

Near-random behavior in strategic games

Tal Boger¹, Dorsa Amir², Sami R. Yousif³

¹ *Johns Hopkins University*

² *Duke University*

³ *Ohio State University*

tboger1@jhu.edu, dorsa.amir@duke.edu, yousif.36@osu.edu

Abstract:

Humans struggle to produce random behavior when asked to do so, leading some to suggest that our ability to randomize is attenuated (or even absent). Here, we demonstrate that the capacity for random behavior is much greater than previously acknowledged: Humans can behave in (near) optimally random ways when the context demands it. We created a custom, online version of the game Battleship in which participants competed against one another. On each turn, one player chose a tile on a 3x3 grid to place a ship, while the other player tried to guess the ship's location. Across six preregistered studies, this strategic context led to dramatic increases in random behavior relative to closely matched non-strategic tasks. Such random behavior was also remarkably flexible and optimal depending on context; both players readily incorporated varying information about the game (e.g., that tiles can differ in value) while simultaneously accounting for their opponents' knowledge of that information. A computational model validated these findings and revealed nuanced differences across contexts. Contrary to decades of work finding that humans fail to randomize, we find that strategic games promote highly flexible, near-random behavior.

Data and materials: <https://osf.io/z928n/>

Version: September 28, 2026; in press, *Cognition*.

Main text

In the forests of northern Labrador, a Naskapi man carefully holds the shoulder bone of a caribou over hot coals. As the bone heats, cracks begin to appear. To the untrained eye, these cracks seem meaningless. But to the Naskapi, they form something much more useful: a map, pointing the way to their next hunt (Speck, 1977). Why have such rituals proliferated across a wide set of cultures reliant on hunting? In the mid-20th century, anthropologist O.K. Moore proposed an intriguing idea: Perhaps practices like scapulimancy succeed, counterintuitively, because they are random (Moore, 1957).

Randomness is a foundational feature of biological and cultural systems. It is a primary engine of evolutionary adaptation (Bonner, 2013), inherent in both basic processes like genetic recombination and complex phenomena such as anti-coordination between predator and prey (Driver & Humphries, 1988). Just as randomness plays a vital role in shaping natural and social processes, so too it underlies essential functions in the mind. Appreciating the presence — and absence — of randomness entails separating signal from noise, thus enabling efficient learning, decision-making, perception, memory, and more (e.g., Saffran et al., 1996, 1999; Newman et al., 2010; for review, see Icard, 2021). Indeed, even young children and infants intuitively appreciate the different kinds of causes that can bring about randomness as opposed to order (e.g., intentional action vs. incidental side effect; Newman et al., 2010). The importance of randomness applies not only to its perception but also to its generation — we rely on our ability to produce randomness in settings ranging from sports (e.g., choosing an unpredictable direction for a penalty kick) to international relations (e.g., crafting unpredictable attacks on adversaries). Further, many cognitive models incorporate randomness or noise (e.g., Sanborn et al., 2025) and some theorize that the ability to generate randomness provides several benefits to any organism with bounded memory (e.g., Icard, 2021).

Yet, decades of research suggest humans notoriously struggle with randomness, falling prey to numerous biases when both producing and perceiving random sequences. People expect random sequences to be locally representative (Kahneman & Tversky, 1972); they judge sequences with a large number of ‘alternations’ (e.g., from heads to tails, or vice versa) to be more random; they quickly cycle through all possible options when generating random sequences; and more (for review, see Wagenaar, 1972; Bar-Hillel & Wagenaar, 1991; Nickerson, 2002; Zhao et al., 2014; Boger et al., 2025). Such biases often carry everyday consequences for our decision-making, as in the case of ‘hot hand fallacy’ or the ‘gambler’s fallacy’ (wherein one believes that random events should be sequentially dependent; Gilovich, 1983; Gilovich et al., 1985; Ayton & Fischer, 2004); and they can even influence other aspects of cognition, such as number estimation (Zhao & Yu, 2016).

How is it that randomness is so prevalent in nature and so fundamental to cognition, yet so difficult for humans to accurately produce and perceive? This question is surprisingly tricky, with researchers proposing several different accounts for our apparent failures to randomize. For example, some suggest that humans' inability to produce and perceive randomness stems from a general tendency to *overattribute* meaning, causes, and intentions to patterns. In other words, many of the same abilities that led to our general success as a species — e.g., sophisticated prosocial behavior, causal reasoning, mentalizing, and imitation — may impede our ability to act randomly (e.g., Henrich, 2015). This tendency to infer hidden causal structure applies both when judging randomness (e.g., when a random streak appears patterned, as in the gambler's and hot hand fallacies) and when producing randomness (e.g., when monitoring streak-like patterns in one's own random behavior in ways that abide by the same fallacies). Furthermore, this logic extends naturally to social contexts. In competitive interactions, the same cognitive machinery that supports mentalizing and strategic reasoning may interfere with the production of randomization. For example, recursive modeling of an opponent's mental states, even at shallow depth (cf. Hedden & Zhang, 2002; Camerer et al., 2004), could inadvertently introduce systematic structure into ostensibly random sequences. Consistent with this possibility, players of rock-paper-scissors often exhibit automatic imitation of their opponents' gestures (Cook et al., 2012), suggesting that our social cognitive biases can hamper randomness production by introducing predictable patterns, even when randomness is strategically advantageous. This view, therefore, suggests that the addition of a social opponent, or the demand to represent that opponent's knowledge state, may compound failures of randomization.

Other — not mutually exclusive — hypotheses implicate other cognitive processes in our apparent failures to randomize. For instance, some have suggested that *resource limitations* (e.g., in working memory) may lead to poor randomization. The basic suggestion is that people fail to randomize because they cannot hold all of their previous choices in their minds at once, making it difficult to produce uniform random sequences over long runs (i.e., beyond the limits of working memory). In other words, randomness is informationally dense (in some sense, definitionally), making random sequences hard to produce and perceive — whether because of limits in our ability to hold information in mind, or even limits in our executive function (Baddeley, 1966). Indeed, prior work documents strong links between memory capacity and randomization ability (Biesaga & Nowak, 2024; see also Icard, 2021); and some argue that certain 'biases' of randomness reflect limits of memory rather than a poor capacity for behaving randomly per se (e.g., Kareev, 1992; Falk & Konold, 1997; Hahn & Warren, 2009; Warren et al., 2018). That said, other researchers propose the *opposite* relationship between randomness and memory: An agent with no memory whatsoever should (in theory) be perfectly random, insofar as they cannot be biased by their prior decisions (Falk & Konold, 1997).

A separate line of work offers various Bayesian or sampling-based theories of our failures to randomize. For example, some researchers account for biases in randomness via appeal to an

underlying statistical inference about the generative process that produces (or, in the case of judgments, produced) a random output (Griffiths & Tenenbaum, 2003; Griffiths et al., 2018). These accounts focus on the posterior odds of random vs. non-random sequences, with the idea being that people generate random-like sequences that maximize these odds under certain hypotheses of what non-random sequences look like. Other accounts appeal to a more general sampling process (which may itself incorporate some bias) that the mind draws from to produce randomness (e.g., Castillo et al., 2026).

The common theme across all of these theories, however, is that humans are poor randomizers. But is this true? An alternative possibility might be that randomness is more context-sensitive than previously appreciated. Perhaps it is not the case that humans are bad randomizers writ large, but rather that humans are *strategic randomizers*: They can produce random (or near-random) behavior when the situation demands it. A clue hinting at the promise of this hypothesis lies in the methodologies typically used to elicit randomness (see Ayton et al., 1989; Nickerson, 2002). A common way to probe random behavior is to simply demand it from people in those terms — i.e., asking naive participants to generate long sequences of random numbers. As an extreme example, one study asked subjects to generate a string of 2,520 random numbers — a task so "difficult and boring" that its own creator admitted having to instruct participants not to "dawdle" (Chapanis, 1995; p. 1350). While such studies have uncovered important and recurrent biases in our attempts to produce randomness, perhaps they fail to simulate the real-world contexts that incentivize truly random behavior. Indeed, a small set of studies employing strategic games found that those contexts facilitate improvements in random behavior (e.g., Rapoport & Budescu, 1992; Budescu & Rapoport, 1994; Mookherjee & Sopher, 1994), but much of this work relies on simple games (e.g., matching pennies, which involves a binary choice), raising questions about how these findings generalize to more naturalistic contexts and how they situate random behavior in cognition more broadly. Is it possible that people behave in nearly optimal, random ways when the context demands it?

The present work: Strategic random behavior

Here, we test people's capacity for *strategic random* behavior — randomness generated to achieve a specific goal in a competitive or uncertain environment.

