UtilitasMathematica

ISSN 0315-3681 Volume 122 (2), 2025

Responsible AI in Real-Money Gaming:

Embedding Regulatory Constraints in Personalization Algorithm

Vatsal Modi[1,*] and Abhijit Chanda[2]

[1]*Carnegie Mellon University, Pittsburgh, PA*

[*]*Corresponding Author Email: vatsalm@alumni.cmu.edu - ORCID ID: 0000-0002-6221-6486*

[2]*Senior Manager, Tredence Inc.. Fremont, CA*

Email: abhijitju252024@gmail.com - ORCID ID: 0009-0002-6976-3922

August 13, 2025

Abstract

The use of AI-powered personalization in real-money gaming has raised significant ethical and regulatory concerns in recent years. These systems are usually designed to maximize user engagement with the platform, which leads to higher monetization, but this sometimes comes at the cost of opaque and potentially discriminatory targeting systems, which risk exploiting user behavior, particularly among the vulnerable population. This study aims to address these issues by developing a regulatory-aligned framework that integrates ethical constraints such as fairness, transparency, and harm mitigation directly into the design of personalization algorithms. Drawing on secondary qualitative sources such as regulatory frameworks, peer-reviewed literature, and documented platform practices, this study analyzes current standards and gaps in industry practice. These insights inform the creation of the Responsible Personalization Framework (RPF), which operationalizes regulatory principles through concrete design patterns. The findings reveal that widely adopted personalization strategies often contradict responsible AI principles. However, the proposed framework, featuring mechanisms like Harm-Aware Personalization and the Offer Parity Rule, offers actionable, regulation-compliant solutions for ethical personalization. Overall, the study demonstrates that embedding ethical and regulatory safeguards at the design stage, instead of post hoc reactive enforcement, enables the development of safer, fairer, and more responsible AI systems in the context of real-money gaming.

Keywords: Responsible AI, Algorithmic Personalization, Real-Money Gaming, Regulatory Compliance, Fairness and Transparency, Ethical Monetization

1 Introduction

The growing influence of artificial intelligence (AI) across high-stakes sectors such as healthcare, finance, and online gaming has made Responsible AI more urgent than ever. Within the realm of real-money gaming, AI-powered personalization systems are increasingly used to tailor user experiences

and maximize revenue. While these systems can enhance gameplay and platform profitability, they also introduce serious ethical and regulatory risks. By closely analyzing player behavior and transaction history, personalization algorithms can lead to exploitative outcomes, particularly among vulnerable users. As regulatory bodies around the world begin to enforce legal frameworks to govern AI development and use, the pressing challenge lies in operationalizing these principles within real-world systems.

Despite the availability of Responsible AI guidelines, many current personalization algorithms still lack essential guardrails. Their decision-making processes remain largely opaque, and in practice, they often reinforce addictive patterns or unfair treatment. Furthermore, even when companies implement ethical or legal constraints, these are typically applied post-deployment, after harm may already have occurred. This reactive approach not only diminishes user trust but also falls short of regulatory expectations.

What's notably missing in the current research landscape is a method for embedding regulatory and ethical constraints directly into the design of personalization algorithms themselves. While post hoc audits and policy frameworks have been explored, very few studies have tackled algorithmic redesign as a way to proactively prevent harm, especially in the dynamic, profit-driven context of real-money gaming.

To address this gap, this study explores how personalization algorithms in real-money gaming can be reimagined to integrate fairness, transparency, and harm prevention from the ground up. The aim is to propose a design framework that helps developers align monetization strategies with legal compliance and player well-being, transforming Responsible AI from a compliance checklist into a core design principle. This paper is intended for researchers, product designers, and policy architects in gaming, AI, and regulatory technology sectors, with the goal of advancing responsible innovation practices in algorithmic personalization.

