

DIGITAL TWIN TO MITIGATE ADVERSE ADDICTIVE GAMBLING BEHAVIOR

Felisa Vázquez-Abad¹, Jason Young², and Silvano Bernabel³

¹Dept. of Computer Science, Hunter College, City University of New York, NY, USA.

²Dept. of Psychology, Hunter College, City University of New York, NY, USA.

³Dept. of Computer Science, Graduate Center, City University of New York, NY, USA.

ABSTRACT

This work develops a simulation engine to create a digital twin that will monitor a gambler's betting behavior when playing games. The digital twin is designed to perform simulations to compare outcomes of different betting strategies, under various assumptions on the psychological profile of the player. With these simulations it then produces recommendations to the player aimed at mitigating adverse outcomes. Our work focuses on efficient simulation and the creation of the corresponding GUI that will become the interface between the player and the digital twin.

1 INTRODUCTION

Gambling has long been a popular recreational activity, providing thrill and entertainment through games of chance. However, it also poses significant risks to financial and psychological well-being particularly for players who lack a disciplined approach or become overly reliant on gambling for emotional or monetary rewards.

In fact, gambling as an online experience is becoming easier to access (including on mobile phones) to more people, and at an earlier age, leading to projected increases in the population exposed to risks for gambling addiction (Sohn 2023). In this study we focus on the American roulette as a case study. Roulette, a staple of casino culture, epitomizes this duality: it is simple to understand but deeply complex in its mathematical underpinnings and behavioral implications. The goal of this research is to create a digital twin that interacts with the player suggesting recommended betting strategies (Barricelli, Casiraghi, and Fogli 2019). While aiming to prevent financial ruin, the recommendations should recognize the behavioral patterns of players in order to be successful.

A pilot study provided a proof of concept for our digital twin. In that work we simulated three gambler's profiles: high, moderate and low risk players. That simulation revealed key aspects of the edge dynamics. Our preliminary study shows that low-risk players sustain final bankrolls of \$5.24 on average, often reaching the maximum 100 rounds, when their initial bankroll is \$100. Moderate-risk players bankrupt before the 100 rounds with (approximate) probability 0.8, while high-risk players bankrupt within 20 rounds on all 30,000 replications that we made of the pilot simulation. This observation, while not surprising, motivates our idea to help gamblers adapt their strategies in order to improve the final outcome. The simulation of gambler's strategies was compared to a simulation of a digital twin that uses a different exit strategy (explained in more detail below), with the aim at stopping the gambler from continuing if the current state has high risk of falling below a certain acceptable loss. Implementing bankroll-based exit mechanisms drastically reduced high-risk player losses (from \$100 expected loss to \$3.33–\$6.67), demonstrating the benefits of adaptive strategies. Following this pilot study we develop here a more detailed model for the digital twin, aiming at keeping losses at a certain “acceptable” level, while acknowledging the nature of addictive gambling.

A risk theoretical approach has been proposed to model gambling behavior (Schnytzer and Westreich 2010), where utility functions model the “satisfaction” of the player, in an attempt to explain how a “rational”

player gambles even when it is known that the expected loss is positive. Barberis (2012) recognizes that the usual economic utility theory is not consistent with observed gambling behavior with respect to risk taking and uses “prospect theory” to model non-concave utility functions, which then determines betting behavior. Building on this research, Blavatskyy (2024) studies various utility models to describe different betting behaviors for just one single round. Although they do not mention the psychology of gambling, it is apparent that different mathematical models of behavior may be related to what we will call the profiles of the players. Hales, Clark, and Winstanley (2023) present a comprehensive study of the “Gambling Disorder”, a psychological condition that has been well identified, but which still eludes definite treatments. Mathematical models have been developed in order to gain insight into the cognitive process that leads to addictive gambling. Among the methods, reinforcement learning, bayesian models and drift diffusion models have been used to model a gambler’s behavior.

None of the above research papers dwell on the question of *mitigation* of the adverse effects of gambling as an addiction, which we propose here. Moving beyond the existing research, in this paper we account for the gambler’s behavior to build a reasonable “agent” in the form of a digital twin that will interact with the player in the form of recommended actions.

2 MODEL FORMULATION

2.1 American Roulette: Model

Figure 1 shows the 38 possible outcomes, assumed equally likely. As well, the various possible *placements* of the bets are illustrated.

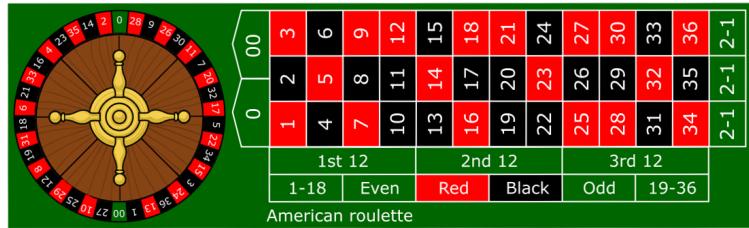


Figure 1: The possible outcomes and placements in American Roulette.

