

People use fast and flat simulation to reason about new games

<https://doi.org/10.1038/s41586-026-10722-1>

Received: 23 September 2025

Accepted: 26 May 2026

Published online: 15 July 2026

Open access

 Check for updates

Katherine M. Collins^{1,2,3,7}✉, Cedegao E. Zhang^{1,7}, Lionel Wong^{1,4,7}, Mauricio Barba da Costa^{1,7}, Graham Todd^{5,7}, Adrian Weller^{3,6}, Samuel J. Cheyette¹, Thomas L. Griffiths² & Joshua B. Tenenbaum¹

Games have long been a microcosm for studying planning and reasoning in both natural and artificial intelligence, often focusing on expert-level or even super-human play^{1–6}. But real life also pushes human intelligence along a different frontier, requiring people to flexibly navigate decision-making problems that they have never thought about before. Here we use novice gameplay to study how people reason about new problem settings. Through a series of large-scale behavioural studies with over 1,000 participants and 121 two-player strategic board games (almost all novel to our participants), we show that people are systematic and adaptively rational in how they play a game for the first time or evaluate a game (for example, how fair or how fun it is likely to be) before they have played it even once. We explain these capacities via a computational cognitive model that we call the ‘Intuitive Gamer’: a model based on mechanisms of fast and flat (depth-limited) goal-directed probabilistic simulation. Our work offers insights into how people rapidly evaluate, act and make suggestions when encountering novel problems, and could inform the design of more flexible and human-like artificial intelligence systems that can determine not just how to solve new tasks but also whether a task is worth thinking about at all.

Games occupy a special space in psychology and computer science^{1–4,6–12} for good reason: games are systems of rules and reward that hold a mirror to the structures, patterns and challenges that reality confronts us with, across familiar and merely possible experiences. Mancala lets us sow and capture abstract resources. Chess lets us strategize over bloodless battlefields. Go lets us encircle and claim territories with silent stone troops. A good game lets us thrill to uncertainties and losses that we might shy from in real life.

Substantial work in cognitive science focuses on the nature and development of expertise, using games as a testbed^{3,6,13}. Human experts, correspondingly, have long served as goalposts in the quest to build intelligent machines, with a particular focus on emulating and ultimately beating humans at their own games, using more game-playing data and raw computational horsepower than any human expert could acquire in their lifetime^{2,4,14,15}.

But people are not experts at most of the problems they encounter throughout their lives, or most systems they are thrown into before they need to decide on an action. Human cognition is flexible enough to consider many potential problems and many potential systems of rule and reward. Thinking about whether a system of rules and reward makes sense, whether to participate, or what actions to take first—across a wide range of novel situations—is arguably at least as important for everyday human cognition (if far less studied) than the cognitive processes that lead a small number of humans to become expert players in any one game.

Here we study the abilities of novice game reasoners across a range of over 100 novel games from a subclass of strategic grid-based board games that are not entirely unlike games they have probably seen before (for example, Tic-Tac-Toe or Connect-4), but vary in their rules and underlying dynamics (see examples in Fig. 1a and Extended Data Table 1). In a series of large-scale behavioural studies, each with hundreds of participants, we assess people’s thinking about novel games in three reasoning settings (Fig. 1b–d): game evaluation (determining how rewarding a game is likely to be before playing); action selection (choosing moves when playing for the first time, with another first-time player); and action prediction (judging the likelihood of moves other first-time players might make without having played directly themselves). Our setting is different from much previous work on games, which usually focuses on one or a small set of games and many rounds of play in that game; here, we consider many games and little (or even no) experience.

Our core contribution is a computational account of how people reason about these novel problems. Expert models of gameplay typically involve deep tree search with potentially thousands of evaluations of possible states (Fig. 2a). It seems unlikely that a novice reasoner is engaging in such intensive search and evaluation before they have extensive—or any—experience. But neither are people likely to be completely unsystematic. We hypothesize that novice thinking occupies a valuable intermediate point between these extremes (Fig. 2b): people are non-random in assessing new problems and run

¹Massachusetts Institute of Technology, Cambridge, MA, USA. ²Princeton University, Princeton, NJ, USA. ³University of Cambridge, Cambridge, UK. ⁴Stanford University, Stanford, CA, USA.

⁵New York University, New York City, NY, USA. ⁶The Alan Turing Institute, London, UK. ⁷These authors contributed equally: Katherine M. Collins, Cedegao E. Zhang, Lionel Wong, Mauricio Barba da Costa, Graham Todd. ✉e-mail: katiemc@mit.edu




Fig. 1 | Our novel game dataset and suite of game tasks. a, Ten example games from our 121 game dataset. Games vary in board sizes and rules, such as what it takes to win and how many pieces any player can place on their opening move. **b–d**, We assess people’s reasoning about novel games through three

behavioural studies designed to test: how people reason about games before they even play a single game (**b**); how people decide what actions to make in their first instance of play (**c**); and how people predict others should play when watching them play (**d**).

computations analogous to those that underlie state-of-the-art artificial intelligence (AI) game systems and cognitive models of expert human gameplay^{4,6,14}—but scaled down to a level that is more realistic for everyday thought.

We instantiate and test this hypothesis in an Intuitive Gamer model that makes judgements and decisions by running mental simulations of gameplay that are (1) fast, (2) flat, (3) goal-directed and (4) probabilistic. The Intuitive Gamer model accounts very well—and significantly better than alternative models—for how people evaluate these games before play. Features computed from the same model can capture people’s judgements about the game’s objective value, as well as more subjective evaluations such as a game’s expected ‘funness’. Action choices simulated by this same model generally capture how novices actually play in games and how they make predictions about unfolding games they have never played themselves.

The Intuitive Gamer model

The capacity to evaluate properties of a new game and decide whether the game is worth engaging with might intuitively feel like it comes before developing a policy for actually playing the game. However, our intuition as modellers is the opposite: game properties such as expected values depend on both the rule specifications of a game and the policies adopted by the players. We hypothesize that this could be true for novice human game reasoners—before taking a single action in the game—if their policies are extremely simple to construct, fast to execute and support probabilistic sampling of reasonable action choices. If so, a policy can be queried just a small number of times to simulate gameplay, from which the game reasoner can draw reasonable inferences about the game’s properties using these simulated traces.

We instantiate this hypothesis with our Intuitive Gamer model. The model consists of two modules: a ‘flat’ game-playing agent, which runs

an extremely depth-limited search to select actions probabilistically on the basis of abstract yet locally evaluable ‘goal-directed’ heuristic functions (Fig. 2c); and a ‘fast’ game reasoner that makes probabilistic inferences about game propositions and their expected values using a small number of simulated games (Fig. 2d).

Player module

A long tradition of modelling human behaviour in games has used some combination of game-state evaluation functions and search over the decision tree of game states^{16,17}, as depicted in Fig. 2b, to construct a policy. Our approach broadly falls into this tradition. However, existing accounts typically assume highly tuned or expert-designed value functions and deep simulation^{2,6,14,18}, which by definition cannot be expected of a novice. Rather, we propose a model that combines a general-purpose abstract value function—which could be easily constructed from understanding the rules of the game—with depth-limited search. Instead of simulating the downstream effects of a move via extensive forward search, the Intuitive Gamer player module uses only a single step of ‘look-ahead’ and selects an action probabilistically by evaluating the resulting board states according to a small collection of general heuristics (Fig. 2c). These heuristics are built on the assumption that players understand and pursue their goals (as defined by the game rules) and attempt to prevent their opponents from doing the same, subject to constraints on computational resources^{19,20}. The Intuitive Gamer player module then is a more general, compute-bounded version of prior game-based agents (Methods).

Reasoning module

Whereas the Intuitive Gamer player model captures action choice during a game, the Intuitive Gamer reasoning module captures judgements about game properties by nesting the player module within a sample-based probabilistic inference procedure (Fig. 2d). For any




Fig. 2 | The Intuitive Gamer model compared with previous models of game reasoning. **a**, Previous work modelling expert gameplay often involves deep tree search to determine what move to make, given a board state S_t (refs. 4, 6). It is unlikely that novice human reasoners conduct such computationally expensive search and state evaluation before deciding whether to engage with the problem at all. **b**, We consider what novice reasoners might be doing instead. Gameplay agents can differ in the amount of compute and expertise brought to bear to reason about any game (differing in search depth and value sophistication). Our proposal is that people reasoning about problems with which they have no experience sit at the lower end of this spectrum, but not the lowest. **c**, The Intuitive Gamer conducts depth-limited ('flat') search with game-general abstract goal-directed value functions that have yet to encode game-specific features. The Intuitive Gamer gameplay agent conducts no more

than a single step of lookahead when deciding what action (A_{t+1}) to take from a given board state (S_t) at a given move turn t . The Intuitive Gamer is goal-directed, assessing whether any action would advance the player's own goal (purple) and how much it might block the progress of the opponent's goal (yellow). The final action is selected probabilistically by sampling from a softmax distribution over the estimated values per action. **d**, To reason about any new game query ψ for a given game description $\mathcal{G}$, we posit that people conduct only a few (k) self-play simulations between the gameplay agent (as depicted in **c**) to answer the query, which could be run to termination or probabilistically stop early. In summary, the Intuitive Gamer model is fast (low k), flat (low search depth), goal-directed (in the value function) and probabilistic (in action selection), and involves mental simulations of gameplay. See Methods for a detailed formalism of the model.

game, the reasoning module infers answers to a query (for example, what is the likelihood that the game will end in a win, loss or draw for a given player?) by simulating k game playouts and iteratively calling the player module at each turn of self-play. Each simulation ends when the game reaches a win or draw state, or may terminate early with some probability. Computations over the simulated gameplay traces then inform the resulting distribution over the query values (for example, game outcomes). We assume simulated games are independent, though this could be relaxed in the future (see 'Discussion'). In keeping with our focus on fast, resource-limited reasoning, only a small number of game simulations are used (which we estimated empirically, see results and Methods) in line with previous evidence that people use a relatively small number of simulations to form beliefs and make decisions^{21–26}.

Alternative models

Our Intuitive Gamer framework naturally extends to model players with more or less competence and motivation. We can think of the Intuitive Gamer as operating under the same mechanisms as more expert models, but scaled down (offering, therefore, a path to scale back up). We can vary parameters in our framework (value function and search depth) to instantiate a version of the kind of cognitive models that have been successful at capturing human experts in a single 'four-in-a-row' game⁶. We refer to this variant as the 'Expert Gamer', which differs from

the Intuitive Gamer by conducting a deeper search (approximately depth-5) and using a more sophisticated value function (Methods). On the other end of the spectrum, we consider unmotivated players who select actions uniformly at random (Fig. 2b), which we refer to as the 'Random Gamer'. We also compare reasoning judgements to Monte Carlo tree search (MCTS)^{4,27,28}, an alternative tree-search approach that is even more computationally intensive and popular in AI systems for strong general game playing (see Methods for contrast with our Intuitive Gamer and other cognitive models). Finally, we compare against models that do not involve explicit structured game simulation, operating either over pre-computed linguistic features from the game description alone (Methods and Supplementary Information section 5) or more flexible linguistic computable models, for example, language models (Supplementary Information section 5).

Resource-rational reasoning

The Intuitive Gamer occupies an interesting point on the Pareto frontier of compute efficiency and game reasoning sophistication. On the one hand, the Intuitive Gamer is orders of magnitude faster than more sophisticated game reasoners (MCTS and the Expert Gamer) in wall clock time and number of board states evaluated (Extended Data Table 2a and Supplementary Information). The Intuitive Gamer is approximately 700× faster than the Expert Gamer in wall clock time, with approximately 500× fewer board evaluations, and is nearly




Fig. 3 | Evaluating games without ever playing. A total of 238 participants judged the expected payoff of games drawn from our 121-game suite. **a**, Expected payoff predicted by humans and by alternative models for each game. Each point represents the payoff for one of the $n = 121$ game stimuli. Error bars show 95% CIs around the mean human estimate and the mean model prediction from $k = 6$ simulations sampled from 20 simulated participants. **b**, Variance explained (R^2) by the full Intuitive Gamer and lesioned variants. Removing flatness, probabilistic simulation or goal-directedness reduces the fit to human payoff judgements. The dashed line indicates the mean and 95% CI split-half R^2 on the human-predicted payoffs, indicating that the full Intuitive Gamer model captures essentially all of the explainable variance that is not due to noise.