2 Literature Review

2.1 The Ethics of Personalization and Monetization in Real-Money Gaming

Personalization has become central to real-money gaming platforms, shaping how users encounter content and make decisions. At its best, personalization can enhance player experience, surface relevant offers, and improve platform usability. Yet, when these systems are tightly coupled with monetization strategies, concerns around fairness, autonomy, and harm become difficult to ignore.

The architecture of recommender and personalization algorithms is typically designed to optimize engagement, often measured in playtime or spending. Cobbe and Singh, 2019 argue that recommender systems can produce harm not only through the content they deliver but through their underlying design, which focuses on improving engagement metrics over user well-being. These algorithms shape player behavior in ways that are not easily observable or contestable, leading to a subtle erosion of user agency over time.

In gaming environments where monetization is the ultimate goal, personalization systems serve as more than just experience enhancers; they function as behavioral steering mechanisms. Packin et al., 2024 examine how gaming and financial platforms use nudges, feedback loops, and gamified incentives to encourage repeated actions such as spending or chasing rewards. Techniques like randomized loot boxes, timed promotions, and dynamic pricing leverage well-documented cognitive biases to influence user behavior, often under the guise of customization. This raises particular concerns in contexts where players may be less equipped to recognize or resist these tactics, such as among younger users or those with lower financial literacy.

Gao et al., 2023 highlight the ethical complexity of this dual-purpose design. While personalization can indeed improve satisfaction and targeting efficiency, its benefits are contingent on whether systems are built with appropriate guardrails. Their analysis stresses the importance of embedding privacy, bias detection, and transparency mechanisms directly into the system architecture, rather than retrofitting ethical considerations after deployment. Without such safeguards, personalization risks functioning as a tool for fine-grained value extraction rather than overall user benefit.

Dastakeer, 2024 further illustrates how personalization can be used to manipulate spending behavior, particularly through techniques such as adaptive pricing and engagement-based incentives. While such strategies may be commercially successful, they introduce asymmetries in how users are treated based on behavioral segmentation and potentially offer fewer protections or more aggressive offers to high-value or susceptible users. In this way, personalization shifts from enhancing relevance to actively shaping outcomes that serve the platforms revenue goals, often in ways that are neither visible nor understandable to players.

Across these perspectives, a clear pattern emerges: personalization systems in real-money gaming are not neutral intermediaries. They play an active role in determining how players engage with the platform, what choices they perceive as available, and how much they ultimately spend. This makes ethical design not just a regulatory issue, but a structural one and requires that fairness, accountability, and harm mitigation be addressed at the point of system development rather than imposed externally.

These concerns cannot be addressed through post hoc oversight alone. The integration of personalization and monetization calls for a shift in design priorities toward systems that actively protect player autonomy, promote transparency in monetization, and ensure that behavioral insights are used responsibly.

2.2 Responsible Personalization Through Constrained Decision Models

AI systems often function under practical constraints like fairness considerations, budgetary restrictions and legal requirements to name a few. Even though these constraints are important, they are typically applied retroactively instead of being embedded into the initial model design.

There are several frameworks developed that incorporate real-world constraints into algorithmic decision-making. For instance, Ivanov and Ben-Porat, 2024 propose represented Markov Decision

Processes (r-MDPs), a formulation that integrates budgetary and policy constraints into personalization strategy. In this setup, users are matched to a small set of representative policies, and the system learns to optimize both the user-to-policy assignments and the policies themselves within the given constraints. This framework provides a useful abstraction for constrained personalization by trying to balance efficiency and societal benefit.

Zhou and Ji, 2022 take a softer constraint approach through kernelized multi-armed bandits. Their framework is built to minimize cumulative harm by allowing for short-term constraint violations. This work focuses more on the value of flexibility and draws our attention to tradeoffs in sequential decision-making.

In contrast, Ding et al., 2013 present a framework that is grounded in strict adherence to budget constraints and deterministic feasibility. While this formulation serves as a clear theoritical baseline, it tends to oversimplify the complexity and unpredictability found in real-world settings.