A betting strategy consists of three parts: the *placement*, the *wager amount* and a *stopping time*.

Placement Strategy and corresponding Payout (s):

1. Single-number: The player chooses one of the 38 possible outcomes. Payout is 36 : 1.
2. Dozens: The player chooses one of the three possible “boxes” of a dozen numbers. Payout is 3 : 1.
3. Binary: The player chooses one of two choices (black or red, even or odd, first half or second half). Payout is 2 : 1.

Call $R(s)$ the payout ratio and $P(s)$ is the corresponding winning probability of placement strategy s , that is:

$$P(s) = \left(\frac{1}{38}, \frac{12}{38}, \frac{18}{38} \right).$$

The house edge is the expected amount of money that the player loses (relative to the bet amount). It is calculated using the expected value formula. Assuming a unit bet of \$1.00 this yields

$$E = (1 - P(s)) - (R(s) - 1)P(s).$$

Direct calculation shows that the house edge for the American roulette with these parameters is always 0.0526, regardless of the value of s for the placement strategy.

Wager Strategy (w): There are a number of websites available today for gamblers to either play the actual games, or to read and practice betting strategies on their own (Hoofe 2025; Grindu 2025; Coyle 2023). These sites provide in-depth breakdown of various wager strategies and explain the dynamics of the game, encouraging players to choose a single strategy and use it throughout the game. Naturally any player may change his/her mind and change strategies at any given round of the game. However we follow here the model where a gambler chooses a wager strategy at the start of the game and uses it all the time, unless the digital twin intervenes with a different recommendation (see below). We consider the most popular five wager strategies, described below.

1. All-In-One: The player wagers the entire bankroll.
2. Martingale: The player defines an initial wager as a fraction of the bankroll. After a loss, the player doubles the amount (if possible). After a win, the player resets the wager amount to the initial one.
3. Fibonacci: The initial wager is \$1. After a loss, the wager amount is the following number to the previous amount in the Fibonacci sequence (if possible). After a win, the wager chosen moves back two places in the Fibonacci sequence (if possible).
4. D'Alembert: The player chooses an initial amount (called the “unit”). After a loss, the wager is increased by one unit (if possible). After a win, the wager is decreased by one unit (if possible).
5. Flat: The player always stakes the same initial amount regardless of previous outcomes (if possible).

Our pilot study used Monte Carlo simulations of the game to establish that aggressive strategies (All-in, and Martingale, with “Single-number”) yield a 0% win rate. Martingale with a “Dozen” bet achieves a 16.67% win rate and gameplay length averaged 30.33 rounds, while Fibonacci improves the win rate (20%) and gameplay length (average 53.27 rounds). D'Alembert performs best with a 50% win rate and average gameplay length of 94.7 rounds, while Flat ensures the longest gameplay length with a comparable win rate. These results justify the order in which we have labeled the strategies.

The effective house edge, considering various wager strategies, can be estimated using the approximation

$$E = \frac{1}{(\sum_{k=1}^N Y_k)} \sum_{k=1}^N (Y_k - W_k),$$

where Y_k are the consecutive wager amounts, W_k the consecutive winnings paid out to players, using a maximum of N rounds for the simulations. In our pilot simulation, high-risk strategies inflate the house edge to 44.93%, while low-risk players maintain a near-theoretical 5.26%, emphasizing the importance of balanced strategies.

Stopping Criterion (τ): The game consists of playing successive rounds until the stoping criterion is met (see description below).

2.2 Markov Model

The stochastic model corresponds to a Markov decision process that describes the bankroll (wealth) of the player. The bankroll X_n denotes the wealth of the gambler at the start of round n . The possible actions are the wager amount Y_n and the placement Z_n at round n . The initial wealth or capital of the player X_0 is known. The transition probabilities are now defined as:

$$p_{ij}(y, s) \stackrel{\text{def}}{=} \mathbb{P}(X_{n+1} = j \mid X_n = i, Y_n = y, Z_n = s)$$

and they can be calculated from the description of the game as follows. Given a wager amount y and a placement strategy s with corresponding payout ratio $R(s)$, it follows that

$$X_{n+1} = \begin{cases} (X_n - y) + R(s)y & \text{w.p. } P(s) \\ X_n - y & \text{w.p. } 1 - P(s), \end{cases}$$

where the notation “w.p.” means “*with probability*”. In addition, the wager amount is related to the previous one. We add a binary indicator $I_n = \mathbf{1}_{\{\text{win}\}}$ for the outcome of the n -th round. Call $F(i)$ the i -th Fibonacci number. In the case of this wager strategy ($w = 3$), $Y_n = F(k_n)$ for some known integer k_n . This yields