We created a custom, online version of the game Battleship to probe strategic random behavior (**Figure 1A**). Players were matched with live opponents, and, on each turn, the first player chose a tile on a 3x3 grid to place a ship, while the other player tried to guess the ship's location. Unlike typical studies of randomness which simply prompt participants to produce random behavior (**Figure 1B**), our task makes randomness instrumentally important — much like it is to animals, people, and cultures in naturalistic settings.




Figure 1. (A) We created a novel, online Battleship game in which participants play against each other live. The first player (the generator) tries to place the ship on a random tile, so as to prevent the second player (the guesser) from locating (‘sinking’) the ship within 3 guesses. Thus, the task requires strategy from both players. (B) This comes in stark contrast to traditional ways of studying randomness, where random behavior has no instrumental goal and subjects are merely instructed to behave randomly. Readers can experience all tasks — including our live, online games — at http://perceptionstudies.github.io/strategic_randomness.

Our approach builds on previous work (particularly in behavioral economics) suggesting that interactive games (e.g., matching pennies) spur improvements in random behavior (Rapoport & Budescu, 1992; Budescu & Rapoport, 1994; Mookherjee & Sopher, 1994). In other words, we are not the first to examine random behavior in games; substantive work examines performance relative to game-theoretic equilibria in rock-paper-scissors (Batzilis et al., 2019; Wong et al., 2021; Brockbank & Vul, 2024), prisoner’s dilemmas (Baker & Rachlin, 2001), and more (with recent work even examining children’s performance in hide-and-seek games; Brocas & Carrillo, 2022).

However, our work goes beyond previous studies of random behavior in strategic games (most proximally, by Rapoport & Budescu, 1992; Budescu & Rapoport, 1994) in several key ways. First, Rapoport and Budescu (1992) and Budescu and Rapoport (1994) test random behavior in simple games with small response spaces (specifically, only 2 or 3 possible responses, compared to our 9 possible responses; this also applies to rock-paper-scissors, as in Batzilis et al., 2019, Wong et al., 2021, or Brockbank & Vul, 2024, and to prisoner’s dilemma games, as in Baker & Rachlin, 2001). While the size of the response space may seem inconsequential, the richness of possible random behavior scales exponentially with the number of possible responses. With 2 possible responses, the main (or even only) emergent properties of a random sequence are its (n -order) probability of alternation and its runs; meanwhile, larger response spaces open the door to much more fine-grained analyses of random behavior. Indeed, Budescu and Rapoport (1994) note that even the jump from 2 possible responses (as in Rapoport & Budescu, 1992) to 3 possible responses results in a “strategically more demanding game”. Second, the non-strategic tasks in Rapoport and Budescu (1992) and Budescu and Rapoport (1994) only emulate typical

single-person subjective randomization experiments (i.e., the kind of task where subjects are simply instructed to generate random numbers; see also Wong et al., 2021) — and other studies (e.g., Batzilis et al., 2019; Baker & Rachlin, 2001) do not employ a non-strategic condition at all. We bridge the gap between a fully strategic game and a classic single-person task with an intermediate game, in which participants play against a non-strategic computer opponent but still receive feedback on their choices. This allows us to examine the ‘gradient’ of strategic random behavior. Third, the use of this richer response space allows us to model complex random behavior in our tasks, permitting analyses of individual strategies for (and components of) random behavior. Specifically, in each study, a computational model of human random behavior validated our key findings and revealed nuanced patterns of behavior.

In addition to these advances, we test the flexibility of strategic random behavior in a variety of ways, exploring whether (and how) participants update their behavior in line with changes in the task environment, such as the value of the tiles and their opponent’s knowledge of those tile values. These changes greatly increase the complexity of the task, insofar as participants must integrate multiple sources of information simultaneously (e.g., about their own prior choices, their opponent’s prior choices, the tile values themselves, and their opponent’s knowledge of those values) in ways that may require recursive theory of mind. Thus, these variants of the task allow us to assess the extent to which, for instance, modeling an opponent’s mental states or expending additional cognitive resources may (or may not) impair random behavior.

To summarize, Study 1 tested strategic random behavior in a live, multi-player Battleship game. Study 2 modified this game, replacing the human guesser with a computer that acts totally randomly; participants played only as the generator, and they were informed their opponent had no memory or strategy. This manipulation offers a ‘bridge’ between the live, two-person game and a non-strategic, one-person game. Study 3 tested the predictability of the non-strategic sequences from Boger et al. (2025) and the sequences collected in Study 2. Here, participants played as the guesser ‘against’ a sequence from these previous tasks. Together, these first three studies revealed a strategic gradient of random behavior, characterizing a difference in random behavior between strategic contexts and closely matched non-strategic contexts.

Three subsequent studies examined the flexibility and optimality of such near-random behavior. In these studies, different tiles provided different rewards for the generator; in other words, one tile might provide twice as many points to the generator as another tile (as long as the generator’s ship survives the turn). In Study 4, both the generator and guesser knew the values of each tile. In Study 5, we hid the tile values from the guessers (and informed the generators about this knowledge asymmetry); and in Study 6, we further masked the tile values from guessers by reassigning (i.e., shuffling) tile values periodically throughout the task. This set of studies probes how effectively participants update their random behavior in line with both varying incentives and varying knowledge states about those incentives.

General methods

Open science practices

We preregistered our sample sizes, exclusion criteria, analysis plans, and task designs for each study. Preregistrations — along with all the code needed to reproduce our tasks and results — are available on our OSF repository (<https://osf.io/z928n/>). Readers can experience all of our tasks exactly as participants did — including the live, online games — at http://perceptionstudies.github.io/strategic_randomness.

All studies

Participants

All participants were recruited via the online platform Prolific. Participants were compensated upon completing the study (with a flat fee based on the estimated time needed to complete the study, at a rate of \$10/hour). No individual could participate in more than one of the tasks reported below. Sample sizes and exclusion criteria were preregistered, and both are reported in the sections that follow for each study.

Computational model

Throughout this work, we use (and adapt) a computational model of random behavior developed by Angelike and Musch (2025). The code we used to run this model is available on our OSF repository.

The model contains 3 parameters: (1) a stay parameter (herein, α), (2) a distance parameter (β) and (3) a cycling parameter (γ). These parameters are fit to each participant's data; they are added sequentially to a uniform random distribution (over the 9 possible choices). Let S be the possible response space (consisting of 9 options); let t be the current trial (thus, for example, $t + 1$ is the next trial); and let c^t be the choice on trial t . Thus, $P(S)^{t+1}$ denotes the probability of each choice in the response space being selected in the next trial.

The model starts with a uniform distribution over the possible choices: $P(S)^{t+1} = U(S) = \frac{1}{9}$.

The stay parameter (α) adjusts the probability of the most recently selected option. It adjusts only the probability of that option. For readability in the final equation presented below, let $\vec{r}$ be a one-hot vector, with the value in the c^t position set to 1, and all other values set to 0. Thus, $\alpha \times \vec{r}$ adjusts the probability of only the most recently selected option.

The distance parameter (β) adjusts the probability of each option in the response space as a function of its distance from the most recently selected choice. Consequently, the first step to incorporating the distance parameter in the model requires computing a distance vector ($\vec{d}$), which encodes the distance between the most recent choice and each other option in the response space. Given that the response format in our tasks is a 3x3 grid and that participants responded with their mouse (which permits diagonal movement), $\vec{d}$ encodes Euclidean distances between the relevant points. The vector $\vec{d}$ is multiplied by the distance parameter β .

Finally, the cycling parameter (γ) adjusts the probability of each choice in the response space as a function of how recently it was selected. Here, a vector $\vec{c}$ encodes the number of trials that have passed since a given response was selected. (So, for example, if the top-left tile in the grid has not been selected for 12 trials, the first element of $\vec{c}$ would be 12.) As Angelike and Musch (2025) note, log-scaling this vector provides the best and most stable approximation of human cycling behavior. The log-scaled vector $\vec{c}$ is multiplied by the cycling parameter γ .

Putting these steps together, we see the following:

$$M: P(S)^{t+1} = \frac{1}{9} + \alpha \times \vec{r} + \beta \times \vec{d} + \gamma \times \log_2(\vec{c})$$

This model was fit to each participant's data using maximum likelihood estimation (implemented via a nonlinear program solver from the MATLAB optimization toolbox). All parameters were bounded between -1 and 1. The result of the above equation was softmaxed, thus transforming $P(S)^{t+1}$ from activations into a probability distribution that summed to 1. Note that softmax is translation invariant, so of course the $\frac{1}{9}$ term could be any arbitrary constant (including 0) and the model would produce the same output. However, here we denote it as $\frac{1}{9}$ to build the intuition that when all parameters equal 0, the model predicts uniform randomness over the 9 possible choices.