A common shortfall across these models is the failure to integrate ethical imperatives, such as fairness or interpretability, directly into the algorithmic core. This research addresses this limitation by focusing on embedding fairness-sensitive utility functions and compliance auditing mechanisms within the optimization framework that drives the personalization logic.

2.3 Responsible AI Frameworks and Regulatory Readiness

A growing number of approaches now aim to bridge the gap between ethical ideals and the realities of regulatory compliance in AI systems. As regulatory frameworks, like the (Regulation (EU) 2024/1689, AI Act, 2024), solidify, the burden of moving past high-level ethical systems and 1building systems that meet both legal standards and societal expectations will fall on the developers of such systems.

To lay the conceptual foundation, there are seven key requirements of trustworthy AI, as outlined by Díaz-Rodríguez et al., 2023. These include human agency and oversight, technical robustness and safety, privacy and data governance, transparency, diversity and non-discrimination, societal and environmental well-being, and accountability. They argue that responsible AI must guide decision-making in system development across the entire lifecycle with these ethical pillars.

The current ethical toolkits are too abstract or disconnected from real-world policy mandates as noted by Constantinides et al., 2024. In response, they propose a method for generating regulation-aligned, role-agnostic guidelines. When these guidelines are embedded into development workflows they promote early-stage ethical reflection and cross-functional accountability.

Finally, from a legal standpoint, Comandè, 2017 warns against procedural compliance models that check boxes without aligning system behavior to regulatory intent. They advocate for proactive governance embedded within algorithmic design. This approach leads to preemptive, rather than reactive, ethical intervention.

These works agree on the urgent need for actionable, regulation-ready frameworks. While Díaz-Rodríguez et al., 2023 emphasize ethical requirements, Comandè, 2017 stresses legal alignment, and

Constantinides et al., 2024 bridge the two in usable tools. My research advances this convergence by integrating regulatory and ethical safeguards directly into the personalization engine, allowing for real-time audits, transparency thresholds, and player protection policies that operate within system logic.

2.4 Synthesis

Although existing research spans personalization, optimization, ethical frameworks, and monetization, several challenges remain unaddressed:

- Technical models rarely embed ethical parameters like fairness or explainability.
- Ethics-oriented frameworks often lack operational specificity, especially for adaptive systems like gaming.
- Most interventions occur post-deployment, after potential harm.
- Regulatory principles are not yet fully integrated into personalization logic.

This study responds by proposing a Responsible Personalization Framework (RPF) tailored to real-money gaming. The framework embeds fairness audits, transparency thresholds, and player feedback loops into the core of algorithmic decision-making.

3 Methodology

This study employs a qualitative, multi-source approach to develop and evaluate a Responsible Personalization Framework (RPF) for real-money gaming platforms. The objective is to translate emerging legal and ethical requirements, such as fairness, transparency, and harm mitigation, into actionable design constraints for AI-driven personalization systems.

3.1 Research Approach

The methodology combines building a conceptual framework with comparing existing approaches and synthesizing findings to evaluate their implications. Three primary research activities inform this process:

1. **Review of regulatory standards** (e.g. Regulation (EU) 2024/1689, AI Act, 2024, Regulation (EU) 2016/679, General Data Protection Regulation (GDPR), 2016, International Organization for Standardization, 2023, and International Organization for Standardization, 2020) to extract enforceable requirements related to AI risk, transparency, and accountability.

2. **Analysis of platform practices** using documented case studies, published audits, and behavioral datasets from the real-money gaming sector (e.g. Zendle et al., 2019 and Macey and Hamari, 2024) to understand personalization strategies and monetization mechanics in use.

3. **Technical review** of recommender system literature to assess how personalization technologies align, or fail to align, with ethical and regulatory standards.

The overall research process is summarized in Figure 1, which illustrates how regulatory, technical, and empirical inputs were synthesized to inform the design and evaluation of the Responsible Personalization Framework.