$$\tilde{Y}_{n+1} = \begin{cases} X_{n+1} & w = 1 \\ 2Y_n(1 - I_n) + Y_0 I_n & w = 2 \\ F(k_n + 1)(1 - I_n) + F(k_n - 2) I_n & w = 3 \\ (Y_n + Y_0)(1 - I_n) + (Y_n - Y_0) I_n & w = 4 \\ Y_0 & w = 5 \end{cases}$$

and then the actual wager must satisfy:

$$Y_{n+1} = \min(\tilde{Y}_{n+1}, X_{n+1})$$

We assume no borrowing is allowed, so if the player has no more money then the game is over. Because by design the probability of winning any round is smaller than 0.5 the main bankroll process $\{X_n\}$ is a super-martingale: $\mathbb{E}[X_{n+1} | X_n] \leq X_n$. In addition, the gambler is ruined when $X_n = 0$, so the bankroll converges to zero w.p.1 (eventual ruin is certain). This is our motivation to introduce a *stopping time* when the gambler cashes the bankroll. The game consists of the consecutive rounds (spins) that yield the process $\{\xi_n\}$, until the game ends. Specifically:

$$\xi_n \stackrel{\text{def}}{=} (X_n, Y_n, Z_n, I_n).$$

The default stopping time is

$$\tau = \min(n \leq N: X_n = 0), \quad (1)$$

where N is the maximum fixed amount of rounds for one player. That is, either the gambler “ruins” losing all the bankroll, or the game ends after N rounds and the player cashes out the amount X_N .

Definition. The *expected loss* incurred by the player is defined by:

$$L = \mathbb{E}[X_0 - X_N]. \quad (2)$$

and it is usually positive.

3 DIGITAL TWIN

A digital twin is proposed in order to provide recommendations to the player. We are currently developing a GUI as an interface that allows players to place bets and follow the normal evolution of roulette playing over a sequence of rounds. The digital twin is therefore “hidden” within a gambling app. The digital twin is capable of performing parallel simulations, given the current state ξ_n , assuming a given gambler profile.

Simulations then help to estimate the expected loss and assess more favorable outcomes under changes in the player’s behavior. To alleviate the possibility of bankruptcy we endow the digital twin with the capability of pre-empting actions by requiring periodic pauses during which the digital twin offers *recommendations* to the gambler. In addition to recommending specific actions, we seek (in future work) to design recommendation strategies that will attempt to modify behavior in order to control adverse addictive conduct. Figure 2 shows the main page of the GUI. After the initial page asks the player for general information, this main page shows the player the amount already lost (*Casino profit*), the current bankroll and the current strategy with corresponding wager amount. The player may change the strategy and the wager amount if desired, however the GUI presents the suggestion using the same strategy. Then the player places the token in the corresponding place that he/she wants to choose (illustrated on the placement zone labeled 2_{nd} 12 in Figure 2). Naturally the GUI does not accept a placement that is not in agreement with the chosen placement strategy.

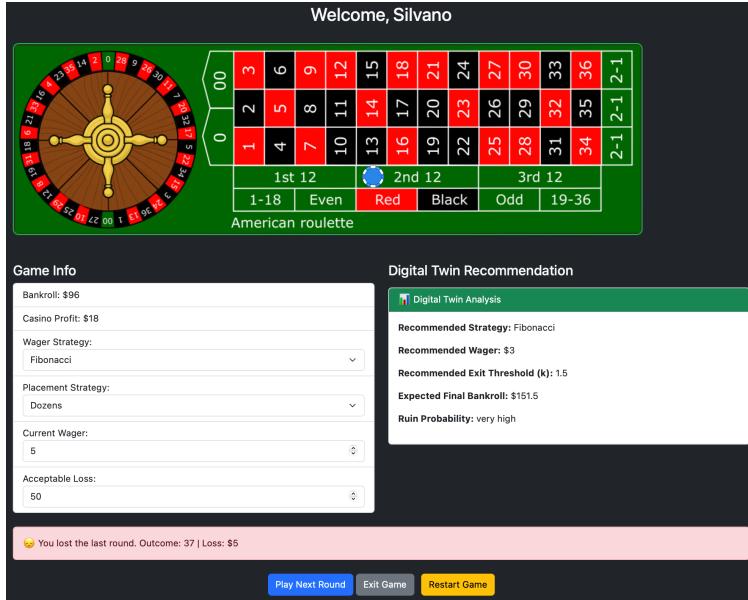


Figure 2: Main page of the American Roulette GUI.

3.1 Recommendations: Early Exit Rule

Although the expected loss is always positive, people are willing to spend money on entertainment, whether it's watching movies, going out for a meal, or playing tennis. Our goal is to suggest a healthy gaming environment that ensures players enjoy the experience as entertainment, without being financially ruined. Recommendations include: evaluating an early exit (cash out), changing the wager amount, or changing the placement. We address here the early exit recommendation. Barberis (2012) presents an exhaustive search numerical solution to assess exit strategies, using a binary tree. Blavatskyy (2024) mentions that "the design of optimal quitting strategy is non-trivial due to problems of [...] computational complexity." In this section we provide analysis of the exit strategies that will reach the desired exit condition. For two of the wager strategies we provide analytical solutions, and we propose simulation methods to find the optimal exit strategy for other wager schemes. Our approach has considerably less computational complexity than exhaustive search.