For Studies 4–6, we introduced a value parameter, δ , which adjusts the probability of each choice in the response space as a function of its value. A value vector $\vec{v}$ encodes the value of each tile on the given trial; $\vec{v}$ is then multiplied by δ . Positive values of δ indicate a propensity to choose the higher value tiles, and negative values of δ indicate a propensity to avoid the higher value tiles. In Studies 4–6, we compare models with this value parameter to models without it.

Exclusions

We apply the same (preregistered) exclusion criteria in every task where participants generate random sequences (Studies 1–2 and 4–6). This includes when we re-analyze data from Experiment 2 of Boger et al. (2025). In each case, we exclude participants who fall outside of the middle 99.99% of uniform random sequences in terms of the frequency of selected choices and the average numerical distance between choices. In other words, to arrive at these exclusion criteria, we simulate several thousand draws of uniformly generated random numbers (with a sequence length equal to the number of trials in the relevant study). We then compute the lower 0.005th and upper 99.995th percentiles for choice frequency (i.e., the lower bound for the least frequently selected choice and the upper bound for the most frequently selected choice) and average numerical distance (Euclidean distance between consecutive choices in the grid).

For a 150-item sequence (as in the data we collected here), the choice frequency criterion excludes sequences where any of the 9 options are selected fewer than 3 times or more than 34 times. The average numerical distance criterion excludes sequences where the average numerical distance is below 1.23 or above 1.71. We excluded participants who failed to meet either of these criteria, as per our preregistrations.

In each study, this resulted in a fair number of exclusions which we further consider here. There are a few important things to note about these exclusion rates and how they relate to our key results. First, these exclusions help ensure that only participants who earnestly complete the task are included in our final sample — and indeed, the exclusion rates are similar to prior work using the same approach (Boger et al., 2025). Second, the exclusion rates on their own are consistent with our hypotheses. Specifically, we hypothesized that live, strategic games spur the strongest random behavior; and indeed, in that task, we observed the lowest exclusion rate (indicating the strongest randomness). Finally, and perhaps most importantly, our Supplementary Materials file demonstrates that the key results hold (and in some cases appear even stronger) when we analyze the data with no exclusions applied. Thus, we feel that although the exclusion rates are high, they fail to explain our key results on their own (especially with regard to the differences across conditions).

Non-strategic data

For our non-strategic data, we use the 2-dimensional random-location-generation data from Experiment 2 of Boger et al. (2025). These data come from participants who were instructed to behave randomly and input responses with their mouse on a 3x3 grid. Thus, the response format is identical to the data we collect here, allowing for direct comparison across tasks.

We re-analyzed the data from Boger et al. (2025) using the exclusion criteria above. However, because those sequences are 250 items long, the exclusion ranges differ slightly. (In this case, the choice frequency criterion excludes sequences where any option is selected fewer than 10 times or more than 50 times; and the average numerical distance criterion excludes sequences where the average numerical distance is not between 1.27 and 1.62.) After exclusions, 119 sequences remained (out of 195 total in the initial data).

Studies 1–3: Random behavior improves along a strategic gradient

Studies 1–3 assessed whether random behavior improves as a function of social context. Here, participants played (a) a live Battleship game against another participant (Study 1), (b) a Battleship game against a random computerized guesser, who was explicitly identified as such (Study 2), and (c) a single-player, non-strategic task (i.e., with no guesser) that uses an identical 3x3 grid response format (from Boger et al., 2025). Study 3 then probed how predictable the sequences from (b) and (c) were to naïve participants by having a new set of participants play as the ‘guessers’ against these sequences.

Methods

Study 1: Live Battleship game

In this game, participants on Prolific were randomly matched with other participants to play against each other. One person was assigned to be the generator, and the other was assigned to be the guesser. The generator placed their ship on one of the tiles in the 3x3 grid on each turn; then, the guesser received 3 attempts to guess the ship’s location. If the guesser hit the ship on any of their three guesses, the turn ended and a new trial began. Guessers were given 3 attempts to identify the ship's location for two main reasons. First, we wanted the game to be as engaging as possible for both players. If the guessers had only one attempt at the correct answer, they would ‘lose’ the vast majority of the time, and thus could become disengaged and stop trying to play the game optimally (which might, in turn, disincentivize the generator from playing optimally). Second, giving the guessers 3 selections allowed for unique strategies, like identifying an entire row/column to target rather than just a single space. Together, these factors ensured that the game remained as competitive as possible.

Point totals were determined such that equilibrium behavior (i.e., both players playing randomly) would result in roughly equal points for both players: Each time the generator’s ship survived, the generator earned +1 point (no matter where the ship was placed). The guesser earned points as a function of how many guesses they needed to sink the generator’s ship. For sinking the ship on the first guess, the guesser earned +4 points; for sinking the ship on the second guess, the

guesser earned +3 points; and for sinking the ship on the third and final guess, the guesser earned +2 points. (This scoring system also applied to Studies 2–3, but differed for Studies 4–6.)

If either player exited the study early, the game ended for the other player too. Thus, we preregistered that we would recruit participants until we reached 100 dyads of full-game participants (pre-exclusion). 29 dyads of participants were excluded due to the preregistered exclusion criteria described above, leaving 71 generator/guesser pairs (see Supplementary Materials for data analysis with no exclusions applied).

Study 2: Computerized Battleship game

In this study, we replaced the live guesser with a computer. This computer had no memory of the generator's previous choices, and chose randomly on each turn; participants were informed of this. Thus, all participants in this task served as the generator. Otherwise, the task looked identical to the live game — i.e., the user interface was the same — with the only other difference being that the computer received just one guess (as opposed to three). We recruited 100 participants, 43 of whom were excluded by the aforementioned criteria.

In theory, one ‘optimal’ strategy in this task might be to select the same tile on each turn, given the ease of doing so. However, only 1 out of 100 participants did this. We think participants largely avoided doing so because (a) they were still instructed to behave randomly and (b) it is possible that, despite being told the opponent was a strategy-free computer, participants still attempted to play seriously (whether because they assumed the computer was intelligent in some way, or out of an intrinsic desire to complete the task correctly).

Study 3: Guessing games

In this study, participants played as the guesser ‘against’ sequences previously generated by participants in the computerized Battleship game (Study 2) and in the non-strategic task (from Boger et al., 2025). Specifically, we selected 10 post-exclusion sequences from each of the computerized and non-strategic tasks at random. Then, participants in this task were assigned to play against one of those randomly selected sequences for the entirety of the task. Otherwise, the task looked identical to the live game. Participants were informed that they were playing against a ‘rehearsed’ version of another real person’s sequence. For each version of the task (guessing vs. computerized sequences and guessing vs. non-strategic sequences), we recruited 100 participants. We excluded participants who did not submit a complete dataset; this excluded 1 participant in the guessing vs. non-strategic sequences task (and 0 in the guessing vs. computerized sequences task).

Results

We first examined the randomness of generators' behaviors by assessing three commonly studied emergent properties in the generated sequences (**Figure 2A**): frequency of repeats, average numerical distance, and frequency of direction switches (see Boger et al., 2025). These properties are common proxies for random behavior (see Towse & Neil, 1998; Angelike & Musch, 2024).

When considering the frequency of repeats — perhaps the most well-documented metric of randomness — we observed significantly closer-to-random behavior in the live game (where participants repeated choices 8.77% of the time) than in the game with the computer opponent (6.53% repeat frequency; $t(125.19) = 2.36, p = 0.02$), which in turn had significantly closer-to-random behavior than the non-strategic task (3.48% repeat frequency; $t(80.65) = 4.18, p < 0.001$). In other words, the strategic context significantly reduced the classic over-alternation bias, wherein people repeat choices less often than expected in random sequences. We did not observe significant differences across the three tasks in terms of their average numerical distance (live mean of 1.50 vs. computerized mean = 1.46, $t(120.95) = 1.85, p = 0.07$; computerized mean of 1.46 vs. non-strategic mean of 1.48, $t(87.55) = 1.31, p = 0.19$). For direction switches, we observed a significant difference between the strategic game (mean direction switch frequency of 66.40%) and the computerized task (mean of 63.05%; $t(93.68) = 2.33, p = 0.02$), which in turn did not significantly differ from the non-strategic task (mean of 64.8%; $t(107.87) = 1.20, p = 0.23$). In either case, note that direction switches contain computational oddities when computed in a 3x3 grid (see discussion in Boger et al., 2025). Importantly, these metrics offer a small subset of the essentially infinite space of ways to analyze the randomness of a sequence (see, e.g., Angelike & Musch, 2024). For example, we analyzed the Lempel-Ziv compressibility of the sequences. For this metric, too, we observed numerically closer-to-random sequences in the live game than in the computerized game ($t(117.41) = 1.43, p = 0.16$); the computerized game displayed significantly closer-to-random sequences than the non-strategic task ($t(112.13) = 20.92, p < 0.001$). That said, we do not claim that every possible metric of randomness will favor the strategic games; it remains entirely possible that other, subtle patterns of random behavior remain relatively constant across our conditions.