40,000 $\times$ faster than MCTS, with almost 10,000 $\times$ fewer node evaluations, as computed under self-play game simulations for the respective models. On the other hand, although not as close to game-theoretic optimal play as the Expert Gamer or MCTS, game evaluations under the Intuitive Gamer are still well-correlated with game-theoretic optimal analyses and far more aligned with game-theoretic optimal play than random play (Extended Data Table 2b and Methods). The Intuitive Gamer's mechanisms of fast, flat goal-directed probabilistic simulation thus constitute a highly efficient and resource-rational approach to game reasoning, which we hypothesize people could plausibly and productively engage when encountering new games.

Games to study novice game reasoning

To explore how people reason about new games within the bounds of a behavioural study, we need a broad collection of games that are novel without requiring lengthy demonstrations or explanations. To that end, we implement a range of two-player, grid-based strategy games derived from classic $M \times N \times K$ games (examples include tic-tac-toe and gomoku, in which players take turns placing pieces on an $M \times N$ grid trying to make K pieces in a row—a subgenre of games frequently studied in previous work^{6,29,30}). Our set of 121 distinct games covers a variety of environment specifications (for example, the size and shape of the board), transition dynamics (for example, the number of moves a player makes on their turn) and win conditions (for example, whether completing a line results in a win or a loss). In addition, games vary in their duration and balance under optimal play (Methods). In each case, however, the core mechanic of the game (that is, placing a piece into an empty grid cell) remains both consistent and familiar. See examples in Fig. 1a and Extended Data Table 1.

Reasoning about games before any play

We first assess how people evaluate a game (Fig. 1b), from 'just thinking' about the game alone—before any play. We consider two kinds of

c, Fit to the variance of human payoff judgements as a function of the number of simulations k . A small number of simulations ($k \approx 5-7$; we use $k = 6$) best captures the variability in participants' judgements. Fit is measured by the root mean squared error (RMSE) between human and model variance by game and by Wasserstein distance between their variance distributions across games. Full scatterplots are shown in Extended Data Fig. 1, and details of our sample complexity analysis are in the Methods. **d**, Expected payoff predicted by humans and by game-theoretic optimal analysis for the 78 games (out of 121 games) for which an optimal payoff could be computed. Error bars depict 95% CIs around the bootstrapped mean human prediction per game.

game evaluation: how people assess the expected outcomes of a game and a more subjective assessment of whether the game is likely to be engaging to play (its funniness).

Is the game likely to be fair?

We recruited 238 participants to evaluate the expected outcomes of the game in a 'zero-shot, zero-experience' experiment. Participants 'just thought' about the games, estimating both the likelihood of a draw and the likelihood that the first player would win if the game did not end in a draw for games solely on the basis of a natural language description of the rules and a depiction of a blank board. These estimates define an expected payoff of the game from the perspective of the first player (Methods), which in turn encodes an intuitive notion of fairness (that is, whether a game is biased towards a particular player).

We assessed how each characteristic of the Intuitive Gamer model (fastness, flatness, goal-directedness and probabilistic simulation) contributes to people's payoff predictions by comparing with alternative models that increase or decrease the sophistication of each component (Fig. 3a). The Intuitive Gamer model correlates well with human estimates ($R^2 = 0.81$ (95% confidence interval (CI), 0.77, 0.85)), matching the total explainable variance from the human data as estimated by split-half correlations ($R^2 = 0.82$ (95% CI, 0.77, 0.86)). The Intuitive Gamer captures human judgements significantly better than alternative models that vary in sophistication. It outperforms the Random Gamer model ($R^2 = 0.47$ (95% CI, 0.41, 0.54)), Expert Gamer model ($R^2 = 0.65$ (95% CI, 0.60, 0.70)), MCTS baseline ($R^2 = 0.60$ (95% CI, 0.55, 0.64)) and a variant of the Intuitive Gamer that is not probabilistic ($R^2 = 0.53$ (95% CI, 0.48, 0.57)).

To more directly probe each component of the Intuitive Gamer model, we conducted a series of experiments in which we varied key model components singly and in combination (see Supplementary Information section 1 for a complete summary). The base model can be parameterized by a series of binary choices: whether




Fig. 4 | Evaluating whether games are likely to be fun before ever playing them. A total of 246 participants judged whether they thought the games would be fun. **a**, Representative games rated by people as more or less fun, with position on the spectrum arranged approximately according to the mean of participants' funness ratings. **b–d**, These funness ratings are related to several game features that can be read out under Intuitive Gamer model simulations: game balance (**b**); predicted advantage over a random agent given the Intuitive Gamer model's gameplay (reward for thinking; **c**); and expected game length, as fit with a quadratic term to account for a nonlinear relationship to funness (**d**). **e**, A regression model with all features predicted under the Intuitive Gamer

model (balance, reward for thinking, and length) captures much of the explainable variance in the human funness ratings ($R^2 = 0.57$ (95% CI, 0.51, 0.63) compared with split-half $R^2 = 0.60$ (95% CI, 0.51, 0.68)). Error bars depict 95% CIs around the human mean funness per game; error bars for the model depict 95% CI over the predicted mean for each of the fit models (models are fit to each bootstrap subsample). Black circles show games from **a**, highlighting that funness judgements under the Intuitive Gamer generally align with the ordering of funness of games judged by people. A full listing of the games and their average human funness scores are in Supplementary Information section 8.

to be goal-directed in value assessments; whether to be probabilistic; and whether to be flat. Modifying any of these factors relative to the full model impairs the fit to human payoff judgements (Fig. 3b). 'Fastness' can be further varied by modulating the number of simulations the Intuitive Gamer reasoning module takes when assessing payoff. A small number of simulations (k more than one, but less than ten) best captures the variance in human prediction (Fig. 3c and Extended Data Fig. 1; see complexity analysis details in Methods). These comparisons suggest that human payoff judgements of novel games are well-modelled by efficient and compute-limited simulation rather than purely random exploration of game states or deeper search and computationally intensive simulation. For simplicity, this estimate assumes that all simulations are independent and run to the end of the game; allowing simulations to probabilistically terminate early (Methods) yields similar fits in the expected payoff evaluations and number of mental simulations per participant (Extended Data Fig. 2).

We were able to estimate the optimal game-theoretic payoff for many (78 of 121; Methods) of the games, allowing us to compare people's subjective judgements with a fully objective game evaluation. Human judgements are reasonable relative to the game-theoretic optimal ($R^2 = 0.62$ (95% CI, 0.57, 0.67)). Although human judgements generally track the direction of the game-theoretic optimal (for example,

estimating a payoff greater than zero when the first player should definitely win), the Intuitive Gamer provides a superior qualitative and quantitative fit to human judgements.

Is the game likely to be fun?

We ran a parallel 'zero-shot zero-experience' study as described above on a new group of 246 participants who were asked to make a subjective estimate of how fun a game is likely to be, before ever playing, to test whether the Intuitive Gamer's reasoning process can also account for how people make more subjective evaluations about new problems. Games varied widely in their judged funness (Fig. 4a and Supplementary Information section 8), and funness judgements generally varied more across individuals than judgements of expected payoff, which is reasonable given that funness is both more vague and a more subjective evaluation. Games that were judged most fun tended to have larger boards (often 10×10) and win conditions connecting a moderate number of pieces in a row (often 4 or 5), but there is no simple relation between these features and funness: many 10×10 games were judged below average (Extended Data Fig. 3b), and the single most fun game was a 5×5 board with a misère rule, where 3 in a row loses.

We hypothesized that several features of a game that can be read off easily from our model's fast and flat probabilistic simulations—how balanced a game is, how much a game rewards thinking, and how long

a game is—could quantitatively explain much (or even most) of the variance in people’s funniness judgements. Our measure of balance captures the notion that people prefer games they expect will be decisive (unlikely to end in a draw), where both players have an equal chance to win. Our thinking reward captures the relative advantage of playing strategically (‘thinking’) versus moving haphazardly (‘without thinking’). The effect of game length is modelled as an inverted U-shape of expected number of moves to a draw or win, reflecting the intuition that longer and more involved games tend to be more fun, unless they extend unnecessarily. (For details of how these game features are computed, see Methods.)

Each of these three features computed under the Intuitive Gamer model captures significant variance in human funniness judgements (Fig. 4b–d and Extended Data Fig. 4a,b). A simple regression model that combines these features (including a quadratic game length term) attains $R^2 = 0.57$ (95% CI: 0.51, 0.63) (Fig. 4e), which approaches the total variance explainable in the human data based on bootstrapped split-half R^2 from participants over all games ($R^2 = 0.60$ (95% CI: 0.51, 0.68)). This model also quantitatively explains the qualitative trends we observed in funniness judgements, capturing the nonlinear relation between funniness, board size and the number K in a row to win (Extended Data Fig. 3).

Finally, a generalization test (Methods) showed that these fits are not sensitive to the choice of games, and that game features estimated from Intuitive Gamer simulations predict funniness judgements better than those from the Random or Expert Gamer models, using more or fewer compute resources (depth of thinking), or models that relied exclusively on surface-level game features (Extended Data Fig. 4c and Supplementary Information section 5.7).

Taken together, these two experiments in which participants ‘just think’ to rate both objective measures of a game (for example, its payoff or fairness) and more subjective evaluations (for example, its enjoyability) lend support to our hypothesis that these kinds of novice judgements are best captured by a reasoning module that is goal-directed but limited in its computational cost (that is, fast and flat) compared with alternatives that vary in terms of computational demands and game reasoning sophistication.

Decision-making in first-time gameplay

Relative to the flat, single-step thinking that our Intuitive Gamer model posits people use to evaluate new games, conventional models of how people actually play games⁶ perform much deeper tree search, simulating multiple future turns by both players to evaluate potential moves before choosing the best move at each step of a game. But these models have been evaluated only on players who are either experts in a game or have at least played a game many times. Here, we consider what human players do when they are playing a game for the very first time. That is, we ask whether they search deeply, as previous models might suggest, or whether they use only a fast and flat decision mechanism, as our Intuitive Gamer model does—and as our results suggest they do when merely imagining how a game might unfold before playing it.

What moves do people make?

To test whether the Intuitive Gamer player module captures how people actually play a game for the first time, we recruited 302 participants to play each other in a subset of 40 of the 121 game variations, plus tic-tac-toe (Fig. 1c). Each participant played only a single round of any given game.

Initial evidence that the Intuitive Gamer model captures not only how people think about new games but also how they play them for the first time comes from comparing the distribution of actual game outcomes (first player win, second player win or draw) in this gameplay experiment with the outcome distributions to be expected if both players

make move decisions according to the model. The model captures most of the variance in actual game payoffs with novice players ($R^2 = 0.72$ (95% CI: 0.67, 0.76)) and predicts these payoffs as well as human estimates of expected game payoffs do, from our earlier game evaluation experiment (Extended Data Fig. 5).

A more direct test is to compare predicted move choices that the Intuitive Gamer model makes at each turn of each game with the actual moves that new players made at the same game steps when they were playing these games for the first time. Because there is a strong element of randomness in people’s move choices as well as the model’s predictions, we evaluated the likelihood of people’s moves under the Intuitive Gamer relative to alternative models that are either more or less sophisticated and computationally intensive, the Expert Gamer and Random Gamer models from our previous study, respectively (Fig. 5a). The Intuitive Gamer model best predicts players’ moves in aggregate, as measured either by average per-move log likelihoods or per-match log likelihoods (summed over all moves within a match), across 1,808 distinct matches with a total of 9,892 moves. All of these differences are highly significant as measured by paired t -tests (mean differences in total per-match log likelihood for the Intuitive Gamer versus Expert model = 2.83, $t_{1807} = 29.4$, $P < 0.001$; versus Random model = 3.73, $t_{1807} = 26.4$, $P < 0.001$; mean differences in per-move log likelihood for the Intuitive Gamer versus Expert model = 0.51, $t_{9891} = 39.9$, $P < 0.001$; versus Random model = 0.68, $t_{9891} = 42.3$, $P < 0.001$).

We also fit these three models of action choice to the distribution of all moves people made in each individual game and the distribution of all moves an individual player made across the different games they played using probabilistic (admixture) models (Supplementary Information). At a per-game level, the Intuitive Gamer model generally, but not always, fits human play better than the Expert or Random models: it captures more than 50% of the probability distribution of moves in 32 out of 41 games, and a plurality of moves in 5 of the remaining games (Fig. 5b). Individual player’s distributions of moves are also best fit by the Intuitive Gamer: it captures more than 50% of the probability distribution of moves for 243 out of 302 players, and a plurality of moves in 16 of the remaining players (Extended Data Fig. 6).