Instead of (1), the digital twin will use an early exit strategy:

$$\tau^{\text{def}} = \min(n \leq N : X_n \in \{0, K(s, w)X_0\}), \quad (3)$$

where $K(s, w) > 0$ is a factor depending on the strategy. That is, if the bankroll attains $K(s, w)$ times the initial value, the player is encouraged to quit the game and cash out. In particular, for $w = 1$ and $w = 5$ we can use stochastic models (Ross 2014; Taylor and Karlin 1998) to estimate the expected loss under this new stopping criterion. We seek to calculate a boundary $K(s, w)$ for each strategy, such that the expected loss is limited to a "reasonable" amount. As mentioned in (Schnytzer and Westreich 2010), some gamblers view this as an "entrance fee" that they are happy to pay in order to have fun.

Theorem 1 Consider the placement strategy s with payout ratio $R = R(s)$ and win probability $P(s)$, under the wager strategy $w = 1$ of "all-in". Let

$$N(k) = \min(n : R^n \geq k),$$

and consider the modified stopping criterion $\tau = \min(X_n \in \{0, kX_0\})$. Then the expected loss of the game is given by

$$L(k) = X_0 - X_0 \left(R^{N(k)} P(s)^{N(k)} \right). \quad (4)$$

Proof. Under this wager strategy $Y_n = X_n$, so each round either results in ruin, or it provides a gain of RX_n . Notice that from the definition of $N(k)$ the game ends at time $\tau \leq N(k)$. If the player succeeds n consecutive winning rounds, then $X_n = R^n X_0$. Thus, it is necessary to have exactly $N(k)$ consecutive winning rounds in order not to bankrupt. This event happens with probability $P(s)^{N(k)}$, which yields the result. \square

Theorem 2 Consider the placement strategy s with payout ratio $R = R(s)$ and win probability $P = P(s)$, under the wager strategy $w = 5$ with a “flat” bet of \$1. For the exit condition B , define the modified stopping criterion :

$$\tau = \min(n : X_n \in \{0, B+1\}). \quad (5)$$

Define the coefficients:

$$\begin{aligned} c_0(0) &= 0, & c_1(0) &= 0, \\ c_0(B) &= 0, & c_1(B) &= 1, \\ c_0(i) &= i, & c_1(i) &= 0 \quad i > B. \end{aligned}$$

Then $v_i = \mathbb{E}[X_\tau | X_0 = i]$ satisfies:

$$v_i = c_0(i) + v_B c_1(i), \quad 0 \leq i \leq B+R, \quad (6)$$

where the tuple $c(i)$ satisfies the recursion:

$$c(i) = \frac{c(i+1) - P c(i+R)}{1-P}; \quad 0 < i < B. \quad (7)$$

Finally, $v_B = \frac{P c_0(1) - c_0(1)}{c_1(R) - P c_1(1)}$.

Proof. Using the definition of the stopping time τ in (5), we obtain the following boundary conditions: $v_0 = 0$; and $v_i = i$, for $i > B$, which establishes (6) for those indices. Given this wager strategy, each spin either decreases the bankroll by 1, which happens w.p. $(1-P)$, or it increases it by $R-1$, w.p. P , which yields the following recursion, using first step analysis (see Chapter III of Taylor and Karlin (1998)):

$$v_i = \mathbb{E}[X_\tau | X_0 = i + (R-1)] P + \mathbb{E}[X_\tau | X_0 = i-1] (1-P); \quad i = 1, \dots, B.$$

This yields:

$$v_i = P \times v_{i+(R-1)} + (1-P) \times v_{i-1}; \quad i = 1, \dots, B. \quad (8)$$

The proof now follows applying mathematical induction, where the induction hypothesis is (6), which is satisfied by the base case when $i \geq B$. From (8) it now follows that for $1 < i < B$:

$$v_{i-1} = \frac{v_i - P v_{i+(R-1)}}{1-P} = \frac{c_0(i) + v_B c_1(i) - P(c_0(i+R-1) + v_B c_1(i+R-1))}{1-P},$$

which establishes (6) for all $i = 0, \dots, B+R$. Simple algebra now verifies (7). The end of the proof uses the boundary condition $v_0 =$, so that

$$v_1 = (1-P) v_0 + P v_R = P v_R \implies c_0(1) + v_B c_1(1) = P(c_0(R) + v_B c_1(R)),$$

which yields $v_B = \frac{P c_0(1) - c_0(1)}{c_1(R) - P c_1(1)}$, as claimed. \square

Algorithm 1 Calculation of v_i as a function of B , using Theorem 2.