Next, we analyzed human guessing performance on the three tasks — i.e., from the live guessers in Study 1, and from the guessers who played against the computerized sequences and non-strategic sequences in Study 3. Guessers performed significantly above-chance against both the computerized sequences (35.36% vs. chance-level performance of 33.33%; $t(99) = 3.96, p < 0.001$) and the non-strategic sequences (35.32%; $t(98) = 4.26, p < 0.001$). However, in the live game, guessers performed indistinguishably from chance (33.63%; $t(70) = 0.39, p = 0.70$; **Figure 2B**). Thus, in a strategic context, people generated sequences that were more random not only in terms of their emergent properties but also in terms of their predictability to opponents.

To decompose the observed patterns into interpretable behavioral signatures, we implemented a computational model from Angelike & Musch (2025; itself a variation of a prior model, Boger et al., 2025). (For additional definitions and details, see *Computational model* above.) This modeling approach allows us to identify the latent, constituent dimensions of (non-)random behavior across an entire individual's sequence — rather than looking only at correspondences between a selection and its immediate neighbors. The model contains three parameters: (1) a 'stay' parameter, which adjusts the probability of choosing the most recently selected choice; (2) a 'distance' parameter, which adjusts the probability of choosing nearby (or far-away) choices in the response space; and (3) a 'cycling' parameter, which adjusts the probability of choosing an option as a function of when it was last selected. Uniform random behavior requires that all parameters are equal to 0, making every choice in the response space equally likely.

The model revealed several nuanced signatures of near-random behavior, beyond what the model-free metrics revealed (**Figure 2C**). There was no difference in the stay parameter across the three tasks (live mean of -0.21 vs. computerized mean of -0.36, $t(112.5) = 1.26$, $p = 0.21$; computerized mean of -0.36 vs. non-strategic mean of -0.35, $t(107.01) = 0.14$, $p = 0.89$). This may initially be surprising given the difference in the frequency of repeats across the tasks. However, this is because the parameters of the model relate to one another in interesting ways. For example, repeating the same choice very often will result in not only a high stay parameter but also a low distance parameter. The burden of modeling a given individual's sequence is thus shared by multiple parameters, making it possible to observe an apparent mismatch between an emergent property and the model parameter that seemingly captures the same behavior. (In other words, emergent properties are computed independently, whereas model parameters interact in predicting outputs.) The other two parameters, however, revealed significantly closer-to-random behavior in the live game (distance live mean of 0.05 vs. computerized mean of -0.19, $t(94.02) = 3.77$, $p < 0.001$; distance computerized mean of -0.19 vs. non-strategic mean of -0.33, $t(97.48) = 2.36$, $p = 0.02$; cycling live mean of 0.04 vs. computerized mean of 0.20, $t(125.19) = 4.63$, $p < 0.001$; cycling computerized mean of 0.20 vs. non-strategic mean of 0.62, $t(154.42) = 12.66$, $p < 0.001$). Additionally, the model's accuracy in predicting participants' next choices revealed differences absent in the human guessing task. The model's predictions were above-chance in all three tasks (all $ps < .001$ relative to chance-level accuracy of 33.33%), but were highest in the non-strategic task (i.e., sequences in the non-strategic task were most predictable; non-strategic mean of 59.00% vs. computerized mean of 45.58%, $t(139.51) = 10.31$, $p < 0.001$; computerized mean of 45.58% vs. live mean of 43.06%, $t(106.91) = 2.11$, $p = 0.04$; **Figure 2D**).




Figure 2. (A) Emergent properties in the random sequences generated by humans in the live Battleship game against an opponent ('Live'), an identical Battleship game where the real opponent is replaced with a computer ('Computer'), and a non-strategic task (from Boger et al., 2025; 'Non-strategic'). (B) Guesser top-3 accuracy (i.e., guesser accuracy when given 3 chances to sink the ship) on each sequence type. (C) Parameters from a model fit to the sequences generated by participants in each of the three tasks. (D) Model top-3 accuracy on each sequence type. Notice that (in general) behavior is closest to random in the live game and farthest from it in the non-strategic task. Dotted lines indicate the expected value in a uniform random sequence (e.g., a uniform random sequence of 9 options will repeat choices around 11.11% of the time); solid lines indicate the mean of each distribution.

These studies demonstrate that people behave more randomly in strategic games than equivalent, non-strategic tasks. But is it fair to say that people's behavior is 'near-random'? It is possible, of course, that despite behaving 'more randomly', people still deviate from uniform randomness in substantial and systematic ways. To address this possibility, we compared the metrics and model parameters described above to uniform randomness. First, we did this with frequentist tests, which revealed departures from randomness in terms of the frequency of repeats (8.77% vs. 11.11% expected; $t(70) = 3.43, p < 0.01$) and average numerical distance (1.50 vs. 1.45 expected; $t(70) = 2.87, p < 0.01$), but not in terms of the number of direction switches (99.61 vs. 98.35 expected, or 66.40% vs. 65.57%; $t(70) = 1.13, p = 0.26$; though note the computational instability of direction switches in 3x3 grids, as in Boger et al., 2025). Note, however, that the 'departures' from randomness in the former cases are not drastic; to better understand this, we asked what percent of uniform random sequences (out of 10,000 generated sequences) contain values more extreme (i.e., less 'random', in a sense, regardless of direction) than what we observed in our data. Perfectly random values (e.g., if people repeated at a frequency of exactly

1/9) would result in 100% of uniform random sequences being more extreme, because one cannot get any closer to the ‘true’ random value. We found that 43.06% of randomly generated sequences had a frequency of repeats more extreme than what we observed in the live game — and these values were 49.68% and 87.38% for average numerical distance and direction switches, respectively. Finally, in terms of the model parameters, only the stay parameter significantly differed from uniform randomness (which is 0 for all parameters; mean = -0.21, $t(70) = 2.70$, $p < 0.01$); the other two parameters did not significantly differ from uniform randomness (distance mean = 0.05, $t(70) = 1.47$, $p = 0.15$; cycling mean = 0.04, $t(70) = 1.79$, $p = 0.08$). Thus, there is legitimate evidence to license the claim that the live game leads not only to ‘more random’ behavior but also specifically to ‘near-random’ behavior.

As an additional exploratory analysis, we conducted a model-selection approach to classify the strategies used by individual participants. We examined which sub-model (i.e., our full three-parameter model compared to each decomposition of just one parameter, and even compared to a uniform random, no-parameter model) best fit each person’s data (according to BIC). This revealed further evidence for a strategic gradient of random behavior. In the live game, only 12.7% were best fit by our full, three-parameter model, and 32.4% of participants were best fit by a uniform random distribution (i.e., a no-parameter model). By comparison, in the computerized game, 33.3% of participants were best fit by our full model, and 28.1% were best fit by a uniform random distribution; and in the non-strategic task, 51.3% of participants were best fit by the full model, and only 1 participant (0.8%) was best fit by a uniform random distribution. These analyses provide further evidence that participants in the live, strategic game behaved in a way that was indeed ‘near-random’.

Study 4: Random behavior with differing tile values

Studies 1–3 demonstrate that near-random behavior is possible when the context demands it. But how flexible is such behavior? Some theories of random behavior (e.g., those that posit that randomization failures stem from either an overattribution of mental states or limitations on cognitive resources) may lead us to predict that random behavior should worsen in more complex, strategic tasks — specifically, ones that require representing minds not only in a simple way (i.e., the presence of an opponent) but also in a richer context (i.e., representing specific belief states).

Study 4 used the exact same design as in Study 1, but here the tiles differed in their point values. This increased the strategic depth of the game, as now there was another factor (differing point values) that might motivate players to choose certain tiles. Both generators and guessers were informed about the presence and location of these point values (**Figure 3A**). Though optimal behavior here entails randomly choosing among all tiles — regardless of value — it is not clear whether each player will stick to this optimal strategy. At the extremes, we imagined two

possibilities: (1) Participants might behave non-optimally, perhaps because the cognitive demands were too great and the allure of the salient point values would be significant, or (2) Participants might continue to behave in a near-optimal way despite the additional demands. This design (along with the ones that follow in Studies 5–6) thus tests random behavior in strategically more demanding contexts, where each player must detect and respond to differences in their opponent’s tile value preferences (and it is not clear in the first place that opponents *must* behave randomly to succeed).

