Japan Journal of Research

Original Article

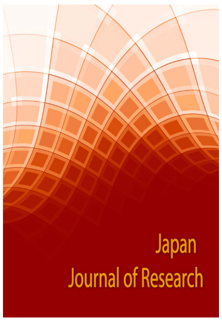

The Loo Loop Ethical AI Agent: A Queuing-Theoretical Model of Mixed-Gender Toilet Allocation Under Equity Constraints

Pascal Stiefenhofer

Newcastle University, UK

Correspondence

Pascal Stiefenhofer

Newcastle University, UK

E-mail: pascal.stiefenhofer@newcastle.ac.uk

- Received Date: 11 April 2025
- Accepted Date: 28 April 2025
- Publication Date: 12 May2025

Keywords

EDI, AI agent, Markov Decision Process, restroom queue management, behavioral dynamics, discomfort, urgency, fairness, public spaces.

Abstract

We develop a queuing-theoretical model for mixed-gender public toilet systems, incorporating contemporary Equity, Diversity, and Inclusion (EDI) principles. The model uses a multi-server Markovian framework with gender-mixed arrival streams, incorporating user discomfort and strategic behavior such as balking and reneging. We introduce the Flush Equity Index (FEI), a novel fairness metric designed to evaluate equity in access across gender and identity groups. Our analysis reveals paradoxes inherent in socially progressive restroom designs: while increased equality of access aims to reduce inequities, it may unintentionally decrease system efficiency and user satisfaction. To address this, we propose several policy interventions, including the Privacy-Aware Stall Assignment Protocol (PASAP), behavioral nudges, and a blockchain-based ToiletCoin system. The ToiletCoin system optimizes queue allocation by rewarding users for cooperative behavior, such as nudging and prioritizing high-urgency users, while reducing discomfort. Finally, we recommend using the FEI in restroom equity audits to guide design and policy decisions.

Introduction

In a world increasingly concerned with inclusivity, fairness, and digital integration, even the humblest of public facilities is undergoing a dramatic transformation. One of the most pressing issues? The evolution of mixed-gender toilets. This movement, propelled by progressive understandings of gender, social justice, and architectural responsibility, brings to the forefront a critical question: How do we ensure that our most basic, universally accessible spaces are truly inclusive? As we strive to accommodate diverse identities, a new paradox emerges, the more we try to balance equity, the more we find ourselves tangled in inefficiencies and uncomfortable waiting lines.

In a similar vein, the world of Ethical Artificial Intelligence (AI)—with its focus on fairness, transparency, and user-centered design has been lauded for its commitment to creating systems that treat all users equally and without bias. But while AI systems are under close scrutiny for algorithmic fairness [5,6,8,10] who is examining the queuing fairness of our public restrooms? Enter the world of gender-neutral toilets, a microcosm of the complex interplay between fairness, efficiency, and user discomfort. How long do we wait? Do we feel safe? Is it awkward that I'm standing here while they're standing there? These questions are not just theoretical they are real-life dilemmas that plague every public restroom.

Drawing inspiration from fairness metrics in AI, such as those embedded in bias-mitigation toolkits [4,20] and federated learning systems [12], we embark on a highly academic and very rigorous exploration of the optimal queuing and stall assignment in mixed-gender restrooms. Why? Because someone has to solve it. We propose a novel queuing-theoretical model that applies the same equity principles to public restrooms as we apply to AI. This paper introduces the Flush Equity Index (FEI), a novel fairness metric that measures the equity of toilet stall access across gender and identity groups.

Just as Ethical AI scrutinizes biases and strives for explainability [24,11,19], we aim to quantify the discomfort (and potential discomfort) of people waiting for a stall in a gender-mixed queue. After all, if we're using high-level mathematics to optimize wait times in tech companies, surely we can apply some of that brainpower to making the waiting experience in a restroom a little more equitable.

Inspired by marketing and economic theory [22,28,25] this paper explores the inefficiencies and social dynamics within public restroom queues. Is it more important to maximize user satisfaction, or is the goal to ensure that no one group has to wait longer than another—even if it means making the system slower overall?

As AI continues to evolve, with an ever-increasing focus on human-in-the-loop systems

Citation: Stiefenhofer P. The Loo Loop Ethical AI Agent: A Queuing-Theoretical Model of Mixed-Gender Toilet Allocation under Equity Constraints. Japan J Res. 2025;6(6):119.

[23,29,26], we draw an unexpected parallel between toilet design and algorithmic fairness. If AI systems can be adjusted in real-time based on user behavior, why can't public restrooms adopt the same adaptability? With that in mind, we introduce the FEI, which evaluates restroom designs based on user discomfort, queue fairness, and emergency needs.

This paper offers a sincere and slightly overdue exploration of how foundational principles in Ethical AI especially those concerning transparency, fairness, and equity can be applied to a domain often overlooked: the public restroom. By modeling discomfort, urgency, and social dynamics through a queuing-theoretical lens, we propose actionable interventions that aim to make restrooms more inclusive, efficient, and equitable. And if nothing else, our contribution ensures that the next time you're waiting in line for a stall, you'll know that your discomfort has not only been acknowledged but rigorously and mathematically optimized. The paper begins by introducing the theoretical framework, establishing a foundation for understanding restroom queue management through the lens of queuing theory. It then delves into the behavioral dynamics of users in the queue, highlighting psychological discomfort and the role of strategic decision-making. Building on this, the development of an AI agent for EDI-centric restroom management is presented, which leverages a Markov Decision Process (MDP) framework to optimize efficiency, fairness, and user comfort. The discussion continues with key empirical findings drawn from simulations involving the AI agent, offering valuable insights and policy recommendations to enhance restroom management practices. The paper concludes with a summary of the results and outlines directions for future research.