Read input parameters s, B .
 $P = P(s), R = R(s)$
Initialize:
for ($i = B + 1$ to $B + R - 1$) **do**
 $c_1(i) = 0, c_0(i) = i$
 $c_1(0) = c_0(0) = 0; c_1(B) = 1, c_0(B) = 0.$
Calculate coefficients:
for ($i = B - 1$ **down to** $i = 1$) **do**
 $c(i) = \frac{c(i+1) - P \times c(i+R)}{1 - P}$
 $v_B = \frac{P \times c_0(R) - c_0(1)}{(c_1(1) - P \times c_1(R))}$
for ($i = 1$ to B) **do**
 $v(i) = c_1(i) \times v_B + c_0(i)$
return v

Using the above theorems it is possible to evaluate numerically the bounds $K(1, s)$ and $K(5, s)$ quickly. Table 1 shows the results for the expected loss as a function of $k = K(1, s)$. Table 2 provides the results for the expected loss for $w = 5$ as a function of k , using Algorithm 1 with $B + 1 = kX_0$. Figure 3 shows the results of 10,000 replications of the simulation using (3) with $N = 100$. The plot illustrates the form of the expected loss as a function of the early exit rule for $s = 2$.

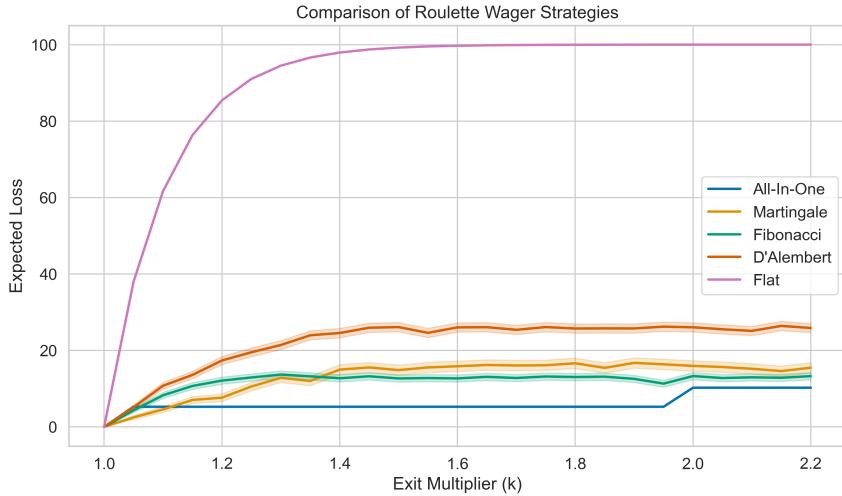
Table 1: Expected loss of theorem 1 ($w = 1$), for $X_0 = 100$.

k	2.00	3.00	4.00	5.00	6.00
s=1	5.2631	5.2631	5.2631	5.2631	5.2631
s=2	5.2631	5.2631	10.2493	10.2493	10.2493
s=3	5.2631	10.2493	10.2493	14.9730	14.9730

Table 2: Expected loss of theorem 2 ($w = 5$), for $X_0 = 100$.

k	1.1	1.2	1.3	1.4	1.5
s=1	3.8259	4.8918	6.4442	8.1816	9.5478
s=2	36.413	59.2803	74.091	83.6072	89.679
s=3	61.646	85.4111	94.4893	97.9307	99.227

The initial page of the GUI prompts the player to fill out info such as initial amount of money, chosen strategy, and if there is a reasonable loss that the player is ready to accept (see Figure 2). With this, the digital twin creates the first simulations to calculate the corresponding exit rule $K(s, w)$. Because the loss is monotone non decreasing, search for the value can be achieved fast using binary search. For $w = 1$ or $w = 5$ the values of the loss function $L(k)$ can be calculated very efficiently for each of the values of k used during the search. For other strategies analytic solutions are elusive and we recur to simulations in order to estimate the corresponding bounds. These simulations have been designed using common random numbers combined with stochastic binary search in order to find the value of k that attains the desired expected loss. Such search method follows the ideas developed in Chakraborty, Das, and Magdon-Ismail (2011) and Vázquez-Abad and Fenn (2016).

Figure 3: The simulation result for $s = 2$ with 95% confidence intervals.

4 MODELING BEHAVIOR

In the previous section we developed the “optimal” exit strategy that ensures a certain acceptable level for the expected loss. However, addictive behavior may cause the player not to follow the recommendations of the digital twin. This requires another level of *learning* in order to adjust the recommendations to the “most likely option” to be accepted. That is, the model becomes a model of optimization (to decrease losses) under constraints (psychological profile).