Post hoc exploratory analyses into intermediate stages of gameplay suggest that the Intuitive Gamer model best fits early and mid-stage play, although in later-stage play, greater depth variants are more comparable (Supplementary Information). This could be due to some people being more motivated to think more deeply at the end of the game or finding it easier to think more deeply with a smaller space of moves to consider (as there are fewer legal moves at the end of the game).

In addition, as in the game evaluation experiments, both aspects of goal-directedness in the Intuitive Gamer value function (offensive progress towards your goal and defensive assessment of how to block your opponent’s progress) matter for capturing human behaviour in novel games (Extended Data Fig. 7). Overall, our results support the hypothesis that when playing a novel game for the first time, people adopt a similar fast, flat and goal-directed approach to choosing their next move as they appear to do in mentally simulating games before they start to play.

What move should someone else make?

In the previous experiment, as in almost all studies of gameplay, participants made only a single choice for each move, and they played each game only once. Yet our model posits that people think about new problems by running probabilistic mental simulations, letting them form graded expectations about the value of a range of possible actions. To assess how well the Intuitive Gamer player module captures people’s probabilistic expectations, we recruited a new group of participants to predict distributions of likely actions when watching videos of other novices playing these games (Methods and Fig. 1d).




Fig. 5 | Modelling people's actions and distribution over predicted actions in the first encounter with a new game. **a**, The average log-likelihood of human moves under three variants of the game player module. Data come from 302 participants playing a single round each of 5 novel games (along with tic-tac-toe); across participants, 40 of our novel games were tested. The Intuitive Gamer model captures players' moves significantly better than alternatives. Error bars depict 95% bootstrapped CIs. **b**, The estimated contribution of each model in a probabilistic mixture (admixture model; Methods) fit to participants' moves for each game. The Intuitive Gamer model is the component that best explains participants' moves in approximately 90% of games (37 of 41); games in which people would not be novices (tic-tac-toe) or some particularly atypical games (for example, misère games, where the first person to get K in a row loses, listed with an 'L' as in '7 × 7 4 L') are comparatively less well captured by the Intuitive Gamer compared with alternative models. Game descriptions are lemmatized to show board size (for example, 7 × 7) and

key features (for example, HV indicates only horizontal and vertical wins are allowed; P22p means player 2 can move twice on their first turn; the number after board size indicates the number in a row needed to win). A full description of lemmatized game codes are in Methods. **c**, TVD (lower is better) between human and model-predicted distributions over next moves for the same three models. Judgements come from a new group of 314 participants who estimated where a participant should move next in frozen videos from the gameplay experiment. The dashed line is the split-half human TVD (effectively the noise ceiling). **d**, Mean TVD by game. The Intuitive Gamer generally best matches human judgement distributions, with similar failures to the play data (for example, for misère games). Error bars depict standard deviation over boards for each game. **e**, Decisions of whether to accept or reject a draw when requested are captured by the expected value of the player's current position, but participants tend to be willing to reject a draw and play out a game with a slightly lower expected value if the game is judged to be fun.

The Intuitive Gamer model generally better captures the distribution of moves predicted by people compared with alternative models across the 249 game boards for each match and game stage, as measured by the total variance distance (TVD) between model and people's predicted distribution, and is near the expected noise ceiling (computed via split-half) human judgements for most game boards (Fig. 5c). Differences in distance (where lower means closer to people's predicted action distribution) are highly statistically significant as measured by paired t -tests in the TVD per game board for the Intuitive Gamer versus Expert model = -0.15 , $t_{248} = -15.0$, $P < 0.001$ and versus Random model = -0.09 , $t_{248} = -7.5$, $P < 0.001$. Exceptions for poor fits align with those found in the play experiments (for example, on misère games; see Fig. 5d and example boards in Fig. 6). Fits are robust to choice of distributional measure (Supplementary Information). People assign similar log probability on the move actually played by the human player to that of the Intuitive Gamer, further highlighting alignment of the Intuitive Gamer to people's probabilistic judgements (Supplementary Information).

When is a game worth continuing?

Game evaluation is important not only when deciding whether a game is worth playing but also for assessing—during play—whether a game is worth continuing to engage with. In our gameplay study, participants had the option to request a draw at any point in play. If they received a draw request from another player, they had the choice of whether to accept the draw and immediately terminate play or reject the draw and keep playing. Over the 1,808 matches participants played, players made

142 draw requests, of which 83 were accepted and 59 were rejected. We next considered how people decide whether or not to keep playing when they receive a draw request.

We modelled these choices as probabilistic value-based decisions trading off a player's expected reward of winning with the expected cost of continuing to play, both of which can be computed under the Intuitive Gamer reasoning module conditioning on the current board state (rather than a blank opening board, as in our earlier game evaluation studies; Methods). In exploratory analyses, we also hypothesized that a player might be willing to keep playing a game even if the expected cost of continuing outweighs the potential objective monetary reward, if playing provides an alternative subjective reward, for example, if it is likely to be fun. A logistic regression model fit to participants' draw request decisions using these three features (Methods) captures the intuition that players' game-time evaluations are probabilistically rational under limited resources (time) and sensitive to subjective factors such as the engagement of games in addition to their objective value (Fig. 5e). Ablations of the Intuitive Gamer evaluation function lead to significantly worse qualitative and quantitative fits to human draw request choices (Extended Data Fig. 8a,b).

Discussion

Systems of rules and rewards govern many aspects of human behaviour—from jobs and institutions to games and geopolitics. A core question is, how can people think reasonably about any such system or problem for the first time, so flexibly and quickly? We formalized and




Fig. 6 | Example human- and model-predicted distributions over the next action in real games. a–j. Each panel shows a representative state taken from a different game between two human players, along with the actual move one player took at the next step and the predictions made by other humans and two models for how that move should have been made. Filled black and white circles show the moves both players took prior to this point in the game. The one open circle on each board indicates where the human participant whose turn it was on this move (white or black) actually played. The top, middle and bottom boards show the predicted next-move distributions from human participants in the watch-and-predict study, the Intuitive Gamer and the Expert Gamer, respectively. The purple shading gives the predicted probability of choosing each empty square. Humans often distribute probability over several plausible

moves (a–e), although sometimes they strongly favour a single move (f,g); the Intuitive Gamer often captures this pattern. The Expert Gamer model fits human predictions less well, in some cases because it allocates high probability to a small set of high-utility moves that are often counter-intuitive to naive players (a–c,e); and in other cases by allocating diffuse probability across many states when deep search indicates a sure loss under perfect play (d,f). The Intuitive Gamer does not always fit best: the Expert Gamer can sometimes better match human confidence (g), and the Intuitive Gamer can be too sharp in its predictions (h). In some cases, both models are sharper than humans, including crowded boards where people may miss a winning or blocking move (i) and misère variants (j). The full set of predictions (for 249 moves across 41 distinct games) is included in Supplementary Information.

empirically evaluated a computational model of an Intuitive Gamer, which draws on fast and flat goal-directed probabilistic mental simulations to make rich quantitative predictions about how novice players judge the objective value of a previously unseen game, as well as to make more subjective evaluations. Our model also captures how people act within these games, and does so better than alternative models with more and less computational power. The same model generalizes across these different tasks and the many different games we study.

Our work differs from most previous studies of gameplay in cognitive science and AI, which focus on how expert play is achieved, and typically consider only expertise in a single game. By contrast, we study a class of 121 strategic games with varying dynamics and highlight the phase of ‘pre-expertise’: the critical first moments before people can rely on the benefits of repeated experience, but when they must nonetheless take reasonable actions and even decide whether to play at all—or to continue playing. Studying pre-expertise across varied domains is important not just for our understanding of cognition generally, but also to construct more accurate models of human behaviour in game

theory under novel economic mechanisms^{31,32} or alongside unfamiliar agents such as AI systems^{33,34}. The Intuitive Gamer approach could inform the design of more efficient AI systems that make reasonable judgements in new multi-agent settings despite using much less compute than conventional expert-level models.

Our work also differs from prior studies of funness in games, which have focused on identifying factors that lead people to experience a game as fun. By contrast, we have focused on modelling the intuitive theories people have about how they would expect to experience a game as fun, before playing it for the first time. Features identified in the previous literature, such as strategic depth³⁵ and competence³⁶ reflecting the need for challenge, or learning progress³⁷ and thinking progress³⁸, do align with those in our model, especially the reward for thinking component. Future work could study whether other aspects of motivation identified in previous gaming literature—autonomy, social connectedness, immersion and ease of control³⁶—also contribute to people’s intuitions about funness, and whether these could also be modelled using a generalization of our Intuitive Gamer model.

As a model of human reasoning, our work at present is limited to a subset of perfect information, two-player competitive board games, and these games reflect novelty within a genre that many people are familiar with rather than completely novel experience. It is an open question how the heuristics our model uses could generalize to more complex games in the same class, such as Go or chess, and to entirely new classes of games. Many extensions are possible, such as generalizing the contiguous-line features that we consider towards connectivity or territory metrics that reflect partial or graded progress towards a game's win conditions. Natural language could be used to suggest short programs expressing such heuristics, as in verifiable code-based rewards in reinforcement learning³⁹. Although we hypothesize that key features of the Intuitive Gamer reasoning module—being fast, with a few probabilistic simulations of play, where play is assumed to involve goal-directed agents—would extend to other game classes, it is an open question how 'flat' reasoning might be and what more general heuristics might capture intuitive reasoning in the wider range of games that characterize human social life. Understanding peoples' general game reasoning on a wider class of games beyond competitive settings, for example, cooperative or multi-agent contexts, is necessary to move towards a general model of people's 'intuitive theory of games', analogous to the computational models of intuitive theories of physics or intuitive theories of mind that cognitive scientists have constructed^{40,41}.

In addition, there is a need for more fine-grained process-level and individual-level accounts of how people reason about novel games. More work should study how people judge the expected outcome of a game if they stop imagining play before reaching a win or draw, if their simulations are temporally dependent, or when they do not simulate at all—and how models can best capture these complexities. Our model also does not address how people may learn and adapt such heuristics within or across games, limitations that suggest important directions for future work. How mental computation for reasoning about familiar games may shape thinking about a new game⁴², especially in high-stakes situations²⁶ or when time is limited or thinking is costly^{19,20}, and how in-game experience^{43,44} and preferences for different kinds of play and strategy⁴⁵ influence game reasoning, will also probably be important for modelling variability in peoples' thinking.

Finally, people also make new games and new problems, either by modifying old ones to make them more challenging or more fun, or by creating entirely novel goals and pursuits—which has led to some of humanity's greatest discoveries and innovations³⁸. In ongoing work, we are exploring whether models based on fast and flat simulation can also explain the moves people make in these 'meta-games' of game-creation and modification.

Science itself is such a 'meta-game': a family of games against nature⁴⁶, where scientists set the rules that lead to the most productive research⁴⁷. Similarly, mathematicians may explore conjectures and proof strategies by imagining adding or relaxing the logical rules of how a problem is formulated⁴⁸. Suppose, more speculatively, that these and other creative activities at the frontiers of human knowledge do in fact share something deeply in common with intuitive reasoning about games, or even children's play. Then, perhaps mechanisms such as the fast probabilistic mental simulations studied here could help to explain aspects of how researchers choose to approach a problem, or which problems to think about in the first place. Whether scientific thinking can be meaningfully captured by such models is very much an open question, but one we are now better positioned to investigate. At the very least, it feels like a game worth playing.

Online content

Any methods, additional references, Nature Portfolio reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions

and competing interests; and statements of data and code availability are available at <https://doi.org/10.1038/s41586-026-10722-1>.

