The rise of the gig economy: fact or fiction? *Journal of Economic Perspectives*, 33(3), 3-28.
3. Deci, E. L., & Ryan, R. M. (2000). The "what" and "why" of goal pursuits: Human needs and the self-determination of behavior. *Psychological Inquiry*, 11(4), 227-268.
4. 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.
5. Fang, A., Kulkarni, N., & Hsieh, G. (2022). "I just want to feel safe": Understanding the experience of income insecurity among on-demand workers. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW1), Article 79, 1-27.
6. Tong, S., Jia, N., Luo, X., & Fang, Z. (2021). The Janus face of artificial intelligence feedback: Performance-enhancing feedback that undermines learning. *MIS Quarterly*, 45(3), 1453-1488.
7. Bucher, E., Fieseler, C., & Lutz, C. (2020). Mattering in digital labor. *New Media & Society*, 22(10-11), 1912-1931.
8. Abeler, J., Falk, A., Goette, L., & Huffman, D. (2011). Reference points and effort provision. *American Economic Review*, 101(2), 470-492.
9. Diakopoulos, N. (2016). Accountability in algorithmic decision making. *Communications of the ACM*, 59(2), 56-62.
10. 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.
11. Shute, V. J. (2008). Focus on formative feedback. *Review of Educational Research*, 78(1), 153-189.
12. Rosenblat, A., & Stark, L. (2016). Algorithmic labor and information asymmetries: A case study of Uber's drivers. *International Journal of Communication*, 10, 3758-3784.
13. 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.
14. 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.

15. Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2), 2053951716679679.
16. Schor, J. B., & Attwood-Charles, W. (2017). The "sharing" economy: labor, inequality, and sociability on for-profit platforms. *Sociology Compass*, 11(8), e12493.
17. Irani, L. (2015). The cultural work of microwork. *New Media & Society*, 17(5), 720-739.
18. Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. *California Law Review*, 104(3), 671-732.
19. Hawlitschek, F., Teubner, T., & Weinhardt, C. (2016). Trust in the sharing economy. *Die Unternehmung*, 70(1), 26-44.
20. Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., & Pedreschi, D. (2018). A survey of methods for explaining black box models. *ACM Computing Surveys*, 51(5), Article 93, 1-42.
21. Kalleberg, A. L. (2009). Precarious work, insecure workers: Employment relations in transition. *American Sociological Review*, 74(1), 1-22.