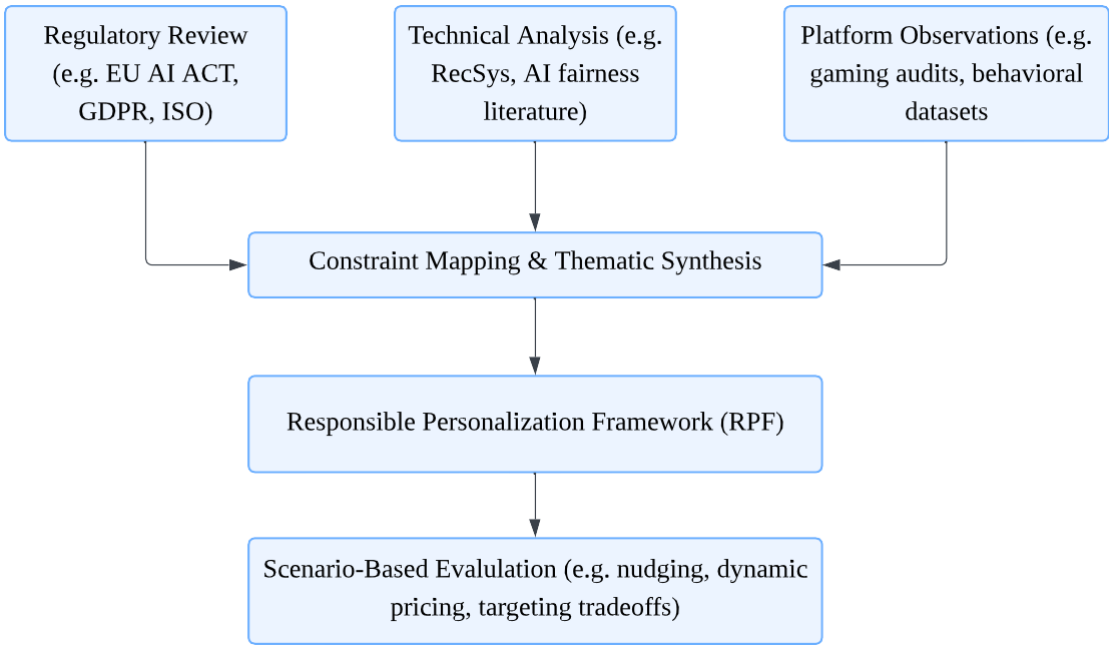

Figure 1: Research design and framework development process

3.2 Data Collection and Analysis

A systematic document review was conducted to extract regulatory principles and map them to features of algorithmic personalization, including targeting logic, ranking functions, and feedback loops. These mappings informed the identification of risk flags and constraint categories relevant to gaming platforms.

To analyze platform behavior, digital ethnography and content analysis were applied to public documentation, empirical studies, and audit reports. An iterative coding process was used to identify recurring design patterns that could pose ethical or regulatory concerns.

Thematic synthesis was employed to consolidate insights across legal, technical, and empirical domains. Where available, open-source algorithmic frameworks were examined to illustrate how specific configurations may or may not conform to expected norms of fairness, user agency, and transparency.

3.3 Framework Development and Evaluation

Insights from the regulatory and platform analysis were translated into a set of constraint-based design principles, forming the basis of the Responsible Personalization Framework (RPF). To evaluate the framework, a series of scenario-based simulations were constructed. These hypothetical cases model common monetization strategies, such as dynamic pricing, behavioral nudging, and reward targeting, and assess their compatibility with the ethical standards defined in the RPF.

Each scenario was evaluated based on its alignment with legal, ethical, and operational expectations for high-risk AI systems. The analysis emphasizes not only theoretical compliance but also practical applicability within platform ecosystems.

4 Results

The Responsible Personalization Framework (RPF) was developed to operationalize fairness, transparency, and harm mitigation principles within AI-driven personalization systems for real-money gaming. It integrates regulatory obligations, such as those outlined in the Regulation (EU) 2024/1689, AI Act, 2024 and ISO standards (International Organization for Standardization, 2020, International Organization for Standardization, 2023), directly into the design and delivery of personalized experiences.