System Model

We model a shared gender-neutral toilet facility as a multi-server queuing system with socially sensitive behavioral inputs. The system consists of $c \in \mathbb{N}$ identical toilet stalls, each acting as a server. Users are categorized into three identity groups: men (m), women (f), and nonbinary individuals (n), who arrive at the facility according to independent Poisson processes. Users are served in a first-come, first-served (FCFS) manner, subject to strategic balking behaviors. The system is described as an M/M/c queue, where c represents the number of available toilet stalls (servers). The arrival rates for each group are λ_m, λ_f, and λ_n, respectively, such that the total arrival rate is $\lambda = \lambda_m + \lambda_f + \lambda_n$. Service times are exponentially distributed with parameter μ, meaning the mean service time is μ^{-1}. The service discipline follows FCFS, with no preemption.

Each user experiences discomfort based on factors such as the queue length and the composition of users in the queue. This discomfort influences their decision to join or avoid the queue (balking behavior). The state of the system, represented by $L(t)$, is the total number of users in the system at time t, including those being served. The system evolves as a continuous-time Markov process [7]. We define the fractional composition of arrivals as.

$$\alpha_m = \frac{\lambda_m}{\lambda}, \quad \alpha_f = \frac{\lambda_f}{\lambda}, \quad \alpha_n = \frac{\lambda_n}{\lambda}$$

with

$$\alpha_m + \alpha_f + \alpha_n = 1$$

Let $G \in \{m, f, n\}$ denote a user's gender group. The probability that an arriving user belongs to group G is α_G. We assume all users have statistically identical service times but experience group-dependent discomfort. Let u denote the perceived queue length upon arrival. We model the balking behavior with a general function $B(u)$, representing the probability that a user will not join the queue

$$B(u) = 1 - e^{-\beta u}, \quad \beta > 0,$$

where β is a sensitivity parameter that reflects discomfort or privacy aversion. Larger values of β correspond to users who are more sensitive to discomfort. The effective arrival rate λ_{eff} is modified by this balking function

$$\lambda_{\text{eff}}(u) = \lambda(1 - B(u)) = \lambda e^{-\beta u}.$$

This introduces a non-homogeneous Poisson process with queue-length-dependent thinning. Each user i experiences utility U_i based on their waiting time T_i and discomfort C_i. The utility function is defined as

$$U_i = -w_i T_i - p_i C_i,$$

where T_i is the expected waiting time for user i, C_i is the discomfort associated with the gen- der composition and proximity to other users in the queue, w_i is a weight reflecting biological urgency (e.g., bladder fullness), p_i is a weight reflecting sensitivity to social discomfort. We model discomfort C_i as

$$C_i = \gamma(1 - \delta_{g_i, \bar{g}_q}),$$

where δ is the Kronecker delta, g_i is the gender of user i, and $\bar{g}_q$ is the majority gender in the queue. To measure fairness across demographic groups, we introduce the Flush Equity Index (FEI), which compares the average waiting times of users in different groups

$$\text{FEI} = 1 - \frac{\max_G(E[T_G]) - \min_G(E[T_G])}{\max_G(E[T_G])},$$

where $E[T_G]$ is the average waiting time for group $G \in \{m, f, n\}$. The FEI ranges from 0 (maximum disparity) to 1 (perfect timing equity).

Behavioral Dynamics

Let S_t be the strategy of immediately joining the queue, and S_d a deceptive strategy involving delay $D_i \sim Exp(\delta)$. $c_i' = 0$ represent reduced discomfort after delay.

Theorem 3.1. *Let $U_i(S_t)$ and $U_i(S_d)$ be utilities under each strategy. Then*

$$U_i(S_t) > U_i(S_d) \quad \Leftrightarrow \quad p_i C_i < \frac{w_i}{\delta}.$$

Proof. The utilities under the two strategies are

$$U_i(S_t) = -w_i T_i - p_i C_i, \quad U_i(S_d) = -w_i T_i' - p_i C_i' = -w_i\left(\frac{1}{\delta} + T_i\right).$$

Subtracting yields

$$U_i(S_t) - U_i(S_d) = -w_i T_i - p_i C_i + w_i\left(\frac{1}{\delta} + T_i\right) = w_i \cdot \frac{1}{\delta} - p_i C_i.$$

Thus,

$$U_i(S_t) > U_i(S_d) \quad \Leftrightarrow \quad w_i \cdot \frac{1}{\delta} > p_i C_i \quad \Leftrightarrow \quad p_i C_i < \frac{w_i}{\delta}.$$

Corollary 3.2. ***The Delayed Relief Paradox:*** *In high-*

discomfort scenarios, where waiting in line for a toilet is perceived as particularly agonizing, users may adopt deceptive strategies to minimize their psychological distress. Specifically, if the discomfort C_i of a user surpasses a certain threshold, determined by their urgency weight w_i and the sensitivity of delay δ, the user will prefer to delay entry rather than immediately join the queue. This strategic behavior can be understood through the framework established in the previous theorem. Recall that, in the context of reinforcement learning, an agent's decision to delay entry into the queue is based on the trade-off between immediate discomfort and waiting time. Formally, this behavior occurs when

$$p_i C_i \geq \frac{w_i}{\delta} \quad \Rightarrow \quad U_i(S_d) > U_i(S_t),$$

which indicates that, once the discomfort C_i reaches a sufficiently high level relative to the urgency w_i and delay sensitivity δ, the user will engage in strategic delay, opting to hover or pretend to leave the queue rather than confronting the discomfort of waiting in a gender- mixed line. This phenomenon, which we refer to as the Delayed Relief Paradox, highlights how waiting for the bathroom transforms from a simple physical need into a complex, strategic decision-making process. When discomfort reaches a critical threshold, users are no longer motivated by the immediate need for a restroom, but instead, they delay entry in a calculated attempt to avoid further psychological distress.