Methods

This study used the live game from Study 1 with one key change: Here, certain tiles were more valuable than others. Tile values affected the generator’s points; e.g., if a generator placed their ship on a tile with value 3, they earned +3 points if the ship survived the turn. Of the 9 tiles in the grid, 3 had value 2, 3 had value 3, and 3 had value 4. We created 4 different conditions of tile arrangements. Thus, generators earned either +2, +3, or +4 points when their ship survived, leading to an average expected point total of +2 on each turn (as the mean tile value was 3, and the ship survives approximately $\frac{2}{3}$ of the time given that the guesser provides 3 guesses). We adjusted the scoring for guessers to match this expected value; now, guessers earned +8 points for sinking the ship on their first guess, +6 points for doing so on their second guess, and +4 points for doing so on their third. This also amounts to an expected value of +2 (as, on average, guessers sink the ship on their second guess, and will do so $\frac{1}{3}$ of the time). Importantly, the value of the tile influenced the generator’s score but *not* the guesser’s score (i.e., the multipliers did not apply to the guesser); the guesser was strictly incentivized to choose tiles as accurately as possible, regardless of point value.

In this game, both the generator and the guesser were informed of the tile values (depicted both in text on each tile and via the tiles’ coloring). We recruited 100 dyads of full-game participants (pre-exclusion) — 25 in each of 4 grid arrangement conditions. 23 dyads were excluded due to the exclusion criteria described above, leaving 77 generator/guesser pairs.

Results

To analyze random behavior in Studies 4–6, we modified the above model to include a fourth parameter: a ‘value’ parameter, which adjusts the probability of choosing a tile depending on its value (see *Computational model*). A high value parameter indicates a propensity to select the high-value tiles, while a low value parameter indicates a propensity to avoid those tiles. A value parameter of 0 indicates no such preference (as if the participant ignored the tile values altogether).

In Study 4, we found that the value parameter did not significantly differ from 0 ($M = 0.07$, $t(76) = 1.79$, $p = 0.07$; **Figure 3B**). (This is also reflected by the fact that adding the value parameter to the model only decreased BIC by 3.) In other words, generators largely — and rationally — ignored the value of the tiles, perhaps understanding that guessers would defend those tiles.¹ Evidence for this possibility arose in two key ways. First, generators indeed behaved randomly, to a similar extent as in Study 1 (where tiles did not differ in value). For example, the mean stay, distance, and cycling parameters were quite similar across the two studies (Study 1: -0.21, 0.05, 0.04, respectively; Study 4: -0.16, 0.05, 0.05, respectively). Second, the final score of the game reflected competitive balance between the two players (with generators losing by an average of 7.5 points, or 1.3% of the average total score; $t(76) = 0.79$, $p = 0.43$). Thus, generators flexibly incorporated this information about tile values and continued to behave randomly in this study. This alone is a surprising conclusion, as one may have expected generators to (mistakenly) exploit the high-value tiles and abandon their optimal random strategies — but that is not what we observed.

Study 5: Random behavior with a knowledge asymmetry

In Study 5, we varied players' knowledge about tile values. Specifically, in this study, the guessers had no access to the tile values; they received no visual indicators as to where the generators' high-value tiles were, or even that such tiles existed (**Figure 3C**). Generators, meanwhile, had access to the tile values and were informed that the guessers *did not* know about those values. Notably, despite the informational advantage possessed by the generators, they hold no true advantage in the long run if both players behave optimally. This is because, if generators exploit their advantage, then they will reveal the nature of the advantage and the position of the high-value tiles, at which point guessers can use this knowledge to better defend the relevant tiles. Thus, the generator should still behave in a near-random way, even with this significant knowledge advantage. How would this knowledge asymmetry affect behavior? We predicted that generators would initially select higher-value tiles before adopting a more random strategy, at which point they would perform no better than their less-informed adversaries.

Methods

This task was identical to Study 4, except that here only the generators were informed about the high-value tiles. The guessers were not explicitly instructed about the fact that certain tiles gave the generators more points than others; and generators were told that guessers did not know

¹ Along with the analyses reported in the main text, we preregistered an analysis exploring the relationship between the generator's value parameter and the final score of the game. In a pilot sample, we observed a quadratic (inverted-U) relationship between the two, such that the best-performing generators were the ones who relatively *ignored* tile value in making their choices. Though our preregistered results provide numerical evidence for this trend, the quadratic analysis we preregistered did not reach significance. Still, for clarity, we report it here: In a linear model predicting final score by the generator's value parameter (and its quadratic term), the quadratic coefficient was not significant ($p = 0.19$; for the linear term, $p = 0.35$).

where their high-value tiles were located. Otherwise, this study was identical to Study 4 (even using the same 4 grid arrangement conditions).

As before, we recruited 100 dyads of participants (pre-exclusion) who completed the game (25 in each of the 4 grid arrangement conditions). 45 dyads were excluded, leaving 55 dyads total.

Results

Consistent with our prediction, the value parameter differed significantly from 0 ($M = 0.37$, $t(54) = 9.21$, $p < 0.001$; and the addition of the value parameter improved model fit, $\Delta\text{BIC} = 12$; **Figure 3D**).² This effect was driven by an early but rapidly dissipating preference for high-value tiles. In the first 25 trials, the generators' mean tile value selected was significantly higher than the mean tile value selected across the guessers' three guesses ($M = 0.21$, $t(54) = 6.85$, $p < 0.001$) — a difference that weakened throughout the rest of the task (value difference in first 25 trials vs. value difference in the rest of the task: $M = 0.16$, $t(54) = 4.71$, $p < 0.001$). This points to a natural ebb and flow between the players' behavior: The generators started by exploiting the high-value tiles (successfully representing their opponents' ignorance of those values), but over time, the guessers learned where those high-value tiles were located, and the players reached equilibrium. Despite the knowledge gap, the game's final score reflected nearly equal performance between the generators and guessers (with generators winning only by an average of 6.1 points, or 1.0% of the average total score; $t(54) = 0.55$, $p = 0.58$).

The relative equilibrium of the game suggests that the players optimally adapted to the task structure. Importantly, this held true not only for the generators but also for the guessers. While the guessers need not themselves act randomly, they have an equally important task: to detect and respond to subtle deviations in random behavior. In many real-world contexts, detecting subtle signals hidden in noise confers strategic advantages. It would thus be highly consequential if, for instance, people were better at generating than perceiving randomness (or vice versa) in ways that created large asymmetries in these adversarial tasks. These results tell a different story: They suggest not only that people are capable of behaving in a near-optimal random way but also that people are capable of perceiving (deviations from) random behavior with a high degree of acuity. In other words, the equality in the final game scores (along with other metrics pointing to near-random behavior) suggests that both parties played the game nearly optimally.

² As above, we preregistered an analysis exploring the relationship between the generator's value parameter and the final score of the game. We made no predictions here, and included it merely for consistency with Study 4. A linear model predicting final score by the generator's value parameter (and its quadratic) revealed neither a linear ($p = 0.94$) nor a quadratic ($p = 0.85$) relation between the two.

Study 6: Random behavior with an insurmountable knowledge asymmetry

In Study 5, generators had access to knowledge that the guessers did not. But that advantage could only be exploited for a short time before the guessers adapted. What if generators received an even larger knowledge advantage? Would they then exploit the high-value tiles? For instance, if the value of the tiles changed periodically, then guessers would struggle to adapt — allowing generators to exploit the high-value tiles even more often. Even when both players play such a game optimally, then, the generators nevertheless hold a decisive advantage.

Study 6 asked exactly this question, modifying Study 5 such that the tile values shuffled every 15 turns (**Figure 3E**). This severely impaired the guesser's ability to infer the location of the high-value tiles. As before, the tile values were hidden from the guesser; but the shuffling ensured that the guesser could not learn any patterns in the generator's behavior regarding tile value. We predicted that this manipulation would lead generators to exploit the high-value tiles and largely outscore their opponents.

Methods

In Study 6, we added one feature on top of Study 5. Here, the tile values shuffled every 15 turns, resulting in 10 arrangements total throughout the game. The grids shuffled in a preset order, such that every participant in this study viewed the same grids in the same order. Generators were informed of changes to the grid via a variety of visual markers. As before, only the generators knew the tile values; but the shuffling ensured that guessers could not learn which tiles were most beneficial to the generator (i.e., by observing their placement).

We recruited 100 dyads, 20 of whom were excluded as per our preregistered exclusion criteria.

Results

As predicted, this manipulation led to dramatic shifts in the generator's behavior: Knowing that their opponent could not see — or even, with time, infer — the location of the high-value tiles led the generators to select those tiles with much greater frequency. The value parameter significantly differed from 0 ($M = 0.70$, $t(79) = 14.89$, $p < 0.001$; **Figure 3F**),³ with many people reaching a value parameter of 1 (the upper bound for the parameter). (Consequently, the addition of the value parameter improved the model fit substantially; $\Delta\text{BIC} = 48$.) Unlike in Study 5, however, guessers failed to adapt to this behavior. Because the tile values shuffled, generators exploited the high-value tiles without consequence, leading them to win by a large margin in this game (winning by an average of 51.8 points, or 8.0% of the average final score; $t(79) = 5.29$, $p <$

³ A linear model predicting final score by the generator's value parameter (and its quadratic) revealed neither a linear ($p = 0.18$) nor a quadratic ($p = 0.99$) relation between the two. As above, we performed this analysis for consistency with our earlier Studies.