- Chase, W. G. & Simon, H. A. in *Visual Information Processing* (ed. Chase, W. G.) 215–281 (Elsevier, 1973).
- Campbell, M., Hoane, A. J. Jr & Hsu, F.-h. Deep Blue. *Artif. Intell.* **134**, 57–83 (2002).
- Gobet, F., Retschitzki, J. & de Voogt, A. *Moves in Mind: the Psychology of Board Games* (Psychology Press, 2004).
- Silver, D. et al. Mastering the game of Go with deep neural networks and tree search. *Nature* **529**, 484–489 (2016).
- Silver, D. et al. Mastering the game of go without human knowledge. *Nature* **550**, 354–359 (2017).
- van Opheusden, B. et al. Expertise increases planning depth in human gameplay. *Nature* **618**, 1000–1005 (2023).
- Cleveland, A. A. The psychology of chess and of learning to play it. *Am. J. Psychol.* **18**, 269–308 (1907).
- Shannon, C. E. XXII. Programming a computer for playing chess. *Lond. Edinb. Dublin Philos. Mag. J. Sci.* **41**, 256–275 (1950).
- Newell, A., Shaw, J. C. & Simon, H. A. Chess-playing programs and the problem of complexity. *IBM J. Res. Dev.* **2**, 320–335 (1958).
- Mnih, V. et al. Human-level control through deep reinforcement learning. *Nature* **518**, 529–533 (2015).
- Tsivdidi, P. et al. Human-level learning of complex novel tasks as theory-based modelling, exploration and planning. *Philos. Trans. A Math. Phys. Eng. Sci.* **384**, 20240529 (2026).
- Yannakakis, G. N. & Togelius, J. *Artificial Intelligence and Games* (Springer, 2018).
- Simon, H. & Chase, W. in *Computer Chess Compendium* (ed. Levy, D.) 175–188 (Springer, 1988).
- Silver, D. et al. Mastering chess and shogi by self-play with a general reinforcement learning algorithm. Preprint at <https://arxiv.org/abs/1712.01815> (2017).
- FAIR et al. Human-level play in the game of diplomacy by combining language models with strategic reasoning. *Science* **378**, 1067–1074 (2022).
- Newell, A. & Simon, H. A. *Human Problem Solving* (Prentice-Hall, 1972).
- Charness, N. in *Toward a General Theory of Expertise: Prospects and Limits* (eds Ericsson, K. A. & Smith, J.) 39–63 (Cambridge Univ. Press, 1991).
- Kocsis, L. & Szepesvari, C. Bandit based Monte-Carlo planning. In *European Conference on Machine Learning* 282–293 (Springer, 2006).
- Gigerenzer, G. in *Blackwell Handbook of Judgment and Decision Making* (eds Koehler, D. J. & Harvey, N.) 62–88 (Blackwell, 2004).
- Lieder, F. & Griffiths, T. L. Resource-rational analysis: understanding human cognition as the optimal use of limited computational resources. *Behav. Brain Sci.* **43**, e1 (2020).
- Griffiths, T. L., Sanborn, A. N., Canini, K. R. & Navarro, D. J. in *The Probabilistic Mind: Prospects for Bayesian Cognitive Science* (eds Chater, N. & Oaksford, M.) 303–328 (Oxford Univ. Press, 2008).
- Sanborn, A. N., Griffiths, T. L. & Navarro, D. J. Rational approximations to rational models: alternative algorithms for category learning. *Psychol. Rev.* **117**, 1144 (2010).
- Vul, E., Goodman, N., Griffiths, T. L. & Tenenbaum, J. B. One and done? Optimal decisions from very few samples. *Cogn. Sci.* **38**, 599–637 (2014).
- Icard, T. Subjective probability as sampling propensity. *Rev. Philos. Psychol.* **7**, 863–903 (2016).
- Zhu, J.-Q., Sanborn, A. N. & Chater, N. The Bayesian sampler: generic Bayesian inference causes incoherence in human probability judgments. *Psychol. Rev.* **127**, 719 (2020).
- Klein, G. A. et al. in *Decision Making in Action: Models and Methods* (eds Klein, G. A., Orasanu, J., Calderwood, R. & Zsombok, C. E.) 138–147 (Ablex, 1993).
- Coulom, R. Efficient selectivity and backup operators in monte-carlo tree search. In *International Conference on Computers and Games* 72–83 (Springer, 2006).
- Genesereth, M. & Thielscher, M. *General Game Playing Synthesis Lectures on Artificial Intelligence and Machine Learning* (Morgan & Claypool, 2014).
- Crowley, K. & Siegler, R. S. Flexible strategy use in young children's tic-tac-toe. *Cogn. Sci.* **17**, 531–561 (1993).
- Amir, O., Tyomkin, L. & Hart, Y. Adaptive search space pruning in complex strategic problems. *PLoS Comput. Biol.* **18**, e1010358 (2022).
- Myerson, R. B. Mechanism design by an informed principal. *Econometrica* **51**, 1767–1797 (1983).
- Koster, R. et al. Human-centred mechanism design with democratic ai. *Nat. Hum. Behav.* **6**, 1398–1407 (2022).
- Carroll, M. et al. On the utility of learning about humans for human–AI coordination. *Adv. Neural Inf. Process. Syst.* **32**, 5174–5185 (2019).
- Collins, K. M. et al. Building machines that learn and think with people. *Nat. Hum. Behav.* **8**, 1851–1863 (2024).
- Lantz, F., Isaksen, A., Jaffe, A., Nealen, A. & Togelius, J. Depth in strategic games. In *The AAAI Workshop on What's Next for AI in Games?* <https://cdn.aaai.org/ocs/ws/ws0340/15142-68464-1-PB.pdf> (AAAI, 2017).
- Ryan, R. M., Rigby, C. S. & Przybylski, A. The motivational pull of video games: a self-determination theory approach. *Motiv. Emot.* **30**, 344–360 (2006).
- Brändle, F., Wu, C. M. & Schulz, E. Leveling up fun: learning progress, expectations, and success influence enjoyment in video games. *Sci. Rep.* **15**, 34153 (2025).
- Chu, J., Tenenbaum, J. B. & Schulz, L. E. In praise of folly: flexible goals and human cognition. *Trends Cogn. Sci.* **28**, 628–642 (2023).
- Lehrach, W. et al. Code world models for general game playing. In *International Conference on Learning Representations (ICLR)* (2026).
- Baker, C. L., Jara-Ettinger, J., Saxe, R. & Tenenbaum, J. B. Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nat. Hum. Behav.* **1**, 0064 (2017).

41. Ullman, T. D., Spelke, E., Battaglia, P. & Tenenbaum, J. B. Mind games: game engines as an architecture for intuitive physics. *Trends Cogn. Sci.* **21**, 649–665 (2017).
42. Dasgupta, I., Schulz, E., Goodman, N. D. & Gershman, S. J. Remembrance of inferences past: amortization in human hypothesis generation. *Cognition* **178**, 67–81 (2018).
43. Kolb, D. A. *Experiential Learning: Experience as the Source of Learning and Development* (FT Press, 2014).
44. Rutledge, R. B., Skandali, N., Dayan, P. & Dolan, R. J. A computational and neural model of momentary subjective well-being. *Proc. Natl Acad. Sci. USA* **111**, 12252–12257 (2014).
45. Nguyen, C. T. Games and the art of agency. *Philos. Rev.* **128**, 423–462 (2019).
46. Hintikka, J. On the logic of an interrogative model of scientific inquiry. *Synthese* **47**, 69–83 (1981).
47. Laudan, L. *Progress and its Problems: Towards A Theory of Scientific Growth* Vol. 282 (Univ. California Press, 1978).
48. Thurston, W. P. On proof and progress in mathematics. *For Learn. Math.* **15**, 29–37 (1995).

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2026

Methods

Game construction

We manually constructed the 121 two-player competitive strategy game variants played on $M \times N$ grids, where M is the number of rows and N is the number of columns. One goal in creating these games was to ensure there is enough diversity in board size and rule structure as well as systematic variance. To that end, we designed a series of variations on square boards: on 10×10 boards, we have K -in-a-row for K from 2 to 10; for 3-in-a-row, we have M by $N = M$ boards from 3 to 10; we also include a few 5×5 boards with varying K . We additionally included a few other square boards varying in complexity, as well as rectangular boards ranging in size from 1×5 to 5×10 , and we have integer $2 \leq n \leq 6$ for K -in-a-row. To assess how people reason about games that are not physically realizable, we included three ‘infinitely’ sized games with $K = 3, 5$ and 10 for K -in-a-row. These categories give us 41 games with typical ‘ $M-N-K$ rules’ (for example, where the players take turns and have the same objective: ‘make K in a row, where horizontal, vertical, or diagonal all count’). This set includes the standard tic-tac-toe as well as $4 \times 9, 4$ in a row wins from ref. 6. We then created a number of games with varied game rules. We kept the selection of board sizes and K -in-a-row fixed across categories (ten games within each category). The selection included $3 \times 3, 4 \times 4, 5 \times 5$ and 10×10 boards and $n \in \{3, 4, 5, 10\}$. We also designed games with more atypical rules, varying the winning conditions (for example, K -in-a-row loses and diagonal connections do not count as wins) and first-mover dynamics (for example, player 1 can place two pieces on their first turn). These categories give us an additional 80 games, totalling the 121 in our dataset. We manually implemented an automated win condition checker that permits flexible game assessment over all game types.

Game name codes. Games are expressed in abbreviated form throughout the paper. Games are described by their board size (rows $\times$ columns) and the number K in a row to win. For example, ‘ $4 \times 4, 3$ ’ means the game is played on a 4×4 grid and the first person to get 3 pieces in a row wins. Unless otherwise stated, horizontal, vertical and diagonal all count. If only horizontal and vertical 3 in a row count, the game would be written as ‘ $4 \times 4, 3$ HV’. If only diagonal counts, then the game will be written as ‘ $4 \times 4, 3$ D’. If the constraint only applies to one player (for example, P1 can only win horizontally and vertically, and P2 can win any way), that is represented as ‘ $4 \times 4, 3$ (P1 HV)’. If one player can go twice on their turn, for example, the second player (P2) can play twice, that is written as ‘ $4 \times 4, 3$ (P2 2p)’. A misère game (where first to K in a row loses) is written with an ‘L’, for example, ‘ $4 \times 4, 3$ L’. In our dataset, multiple rule modifications cannot co-occur, so each game can be expressed in this abbreviated way.

The full game natural language game descriptions were provided to participants (for example, as shown in Extended Data Table 1). The game codes are used only for ease of presentation in this paper.

Human experiments

All human experiments were conducted under prior approval from the institutional review board at the Massachusetts Institute of Technology through the Computational Cognitive Science Lab. All participants provided informed consent.

Zero-shot outcome evaluation experiment. We recruited 238 participants from Prolific⁴⁹ to judge novel games. Each participant was randomly presented with 10 games sampled from our 121 diverse game stimuli, as well as regular tic-tac-toe (won by making 3 in a row on a 3×3 board) to set baselines for game judgements. We collected approximately 20 judgements per game stimulus for each game reasoning query. Participants were paid at a base rate of US\$12.50 per hour with an optional bonus up to US\$15 per hour; the full experiment approximately took 25 minutes.

Participants were instructed to evaluate likely game outcomes. Specifically, participants produced judgements on a continuous 0 to 100 probability scale to predict the likelihood of a first player win (“If the game does not end in a draw, assuming both players play reasonably, how likely is it that the first player is going to win (not draw)?”) and a draw (“Assuming both players play reasonably, how likely is the game to end in a draw?”). Judgements were made using sliders. Both game outcome question sliders appeared on the same page.

Participants produced judgements about each game based on a linguistic game specification. We additionally provided participants with an interactive scratchpad board that they were told they could, but were not required to, use to inform their judgements. The scratchpad was automatically sized to the board of the game specified; infinite boards were presented on a 13×13 grid with dashed lines indicating the board could continue. The scratchpad permitted automatically placing pieces of different colours (‘red’ and ‘blue’ to simulate different players); participants could force the next play to be made by the same player (for example, colour red twice in a row) by pressing the spacebar. Two buttons appeared below the scratchpad, permitting the user to either undo their last move or clear the screen to begin a new ‘game’. See Supplementary Information section 4 for examples of the experiment interface. Participants were required to consider each game for at least 60 s before being allowed to make their game judgements. Outliers were determined as the 10% of the judgements farthest from the mean judgements of other participants (in terms of the summed distance from the two queries) and filtered out for each game.

After answering all game reasoning queries, we additionally asked participants to create a new grid-based game variant that they would find fun. Participants wrote a linguistic game specification, describing the board size and win conditions. As in the game judgement queries, participants were again provided an optional scratchpad and required to spend 60 s before submitting a response. The scratchpad enabled participants to try out the game they intended to create. After specifying a game, participants were asked to answer the same game reasoning query (either game outcomes or game fun) about their own game. These game-generation responses were not studied here and are actively being explored in a follow-up study. Example screenshots of experiment interfaces are included in Supplementary Information section 4.

Zero-shot funniness evaluation experiment. We repeated the same methodology above with a new group of 257 participants, replacing the game outcome questions with a question about game funniness. Participants instead assessed the expected funniness of the game (“How fun is this game?”) on a confidence scale spanning 0 (the least fun of this class of game) to 100 (the most fun of this class of game). 11 participants were filtered due to having provided non-effortful or AI-generated games in the creation stage, resulting in a total of 246 participants, and the same outlier filtering was applied per trial as in the outcome evaluation.