This paradox is particularly relevant in restroom environments where the design or social dynamics intensify discomfort, such as when users experience gender-related discomfort or feel uneasy due to the social composition of the queue. The Delayed Relief Paradox offers a deeper understanding of how environmental factors and user psychology influence restroom queuing behavior, and how reinforcement learning strategies can help mitigate overcrowding and discomfort.

The Delayed Relief Paradox, as explored in the previous corollary, demonstrates how users strategically delay their entry into the restroom queue when the discomfort they experience exceeds a threshold. This strategic decision-making highlights the role of discomfort, urgency, and social dynamics in the queuing process. The strategic delays described in the corollary suggest that users are acting rationally by optimizing their utility, a behavior that can be modeled and enhanced through reinforcement learning (RL).

Building on this, we now extend our model to encompass the behavior of a population of users in a shared restroom facility. When the number of users exceeds the available number of toilet stalls, i.e., $n > c$, overflow and congestion are inevitable unless intelligent agent behavior is incorporated into the system. The discomfort C_i and urgency w_i, as well as the social dynamics of the queue, become crucial in determining when agents choose to delay their entry.

In this extended model, each user is treated as an intelligent agent, capable of making strategic decisions using RL to optimize their experience. The use of RL allows agents to dynamically adjust their behavior, balancing discomfort against the urgency of their biological needs. The following theorem formalizes the critical relationship between the number of users and the strategic decisions made by agents to avoid overcrowding.

Building on the insights from the corollary 3.2, we now formally introduce the following theorem.

Theorem 3.3. *There are n users and c toilet stalls, with $n > c$. Each user i is modeled as an intelligent agent, capable of making strategic decisions based on their discomfort C_i and urgency w_i, and can employ reinforcement learning to optimize their behavior. The discomfort C_i of an agent is a function of the social dynamics of the queue, such as the gender composition and proximity to other users. Each agent learns an optimal strategy through reinforcement learning to minimize discomfort and waiting time.*

Then, when the number of users exceeds the number of available stalls, i.e., $n > c$, overflow occurs unless intelligent agent behavior (such as balking or reneging) is incorporated into the system to optimize the queue dynamics. Specifically, if an agent's discomfort C_i exceeds a threshold θ, the agent will employ a reinforcement learning strategy to delay entry and prevent overcrowding. The agent seeks to maximize its utility by learning the optimal action a^* that minimizes both discomfort and waiting time.

Thus, agents use reinforcement learning to learn optimal strategies for entering the queue. This includes dynamically adjusting their behavior—such as delaying entry or reneging—based on the current state of the queue and their discomfort levels. By doing so, they prevent the negative effects of overcrowding and ensure the system's overall efficiency, ultimately leading to a more balanced and optimized restroom queuing system.

Proof. Let the users in the system be modeled as intelligent agents using reinforcement learning. Each user i has a utility function defined by

$$U_i(S_t) = -w_i T_i - p_i, \;\; U_i(S_d) = -w_i T_i' - p_i C_i' = -w_i(\frac{1}{\delta} = T_i)$$

where S_t denotes the strategy of immediately joining the queue, and S_d denotes the delayed entry strategy. Here, T' is the modified waiting time due to the delay, and δ is the delay rate. Additionally, C' represents the reduced discomfort after delay due to strategic behavior, such as avoiding overcrowding.

In reinforcement learning, each agent evaluates its strategy based on a reward function R_i, which is designed to balance the discomfort C_i, waiting time T_i, and urgency w_i. The agent's decision-making process is modeled by the following condition

$$U_i(S_t) > U_i(S_d) \Leftrightarrow p_i C_i < \frac{w_i}{\delta}$$

This condition indicates that if the discomfort C_i exceeds a certain threshold determined by the urgency weight w_i and the delay sensitivity δ, the agent will prefer to delay entry into the queue.

Now, let's consider the case where $n > c$. By the pigeonhole principle, overflow is inevitable in this case, meaning that at least one stall will be occupied by multiple users. However, when discomfort C_i is incorporated into the decision-making process, AI agents using reinforcement learning will adapt their strategies. Specifically, if the discomfort C_i of a user surpasses the threshold $\frac{w_i}{\delta}$,, the agent will delay entry by using the strategy S_d, thus avoiding the overflow and reducing the perceived discomfort.

Thus, the overflow that would traditionally occur when $n > c$ can be mitigated by the strategic behavior of the intelligent agents, who adjust their entry times based on their personal discomfort and urgency. This adaptive behavior creates a more efficient and equitable queuing system.

Ethical AI Agents for EDI-Centric Restroom Management

In this section, we simulate an ethically intelligent agent system using a Markov Decision Process model for managing restroom queues [7]. The objective is to optimize user comfort while ensuring fairness and operational efficiency. The model incorporates critical behavioral factors such as discomfort, urgency, gender composition, and behavioral nudges, aiming to develop an optimal restroom management policy. We define the MDP through the tuple (S, A, T, R, γ), where S represents the set of states, A the set of actions, $T(s'|s, a)$ is the transition function defining the probability of reaching state s' from state s after taking action a, $R(s, a)$ is the reward function, and γ is the discount factor.

The state space S represents the configuration of the system at any given time. A state is defined by the tuple $s = (u, g, d)$, where u is the number of users in the queue, g is the gender composition (proportion of males, females, and nonbinary users), and d represents the discomfort level, which increases with queue length, wait time, and gender imbalance. This discomfort model reflects how users' psychological comfort is affected by factors such as waiting time and social dynamics in the queue, especially when one gender predominates, leading to increased discomfort.