Figure 2 presents a high-level architectural model of the RPF. It illustrates how ethical and regulatory constraints are embedded into the personalization pipeline via a dedicated constraint engine, which sits between preprocessing and algorithmic decision-making. The architecture also incorporates feedback loops through an audit module and user interface layer to support real-time accountability and compliance monitoring.

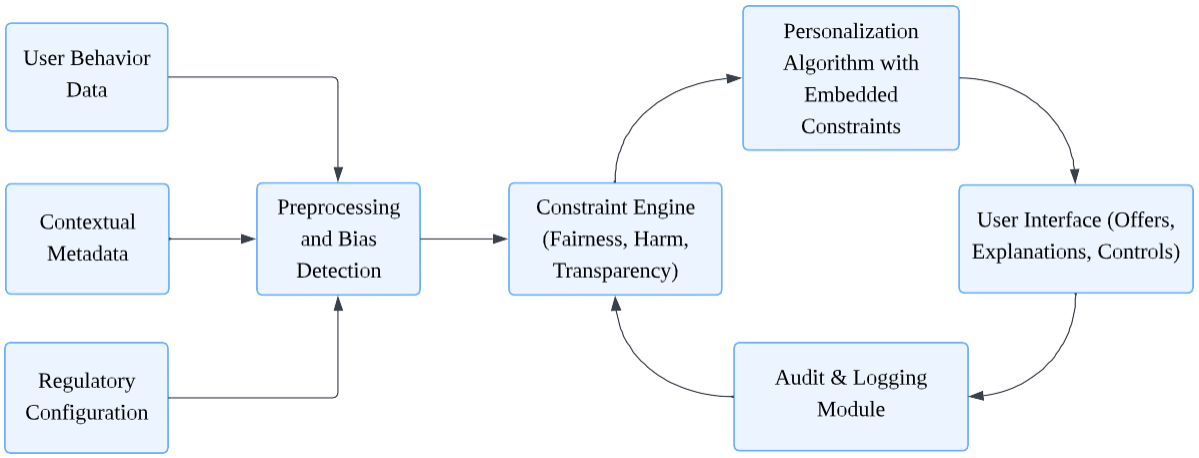

Figure 2: Responsible Personalization Framework Architecture

The following sections break down the results of applying this framework: identifying problematic platform practices, translating regulatory mandates into system-level constraints, and evaluating personalization strategies against responsible AI principles.

4.1 Key Regulatory Constraints for Personalization in Real-Money Gaming

A detailed analysis of leading regulatory sources, namely the Regulation (EU) 2024/1689, AI Act, 2024, International Organization for Standardization, 2020, and Regulation (EU) 2016/679, General Data Protection Regulation (GDPR), 2016 revealed three central obligations directly applicable to AI-based personalization systems in the real-money gaming context:

- **Transparency and Explainability:** Systems must provide clear disclosure regarding the purpose, mechanisms, and potential outcomes of algorithmic decisions (Regulation (EU) 2024/1689, AI Act, 2024, Articles 13 and 69).
- **Harm and Risk Mitigation:** AI systems must incorporate safeguards to prevent financial exploitation, psychological harm, or manipulation, particularly for systems that drive behavior through personalization or gamification (Regulation (EU) 2024/1689, AI Act, 2024 Articles 9 and 10).
- **User Autonomy and Oversight:** Personalization systems are expected to preserve human agency, offering opt-outs and meaningful user control while avoiding covert behavioral nudging or manipulative design (Regulation (EU) 2024/1689, AI Act, 2024 Articles 5 and 12).

Table 1: Common Personalization Strategies and Their Ethical Violations

Strategy	Violation Type	Ethical Concern
Loot box outcome targeting	Lack of transparency	Simulated gambling, unexplained odds
Dynamic pricing per user group	Discriminatory personalization	Exploits behavioral segmentation
Time-based reward throttling	Opaque nudging	Risk of addiction, loss of user control
High-risk user reactivation AI	Exploitative targeting	Violates harm mitigation and fairness norms

4.2 Misalignment with Responsible AI Standards

Content analysis of industry audits and academic studies (e.g., Zendle et al., 2019, Macey and Hamari, 2024) uncovered several personalization strategies widely used in real-money gaming that do not align with emerging regulatory and ethical standards. Table 1 summarizes key tactics, their associated violations, and ethical concerns.