0.001). In other words, this manipulation demonstrated flexible behavior by generators: They readily exploited the high-value tiles when it was clear the guessers could not defend those tiles consistently. These results also help contextualize the prior findings, insofar as they reveal that generators readily exploit a genuine advantage when it exists. It is not for a lack of trying that the generators failed to convert an advantage in the prior experiments. Rather, it is because the guessers successfully perceived and responded to the generators' actions.




Figure 3. (A) In Study 4, tiles differed in value, and both players knew the value of all tiles. (B) Generators did not have a significant preference for selecting the high-value tiles; and guessers were relatively stable in defending those tiles throughout the game. (C) Study 5 altered this design to hide tile values from the guesser. (D) This led to increases in the value parameter (i.e., this manipulation led generators to select high-value tiles more frequently); but this preference evolved throughout the course of the game. Early in the task, the generators selected high-value tiles often, until the guessers adjusted and the mean preference for tile value leveled out. (E) In Study 6, the tile values shuffled every 15 turns, such that the guesser could not learn the tile values based on the generator's choices. (F) As expected given their strategic advantage, generators chose the high-value tiles very frequently; and this preference remained stable throughout the game.

Discussion

Despite the importance of randomness to many aspects of cognition, humans often struggle to produce and perceive it. Why? Here, we proposed that humans are not unequivocally bad randomizers, but rather *strategic randomizers*, able to produce near-random behavior in ecologically valid contexts that demand it. The present work supports this possibility. We demonstrate that people achieve near-random behavior in live, strategic games, much more so than when asked to simply produce random behavior (as typical tasks require). In other words, people hold the capacity for near-random behavior, but perhaps only selectively invest the resources to achieve it. Moreover, we find that the nature of the opponent influences people's ability to achieve such near-random behavior, as performance improves when participants play against real, human opponents (as opposed to a computerized opponent). When random behavior is strategically valuable in this way, the best model of human randomness is no model at all; people play the game in near-optimal fashion. This impressive ability to behave randomly comes in stark contrast with an extensive literature that characterizes significant departures from random behavior in standard laboratory tasks (e.g., Kahneman & Tversky, 1972; Bar-Hillel & Wagenaar, 1991; Nickerson, 2002; Boger et al., 2025). At the same time, our findings build upon previous work using strategic games to probe random behavior (e.g., Rapoport & Budescu, 1992; Budescu & Rapoport, 1994), both in humans and in other animals. For example, chimpanzees behave in a near-optimal fashion in competitive games — matching equilibrium game theory predictions — suggesting that the capacity for near-random behavior has deep evolutionary roots (Martin et al., 2014).

In addition to demonstrating that humans are successful at strategic randomness, we probed the role of mentalizing in this process. Does the representation of another person's knowledge state and beliefs hinder random behavior, such that we get caught in imitative loops? Or are we able to use this information strategically, producing behavior that tracks those incentives? Here, we find support for the latter possibility, insofar as people flexibly incorporate relevant strategic information into their random actions. When we varied the value of individual tiles in the game (Studies 4–6), participants adapted their behavior accordingly: They acted randomly when it was beneficial to do so, but exploited high-value opportunities when they were available. This shows a sophisticated integration of others' knowledge states that improved (rather than impaired) their payoffs. Altogether, these results provide a much more nuanced view of human randomization abilities and suggest there is more to explore regarding how (and when) people engage in optimal random behavior.

A multitude of possible mechanisms for achieving near-random behavior

How do participants achieve near-random strategic behavior? Given the difficulty of defining randomness, answers to this question are elusive. Consider, for example, one popular theory that

explains biases of randomness by appealing to limitations in working memory. This theory faces a unique challenge (as described in, e.g., Falk & Konold, 1997): On the one hand, randomness can be achieved with zero memory. An agent with no memory avoids classic biases (e.g., over-alternation) simply because they do not have a ‘previous response’ in mind to compare against possible next responses. On the other hand, an agent with infinite memory could (in theory) produce uniform randomness by considering *all* of their prior responses to select a next response that brings the sequence closest to uniform randomness (assuming they properly understand what it means to be random). Other accounts of randomness face a similar problem. For example, in our tasks, one might hypothesize that participants employ k -level reasoning in the live game, which happens to produce near-random behavior (i.e., without appeal to a randomness generator that differs across tasks). In other words, perhaps an optimal level of k , wherein an agent models their opponent at the correct depth, facilitates near-random behavior in our tasks. However, here too, extreme values of k in either direction (whether 0 or infinite) seem to result in the same behavior. In this account, then, it remains optimal for participants to simply operate at $k = 0$ (instead of high values of k) given the cognitive cost. Of course, as it relates to the present work, the relevant question is whether intermediate values of k successfully account for increases in randomness; we believe this is an interesting question for future work. These different realizations of near-random behavior may even entail different processes altogether. For example, one hypothesis is that an underlying generator or sampling algorithm (e.g., Castillo et al., 2026) gets ‘tuned’ by strategic contexts through these mechanisms. Another hypothesis, however, is that the different strategic contexts recruit entirely different kinds of generators from non-strategic tasks (e.g., Budescu & Rapoport, 1994). Thus, multiple possible mechanisms may contribute to strategic randomization — perhaps even in tandem, making it difficult to infer mechanism from behavior.

One other possibility centers around features of our task unrelated to strategic randomness. For example, the live and computerized games offer participants feedback (by way of the ship surviving or sinking), while the non-strategic task does not. Thus, it may be that feedback itself produces near-random behavior regardless of context. However, we think this is unlikely, for two reasons. First, though prior work demonstrates improvement in randomization following feedback (Neuringer, 1986), such improvements fade quickly and only apply to the specific sequence features on which participants received feedback. Second, other work has shown that benefits in randomization that arise during a strategic task seem not to transfer to identical non-strategic tasks (Wong et al., 2021). That said, even if feedback fails to fully explain the differences we observe across conditions, it still seems plausible that feedback plays at least some role in tuning random behavior. This is, in itself, non-trivial insofar as other (sometimes heavy-handed) task manipulations seem not to affect randomization performance. For example, even with added financial incentives, people struggle to behave randomly (Budescu & Rapoport, 1994), suggesting that improvements in randomization go beyond mere effort.

In each case, we believe an important distinction arises between process and product (or, as Falk & Konold, 1997 put it, “process and pattern”). Our work provides evidence for near-random *products* when the context demands it. However, we remain open to a multitude of specific processes that lead to that product. We think this is a difficult yet exciting problem, as many processes can result in the same ultimate product. For example, as we mentioned before, an agent with perfect memory and another agent with no memory whatsoever would both be capable of perfectly random behavior; thus, the product alone provides insufficient information from which to infer the process. This remains true even in literal random devices (e.g., as used by cybersecurity companies); many different possible devices, each with wildly different implementations, may nevertheless produce similarly strong randomness. We believe elucidating the cognitive processes underlying the production of randomness is a particularly promising avenue for future work.

Strategic games as a window into cognition and randomness

Strategic contexts differ fundamentally from non-strategic contexts in that the ‘optimal’ strategy always depends on the opponent’s behavior. Even in a game like chess, a seemingly imperfect or even weak move can be effective if it tricks the opponent. In this way, strategic games implicate many other cognitive capacities like theory of mind (Kleiman-Weiner et al., 2025; see also Li et al., 2025; Mazor et al., 2026; for review, see Allen et al., 2024) and working memory (which, as mentioned above, is especially important in the study of random behavior; see Hahn & Warren, 2009; Warren et al., 2018). For example, recent work has used complex, collaborative games as a natural testing ground for the cognitive mechanisms underlying multi-agent collaboration (Wu et al., 2021; see also Towse et al., 2016 for collaborative generation of random sequences). Other, simpler games — including matching pennies (Rapoport & Budescu, 1992) and rock-paper-scissors (Batzilis et al., 2019; Wong et al., 2021; Brockbank & Vul, 2024) — have been used to study randomness given their clearly defined ‘equilibrium’ predictions. The use of optimal computer opponents (e.g., Wong et al., 2021; Brockbank & Vul, 2024) provides an especially interesting upper bound for random behavior, as the only way to defeat such opponents is to behave randomly with respect to any number of n -order transition probabilities. The studies reported here combine random behavior, complex environments, and naturalistic games, leading to unique, emergent gameplay that involves the interplay of many cognitive processes at once.