Zero-shot human–human play experiment. We recruited 302 participants to play these novel games in a pre-registered experiment. We selected a subset of 40 games from the full set of 121 to span a representative range of the gameplay variations (board shapes, board sizes and win rules) in the original dataset while generally favouring games that would not take very long to play in a live experiment. We randomly constructed eight batches of five of these games. Each participant additionally played one round of tic-tac-toe. The order of games was shuffled for each new set of participants.

Each participant was automatically paired with another player. We developed our interface using Empirica⁵⁰, which supports synchronous human–human pairing. Participants played one round from five different games. Players were informed they would get a bonus of US\$0.50 for every win. Participants had to spend at least 5 s reading the game

description before they began. We appended ‘Horizontal, vertical and diagonal all count’ to all game descriptions where any direction was allowed after we noticed some participants in pilots were confused as to which line directions would result in a win. Players were randomly assigned to move either first or second and a corresponding piece colour (red or blue). Players took turns making moves on the synchronous game interface. Players had no time limit on their turn.

Players were also allowed to request a draw or decide to surrender using buttons at the bottom of the interface. If a player surrendered, the game ended immediately (and that player lost). If a player requested a draw, the other player was allowed to either accept the draw (after which the game ended immediately and no player won) or reject the draw (leading the game to continue being played). Draw requests appeared as a popup banner for the other player. We include screenshots of the interface in Supplementary Information section 4.

The match ended when either a player won, a player surrendered, the board filled up completely (draw), or the players agreed to a draw. Both participants were informed about the game outcome. After each match, players made a judgement about either the expected outcomes of that game overall (with a new set of reasonable players) or the game’s funness (in a new match against a new player). Each pair of players was randomly assigned to either the outcome or funness rating condition. Judgements were made on a slider. Players were also presented with a ‘frozen’ version of the match on an example board with which they could replay all of the moves they and their opponent had made. Players also indicated how skilled they thought their opponent was at this game (“Out of 100 other random new players, where do you think the opponent you just played would rank in skill for this game?”). After the judgements were made, the players continued to the next, new game. At the end of the study, they filled out a text-based survey providing general information on their strategy and how fun they found the experiment. We filtered out 18 participants who did not pass our quality control (that is, they provided judgements that were ‘standard’ values (near 0, 50 or 100) on 80% or more of judgements) for a total of 284 subjects.

Watching and predicting play experiment. We recruited a new set of 314 participants, in a pre-registered experiment, to reason about the games zero-shot from only indirect experience: watching two other agents play. We selected a subset of 20 of the games from the previous human–human play study to ensure representation across game rules and dynamics. We also included tic-tac-toe (totalling 21 games). Participants watched a series of videos of other agents’ gameplay. Each video involved two humans playing each other, sourced from our live human–human gameplay experiment. We sampled four human–human played matches randomly from each of the 21 games (owing to a randomized batching error, only three unique matches were sampled for tic-tac-toe; hence, 249 game boards over the matches from 21 games and three stages per game), after filtering out any matches that ended preemptively from a draw request or surrender. For each match, we sampled three specific boards to be evaluated corresponding to the beginning, middle and end of the match. For the beginning and end boards, we randomly selected either the third or fourth move and the second-to-last or third-to-last move, respectively. For the middle board, we selected the median move. We filtered out any match that ended before eight moves. Participants watched one match from five different games, plus tic-tac-toe.

Before each match, participants were informed of the game rules and required to think about the rules for 5 s before the video began. We again appended “Horizontal, vertical, and diagonal all count” to all game descriptions where any direction was allowed. Videos played forward at a fixed rate, as in ref. 6. We chose 2 s per move to give viewers enough time to process each move without taking too long overall. Each video was stopped at the three time points described above. At

each stopping point, participants indicated their belief over where they thought the acting player should move next. Participants were given five clicks, which they could spread across the legal moves on the board to indicate their confidence that the player should move there. We chose five clicks to balance granularity of the elicited belief distribution against the burden on the participants. After each click, the opacity of the cell increased to indicate higher confidence. Participants were informed of the number of clicks they had left and could reset their clicks by clicking on a button below the interactive board.

After watching each video and indicating where they thought a player should move at each of the three timepoints, participants were then shown the remainder of the game as a board snapshot (cells indicated where players had moved and the order of play). Participants then answered either the same game outcome or funness judgements about the game overall, as described above. Judgements were made on a slider. Participants also indicated how skilled they thought each of the players was. We filtered out 10 participants who did not pass our quality control, leaving us with a total of 304 valid participants.

Problem formulation

Formally, one can think of a system, problem, or game G in terms of a space of feasible states $\mathcal{S}$; possible actions $\mathcal{A}$; rules $\mathcal{T}$ specifying valid actions and state transitions, and governing the overall dynamics; and goal functions mapping from states $\mathcal{S}$ to possible rewards $\mathcal{R}$. We are interested in how a reasoner infers properties of a new game $\psi(G)$ (for example, whether a game is likely to end in a draw or have a bias towards a particular player; whether the game is likely to be engaging and fun) as well as a policy $\pi_G(a_t|s_t)$ for how to play (choosing actions given the current state at time t , to achieve their goal, for example, to win), without experience of actual traces of gameplay and relying instead on simulated or imagined traces. Our aim is to model how people approximate $\psi(G)$ and π_G in a way that can (1) take in any G as input, and (2) do so with a limited compute budget and no direct experience with game G .

Intuitive Gamer model specification

We first describe a formal account of the Intuitive Gamer player module, and then describe the reasoning module.

Intuitive Gamer player module. Formally, at a given board state ($s_t \in \mathcal{S}$), the Intuitive Gamer player module scores all legal next actions ($a_t \in \mathcal{A}$) according to a measure of immediate progress made towards a player’s own goal (U_{self}) and a measure of progress blocked (U_{opp}) towards the opponent’s goal (Fig. 2c). Progress is based on the extent to which an action connects more contiguous pieces towards a winning K -in-a-row configuration. These two utilities are designed to capture the general intuition that game players aim to make progress towards their goal while preventing their opponents from doing the same. In addition, we can consider other easily computable heuristic functions that may bear on the value of an action (U_{aux} , for ‘auxiliary’). In our experiments, we consider an auxiliary heuristic based on proximity to the centre of the board. It encodes a ‘centre bias’—a preference for making moves near the centre of the board, which allows a piece to participate in more winning terminal configurations. These heuristics are drawn from features used by previous studies in similar games^{6,29,30}, although we generalize them to our broader class of strategic games. The final heuristic value assigned to a given action on a particular board, $\tilde{V}(s_t, a_t)$, is a sum of the three utility components described above:

$$\tilde{V}(s_t, a_t) = U_{\text{self}}(s_t, a_t) + U_{\text{opp}}(s_t, a_t) + U_{\text{aux}}(s_t, a_t). \quad (1)$$

Our approach to action evaluation and choice reduces computational cost, as it is strongly local in both space and time: it considers only the progress immediately stemming from a specific action and does not account for either past or potential future returns. By contrast, it is

Article

common for value functions to be action-agnostic (that is, to depend only on s_i (ref. 6)) and to explicitly capture some notion of the game's terminal rewards (for example, via a learned value network⁴). Intuitively, each of these alternatives represents a substantial increase in cognitive load: the former requires scanning over the entire board to evaluate any action, and the latter requires mental simulation all the way to the end of the game (or access to a function that is derived from such simulations) before deciding what action to take.

We next describe our specific implementation of each utility function. As mentioned, we came up with these heuristics via prior literature grounding, as well as expert intuitions. Concepts of progress and blocking are present in the classic study of ref. 29, and similarly progress patterns and central tendencies are used in the recent ref. 6, on which we base our expert model. More generally, features based on connections (or lack thereof), implemented differently based on games, is pretty common in game-playing AI. Most of the authors are expert cognitive modellers and some are expert board game players (and played many connect- N style games), so the combination of progress, blocking and centring is effectively the first modelling hypothesis that we came up with and judged to be plausible. More details on the lesioning of these heuristics are included in Supplementary Information section 3.3.

The first utility (U_{self}) computes a measure of intermediate 'goal progress' of the active player, based on a feature n_1 defined as the largest line of contiguous pieces created by the action that could be extended to result in a win. This means that actions which extend a line in an illegal direction (for example, a diagonal line in a game with only horizontal or vertical wins) or that are already blocked by an opponent's piece or the edge of the board do not contribute to goal progress. If the action would result in a win for the active player (that is, $n_1 = K$) then an additional 1 is added to the feature to magnify the value of winning actions. We note that our choice of feature only rewards actions that form contiguous lines of pieces—any piece that is not immediately adjacent to a previously placed piece from the active player has $n_1 = 0$. Although other formulations (such as rewarding actions that form the ends of a line that is empty in the middle or detecting particular patterns of pieces on the board) are sensible, we choose a simple and easy-to-calculate function that reflects a straightforward intuition (that is, making K in a row by first making $K - 1$ in a row, $K - 2$ in a row and so on).

The second utility (U_{opp}) computes a measure of 'progress blocked' for the opponent, based on a feature n_2 that is largely symmetric with the goal progress feature above: it is computed exactly as the goal progress the opponent would obtain if they made the same move (that is, with respect to the opponent's allowed winning directions). We subtract 0.5 from n_2 to reflect people's tendency to weigh offence over defence²⁹, so blocking the opponent's $\hat{K}$ in a row is not as good as making $\hat{K}$ in a row for oneself (but is better than making $\hat{K} - 1$ in a row). However, if n_2 equals the winning K , we do not subtract 0.5 in order to account for the importance of blocking winning threats. As above, there are other reasonable formulations of this feature that we leave to future work.

Finally, the third utility (U_{aux}) is computed as the normalized Euclidean distance between the position of the action and the centre of the board, $\xi \in [0, 1]$. This reflects the intuition, applicable across our family of games, that people often place pieces around the centre of the board. This tendency may in part be explained by the fact that placing pieces closer to the centre of the board allows that piece to participate in the most possible winning terminal states for any K -in-a-row win condition. This measure generalizes to arbitrary rectangular boards. For example, on a 4×6 board, intuitively the 'middle point' is the centre, and the middle four cells are all closest to the centre ($\langle 2, 3 \rangle$, $\langle 2, 4 \rangle$, $\langle 3, 3 \rangle$, $\langle 3, 4 \rangle$).

We assume each utility function is represented as an exponentiation of the respective feature:

$$\tilde{V}(s, a) = w_{\text{connect}} \times 2^{n_1} + w_{\text{block}} \times 2^{n_2} + w_{\text{centre}} \times 2^{(1-\xi)}. \quad (2)$$

The choice to exponentiate some measure of progress for heuristic functions is common in gameplay modelling (for example, refs. 51–53). We chose base two on the basis of light initial exploration (prior to collecting any human gameplay data), under the goal of keeping our intuitive model simple. Future work should explore other modelling choices to capture how people might represent and combine multiple utility functions, including the role of learning in how people might come to flexibly synthesize these functions.

Moves are selected by following Boltzmann rationality^{54–56}, sampling actions from a softmax function over their estimated value (based on the above goal-directed heuristics). We assume that players have already developed the capacity to account for multiple goals simultaneously, unlike potentially even more naive child-like game reasoners²⁹. Concretely, the probability of choosing action $\hat{a}$ at state s is given by:

$$P(\hat{a}|s) = \frac{\exp(\tilde{V}(s, \hat{a})/\tau)}{\sum_a \exp(\tilde{V}(s, a)/\tau)}. \quad (3)$$

We fix temperature (τ) at 1 for our game reasoning and action experiments (with the exception of marginalizing over τ only for the admixture analyses in the 'play' experiment). We set the weights of each component (w) to 1 for all experiments. We use the same settings for the Expert Gamer model (as it uses the same value function and softmax-based action selection). We find that equal weights is a reasonable fit for the Intuitive Gamer model to human payoff predictions (Supplementary Information section 3.3). When lesioning a component of the value function, we set its weight to zero.