4.3 Codified Design Patterns for Regulatory-Aligned Personalization

To address these ethical risks, regulatory mandates were translated into design-level constraints that could be embedded within real-time personalization systems. These constraints were formalized as design patterns, shown in Table 2.

Table 2: Design Patterns for Regulatory-Aligned Personalization

Pattern Name	Regulatory Principle	Actionable Constraint
Explain Before Act	Transparency (Art. 13)†	Disclose logic behind targeted offers and personalization triggers
Throttle by Consent	User Autonomy (Art. 5)†	Allow user-set pacing limits on reward or engagement nudges
Harm-Aware Personalization	Risk Mitigation (Art. 9, 10)†	Suspend personalization for users exhibiting harmful patterns
Offer Parity Rule	Fairness & Non-discrimination	Ensure equal offers for users with similar profiles and behavior

† Regulatory principles referenced are from the Regulation (EU) 2024/1689, AI Act, 2024

4.4 Gaps in Industry Practice and Self-Regulation

Publicly available reports and regulatory summaries reveal significant shortcomings in how gaming companies approach responsible personalization:

- **Inconsistency:** Ethical evaluations, where they exist, are sporadic and lack standardized enforcement mechanisms.
- **Reactivity:** Platforms typically update systems in response to external scrutiny or media attention, rather than through proactive audits or risk forecasting.
- **Opacity:** Players are rarely informed about how personalization works or how their data is used to influence in-game offers, undermining informed consent and user control.

These gaps highlight the need for embedded ethical mechanisms rather than post hoc compliance.

4.5 Illustrative Examples of Compliance Across Strategies

A comparative classification of personalization strategies based on transparency, fairness, autonomy, risk mitigation, and regulatory compliance is presented in Table 3.

These evaluations reinforce the frameworks utility as both a compliance and design tool, enabling system designers to quickly assess the ethical and regulatory acceptability of personalization strategies in real-money gaming contexts.

Table 3: Compliance Assessment of Personalization Tactics

Tactic	**Transparency**	**Fairness**	**Autonomy**	**Risk Mitigation**	**Compliance Status**
Personalized loot odds	C	C	C	C	C Unacceptable
Time-based bonus pacing	C	✓	C	C	¼High-Risk
Avatar-based cosmetic targeting	✓	✓	✓	✓	✓ Compliant

5 Discussion

5.1 Personalization Conflicts with Responsible AI Principles

This study demonstrates that personalization systems currently deployed in real-money gaming environments are fundamentally misaligned with key principles of responsible AI. In particular, widespread platform practices conflict with standards related to fairness, transparency, user autonomy, and harm mitigation as outlined in the Regulation (EU) 2024/1689, AI Act, 2024 and supported by ISO frameworks (International Organization for Standardization, 2020, International Organization for Standardization, 2023). The lack of embedded ethical safeguards at the design level often enables manipulative personalization tactics, increasing the risk of regulatory non-compliance and erosion of user trust.

5.2 Reframing Personalization as an Ethical Design Challenge

These findings reinforce the idea that personalization, particularly in high-risk domains such as real-money gaming, should not be treated purely as a technical optimization problem. Rather, it represents a deeply ethical and regulatory design challenge. Algorithms that adjust dynamically to user vulnerabilities or behavioral triggers, without offering transparency or control, risk breaching both ethical norms and legal boundaries. This research underscores the necessity of treating responsible AI constraints as core functional requirements within system architecture, rather than as optional compliance layers added post hoc.