The psychology of gambling has been described in various scholarly treatises, and it encompasses a variety of different mindsets. People gamble for different reasons: enjoyment, excitement, socializing, to impress the crowd, etc. Depending on the person, he/she may be more (or less) likely to exit the game when the digital twin prompts the recommendation. Note that, in this research, we only address those who gamble for entertainment. We do not address pathological gamblers: we believe that those who are deeply affected should seek professional treatment, rather than computer-guided apps, no matter how “intelligent” they are. People who show high addictive behavior will not be very likely to accept an early exit, for example.

Most individuals who see gambling as an entertaining diversion will be more likely to accept the digital twin’s recommendations. Acceptance of a suggested reduced wager or a lower risk placement may have higher probability of acceptance by a casual gambler than by the more addictive ones. The digital twin initially calculates the early exit strategy and proposes this as a recommendation. If the player does not accept it, then the digital twin conducts fast simulations considering alternative recommendations. As well, the estimated user profile is updated.

Hales, Clark, and Winstanley (2023) discusses reasons why people keep gambling. Among these, we will consider:

- Some people believe that they can actually beat the odds and win money. A key goal of the digital twin will be to convince players that a) their initially planned exiting strategy will lead to loss, and b) the alternative strategy can still involve play (“staying in the game”), while mitigating losses. This approach directly addresses what is known as the “gambler’s fallacy” (Kong, Granic, Lambert, and Teo 2020; Rao and Hastie 2023), whereby players mistakenly believe their past gambling outcomes have implications for future outcomes.
- People may keep gambling even after losses because they crave special attention and other rewards that they receive from other people. To address this, another goal of the digital twin will be to

convince players to reframe their external perceptions of others' expectations of their performance (i.e., social incentives, see Russell, Langham, and Hing (2018)).

- Finally, after losing, some players may believe that they have no choice but to keep playing in order to recover their financial loss. This pattern, known as the “sunk cost fallacy,” will require the digital twin to convince players to reframe their internal perceptions by pulling back the lens to confront the implications of their losses (Doerflinger, Martiny-Huenger, and Gollwitzer 2023).

Based on this model, the digital twin recommendations follow two guidelines:

- The actual recommendation rank: (a) strong (early exit), (b) moderate (change placement strategy), or (c) mild (decrease wager amount or change the wager strategy).
- The message that is sent along with the recommendation. These messages will be based on long-standing psychological research (Priester and Petty 1995) on the factors that make them more likely to influence individuals, particularly when they come from a source that has either credibility (i.e., seen as trustworthy and knowledgeable), attractiveness (i.e., through adherence to a social norm) or power (i.e., able to administer rewards or punishments based on message adherence (Kelman 1958)). For the digital twin, this will involve describing its recommendations as coming from a source that reflects either (a) credibility (by sending the player results from the simulation scenarios and perhaps statistical data to support the recommendation), (b) a social norm (by sending the player messages that will reinforce social recognition from other players for more responsible behavior), or (c) reward/punishment (by sending the player warnings about mounting debt and information about the difficulty of paying off debt).

We propose to model categories of player's profiles that have labels A , denoting their level of “addiction” (as measured by the Problem Gambling Severity Index, described in the next paragraph) and M , denoting the type of message that they will respond positively to. Given a category (A, M) , the probability of acceptance of a recommendation/message pair is used by the digital twin in order to minimize loss, subject to the estimated category of player. The digital twin keeps an estimate of the probability of acceptance of the pair of recommendation and message that it will send. When the player responds by either accepting or rejecting recommendations, the digital twin will update its estimated guess of the player's profile.

Some empirical pre-testing will be needed to gauge the effectiveness of the messages presented by the digital twin, particularly to ensure they psychologically have the intended impact of credibility, normative influence, and power. An experimental design would first ask participants to complete the Problem Gambling Severity Index (Ferris and Wynne 2001) to identify the degree of any gambling addiction issues. Participants would then be randomly assigned to conditions varying in the amount of gambling losses incurred over multiple trials, with the digital twin triggered to intervene after a set level of gambling losses. This experimental approach will help to identify a ground truth benchmark to estimate the acceptance probabilities given the degree of losses incurred. A unique psychological contribution of this proposed study (apart from the use of the digital twin) will be examining the interaction between level of addictiveness and gambling loss, on one hand, and type and source of persuasive message to see where each has maximum impact. Because this goal will require much experimentation and additional sources of support, it falls outside the scope of this first paper, which focuses on producing the simulation engine and corresponding GUI that will enable the app to interact with the players.