Intuitive Gamer reasoning module. The Intuitive Gamer player module is queried as part of the game reasoning module. The Intuitive Gamer reasoning module nests player module-based simulations to compute a series of gameplay traces ($\{(s_0^0, s_1^0, \dots, s_T^0), \dots, (s_0^k, s_1^k, \dots, s_T^k)\}$). These simulations involve self-play between the same player module type (with the exception of the funniness reward for thinking computation, see below), where each agent is trying to make progress towards their own goal (which may be different, for example, player 1 trying to make 4 in a row only diagonally, while player 2 can win in any direction). From these simulations, a probabilistic judgement of the expected outcomes can be made. That is, for each game G , k game simulations are sampled, producing a set of k outcomes $o \in \{-1, 0, 1\}$ encoding whether the first player won (1), lost (-1), or the game ended in a draw (0). From these outcomes, we can compute a payoff of a game G :

$$\psi(G) = (1)P(\text{win}|\neg\text{draw}) \cdot P(\neg\text{draw}) + (-1) \cdot P(\text{lose}|\neg\text{draw}) \cdot P(\neg\text{draw}) + (0) \cdot P(\text{draw}). \quad (4)$$

$$\psi(G) = P(\neg\text{draw}) \cdot [P(\text{win}|\neg\text{draw}) - P(\text{lose}|\neg\text{draw})]. \quad (5)$$

Our game reasoning module is constructed to represent the reasoning of any one participant. Therefore, we draw k simulated matches for $N = 20$ simulated participants (as each game has, generally, 20 participant responses; see 'Analysis methods' below for details on selecting k).

Partial game simulations. Our primary game reasoning module assumes that mental simulations are conducted until the end of the game is reached: either a player attains the win condition, or all open board positions are filled and the game is called a draw. While many of our games do not take many moves to reach an end, people might not mentally simulate all the way to the end of the game in each of their k simulations. We conduct a preliminary exploration into the impact of partial game simulations in fitting peoples' game fairness evaluations by modelling the possibility that people halt any one of their k simulations early. We sample a 'stopping time' uniformly from 1 to the size of the board and end the game after that many turns. To determine the outcome of

games that are stopped early, we apply a simple rule and treat each of them as a draw (allowing us to explore encoding a weak prior towards games ending in draws; Supplementary Information section 5.4). We repeat the same exploration of variance (see ‘Analysis methods’ below) and find that the optimal number of samples is similarly more than 1 and less than 10, but may be closer to $k = 4$ (Extended Data Fig. 2). We leave further analysis of the distributions over stopping times and how people assign outcomes for partial games to future work.

Funniness model. To estimate a game’s funniness, we consider three features derived from playouts under the Intuitive Gamer player module: balance, reward for thinking and game length. We define a game’s balance as the difference between the probability distribution of observed outcomes and an ideal outcome distribution relative to a game where exactly half of the games end in wins and losses and none end in draws (measured by the Earth Mover’s Distance⁵⁷). As described in the main text, this feature captures the notion that players prefer decisive games (that is, those that do not end in draws) that are also balanced across players^{58–60}. We measure game length as the expected number of moves until a game ends in a draw or win, computed from the same Intuitive Gamer simulations. The effect of game length on funniness is modelled using a quadratic function, based on the intuition that the most fun games will be neither too short nor too long. We define a game’s reward for thinking as the proportion of simulated games won by the Intuitive Gamer model playing against the Random Gamer player module, that is, a player that chooses actions uniformly at random from any legal move. This feature captures the intuition that players prefer games that they expect will challenge their thinking^{35,38} and reward them for at least some strategic thinking, rather than just responding arbitrarily or without any strategy.

We use 120 game simulations for each feature (to align with the $k = 6$ simulations for each of 20 participants), randomizing whether the Intuitive Gamer plays first against the random agent in the simulations used to assess reward for thinking. Game simulations are run to the end of the game (where either a player wins or the game ends in a draw); future work can explore variants of the funniness model in which features are computed under partial game simulations. We use the same features and number of simulations when comparing to alternative models (see below in ‘Fitting regression models to funniness judgements’).

Alternative models

We next detail several alternative models.

Expert Gamer. We compare our Intuitive Gamer model against an Expert Gamer model that differs along two key dimensions: (1) sophistication of the value function, and (2) depth of search. This model is closely based on the model of human expert play for 4-in-a-row on a 4×9 board in ref. 6, which empirically estimates tree search depth from human players after hours of continuous gameplay experience. As in ref. 6, the depth is controlled by a probabilistic ‘stopping parameter’ governing how many the search tree is expanded; we run all simulations by expanding the search tree a fixed number ($k_{\text{iterations}} = 636$) of iterations, where $k_{\text{iterations}} = 636$ is the empirical mean value of this parameter estimated from the ref. 6 data. This setting corresponds to approximately depth-5 search. On each iteration, the Expert Gamer selects a node (corresponding to a board state). We then expand this node, considering all legal actions, wherein we compute the value of each action as outlined below. The Expert Gamer model then conducts best-first search over the game tree, using its state value function defined above. Specifically, it probabilistically expands nodes in the search tree by repeatedly sampling actions from a softmax policy governed by temperature τ over its current action estimates based on these computations, expanding unexplored states in the search tree, and backpropagating utilities estimated at future states. As in ref. 6, any node that ends in a definite win or loss is assigned $\pm\sigma$ based on

whether the move would result in a win for the current player (+) or loss (−). We set $\sigma = 1,000$.

We next detail the Expert Gamer value function. Following ref. 6, we compute the value of any move by looking at features over the entire board state, rather than locally circumspect around the possible move position in question (like the Intuitive Gamer model). This is more computationally intensive, particularly for larger boards. Specifically, for any open legal position p , we imagine a state s' that has that position played by the current player. We then sweep over all played positions with that board state. For each played position p' , we compute the same novice, partial-progress value function. We then define the value of that state s' as the difference in cumulative value from the played positions by the current player minus the opponent. We set the value function feature weights (w_{connect} , w_{block} , w_{centre}) using the fit values from the IntuitiveGgamer model. It is worth noting that the Expert Gamer is itself an important contribution—it is a more general version of a relatively deep goal-directed model. We leave validation of how well the Expert Gamer captures human expert reasoning and play for future work.

Random Gamer. The Random Gamer player selects actions uniformly over the space of legal moves. Games are simulated to the end (for example, until either player wins or the game ends in a draw).

Monte Carlo tree search. We additionally implement a standard upper-confidence-bound MCTS^{27,28} agent to act as an approximate gameplay ‘oracle’ (pseudocode can be found in Supplementary Information section 2). Unlike the Intuitive Gamer and Expert Gamer models, MCTS does not use game-specific heuristic features and instead estimates intermediate utilities by expanding a search tree guided by repeated random rollouts to terminal states. MCTS algorithms are commonly used to approximate optimal gameplay in arbitrary games²⁸, as they are empirically both efficient and accurate. Specifically, the MCTS implementation we consider here uses a large number of tree expansions (10,000) to model the behaviour of a near-perfect player in our novel games. Owing to computational costs, we estimate the expected payoffs using 50 simulations per game; these samples are then bootstrapped in their fit to people, as with the other models.

Comparison of MCTS and Intuitive Gamer. We briefly clarify the relation between MCTS—the planning algorithm behind many AI systems that achieve expert-level gameplay such as AlphaGo and AlphaZero—and the contribution the Intuitive Gamer model offers as a model of human reasoning. One could potentially view the Intuitive Gamer’s player module as a particularly restricted or shallow variant of MCTS (as MCTS is also compatible with constraining computational resources and adopting heuristic functions). However, for several reasons we think it is important to distinguish between these classes of models—recognizing that this is in part a matter of scientific emphasis and interpretation rather than algorithmic innovation.

First, framing MCTS so generally as to include the Intuitive Gamer’s play would also include almost any kind of stochastic tree search, including other models in our comparison set, for example, the Random Gamer. More deeply, it would leave out core features of MCTS that have made it so powerful in AI systems as well as the most distinctive features of the Intuitive Gamer as a cognitive model. MCTS made such an impact in AI gameplay and reinforcement learning^{4,14,27,28} precisely because it could achieve very strong play without the need for heuristics to guide search. Instead, it uses extensive inference-time computation and sophisticated tree traversal arithmetic (backtracking, node visit counting) to effectively explore a very large, unstructured game tree. This is in contrast to the Intuitive Gamer, which does not require any sophisticated algorithms and uses very minimal computation, although it does rely on a small number of abstract heuristics when assessing the value of next moves. Such a design leads the Intuitive Gamer to be

Article

a highly efficient if less than optimal player, and one which we believe much better captures how people reason (as evidenced in our behavioural action selection and action prediction studies). As described in the ‘Resource-rational reasoning’ section of the main text and captured in Extended Data Table 2b, the Intuitive Gamer is orders of magnitude more efficient than either MCTS (run for 10,000 iterations, as described above) or the Expert Gamer, in terms of wallclock time and number of board states evaluated.

Variations on the full Intuitive Gamer model. We compare the full Intuitive Gamer model against several variants that modulate whether it is flat, goal-directed or probabilistic. The non-flat version of the Intuitive Gamer model performs a deeper search over possible moves when selecting actions. The non-goal-directed version of the Intuitive Gamer model ablates one or more of the value function components (that is, player progress or blocking progress—see Supplementary Information for ablations of the centre-bias component). The non-probabilistic version of the Intuitive Gamer model replaces the softmax element of action selection with a deterministic choice (or equivalently, softmax with temperature zero). We further vary fastness by conducting a sample complexity analysis wherein we modulate the number of simulations drawn (k) from the gamer model, over which the game reasoning engine computes the expected payoff (Extended Data Fig. 1).

Game-theoretic optimal payoffs. Many of the games in our set have known game-theoretic ground-truth outcome values assuming optimal play from both players. We describe how we filter which games of the 121 are estimatable via a game-theoretic optimal payoff. We first include those with known game-theoretic optimal payoffs (either definite player 1 win, definite draw, or definite player 2 win) according to ref. 61. Then, to approximate a larger set of games’ optimal payoffs, we additionally include games for which MCTS converged to $\{-1, 0, 1\}$ in its payoff predictions (noting that MCTS estimates are not guaranteed to be perfect). This process results in a set of 78 out of the full 121 games for which we have estimated game-theoretic payoffs.

Analysis methods

We next provide additional details on experimental analyses.

Comparing payoff predictions. Participants answered two questions in the game outcome reasoning experiment: how likely the first player wins given the game does not end in a draw ($P(\text{P1 wins}|\text{not draw})$), and how likely the game is to end in a draw ($P(\text{draw})$). Together, these responses yield all information we need to compute an expected payoff. We compare the expected payoff from participants against those of models and report the R^2 . Specifically, we bootstrap subsample with replacement from the human participant samples, per game, and compare models against the mean payoff per sample. We also bootstrap subsample k for 20 simulated participants per model from the pool of model simulations. Unless otherwise stated, we run 1,000 bootstrap samples.

Sample complexity analyses. We select k (the number of simulations sampled for each simulated participant) by inspecting the variance of the payoff predictions over the simulated set of 20 participants against the empirical variance observed in the human data. We measure the root mean squared error between the model- and human-predicted variance for each game, as well as the Wasserstein distance between the distribution over model and human variances (Fig. 3c). We compute the latter to better capture the poor match of high k to people’s variances (that is, with increased k , variance collapses to zero). We find that that generally approximately $k = 5\text{--}7$ full game simulations reasonably well-captures human variance (Fig. 3c and Extended Data Fig. 1) under both measures. We report all main results with $k = 6$.

Fitting regression models to funniness judgements. We fit a linear regression model to these features using the lmer package in R. Features are normalized to zero-mean and unit standard deviation. We fit the models to 1,000 bootstrapped subsamples of the human data for all games. Models are tasked with predicting the mean of participants’ funniness judgements for each game. We compare models via an ANOVA test and AIC in Supplementary Information section 5.7.

In addition, we ran a generalization test wherein we fit regression models on 50% (59) of the 118 games studied (infinite boards were removed) and tested each model on the held-out 50% of games. We report both results in Extended Data Fig. 4 and include additional details in Supplementary Information section 5.7.

Assessing contribution of non-simulation-based features to funniness judgements. In exploratory analyses, we assess the potential role of non-simulation-based features that can be read off of the game description alone. We compute a series of binary game traits that capture ways in which a game may differ from the base tic-tac-toe. For instance, a game may not be a 3×3 board; the game may end with $K \neq 3$ pieces in a row to win; the second player may have a different win condition than the first player. We consider the following binary features: if the game has constrained win conditions (for example, no diagonals allowed); if the game has asymmetric win conditions (between players 1 and 2); if the game has asymmetric play dynamics (for example, player 2 can place two pieces on their first turn); if the board is not square; $K \neq 3$ pieces in a row to win; if the board is larger than 3×3 ; if the game ends when the first player to achieve K in a row loses (that is, misère variants). We combine the binary features into a single aggregated measure of ‘approximate novelty’, which is the sum of the number of binary features present in any one game. We also explore the incorporation of board size (encoded as the number of rows $\times$ number of columns) as part of non-simulation-based game features that reasoners may consider when assessing the funniness of games. We assess how well participants’ funniness judgements can be explained by these features alone by fitting bootstrapped sets of linear regression models to these features in the same 50/50 train/test splits over the games. In addition, we explore the relative benefit of adding all binary features or any of the aggregate features to the simulation-based model over the full set of games (see Supplementary Information section 5.7 for more details and additional analyses).