5.3 Alignment with Prior Research and Novel Contributions

The conclusions of this study are consistent with existing literature that highlights behavioral and regulatory concerns surrounding algorithmic personalization. Zendle et al., 2019 and Macey and Hamari, 2024 exposed the psychological and behavioral risks associated with loot boxes and nudging mechanisms. Dastakeer, 2024 further observed rising user distrust in AI-driven monetization

systems. Meanwhile, Constantinides et al., 2024 advocated for aligning AI system design with regulation to foster responsible development practices.

Building on these prior works, this research contributes an original set of actionable, design-level controls, such as *Explain Before Act* and the *Offer Parity Rule*, that directly map to specific articles within the Regulation (EU) 2024/1689, AI Act, 2024. Unlike most theoretical proposals, these mechanisms are designed for integration into system development workflows, thereby offering operational pathways to meet ethical and regulatory standards.

5.4 Implications for Responsible AI Implementation

The integration of regulatory principles into personalization frameworks offers several critical benefits. First, it has the potential to rebuild user trust by promoting transparency and fairness. Second, it reduces legal exposure by aligning system behavior with binding obligations under emerging AI laws. Third, it signals a shift in industry perspective, from retroactive compliance strategies toward proactive ethics-by-design approaches.

These findings also reflect a broader evolution in the field of AI development. As personalization becomes increasingly embedded in high-stakes decision-making systems, it is essential to evaluate not just its efficiency or profitability, but also its performance against ethical indicators such as fairness, autonomy, and harm prevention.

5.5 Limitations

While this study offers conceptually grounded and regulation-aligned framework for ethical-personalization, there are a few limitation that point to future direction for validation and refinement. The Responsible Personalization Framework proposed here has not yet been tested in live production environments, and thus, its real-world feasibility and commercial impact remain unverified. Moreover, the business tradeoffs, such as potential reductions in revenue from constrained personalization, have not been analyzed.

Additionally, most of the regulatory references originate from the European context. Given the jurisdictional specificity of laws like the Regulation (EU) 2024/1689, AI Act, 2024 and Regulation (EU) 2016/679, General Data Protection Regulation (GDPR), 2016, the generalizability of this framework to global markets with differing legal environments remains uncertain. Future iterations of this framework should consider adaptations for broader regulatory compatibility.

5.6 Future Directions

To build on the findings of this research, several directions for future work are proposed:

- **Simulations:** Run controlled simulations of personalization frameworks with and without em-

bedded ethical constraints using open-source AI infrastructure.

- **Industry Pilots:** Partner with gaming platforms to implement and evaluate the Responsible Personalization Framework in sandbox or production settings.
- **Ethical KPIs:** Develop measurable performance indicators for ethical AI dimensions such as transparency, user control, and harm avoidance.
- **Cross-domain Extension:** Expand the scope of the framework to cover other high-risk personalization domains such as finance, insurance, and digital health.

By pursuing these directions, future research can continue to operationalize responsible AI principles and ensure they are embedded into the systems that increasingly shape human behavior and well-being.

6 Conclusion

This research has shown that current personalization frameworks used in real-money gaming frequently diverge from key principles of Responsible AI, specifically those related to fairness, transparency, and harm mitigation. Through a comprehensive analysis of regulatory texts, technical standards, and real-world personalization practices, the study has introduced a Responsible Personalization Framework. This framework translates abstract legal mandates into actionable algorithmic patterns such as *Explain Before Act*, *Throttle by Consent*, *Harm-Aware Personalization*, and the *Offer Parity Rule*. These patterns offer a clear pathway for embedding ethical and legal safeguards directly into system design.

The importance of this work lies in the unique positioning of real-money gaming at the intersection of AI, behavioral design, and human vulnerability. In the absence of built-in protections, personalization systems risk manipulating users, particularly those who may be predisposed to addictive behaviors or financial instability. Such outcomes not only raise serious ethical concerns but also risk breaching emerging regulatory frameworks like the Regulation (EU) 2024/1689, AI Act, 2024. The proposed framework addresses these challenges by enabling a proactive, design-level strategy for compliance and user protection, helping the industry move beyond reactive governance.