5 SIMULATION ENGINE

The main simulation engine for each round is simply a Bernoulli trial, given the strategy. The Markov process $\{\xi_n\}$ is simulated to estimate the final bankroll. Because simulations are used by the digital twin, various scenarios must be considered in order to estimate future options for the recommendations. To achieve efficiency in the comparison, we use common random numbers and parallel computations. This

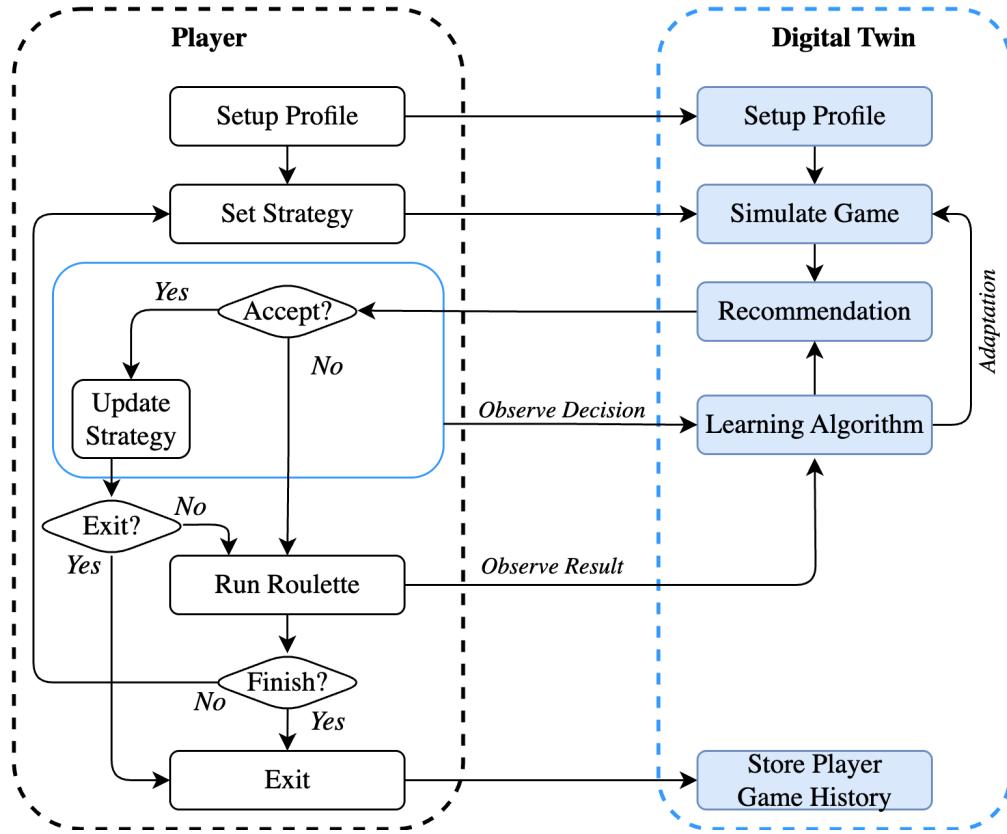


Figure 4: Activity diagram of the American Roulette.

way the digital twin, with an estimated probability of acceptance from the player's estimated profile, can adapt consecutive recommendations/messages. Figure 4 shows the workflow of the simulation engine.

6 CONCLUDING REMARKS

This research presents the development of a simulation engine to empower a digital twin to help mitigate addictive behavior. The benchmark model used is that of American roulette games, but the main concepts can be adapted to other games or addictions. In future we plan to conduct controlled psychological experiments to determine the appropriate model for the player's profiles and test the effectiveness of the proposed digital twin. At that time the digital twin should learn the player's profile. We will compare three methods for learning the category: (a) multi-armed bandits (reinforcement learning), (b) recommender systems methods, and (c) classifier methods.

7 ACKNOWLEDGEMENTS

The pilot simulation study was carried out as part of a graduate course project. We thank the team members Eungkyung Lee, Rida Sohail, Dylan Domínguez and Perla Escaño for their contribution to that work.

REFERENCES

Barberis, N. 2012. "A Model of Casino Gambling". *Management Science* 58(1):35–51.
 Barricelli, B. R., E. Casiraghi, and D. Fogli. 2019. "A Survey on Digital Twin: Definitions, Characteristics, Applications, and Design Implications". *IEEE access* 7:167653–167671.
 Blavatskyy, P. 2024. "A Puzzle of Roulette Gambling". *Journal of Risk and Uncertainty* 69(2):219–234.

Chakraborty, M., S. Das, and M. Magdon-Ismail. 2011. "Near-Optimal Target Learning With Stochastic Binary Signals". In *Proceedings of the 2011 Conference on Uncertainty in Artificial Intelligence*, edited by F. G. Cozman and A. Pfeffer, 69–76. Corvallis, Oregon: Association for Uncertainty in Artificial Intelligence Press.

Coyle, S. 2023. "How to Win at Roulette". <https://www.pokernews.com/casino/roulette/win-at-roulette.htm#martingale-roulette>. Accessed 2024-12-16.

Doerflinger, J. T., T. Martiny-Huenger, and P. M. Gollwitzer. 2023. "Exploring the Determinants of Reinvestment Decisions: Sense of Personal Responsibility, Preferences, and Loss Framing". *Frontiers in Psychology* 13:1025181.