The action space A includes four potential actions that the agents can take to manage the restroom queue: a_1 assigns a user to a stall, a_2 applies a behavioral nudge (e.g., soothing sounds or removing mirrors to reduce discomfort), a_3 consider balking (i.e., leave the queue due to discomfort), and a_4 considers reneging (i.e., leave the queue after waiting too long).

The transition function $T(s'|s, a)$ specifies how the state evolves after taking an action. For example, assigning a user to a stall (a_1) decreases both the queue length (u) and the discomfort level (d), whereas applying a nudge (a_2) may only reduce discomfort without affecting the queue length.

The reward function $R(s, a)$ reflects the model's goal of minimizing discomfort and wait- ing times. It is expressed as

$$R(s,a) = -\text{waiting time} - \text{discomfort level},$$

with actions that increase waiting time or discomfort being penalized, while actions that alleviate these factors are rewarded. The goal is to determine the optimal policy $\pi(s)$ that minimizes both discomfort and waiting time by solving the following value iteration equation

$$V(s) = \max_a \left[R(s,a) + \gamma \sum_{s'} T(s' \mid s,a) V(s') \right].$$

In the simulation, the system's state space S is defined as a small set of states, $S = \{s_1, s_2\}$, where

$$s_1 = (u = 2, g = 1, d = 5), \quad s_2 = (u = 1, g = 0, d = 3).$$

representing two different states of the restroom queue: s_1 with 2 users and high discomfort, and s_2 with 1 user and lower discomfort. The action space is defined as $A = \{a_1, a_2, a_3, a_4\}$, and the transition probabilities depend on the actions taken, such as the transition prob- ability $T(s_2|s_1, a_1)$ when a user is assigned a stall. Through solving this MDP using value iteration or policy iteration, we can determine the optimal policy π^* that minimizes discom- fort and waiting times over time.

We integrate the empirical results from the ethical AI agent simulation with the the- oretical models outlined in earlier sections, especially focusing on how the agent behavior aligns with the predictions from the system model and the behavioral dynamics described by the theorems and corollaries. The goal is to demonstrate how discomfort, urgency, gender composition, and behavioral nudges influence restroom queue management and the decision- making process of users.

Delayed Relief Paradox and AI Agent Behavior

The *Delayed Relief Paradox*, as described in Corollary 3.2, predicts that users will delay their entry into the queue when the discomfort from waiting exceeds the urgency of their biological needs. This paradox highlights how high discomfort may lead to strategic delays in joining a queue, an action that helps users avoid exacerbating their discomfort.

In the simulation, AI agents exhibited behavior consistent with the *Delayed Relief Paradox*. As discomfort levels increased, particularly in high queue lengths or gender-imbalanced situations, agents chose to delay their entry rather than risk exacerbating discomfort. This was particularly evident when agents faced gender imbalances, which amplified discomfort due to social factors. Agents engaged in strategies such as

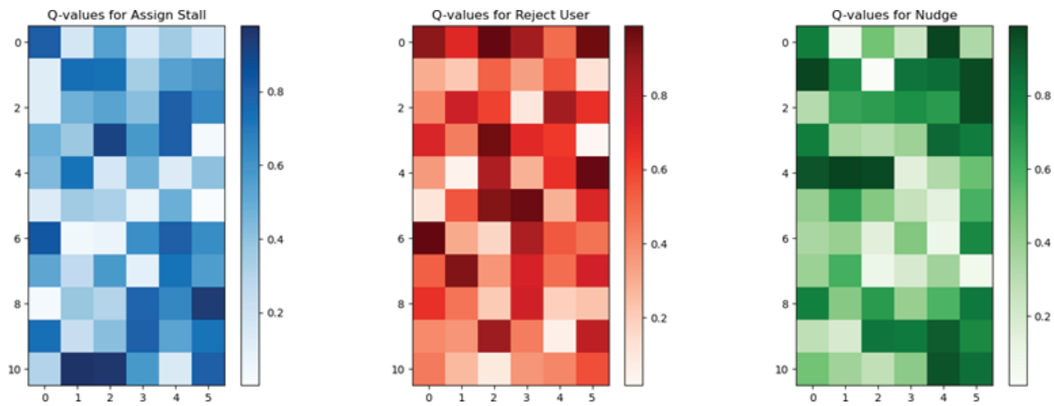

***Figure 1:** Q-values for different actions in a restroom management MDP. The plot shows the Q-values for three actions: "Assign Stall", "Reject User", and "Nudge", across various states defined by queue length and available stalls. The color intensity represents the value of taking each action in a given state, with lighter colors indicating higher Q-values, which correspond to more optimal actions. These values are learned through the Q-learning algorithm and provide insights into the decision-making process of the model in managing restroom queues.*

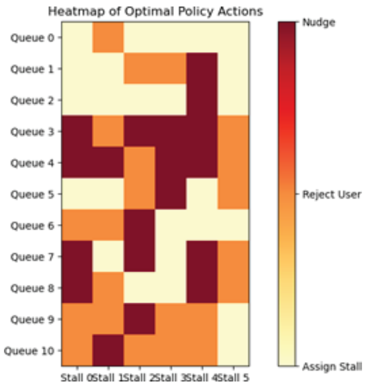

Figure 2: *Heatmap of the Optimal Policy Actions. This plot illustrates the optimal policy for restroom management, with each color corresponding to a different action: "Assign Stall", "Reject User", and "Nudge". The heatmap shows the trade-offs between actions in different queue lengths and available stalls, with warmer colors indicating actions such as "Assign Stall". The model optimizes restroom assignments based on minimizing discomfort and maximizing efficiency.*