This is especially true of Studies 4–6. Across these studies, behavior differed dramatically as a result of differences in knowledge between the two players (and not merely because of the underlying rules of the game or participants’ ability to act randomly within its confines). When *both* players had access to knowledge of the high-value tiles, the guessers preemptively defended those regions, such that optimal behavior involved the generator ignoring the tile values and behaving randomly. And indeed, this is how generators behaved, consequently performing

adequately in the game. When only the *generator* knew the location of the high-value tiles, behavior changed slightly — but primarily at the start of the task, as the guessers quickly adapted. When generators exploited the high-value tiles, the guessers learned the location of those tiles and began to defend them. Thus, optimal behavior in the long run involved integrating information about the tile values with the opponent’s knowledge, rather than relying on the values alone. Finally, when only the generator had knowledge of the high-value tiles, and the location of those tiles fluctuated at regular intervals, the guessers’ inferences were disrupted, as they could not learn which tiles were valuable. Consequently, generators could — and did — exploit this advantage to greatly outperform their competitors. Collectively, then, these latter studies demonstrate not only that people are capable of near-random behavior, but also that they optimally incorporate multiple sources of information (e.g., tile values, opponent’s knowledge and beliefs) into their action plans.

So far, we have focused primarily on *generator* behavior, but the guessers’ behavior is equally revealing. Specifically, the group-level consistency between people’s tendencies as generators and their tendencies as guessers points to the possibility that both behaviors are served by a common process. Indeed, many biases of random generation also appear in judgment tasks. For example, people judge sequences with more alternations to be more random, just like they create more alternations when generating such sequences. Yet the fact that generators and guessers play these games in a way that results in (group-level) equilibrium suggests a deeper relation between generation and perception of randomness. One very simple mechanism by which random generation and random perception could be linked — especially in this task — is via self-modeling, wherein a guesser thinks, “If I were the generator and I placed the ship here last turn, where would I place the ship this turn?” Of course, this crude link would not always work — e.g., because different people generate randomness in different ways (Boger et al., 2025) — but our results point to the existence of a deep (rather than superficial) relation between random action and random perception.

One way of understanding these results could be that the *performance* measured in prior studies of randomization fails to fully capture people’s *competence* (Chomsky, 1965). In prior studies, people may have simply lacked the proper incentive to behave in a truly random way, and thus, not wanting to expend cognitive resources, made use of various heuristics. Some even suggest that strategic tasks may recruit entirely different cognitive processes from typical laboratory tasks (a possibility that is not mutually exclusive and is suggested in Budescu & Rapoport, 1994). That is: Much like social cognition makes all the difference in how people understand logic puzzles such as the Wason selection task, perhaps the social context of these strategic games recruits unique cognitive systems that are more capable of approximating randomness (Cosmides & Tooby, 1992). The social context is importantly different from the non-social context in our tasks, after all, as the former necessarily involves recursive reasoning and theory

of mind. But on this view, counter to the overattribution account, this social context helps rather than hinders randomization ability.

Concluding remarks

For decades, failures of human randomization have been treated as a cognitive flaw — a 'bug' revealed whenever people are asked to be random. Our results suggest a different conclusion. Randomness, for humans, is not a default mode of action but a strategic tool, deployed when advantageous and conserved otherwise. In competitive contexts that mirror the ecological pressures under which randomness is useful — avoiding predictability, tricking adversaries, exploiting asymmetric knowledge, and so on — people implement random behavior, and do so flexibly, intelligently, and almost optimally. Far from being derailed by mentalizing and recursive reasoning, participants leverage these capacities to calibrate when to randomize, when to exploit structure, and when to shift between the two. This reframes classic findings in the literature: What appears as a failure to randomize may instead reflect a mismatch between the task and the function randomness serves in cognition.

Supplementary Materials

This document provides supplementary analyses to “Near-random behavior in strategic games”. Here, following the feedback of a peer reviewer, we discuss the nature of our exclusion criteria and confirm that our results hold even when we do not apply any exclusions to the datasets used in Studies 1–3.

The key takeaway from this document is that our exclusions on their own fail to explain the results we observe. In other words, even with no exclusions applied, the behavior in the live game is closest to random. We include analyses of the excluded data for transparency (and upon the request of a peer reviewer).

Definitions of our exclusion criteria


In each study, we excluded participants who behaved in clearly non-random ways. We operationalized that as follows (in accordance with previous work; Boger et al., 2025): Participants who fell outside of the middle 99.99% of uniform random sequences in terms of the frequency of selected choices and/or the average numerical distance between choices were excluded. Note that this is, in theory, a very lenient exclusion. We arrived at the specific numerical bounds by simulating several thousand draws of uniformly generated random numbers (with a sequence length equal to the number of trials in the relevant study) and then computing the lower 0.005th and upper 99.995th percentiles of the relevant value. Participants needed to pass both criteria (frequency of selected choices and average numerical distance) to be included. Note that the nature of these exclusions makes it important to visualize not just the means but also the distributions of the resultant properties (as, for example, a sample might display random-like repeats by having some participants avoid repeats entirely and other participants repeat constantly). For that reason, we depict the distributions in each case.

This approach excluded 29 (out of 100) participants in Study 1 (live game), 43 (out of 100) participants in Study 2 (computerized game), and 76 (out of 195) participants in the non-strategic dataset (from Boger et al., 2025).

These exclusion numbers are in line with a previous study using the same criteria (Boger et al., 2025). In addition, the exclusion rate was lowest when we expected behavior to be most random — Study 1 (the live game) excluded the fewest participants. At a glance, then, exclusions seem unlikely to explain our results. To be thorough, we also reanalyzed the data with no exclusions applied at all. These analyses are available on our OSF repository (<https://osf.io/z928n/>).

Emergent properties without exclusions

Below are the emergent properties of the sequences produced in each dataset without exclusions. We also compare these values to the values with the pre-registered exclusions applied (i.e., as in our main text), and to the values when analyzing *only* the excluded participants. As before, the dotted line indicates the expected value for that property in a uniform random sequence.




We then performed statistical tests asking whether our key trends hold with no exclusions applied. (We do not do so for the dataset consisting *only* of excluded subjects because that question is unrelated to our key hypotheses.) For each of these three emergent properties, the sequences in the live game show no significant difference from uniform randomness ($p = 0.93$

for frequency of repeats, $p = 0.88$ for average numerical distance, and $p = 0.29$ for frequency of direction switches). Each comparison across datasets, as conducted in the main paper (e.g., comparing the live game's data to the computerized game's data, or comparing the computerized game's data to the non-strategic task's data) is significant (all $ps < 0.05$). Thus, in this sense, our primary results from the main text hold with no exclusions applied (and in some sense, understate how closer behavior is to uniform randomness).

Note that the mean frequency of repeats for the excluded subjects in the non-strategic task appears close to random. This trend, however, is explained by the fact that very few subjects repeat their choices at a high rate, while the rest of the subjects avoid repeats — thereby resulting in a near-random mean. Importantly, participants were excluded for extreme behavior on particular dimensions of randomness (namely, choice frequency and numerical distance). Because the properties of interest are interrelated (and because different forms of non-random behavior can push a given metric in opposite directions), including the excluded participants can sometimes move an aggregate measure closer to the expected value in a uniform random sequence.

Model parameters without exclusions

Below, we report the model parameters for the data without any exclusions applied. This provides an additional way to examine the strategic gradient of random behavior (and helps clarify any potential differences between the tasks).




Here, too, behavior in the live game appears closest to random. For the stay parameter, neither comparison (live game vs. computerized game, and computerized game vs. non-strategic task) revealed a significant difference (both p s > 0.45). In terms of the distance parameter, the live game differed significantly from the computerized game ($p < 0.001$), but the computerized game did not significantly differ from the non-strategic task ($p = 0.10$). Finally, the cycling parameter

did not significantly differ between the live game and the computerized game ($p = 0.22$), but did differ between the computerized game and the non-strategic task ($p < 0.001$).

Model accuracy

Finally, we compare the model's predictive ability across these three datasets. Note that we do not compare the human guessers' performance because the human guessers in Study 3 only played 'against' a randomly selected subset of post-exclusion sequences from the computerized game and non-strategic task. Thus, a 'no exclusions' version of the human guesser accuracy would require additional data collection. However, we can compute the model's top-3 accuracy on the data without exclusions. Those results are below.



Each of the relevant comparisons here (live game vs. computerized game, computerized game vs. non-strategic task) is significant (both $ps < 0.001$). This provides perhaps the strongest evidence for the 'strategic gradient' in randomness, even without any exclusions applied.