Assessing game-time action selections. To assess how well models capture how people play, we condition the model in question on an intermediate board state of the actual human-versus-human game. All moves in the game are considered (encompassing early-, mid- and late-stage play). We computed the model’s predicted distribution over next moves. We assess the fit of this predicted distribution in two ways. First, we compute the log likelihood of the actual players’ moves against those predicted by the model. To handle cases where any model may place at or near-zero probability on any given position, thereby skewing the log likelihood, we incorporate a ‘slop’ parameter α that captures the probability that a player makes a move from the primary model ($1 - \alpha$ that they make a random move on that turn⁶²). We sweep over α between 0.5 and 0.95 in increments of 0.05 and compute the average log likelihood for each model over the settings of α . We show per-game stage and per-game results in Supplementary Information section 6. We conduct a one-sided paired t -test over each individual move as well as the aggregated move likelihood per game, comparing the Intuitive Gamer with the Expert Gamer and Intuitive Gamer with the Random Gamer, respectively.

Fitting admixture models to actions. We then consider two different admixture models: one over all moves made for a game, and one over all moves made for each player. Admixture models captures the notion that a player may play according to either the Intuitive Gamer, Expert

Gamer or Random Gamer model on any turn. We jointly fit the weights of the Intuitive Gamer and Expert Gamer models and take the weight of the Random Gamer model to be the remaining mass that brings the weights to 1, over bootstrapped samples over the entire population of players. As the weights are fit jointly across the models, we need to be particularly careful about the magnitude of the results for the respective models. To that end, we jointly sweep over the temperature used in the action selection (with temperature ranging from 0.5 to 3.0 to produce a distribution over weights). We also fit the admixture over each individual, estimating the relative contribution of each model to the way they play. As each player plays only six games (and one round per game), we fit these weights to all of their moves across all games. We use scikit-learn for fitting, with the SLSQP optimizer.

Computing aggregate human distribution over predicted next action. For our ‘watch-and-predict’ experiments, we have access to finer-grained predictions from any one participant. We again extract distributions over likely next moves, conditioned on the observed intermediate board state from the novice and alternative models; however, we now compare the distributions directly with the suggested distribution of moves made by participants. We construct an aggregate move distribution over the participants, treating each of the five clicks per participant as contributing some mass to the ‘aggregate human’ distribution.

Comparing model- with human-predicted distributions over next actions. We compare distributions using the TVD; we fix temperature at 1 and sweep over α for both models and people. We conduct a similar one-sided paired t -test over the TVD across the three game stage boards seen per match and compare the Intuitive Gamer and Expert Gamer, and the Intuitive Gamer and Random Gamer. We compute split-half TVDs per stage, per match, per game by randomly splitting the participants’ distributions who saw that board, averaging those boards into two aggregate distributions and computing TVD between them. This forms our split-half human estimate, which we use to normalize the other models’ TVD. We demonstrate the robustness of our results to our choice of distributional measure by repeating the analyses with the Jensen–Shannon divergence in Supplementary Information section 6.2. We additionally run an admixture model over both participant- and game-level predicted watch distributions. These are optimized against the Jensen–Shannon divergence.

Comparing model- with human-predicted distributions over next actions to people’s played actions. We secondarily assess how well the distribution over suggested moves made by the watcher captures how people actually played by treating the suggested distribution of clicks by the human watcher akin to the move distributions from the models. We again assess the likelihood of the move played by the active player to the moves suggested by the predictors, the Intuitive Gamer model, and alternative models by sweeping over and averaging out a slippage parameter (α). We conduct these analyses in aggregate (over all board stages, matches and games), as well as at a per game level (over board stages and matches).

Modelling draw requests. We take initial steps to analyse evaluations made during the game by looking at players’ decisions of whether to accept or reject a draw, when offered. Draw requests could take place at any point during the game. For each board where a draw request was made, we run 40 simulations to the end of the game under the Intuitive Gamer agent model. We compute the expected payoff over bootstrapped subsampled sets of $k = 6$ outcomes (simulating a single player; previously computed under one of our $\psi(G)$ game reasoning queries), as well as the expected remaining length of the game ℓ from state s_t . We compute expected payoff here with respect to the

player who is deciding whether or not to accept the draw request. Players receive a bonus payout of US\$0.50 only if they win (nothing if they lose or draw); we can then compute an expected value that is the payoff $\times 0.5$. For each set of bonus-adjusted expected payoffs and expected match length remaining, we fit a logistic regression model to predict whether a player would accept or reject the draw request, based on expected payoff and length, as well as the average predicted game funniness predicted by people in the ‘just think’ game reasoning experiment. We fit the logistic regression model using the glm R package. We also explore the same computations under different agent models varying in depth (for example, the Expert Gamer model) and goal-directedness of the value function (for example, ablating the goal progress component).

We show the decision boundaries in Extended Data Fig. 8a and report the bootstrapped parameter fits in Extended Data Fig. 8b. For plotting, we use a collapsed measure of expected value which we call the ‘expected value of continuing’ (C) from state s_t :

$$C(s_t) = (\psi(G) \times 0.5) - (\beta_{\text{length}}/\beta_{\text{utility}}) \cdot \ell. \quad (6)$$

This measure combines the expected value (bonus-adjusted payoff) with the opportunity cost of continuing (based on expected length remaining).

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Data availability

Anonymized human data game evaluation, action selection and action prediction data are hosted via GitHub at github.com/collinskatie/intuitive-game-reasoning.git.

Code availability

Code is available via GitHub at github.com/collinskatie/intuitive-game-reasoning.git.

49. Palan, S. & Schitter, C. Prolific.ac—a subject pool for online experiments. *J. Behav. Exp. Finance* **17**, 22–27 (2018).
50. Almaatouq, A. et al. Empirica: a virtual lab for high-throughput macro-level experiments. *Behav. Res. Methods* **53**, 2158–2171 (2021).
51. Cao, X. & Lin, Y. UCT-ADP progressive bias algorithm for solving Gomoku. In *2019 IEEE Symposium Series on Computational Intelligence (SSCI)* 50–56 (IEEE, 2019).
52. Sheoran, K., Dhand, G., Dabas, M., Dahiya, N. & Pushparaj, P. Solving connect 4 using optimized minimax and monte carlo tree search. *Adv. Appl. Math. Sci.* **21**, 3303–3313 (2022).
53. Li, H. gobang: JavaScript Gobang AI based on alpha-beta pruning. *GitHub* <https://github.com/lihongxun945/gobang> (2025).
54. Luce, R. D. *Individual Choice Behavior* Vol. 4 (Wiley, 1959).
55. Train, K. E. *Discrete Choice Methods with Simulation* (Cambridge Univ. Press, 2009).
56. Franke, M. & Degen, J. The softmax function: properties, motivation, and interpretation. Preprint at OSF <https://doi.org/10.31234/osf.io/vsw47> (2023).
57. Rubner, Y., Tomasi, C. & Guibas, L. J. A metric for distributions with applications to image databases. In *Sixth International Conference on Computer Vision* 59–66 (IEEE, 1998).
58. Althofer, I. *Computer-aided Game Inventing* Technical Report (Friedrich Schiller Univ., 2003).
59. Browne, C. B. *Automatic Generation and Evaluation of Recombination Games*. PhD thesis, Queensland Univ. Technology (2008).
60. Todd, G. et al. Generating games via evolution and language models. *Adv. Neural Inf. Process. Syst.* **37**, 110723–110745 (2024).
61. Uiterwijk, J. W. Solving strong and weak 4-in-a-row. In *2019 IEEE Conference on Games (CIG)* 1–8 (IEEE, 2019).
62. Nassar, M. R. & Frank, M. J. Taming the beast: extracting generalizable knowledge from computational models of cognition. *Curr. Opin. Behav. Sci.* **11**, 49–54 (2016).

Acknowledgements We thank M. Kleiman-Weiner, L. Schulz, R. Saxe, M. Bernstein, S. Stern, I. Sucholutsky, L. Ying, J. Chu, T. Mills, J. Andreas, T. Brooke-Wilson, J. Fan, Y. Machino, S. Levine, I. Kuperwajs, U. Bhatt, K. Smith, W. J. Ma, N. Daw, G. Poesia and B. Prystawski for comments on the paper or conversations that informed this work. We thank our family and friends for many fun game nights. K.M.C. discloses support from the NSF SBE SPRF (2507321) and the

Article

Cambridge Trust and King's College Cambridge. J.B.T. and C.E.Z. disclose support from AFOSR (FA9550-22-1-0387) and the ONR Science of AI programme (N00014-23-1-2355). J.B.T. additionally discloses support from a Schmidt AI2050 Fellowship and the Siegel Family Quest for Intelligence at MIT. L.W. discloses support from a Stanford HAI Fellowship. G.T. discloses support from the National Science Foundation Graduate Research Fellowship under Grant DGE-2234660. T.L.G. discloses support from the NOMIS Foundation. A.W. discloses support from a Turing AI Fellowship under grant EP/V025279/1, The Alan Turing Institute, and the Leverhulme Trust via CFI. This work is supported (in part) by ELSA - European Lighthouse on Secure and Safe AI funded by the European Union under grant agreement No. 101070617. C.Z., M.B.d.C. and S.J.C. disclose no relevant funding. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Commission.

Author contributions All of the authors contributed to project framing and editing the paper. K.M.C., C.E.Z., L.W. and J.B.T. conceived of and designed the 'just think' experiment. K.M.C., C.E.Z., L.W., M.B.d.C., G.T., T.L.G. and J.B.T. conceived of and designed the play and watch studies. K.M.C., C.E.Z., L.W. and J.B.T. developed the Intuitive Gamer model framework. C.E.Z.

developed the Intuitive Gamer heuristics. K.M.C., C.E.Z., L.W., M.B.d.C. and G.T. implemented the models. C.E.Z. designed the 121 games. K.M.C. implemented the human experiment interfaces and ran human data collection. K.M.C., C.E.Z., M.B.d.C. and G.T. implemented and ran the analyses. K.M.C., C.E.Z., L.W., M.B.d.C., G.T. and J.B.T. drafted the paper. J.B.T. supervised the project and acquired funding.

Competing interests The authors declare no competing interests.


Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41586-026-10722-1>.

Correspondence and requests for materials should be addressed to Katherine M. Collins.


Peer review information *Nature* thanks Christoffer Holmgård, Gary Lupyan and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Peer reviewer reports are available.

Reprints and permissions information is available at <http://www.nature.com/reprints>.




Extended Data Fig. 1 | Sample complexity analysis, based on the variance of human and model payoff predictions with varied compute budget (number of simulations, k) for the Intuitive Gamer model). The variance over each approximately 20 human participant payoff judgments per game and compare against the average variance under simulated sets of $N=20$


simulated participants each drawing k simulations from the model. Model variance is binned into quintiles along the horizontal axis and the average human variance is computed for points in that bin (the black dot overlaid on the scatterplot). Error bars are standard deviation of the variance for points in that bin.



Extended Data Fig. 2 | Partial game simulations in the Intuitive Gamer reasoning module. All main results are reported under the assumption that the game reasoning module runs simulations to the end. While most games do not involve a large number of moves, under the Intuitive Gamer player module, for the game to end—it is possible that people are engaging even cheaper partial simulations when evaluating novel games for the first time. **a**, Exploratory analyses into running the Intuitive Gamer reasoning module under a probabilistic stopping rule equivalently captures the majority of the explainable variance in human payoff judgments as running the Intuitive Gamer reasoning module under full simulations (see Fig. 3a). For direct comparison with the full simulation model, $k = 6$ partial simulations were run to estimate game payoff. The partial simulations were computed as follows. For each simulated match of a game, a stopping time was sampled from a uniform distribution over board size and terminated early if the game did not end before that time. If the game terminated before a winner or draw was called, the game counted as a draw. **b**, To assess the impact of compute budget in the number of simulations used to estimate payoff, the variance of payoff under different number of partial simulations k was again compared with the variance in people's payoff evaluations (as in Fig. 3c and Extended Data Fig. 1). The best fitting number of samples (k) is similarly greater than one and less than ten when using partial simulations, and perhaps even a bit less (e.g., $k = 4$) compared to using full game simulations (see Fig. 3c).