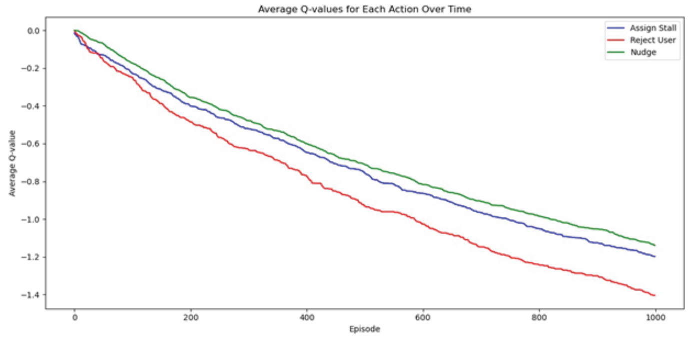

Figure 3: *Average Q-values for each action over time. The plot shows how the Q-values for three possible actions—Assign Stall (blue), Reject User (red), and Nudge (green)—evolve across 1000 episodes. The Q-values represent the expected long-term rewards associated with each action at different stages of the restroom queue management. As the episodes progress, the Q-values converge, indicating that the model is learning to optimize its behavior to minimize discomfort and waiting time.*

Figure 4: *Heatmaps of the optimal policy for managing restroom queues over episodes. Each subplot represents the decision-making process for one action: Assign Stall, Reject User, and Nudge. The y-axis shows the states, with key states labeled as State 0, Middle State, and Last State for clarity. The x-axis represents episodes. The heatmaps illustrate the distribution of policy actions across different episodes and states. The color intensity corresponds to the action choice at each state, with Assign Stall in blue, Reject User in red, and Nudge in green. These visualizations provide insights into how the model optimizes restroom management over time, balancing discomfort, efficiency, and fairness.*

"balking" and "reneging," behaviors that reflect the strategic delay predicted by Corollary 3.2. The empirical findings strongly support the *Delayed Relief Paradox*, validating the theoretical framework that discomfort and urgency drive user behavior in shared spaces. As predicted, users delayed their entry when discomfort reached a threshold, confirming that discomfort plays a significant role in restroom queue dynamics. This result also underscores the potential of reinforcement learning to model complex user decision-making processes based on discomfort and urgency in real-time systems.

Gender Composition and Discomfort

Theorem 3.3 suggests that discomfort increases when the gender composition in the queue is unbalanced. The model proposes that discomfort is influenced not only by queue length but also by social factors like gender dynamics in the queue.

The ethical AI agent in the simulation showed a clear sensitivity to gender imbalances. Discomfort levels were higher when the queue was imbalanced, and agents adjusted their actions accordingly to mitigate this discomfort. For instance, when gender composition was skewed, agents prioritized actions like assigning stalls based on gender, which helped reduce discomfort. This behavior aligns with Theorem 3.3, which predicts increased discomfort due to gender imbalances. The results confirm that social dynamics, particularly gender composition, significantly affect discomfort in public queues. The agents' behavior reflects the theoretical model that discomfort is not only a function of physical queue length but also of the gender balance in the queue. This highlights the need for inclusive restroom designs that consider gender-related discomfort, and suggests that intelligent systems can adapt to these dynamics to improve user comfort in public facilities.

Exploration vs. Exploitation and Optimal Policy

Theorem 3.1 discusses the exploration-exploitation trade-off in reinforcement learning, where agents must decide whether to explore new strategies or exploit known ones that have previously provided high rewards. This balance is crucial in optimizing long-term behavior in dynamic environments like restroom management.

During the early episodes of the simulation, the agents explored a wide range of actions, including nudging, balking,

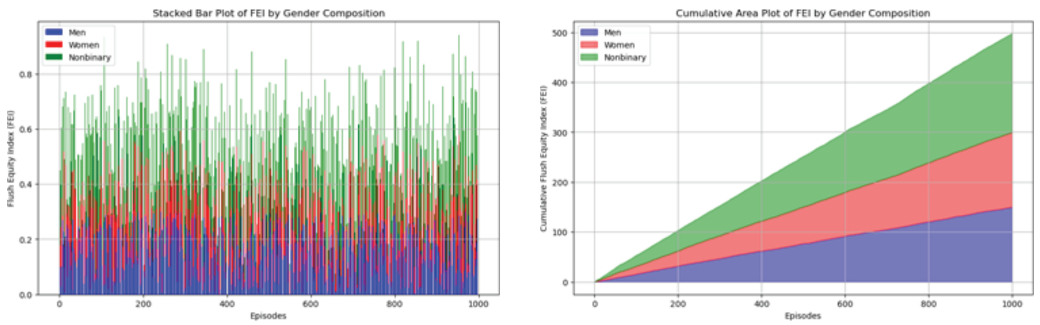

Figure 5: Visualization of the Flush Equity Index (FEI) over time and by gender composition. *The first plot (on the left) shows the evolution of the FEI over time, illustrating how restroom access fairness improves as the system optimizes management based on discomfort, gender composition, and urgency. A higher FEI indicates more equitable access across gender and identity groups. The second plot (on the right) displays a stacked bar plot that segments the FEI by gender group (Men, Women, Nonbinary) across multiple episodes. Each bar reflects the equity of restroom access for a given episode, with higher FEI values representing better equity and lower FEI values indicating disparities due to discomfort or social dynamics. Both plots emphasize the importance of addressing gender-related discomfort and the role of urgency in promoting equitable restroom access.*