Ferris, J. A., and H. J. Wynne. 2001. *The Canadian Problem Gambling Index: User manual*. Canadian Centre on Substance Abuse.

Grindu, V. 2025. "Winning Roulette Strategy: Best Betting Strategies, Tips & Tricks". <https://www.techopedia.com/how-to/roulette-strategy>. Accessed: 31 March 2025.

Hales, C., L. Clark, and C. Winstanley. 2023. "Computational Approaches to Modeling Gambling Behaviour: Opportunities for Understanding Disordered Gambling". *Neuroscience & Biobehavioral Reviews* 147:105083 <https://doi.org/https://doi.org/10.1016/j.neubiorev.2023.105083>.

Hoofe, Liam 2025. "Best Roulette Strategies: How to Win at Roulette". <https://www.gamblingzone.com/uk/roulette/strategy>. Accessed: 31 March 2025.

Kelman, H. C. 1958. "Compliance, Identification, and Internalization Three Processes of Attitude Change". *Journal of conflict resolution* 2(1):51–60.

Kong, Q., G. D. Granic, N. S. Lambert, and C. P. Teo. 2020. "Judgment Error in Lottery Play: When the Hot Hand Meets the Gambler's Fallacy". *Management Science* 66(2):844–862.

Priester, J. R., and R. E. Petty. 1995. "Source Attributions and Persuasion: Perceived Honesty As a Determinant of Message Scrutiny". *Personality and Social Psychology Bulletin* 21(6):637–654.

Rao, K., and R. Hastie. 2023. "Predicting Outcomes In a Sequence of Binary Events: Belief Updating and Gambler's Fallacy Reasoning". *Cognitive Science* 47(1):e13211.

Ross, S. M. 2014. *Introduction to Probability Models*. Academic press.

Russell, A. M., E. Langham, and N. Hing. 2018. "Social Influences Normalize Gambling-Related Harm Among Higher Risk Gamblers". *Journal of Behavioral Addictions* 7(4):1100–1111.

Schnytzer, A., and S. Westreich. 2010. "Attitudes to Risk and Roulette". Technical report, Working Paper.

Sohn, E. 2023. "How Gambling Affects the Brain and Who is Most Vulnerable to Addiction". *APA Monitor* 5:62–63.

Taylor, H. E., and S. Karlin. 1998. *Introduction to Stochastic Modeling*. 3rd ed. Academic Press.

Vázquez-Abad, F. J., and L. Fenn. 2016. "Mixed Optimization for Constrained Resource Allocation, an Application to a Local Bus Service". In *2016 Winter Simulation Conference (WSC)*, 871–882. IEEE.

AUTHOR BIOGRAPHIES

FELISA J. VÁZQUEZ-ABAD is Professor of Computer Science (CUNY). She is interested in stochastic optimization and computer simulation of complex systems under uncertainty. She obtained a B.Sc. in Physics (1983) and a M.Sc. in Statistics and Operations Research (1984) from UNAM (Mexico). In 1989 she obtained a Ph.D. in Applied Mathematics from Brown University. She was postdoctoral researcher at the INRS-Telecom in Montreal, Canada from 1990 to 1993. She was a professor at the University of Montreal (1993-2004), and then a professor at the University of Melbourne until 2009 when she joined the CUNY Faculty. In 2000, she was a recipient of the Jacob Wolfowitz award for advances in the mathematical and management sciences. Her team was a finalist in the MEDSTART competition (2014) and her team got third prize in the CUNY-IBM Watson Competition (2017). She has participated in Grant Selection Committees and has been Associate Editor for IEEE Transactions on Automatic Control, Management Science, and OR Letters, Area Editor of the ACM TOMACS, and web editor of the INFORMS College on Simulation. She actively participates in events and programs to encourage women and minority students to succeed in Computer Science. Her email address is felisav@hunter.cuny.edu.

JASON YOUNG is Professor of Social Psychology at Hunter College (CUNY), where he has been since 1990. His research focuses on the influence of emotion of risky decision-making as applied to a number of real-world contexts, including public perceptions of fear of crime, the impact of math anxiety on financial literacy and behavior, and the role of belief resistance in pedagogy. He obtained his undergraduate degree in political science and psychology at Indiana University, Bloomington, and his PhD in Social Psychology, with a minor in statistics, at the University of Minnesota, Minneapolis. His email address is jason.young@hunter.cuny.edu.

SILVANO BERNABEL received his B.Sc. in Electrical Engineering from the City College of the City University of New York (CUNY) and his M.Sc. in Pure Mathematics from Hunter College, CUNY. He is currently pursuing a Ph.D. in Computer

Vázquez-Abad, Young, and Bernabel

Science at the Graduate Center, CUNY, where his research bridges Artificial Intelligence, Machine Learning, and Computer Vision with simulation-based methods to train intelligent agents. His email address is sbernabel@gradcenter.cuny.edu.