More broadly, this framework demonstrates the feasibility and necessity of building ethical and legal constraints into the technical architecture of AI systems. Although this study focuses on gaming, the framework's underlying logic can be extended to other high-risk domains such as finance, healthcare, and education, sectors where personalization is equally consequential. It also lays the foundation for a new class of "compliance-aware AI architectures" that embed regulatory alignment and social responsibility at their core.

For designers, product leaders, and compliance teams, this research offers a tangible path forward: a way to build personalization engines that not only optimize engagement but also earn trust, resist

misuse, and anticipate future regulation. As personalization becomes a default interface for digital life, embedding responsibility at the level of system logic is no longer optional - it's the new baseline.

7 References

References

Cobbe, J., & Singh, J. (2019). Regulating recommending: Motivations, considerations, and principles [SSRN Scholarly Paper 3371830, Social Science Research Network].

Comandè, G. (2017). Regulating algorithms regulation? first ethico-legal principles, problems, and opportunities of algorithms. In T. Cerquitelli, D. Quercia, & F. Pasquale (Eds.), *Transparent data mining for big and small data* (pp. 169–206). Springer International Publishing. https://doi.org/10.1007/978-3-319-54024-5_8

Constantinides, M., Bogucka, E., Quercia, D., Kallio, S., & Tahaei, M. (2024). Rai guidelines: Method for generating responsible ai guidelines grounded in regulations and usable by (non-)technical roles. *Proceedings of the ACM on Human-Computer Interaction*, *8*(CSCW2), 1–28. https://doi.org/10.1145/3686927

Dastakeer, I. (2024). Innovating ethical monetization for profitable and engaging games.

Díaz-Rodríguez, N., Del Ser, J., Coeckelbergh, M., López De Prado, M., Herrera-Viedma, E., & Herrera, F. (2023). Connecting the dots in trustworthy artificial intelligence: From ai principles, ethics, and key requirements to responsible ai systems and regulation. *Information Fusion*, *99*, 101896. https://doi.org/10.1016/j.inffus.2023.101896

Ding, W., Qin, T., Zhang, X.-D., & Liu, T.-Y. (2013). Multi-armed bandit with budget constraint and variable costs. *Proceedings of the AAAI Conference on Artificial Intelligence*, *27*(1). https://doi.org/10.1609/aaai.v27i1.8637

Gao, B., Wang, Y., Xie, H., Hu, Y., & Hu, Y. (2023). Artificial intelligence in advertising: Advancements, challenges, and ethical considerations in targeting, personalization, content creation, and ad optimization. *SAGE Open*, *13*(4), 21582440231210759. https://doi.org/10.1177/21582440231210759

Towards a Trustworthy AI, 2020.

ISO/IEC 23894:2023 — Information technology Artificial intelligence Guidance on risk management, 2023.

Ivanov, D., & Ben-Porat, O. (2024). Personalized reinforcement learning with a budget of policies. *Proceedings of the AAAI Conference on Artificial Intelligence*, *38*(11). https://doi.org/10.1609/aaai.v38i11.29169

Macey, J., & Hamari, J. (2024). Gamblification: A definition. *New Media & Society*, *26*(4), 2046–2065. https://doi.org/10.1177/14614448221083903

Packin, N. G., Kliger, D., Reichman, A., & Rabinovitz, S. (2024). Hooked and hustled: The predatory allure of gamblified finance [SSRN Scholarly Paper 5090194].

Regulation (EU) 2016/679, General Data Protection Regulation (GDPR), 2016.

Regulation (EU) 2024/1689, AI Act, 2024.

Zendle, D., Meyer, R., & Over, H. (2019). Adolescents and loot boxes: Links with problem gambling and motivations for purchase. *Royal Society Open Science*, *6*(6), 190049. https://doi.org/10.1098/rsos.190049

Zhou, X., & Ji, B. (2022). On kernelized multi-armed bandits with constraints. *Advances in Neural Information Processing Systems*, *35*, 14–26.