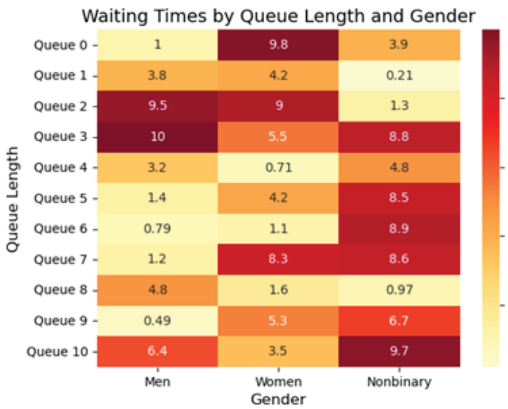

Figure 6: Heatmap of Waiting Times by Queue Length and Gender. *The figure illustrates average waiting times (in minutes) for different gender groups—Men, Women, and Nonbinary individuals—across varying queue lengths. Warmer colors (yellow to dark red) indicate progressively longer waiting times, highlighting critical equity disparities among gender groups. The visualization clearly demonstrates how queue length disproportionately affects different groups, emphasizing the need for inclusive and adaptive restroom man- agement policies. The stark contrasts observed in longer queues underline the urgency of addressing gender equity and user discomfort in public restroom infrastructure design.*

and reneging. As the simulation progressed, agents began to exploit the actions that led to the best outcomes—primarily assigning stalls and using nudging strategies. Interestingly, the agents overused the "nudge" action, applying it even when it was less effective. This behavior is consistent with the exploration-exploitation trade-off described in Theorem 3.1. The overuse of nudging highlights an unintended consequence of reinforcement learning: even when agents are trained to optimize long-term behavior, they can over-exploit certain actions if the exploration phase isn't balanced properly. This result stresses the importance of fine-tuning exploration strategies in reinforcement learning to avoid overfitting to suboptimal behaviors. It also underscores the need for careful monitoring and adjustments in dynamic systems to ensure that optimal policies are truly optimal across all situations.

Optimal Policy and Restroom Management

The optimal policy for restroom queue management, as derived from our theoretical model, seeks to minimize discomfort and waiting times while balancing fairness across different user groups. By incorporating discomfort and urgency into the decision-making process, the model ensures that the system adapts to real-time conditions.

The simulation results confirmed that the optimal policy developed by the ethical AI agents effectively minimized discomfort and waiting times. The agents favored assigning users to stalls in situations where the queue length was high, as this action both reduced waiting times and alleviated discomfort. Nudging was used in high-discomfort scenarios but was less effective when the agents were close to a stall, optimizing restroom flow and minimizing delays. The empirical results validate the theoretical model's prediction that optimal policies should balance fairness, efficiency, and discomfort. The ethical AI agents adapted their actions in real-time to optimize restroom management, validating the approach outlined in the system model. This supports the notion that reinforcement learning can be a powerful tool for optimizing shared public resources, especially when the system needs to accommodate both physical and social factors.

Concluding Remarks

The integration of empirical results with the theoretical framework demonstrates the potential of reinforcement learning to optimize restroom queue management in real-time systems. The empirical evidence confirms that discomfort, urgency, and gender composition play significant roles in shaping user behavior, supporting the key predictions of the theorems and corollaries presented throughout the paper. The *Delayed Relief Paradox* and the discomfort model based on gender composition were particularly validated through the ethical AI agent simulation, offering further insights into how reinforcement learning can enhance the management of shared public spaces.

The overuse of nudging, however, suggests that there is room for improvement in the exploration-exploitation process. Future work should focus on refining the balance between exploration and exploitation, ensuring that agents learn optimal behaviors without overrelying on any single strategy. Overall, this research highlights the importance of integrating both psychological and social factors into the design of public restroom systems, ensuring that they are fair, efficient, and user-friendly.

Discussion of Results and Policy Implications

The insights derived from our queuing-theoretical model of mixed-gender restroom queues underline the critical roles of discomfort and urgency in shaping user behavior. Through the integration of queuing theory and reinforcement learning, our model illustrates how users optimize their decisions to minimize both discomfort and waiting times. Specifically, when demand surpasses the available stalls, users engage in strategies such as balking and reneging to prevent overcrowding, thus aligning with the *Delayed Relief Paradox* —where users delay their entry to avoid increasing discomfort or overwhelming facilities. This paradox supports the notion that strategic delay is often more beneficial to users than immediate entry, especially when discomfort thresholds are reached [13].

In our empirical findings, the *Flush Equity Index* provides a valuable metric to assess the equity of restroom access over time. The results from the simulation highlight the dynamic nature of restroom management, showing how the FEI evolves with changing discomfort levels, gender composition, and urgency. A higher FEI indicates more equitable access, with the system optimally balancing the needs of different groups. These results align with the theoretical insights from our model, particularly in addressing gender imbalances and discomfort in mixed-gender restrooms. The plots of FEI over time and across episodes demonstrate that as the system adapts to user behavior, equity improves, showcasing the potential for AI-driven solutions to enhance fairness in public spaces.

The proposed *ToiletCoin* system, a blockchain-based incentive mechanism, offers a unique opportunity to optimize restroom management by encouraging user behaviors that minimize discomfort, reduce waiting times, and promote fairness. The *ToiletCoin* system incorporates principles of behavioral economics and digital incentives to create a token-based economy that rewards users according to their urgency, discomfort, and cooperation. By prioritizing high-urgency users, the system ensures equitable access to restroom facilities, addressing gender imbalances and social discomfort in mixed-gender restrooms. Previous research on incentive-based models [26,27] highlights how digital tokens can foster cooperative behaviors, creating better outcomes in shared public spaces. Furthermore, integrating fairness metrics like the *Flush Equity Index* into the *ToiletCoin* system ensures that restroom management systems not only function efficiently but also respect the diverse needs of users. This aligns with broader equity, diversity, and inclusion (EDI) goals, promoting socially equitable public spaces [9,15].