References

- Allen, K., Brändle, F., Botvinick, M., Fan, J. E., Gershman, S. J., Gopnik, A., ... & Schulz, E. (2024). Using games to understand the mind. *Nature Human Behaviour*, 8(6), 1035-1043.
- Angelike, T., & Musch, J. (2024). A comparative evaluation of measures to assess randomness in human-generated sequences. *Behavior Research Methods*, 56(7), 7831-7848.
- Angelike, T., & Musch, J. (2025). An improved modeling approach to investigate biases in human random number generation. *PLoS One*, 20(5), e0324870.
- Ayton, P., Hunt, A. J., & Wright, G. (1989). Psychological conceptions of randomness. *Journal of Behavioral Decision Making*, 2(4), 221-238.
- Ayton, P., & Fischer, I. (2004). The hot hand fallacy and the gambler's fallacy: Two faces of subjective randomness?. *Memory & Cognition*, 32(8), 1369-1378.
- Baddeley, A. D. (1966). The capacity for generating information by randomization. *The Quarterly Journal of Experimental Psychology*, 18(2), 119-129.
- Baker, F., & Rachlin, H. (2001). Probability of reciprocation in repeated prisoner's dilemma games. *Journal of Behavioral Decision Making*, 14(1), 51-67.
- Bar-Hillel, M., & Wagenaar, W. A. (1991). The perception of randomness. *Advances in Applied Mathematics*, 12(4), 428-454.
- Batzilis, D., Jaffe, S., Levitt, S., List, J. A., & Picel, J. (2019). Behavior in strategic settings: Evidence from a million rock-paper-scissors games. *Games*, 10(2), 18.
- Biesaga, M., & Nowak, A. (2024). The role of the working memory storage component in a random-like series generation. *PLoS One*, 19(1), e0296731.
- Boger, T., Yousif, S. R., McDougale, S. D., & Rutledge, R. B. (2025). Random behavior is stable across tasks and time. *Journal of Experimental Psychology: General*, 154(6), 1571-1582.
- Bonner, J. T. (2013). *Randomness in evolution*. Princeton University Press.
- Brocas, I., & Carrillo, J. D. (2022). The development of randomization and deceptive behavior in mixed strategy games. *Quantitative Economics*, 13(2), 825-862.
- Brockbank, E., & Vul, E. (2024). Repeated rock, paper, scissors play reveals limits in adaptive sequential behavior. *Cognitive Psychology*, 151, 101654.
- Budescu, D. V., & Rapoport, A. (1994). Subjective randomization in one- and two-person games. *Journal of Behavioral Decision Making*, 7(4), 261-278.
- Camerer, C. F., Ho, T. H., & Chong, J. K. (2004). A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, 119(3), 861-898.
- Castillo, L., León-Villagrà, P., Falbén, J. K., Chater, N., & Sanborn, A. N. (2026). Random generation is what comes to mind in naturalistic settings. *Cognition*, 274, 106554.
- Chapanis, A. (1995). Human production of "random" numbers. *Perceptual and Motor Skills*, 81(3_suppl), 1347-1363.
- Chomsky, N. (1965). *Aspects of the Theory of Syntax*. MIT Press.

- Cook, R., Bird, G., Lünser, G., Huck, S., & Heyes, C. (2012). Automatic imitation in a strategic context: players of rock–paper–scissors imitate opponents' gestures. *Proceedings of the Royal Society B: Biological Sciences*, 279(1729), 780-786.
- Cosmides, L., & Tooby, J. (1992). Cognitive adaptations for social exchange. *The adapted mind: Evolutionary psychology and the generation of culture*, 163, 163-228.
- Driver, P. M., & Humphries, D. A. (1988). *Protean behaviour: The biology of unpredictability*. Clarendon Press.
- Falk, R., & Konold, C. (1997). Making sense of randomness: Implicit encoding as a basis for judgment. *Psychological Review*, 104(2), 301-318.
- Gilovich, T. (1983). Biased evaluation and persistence in gambling. *Journal of Personality and Social Psychology*, 44(6), 1110-1126.
- Gilovich, T., Vallone, R., & Tversky, A. (1985). The hot hand in basketball: On the misperception of random sequences. *Cognitive Psychology*, 17(3), 295-314.
- Griffiths, T. L., & Tenenbaum, J. B. (2003). Probability, algorithmic complexity, and subjective randomness. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 25, No. 25).
- Griffiths, T. L., Daniels, D., Austerweil, J. L., & Tenenbaum, J. B. (2018). Subjective randomness as statistical inference. *Cognitive Psychology*, 103, 85-109.
- Hahn, U., & Warren, P. A. (2009). Perceptions of randomness: why three heads are better than four. *Psychological Review*, 116(2), 454-461.
- Hedden, T., & Zhang, J. (2002). What do you think I think you think?: Strategic reasoning in matrix games. *Cognition*, 85(1), 1-36.
- Henrich, J. (2015). *The secret of our success: How culture is driving human evolution, domesticating our species, and making us smarter*. Princeton University Press.
- Icard, T. (2021). Why be random?. *Mind*, 130(517), 111-139.
- Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology*, 3(3), 430-454.
- Kareev, Y. (1992). Not that bad after all: Generation of random sequences. *Journal of Experimental Psychology: Human Perception and Performance*, 18(4), 1189-1194.
- Kleiman-Weiner, M., Vientós, A., Rand, D. G., & Tenenbaum, J. B. (2025). Evolving general cooperation with a Bayesian theory of mind. *Proceedings of the National Academy of Sciences*, 122(25), e2400993122.
- Li, Z., Tang, N., Gong, S., Zhou, J., Shen, M., & Gao, T. (2025). Territorial Gestalt in the Strategy of Conflicts. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 47).
- Martin, C. F., Bhui, R., Bossaerts, P., Matsuzawa, T., & Camerer, C. (2014). Chimpanzee choice rates in competitive games match equilibrium game theory predictions. *Scientific Reports*, 4(1), 5182.
- Mazor, M., Firestone, C., & Phillips, I. (2026). Pretending not to know reveals a capacity for model-based self-simulation. *Psychological Science*, 37(2), 136-149.

- Mookherjee, D., & Sopher, B. (1994). Learning behavior in an experimental matching pennies game. *Games and Economic Behavior*, 7(1), 62-91.
- Moore, O. K. (1957). Divination-a new perspective. *American Anthropologist*, 59(1), 69-74.
- Neuringer, A. (1986). Can people behave "randomly?": The role of feedback. *Journal of Experimental Psychology: General*, 115(1), 62-75.
- Newman, G. E., Keil, F. C., Kuhlmeier, V. A., & Wynn, K. (2010). Early understandings of the link between agents and order. *Proceedings of the National Academy of Sciences*, 107(40), 17140-17145.
- Nickerson, R. S. (2002). The production and perception of randomness. *Psychological Review*, 109(2), 330.
- Rapoport, A., & Budescu, D. V. (1992). Generation of random series in two-person strictly competitive games. *Journal of Experimental Psychology: General*, 121(3), 352.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, 274(5294), 1926-1928.
- Saffran, J. R., Johnson, E. K., Aslin, R. N., & Newport, E. L. (1999). Statistical learning of tone sequences by human infants and adults. *Cognition*, 70(1), 27-52.
- Sanborn, A. N., Zhu, J. Q., Spicer, J., León-Villagr , P., Castillo, L., Falb n, J. K., ... & Chater, N. (2025). Noise in cognition: bug or feature?. *Perspectives on Psychological Science*, 20(3), 572-589.
- Speck, F. G. (1977). *Naskapi: The savage hunters of the Labrador Peninsula* (Vol. 10). University of Oklahoma Press.
- Towse, J. N., & Neil, D. (1998). Analyzing human random generation behavior: A review of methods used and a computer program for describing performance. *Behavior Research Methods, Instruments, & Computers*, 30(4), 583-591.
- Towse, J. N., Towse, A. S., Saito, S., Machara, Y., & Miyake, A. (2016). Joint cognition: thought contagion and the consequences of cooperation when sharing the task of random sequence generation. *PLoS One*, 11(3), e0151306.
- Wagenaar, W. A. (1972). Generation of random sequences by human subjects: A critical survey of literature. *Psychological Bulletin*, 77(1), 65.
- Warren, P. A., Gostoli, U., Farmer, G. D., El-Deredy, W., & Hahn, U. (2018). A re-examination of "bias" in human randomness perception. *Journal of Experimental Psychology: Human Perception and Performance*, 44(5), 663-680.
- Wong, A., Merholz, G., & Maoz, U. (2021). Characterizing human random-sequence generation in competitive and non-competitive environments using Lempel–Ziv complexity. *Scientific Reports*, 11(1), 20662.
- Wu, S. A., Wang, R. E., Evans, J. A., Tenenbaum, J. B., Parkes, D. C., & Kleiman-Weiner, M. (2021). Too many cooks: Bayesian inference for coordinating multi-agent collaboration. *Topics in Cognitive Science*, 13(2), 414-432.

- Zhao, J., Hahn, U., & Osherson, D. (2014). Perception and identification of random events. *Journal of Experimental Psychology: Human Perception and Performance*, 40(4), 1358-1371.
- Zhao, J., & Yu, R. Q. (2016). Statistical regularities reduce perceived numerosity. *Cognition*, 146, 217-222.