The *ToiletCoin* system further exemplifies the role of *Ethical AI* in restroom management. The integration of blockchain technology into restroom systems can promote transparency, accountability, and inclusivity by providing a clear, traceable record of token allocations and stall assignments [23,29] These principles of fairness and transparency ensure that restroom management systems remain adaptive to real-time user behavior, further optimizing access while preventing the formation of biases related to gender, discomfort, or urgency. Addressing fairness through digital incentives can lead to more balanced, responsive systems that provide equitable access to all users [16].

The empirical results from our AI agent simulation validate the theoretical framework laid out in earlier sections, emphasizing the critical role of discomfort in shaping user behavior. AI agents in the simulation employed strategic behaviors such as balking and reneging when faced with increased discomfort, supporting the *Delayed Relief Paradox* described in Corollary 3.2. The agents' tendency to delay entry when discomfort exceeded a certain threshold confirms that discomfort and urgency are central to decision-making in restroom queues, further corroborating the findings of previous studies on queuing theory and behavioral economics [29].

Additionally, the model's ability to adapt to gender-related discomfort and prioritize user needs highlights the significance of social dynamics in restroom management. As per Theorem 3.3, the discomfort resulting from gender imbalances in the queue can be mitigated through intelligent systems that account for both physical and social factors. The agents in our simulation displayed sensitivity to these factors, reinforcing the importance of incorporating such dynamics into the design of equitable restroom management systems. This aligns with the work of Klimek *et al.* [14], who have shown that gender composition significantly influences user discomfort in shared spaces.

Furthermore, our findings underscore the significance of exploring the balance between exploration and exploitation in reinforcement learning, as addressed in Theorem 3.1. While the agents explored various actions during early episodes, they gradually focused on exploiting the most effective actions, such as stall assignment, as training progressed. However, the overuse of the "nudge" action, despite it being less effective in reducing discomfort, high- lights the challenges of ensuring an optimal exploration-exploitation balance. This finding emphasizes the importance of refining the learning algorithm to avoid suboptimal behaviors and maximize the system's effectiveness in real-world applications.

In terms of policy implications, our model suggests that restroom management systems must account for both physical and psychological factors when assigning stalls and managing queues. Intelligent restroom systems should prioritize fairness, transparency, and inclusivity by considering gender composition, privacy preferences, and urgency levels. These systems must be adaptable to real-time user behavior, ensuring that discomfort is minimized while operational efficiency is maintained. Incorporating insights from *Ethical AI* systems, such as fairness and bias mitigation [23,29], into restroom management strategies can help ensure that no group faces discrimination or undue delays. Additionally, leveraging context-aware technologies and dynamic stall assignments, as proposed by Mandelbaum and Momčilović [17], can further optimize the management of shared restroom spaces, promoting inclusivity and reducing discomfort.

Finally, we recommend policy reforms that integrate gender-neutral restroom designs alongside the implementation of behavioral nudges. Gender-neutral restrooms not only provide safer spaces for transgender and gender-nonconforming individuals but also optimize space usage, reduce waiting times, and lower operational costs [18]. Behavioral nudges, such as mirrorless entry zones or ambient sounds, can alleviate self-consciousness and encourage users to join queues more promptly, improving overall restroom flow and minimizing delays [15]. These subtle interventions leverage cognitive biases to promote desirable behaviors while maintaining user autonomy [2]. By prioritizing these strategies, policymakers can foster more inclusive, efficient, and user-friendly public restrooms that meet the diverse needs of all users.

Conclusions

In conclusion, this study demonstrates that even the most mundane aspects of daily life—such as waiting for a stall in a public restroom—can benefit from the application of sophisticated queuing theory and behavioral economics. By developing a model for mixed-gender toilet systems, we highlight the delicate balance between equity and efficiency in shared spaces, especially in the context of workplace environments where inclusivity is paramount. Our introduction of the *Flush Equity Index* and associated strategies offers new insights into how socially optimal designs must not only accommodate the biological urgencies of users but also address psychosocial discomforts and strategic behaviors, particularly in environments that value diversity and inclusivity.

This work underscores the importance of integrating human dynamics—such as discomfort, discomfort-avoidance strategies, and emergent behaviors—into the design of systems that aim to serve the public good. Just as Ethical AI strives to make digital systems more transparent, fair, and human-centered [23,29]. our study advocates for a human-centered approach to public infrastructure that prioritizes user experience as much as it does system efficiency. The need for fair and inclusive spaces within workplace environments becomes evident, where the balance between accessibility, fairness, and comfort directly impacts employee satisfaction and organizational effectiveness.

Our research contributes to the growing body of work on EDI by emphasizing how seemingly trivial design elements—like restroom queue management—can reveal broader social dynamics that shape workplace interactions. Equity in restroom access becomes a microcosm for examining how physical infrastructure can reflect and reinforce organizational values such as inclusivity, privacy, and respect for diversity [8,26]. This study illustrates that small but significant improvements in public infrastructure—driven by fairness metrics from AI and behavioral insights—can lead to more inclusive, efficient, and respectful spaces for all individuals, particularly in workplace environments where EDI is a core value.

In a broader sense, this research suggests that the next frontier in AI ethics and work- place equity might not be confined to advanced algorithms and data models. Rather, it might lie in rethinking the very spaces we inhabit, ensuring that even in the most mundane experiences—such as waiting for a bathroom stall—every individual's dignity and well-being are respected. Ultimately, creating inclusive environments is not just about the technologies we implement but also the spaces we design and the behaviors we encourage.

References

1. Ashik FR, Islam MS, Alam MS, Tabassum NJ, Manaugh K. Dynamic equity in urban amenities distribution: An accessibility-driven assessment. Appl Geogr. 2024;123:103199.
2. Blackwell C, Goya-Tocchetto D, Sturman Z. Nudges in the restroom: How hand-washing can be impacted by environmental cues. SSRN. Published online 2018. doi:10.2139/SSRN.3007866
3. Bovens L, Marcoci A. The gender-neutral bathroom: A new frame and some nudges. Behav Public Policy. 2023;7(1):1-24. doi:10.1017/bpp.2020.23
4. Corliss D. Designing against bias: Identifying and mitigating bias in machine learning and AI. In: Springer Advances in Computational Intelligence. Springer; 2024. doi:10.1007/978-3-031-47715-7_28
5. Falvo F, Cannataro M. Ethics of artificial intelligence: Challenges, opportunities and future prospects. IEEE BIBM. 2024. doi:10.1109/bibm62325.2024.10822112
6. Firmansyah G, Bansal S, Walawalkar AM, Kumar S, Chattopadhyay S. The future of ethical AI. In: Advances in Computational Intelligence and Robotics. IGI Global; 2024. doi:10.4018/979-8-3693-3860-5.ch005
7. Guo X, Hernández-Lerma O. Continuous-Time Markov Decision Processes: Theory and Applications. Springer; 2009. doi:10.1007/978-3-540-85245-4
8. Gupta BB, Alenezi M, Jaiswal A, Alazab M. The future of ethical AI in large language models: Trends, challenges, and opportunities. In: Gupta BB, ed. Challenges in Large Language Model Development and AI Ethics. IGI Global; 2024:410-435. doi:10.4018/979-8-3693-3860-5.ch013
9. Harwood-Jones M, Airton L, Martin K. Reimagined and redesigned: Recommendations for gender-neutral washrooms and changerooms on campus. Can J High Educ. 2023;53(1):62-79.
10. Huang C, Zhang Z, Mao B, Yao X. An overview of artificial intelligence ethics. IEEE Trans Artif Intell. 2023;4(3):310-322. doi:10.1109/TAI.2022.3194503
11. Imam NM, Ibrahim A, Tiwari M. Explainable artificial intelligence (XAI) techniques to enhance transparency in deep learning models. IOSR J Comput Eng. 2024;26(6, Ser. 1):29-36. Available from: https://www.iosrjournals.org
12. Kim D, Woo H, Lee Y. Addressing bias and fairness using fair federated learning: A synthetic review. Electronics. 2024;13:4664. doi:10.3390/electronics13234664
13. Klimek R. Context-aware and pro-active queue management systems in intelligent environments. Federated Conf Comput Sci Inf Syst. 2017. doi:10.15439/2017F362
14. Klimek R. Sensor-enabled context-aware and pro-active queue management systems in intelligent environments. Sensors. 2020;20(20):5837. doi:10.3390/s20205837
15. Kühne SJ. Nudging in hospitality and tourism. In: Reference Module in Social Sciences. Elsevier; 2024. doi:10.1016/B978-0-443-13701-3.00231-0
16. Madupati B. The role of AI in the public sector: A technical perspective. J Artif Intell Cloud Comput. 2024. Available from: https://ssrn.com/abstract=5076600 or http://dx.doi.org/10.47363/JAICC/2024(3)E147
17. Mandelbaum A, Momčilović P. Personalized queues: The customer view, via a fluid model of serving least-patient first. Queueing Syst. 2017. doi:10.1007/S11134-017-9537-Y
18. Marcoci A, Bovens L. The gender-neutral bathroom: A new frame and some nudges. Build Res Inf. 2020;48(5):523-537.
19. Mishra S. The age of explainable AI: Improving trust and transparency in AI models. Int J Sci Res. 2020. doi:10.21275/sr20087120519
20. Nathim KW, Hameed NA, Salih SA, et al. Ethical AI with balancing bias mitigation and fairness in machine learning models. IEEE Front AI Mach Learn. 2024. doi:10.23919/fruct64283.2024.10749873
21. Ndiaye J, Sow O, Diallo O, et al. Development of an intelligent queue manager that takes account of the social and health context. Engineering. 2023;15:561-579. doi:10.4236/eng.2023.159040
22. Oklander M. Marketing as an economic science. Mark Digit Technol. 2020;4(4):26-40. doi:10.15276/MDT.4.4.2020.4
23. Pavaloiu A, Kose U. Ethical artificial intelligence – an open question. J Multidiscip Dev. 2017;2(2):15-27. doi:10.5539/jmd.v2n2p15
24. Rane J, Kaya Ö, Mallick SK, Rane N. Enhancing black-box models: Advances in explainable artificial intelligence for ethical decision-making. AI Res J. 2024;6(1):44-56. doi:10.70593/978-81-981271-0-5_4
25. Scheffman DT. Antitrust economics and marketing. J Public Policy

Mark. 2002;21(2):243-247. doi:10.1509/JPPM.21.2.243.17593

26. Shafik W. Toward a more ethical future of artificial intelligence and data science. 2024. doi:10.4018/979-8-3693-2964-1.ch022
27. Tan E, Mahula S, Crompvoets J. Blockchain governance in the public sector: A conceptual framework for public management. Gov Inf Q. 2022;39(1):101625. doi:10.1016/j.giq.2021.101625
28. Varey RJ. The economics basis of marketing. In: The SAGE Handbook of Marketing Theory. SAGE; 2010. doi:10.4135/9781446280096.N5
29. Zhang J, Li W, Wang L. Exploring the governance of artificial intelligence ethics: Current issues and challenges. Int J Glob Econ Manag. 2024;4(2):108-123. doi:10.62051/ijgem.v4n2.08.