Extended Data Fig. 3 | Systematic variation in human and model predictions with change in game rules. Human and model judgments follow similar qualitative trends in **a**, payoff and **b**, funness predictions as K varies for a 10×10 board. Human and model average predictions are shown. Error bars depict 95% CIs over the mean (1000 bootstraps). Funness model error bars (**b**) are computed over repeated model fits to 50% splits of the games, holding out the subset of 10×10 games of varied K ($\in \{2, \dots, 10\}$) from each train split. These data reveal systematic and interpretable patterns in the human data (e.g., estimated payoff plateauing towards zero with increased K ; funness peaking at intermediate values of K). In particular, when K is too small (3 or less) the game is unbalanced, as the first player can easily win and the Intuitive Gamer simulates this easily. For too large values of K , the game is hard for either player to win and there is little reward for thinking; the Intuitive Gamer also simulates this.




Extended Data Fig. 4 | Funness model feature importances and generalization test comparing funness under alternate Gamer agents and non-simulation based funness models. a, Funness model feature weights reveal that games that are more balanced are generally considered more fun, as are games that are more rewarding to think about (though they contribute less per unit to the funness than balance). Game length has an inverted-U shaped relationship to fun (as evidenced by the positive linear weight and negative quadratic weight; games that are too short or too long are less fun). Weights of the features over 1000 models fits to bootstrapped mean human funness predictions, over all games. Coefficients are over normalized features (see Methods). **b**, Ablating each linear feature in the primary funness model reveals that all features matter. Higher *F*-statistic means that ablating the feature explains less of the variance in the human data as the full model. Δ AIC is

$AIC_{\text{ablated}} - AIC_{\text{full}}$ (a higher value indicates that ablating the feature impairs fit to human data, relative to the full model). **c**, Simulations under the Intuitive Gamer model better capture human judgments than simulations under models varying the computational sophistication (Random and Expert Gamer) in a generalization test, where the parameters of each model are fit on 50% of the games and tested on the other 50% of games. Funness models based on explicit simulation better capture human judgments than models based only on non-simulation based linguistic features (e.g., binary game traits—whether the game has asymmetric win conditions, non-square boards, etc) and aggregate features based binary traits (counts of how many are “on” to capture “novelty” relative to Tic-Tac-Toe; and board size, encoded as number of rows $\times$ number of columns). Error bars depict 95% CIs of the held-out R^2 over 500 bootstrapped train/test splits.



Extended Data Fig. 5 | Empirical observed payoffs in actual human-human play against human- and model-predicted payoffs. The empirical payoffs (based on the observed human-human play game outcomes) generally follow the payoffs predicted by people and the Intuitive Gamer model in the “just


think” experiment (without any live play). The empirical payoffs do not align with what one would expect under optimal play, nor play under relatively more sophisticated expert models (MCTS or the Expert Gamer) or game agents that use less compute (random).




Extended Data Fig. 6 | Modeling individual game players. A probabilistic mixture model (admixture) is fit for each individual participants' action choices (left) and action predictions (right). Bar heights represent the estimated contribution of each model in the probabilistic mixture. The Intuitive Gamer is

the dominant component for most participants; however, there is variability across individual participants' behavior—some participants make action choices more aligned with a Random Gamer or more sophisticated Expert Gamer model.

Action Selection




Action Prediction



Extended Data Fig. 7 | Heuristic value function ablations. To assess the importance of components of the Intuitive Gamer model to fitting participant data, key features of the full Intuitive Gamer value function $\tilde{V}$ are lesioned by setting the component weight to zero and leaving the others fixed at one. **a**, Aggregate log probability under each variant of the human played moves from the “play” (action selection) experiment; **b**, aggregate TVD comparing the model and human predictor distributions per game stage made in the “watch-and-predict” experiment. Error bars depict bootstrapped 95% CIs over

the mean measure per setting. Ablating the connect weight or the block weight has major effects on how well the model explains human data in all three experimental conditions. However, ablating the center weight only has effects in the play experiment. The lack of impact of the center bias in the “watch-and-predict” data may arise from the center bias only mattering in the first moves (and participants are only asked to provide judgments over the third move onwards for the “watch-and-predict” task).



b

Variant	Avg Log Lik	Utility	Game Length	Funness
Intuitive Gamer	-0.56 (-0.59, -0.54)	-4.16 (-5.03, -3.50)	0.020 (0.015, 0.027)	-0.0130 (-0.0158, -0.0114)
Self-Prog Only	-0.61 (-0.62, -0.60)	-2.15 (-2.32, -1.99)	-0.013 (-0.017, -0.010)	-0.0110 (-0.0119, -0.0097)
Block-Opp Only	-0.64 (-0.66, -0.63)	-5.13 (-8.21, -2.16)	-0.016 (-0.019, -0.015)	-0.0023 (-0.0045, -0.0009)
Random Gamer	-0.64 (-0.66, -0.61)	-2.74 (-3.70, -1.51)	-0.007 (-0.013, 0.002)	-0.0080 (-0.0096, -0.0051)
Expert Gamer	-0.58 (-0.59, -0.57)	-3.18 (-3.54, -2.86)	0.027 (0.024, 0.034)	-0.0168 (-0.0180, -0.0159)

Extended Data Fig. 8 | Modeling people’s decisions about whether to keep playing when they receive a draw request. **a**, Logistic regression models predict whether a player would accept or reject a draw request if received are fit to the expected value of continuing to play under the respective player module variants, and the expected funness of games (as predicted by people in the just-think experiments). Decision surfaces are plotted, averaged over 100 bootstrapped samples of expected value calculations over $k = 6$ simulations. The Intuitive Gamer and Expert Gamer qualitatively encode the intuition that games that have higher expected value from continuing or are expected to be fun are games where players are more open to continuing to play. **b**, Coefficients for draw request models, depicting 95% CIs over bootstrapped subsamples of $k = 6$ game simulations. The actions taken by a player (whether to accept or reject a draw) are assessed under the models by computing the log likelihood of doing the action, given the features computed over the particular board.

The model was trained to predict whether a draw request would be accepted; therefore, negative weights for features mean that a higher value of that feature is more predictive of rejecting a draw (i.e., wanting to continue a game). Log likelihood is computed as the averaged likelihood assigned to each decision and averaged over bootstrapped model fits. Higher likelihood means better fit. Using the expected value of continuing as computed under Intuitive Gamer and deeper Expert Gamer simulations similarly explain people’s decisions of whether to accept or reject a draw. Lesioning either or both of the goal-directed components of the Intuitive Gamer value function substantially impairs fit to people’s decisions, particularly when lesioning the offensive component (i.e., the component specifically driving a player to make progress towards their own goal). This may be because, upon receiving a draw request, a player decides whether to accept or reject from a more offensive stance.

Extended Data Table 1 | Diverse set of novel games

Category	Example(s)	Count
M in a row on square boards	3 pieces in a row wins on a 6×6 board; 7 pieces in a row wins on a 10×10 board	20
M in a row on rectangular boards	3 pieces in a row wins on a 1×5 board; 6 pieces in a row wins on a 5×10 board	18
Infinite boards	3 pieces in a row wins on an infinite board; 10 pieces in a row wins on an infinite board	3
M in a row loses	A player loses if they make 5 pieces in a row, on a 5×5 board.	10
No diagonal wins allowed	4 pieces in a row wins on a 10×10 board, but a player cannot win by making a diagonal row.	10
Only diagonal wins allowed	4 pieces in a row wins on a 5×5 board, but a player can only win by making a diagonal row.	10
First player moves 2 pieces	5 pieces in a row wins on a 10×10 board; the first player can place 2 pieces as their first move.	10
Second player moves 2 pieces	10 pieces in a row wins on a 10×10 board; the second player can place 2 pieces as their first move.	10
First player handicap (P1 no diag)	3 pieces in a row wins on a 3×3 board, but the first player cannot win by making a diagonal row.	10
First player handicap (P1 only diag)	4 pieces in a row wins on a 7×7 board, but the first player can only win by making a diagonal row.	10
Second player needs M-1 to win	The first player needs 4 pieces in a row, but the second player only needs 3 pieces to win on a 5×5 board.	10

Wide variety of 121 two-player grid-based strategy games designed to study how people play and think about games when they are novices. Games vary in the board shape (e.g., square versus rectangular), win conditions (e.g., first to K in a row wins versus loses), and game dynamics (e.g., one player can play two pieces on their turn). Counts of games in each category, as well as one representative example game description for each.

Extended Data Table 2 | Resource-rational game reasoning, comparing compute efficiency and similarity of payoff estimates to estimates of game-theoretic optimal payoff

a

Model	Time per Action (s)	States Evaluated per Action
Intuitive Gamer	0.0089 ± 0.0113	57.28 ± 33.18
Expert Gamer	6.48 ± 9.33	31, 417.80 ± 22, 145.58
MCTS	342.40 ± 380.55	523, 297.22 ± 348, 852.03

b

Reasoner	Acc to Game-Theoretic (GT)	R ² to GT	Abs Diff from GT
Human	0.69 (0.65, 0.72)	0.62 (0.57, 0.67)	0.32 (0.31, 0.33)
Intuitive Gamer	0.75 (0.72, 0.78)	0.69 (0.66, 0.72)	0.25 (0.24, 0.26)
Expert Gamer	0.92 (0.91, 0.92)	0.87 (0.85, 0.88)	0.08 (0.08, 0.09)
MCTS	0.91 (0.90, 0.92)	0.89 (0.88, 0.91)	0.06 (0.06, 0.07)
Random Gamer	0.57 (0.55, 0.59)	0.39 (0.34, 0.44)	0.43 (0.41, 0.44)

a, The Intuitive Gamer is more efficient than more sophisticated game reasoners, both in terms of average time to evaluate before deciding what action to select and average number of game states to evaluate before deciding what action to select. represent mean ± standard deviation over all moves when simulating evaluations for the game reasoning payoff predictions, for a subset of 41 of the 121 games. Additional details on cost comparisons are included in Supplementary Information 3. **b**, Human and model judgments are compared to the 78 of the 121 games where the game-theoretic optimal payoff is estimatable. Accuracy between predicted payoff and the approximate game-theoretic optimal is computed by taking the predicted payoff as “correct” if the game reasoner predicted a payoff > 0.5 and the expected payoff is 1; correct if the predicted payoff is < -0.5 and the expected payoff is -1; correct if the predicted payoff is between -0.5 and 0.5 and the game is expected to end in a draw. R² correlation is computed between the raw predicted payoffs and the game-theoretic optimal values, as well as the average absolute difference between the expected predicted payoff and approximate game-theoretic payoff (lower is closer to “correct”). Bootstrap 95% confidence intervals (CIs) are shown in parentheses, where bootstraps are over bootstrapped samples of participants (with replacement) for people or model simulations.

Reporting Summary

Nature Portfolio wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Portfolio policies, see our [Editorial Policies](#) and the [Editorial Policy Checklist](#).

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a Confirmed

- ☐ ☒ The exact sample size (n) for each experimental group/condition, given as a discrete number and unit of measurement
- ☐ ☒ A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
- ☐ ☒ The statistical test(s) used AND whether they are one- or two-sided
Only common tests should be described solely by name; describe more complex techniques in the Methods section.
- ☐ ☒ A description of all covariates tested
- ☐ ☒ A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
- ☐ ☒ A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
- ☐ ☒ For null hypothesis testing, the test statistic (e.g. F , t , r) with confidence intervals, effect sizes, degrees of freedom and P value noted
Give P values as exact values whenever suitable.
- ☒ ☐ For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
- ☐ ☒ For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
- ☐ ☒ Estimates of effect sizes (e.g. Cohen's d , Pearson's r), indicating how they were calculated

Our web collection on [statistics for biologists](#) contains articles on many of the points above.

Software and code

Policy information about [availability of computer code](#)

Data collection We ran model simulations, which are provided in the Supplementary Materials.

Data analysis Data analysis details are included in the Methods. Code and descriptions of our analysis code are provided in the Supplementary Materials. We also provide details on our code at: <https://github.com/collinskatie/intuitive-game-reasoning>

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Portfolio [guidelines for submitting code & software](#) for further information.

Data

Policy information about [availability of data](#)

All manuscripts must include a [data availability statement](#). This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A description of any restrictions on data availability
- For clinical datasets or third party data, please ensure that the statement adheres to our [policy](#)

All datasets are new. We developed the 121 games and ran new human studies and model simulations (as outlined in Methods). Data is included in Supp Materials and downloadable with our code (see above).

Research involving human participants, their data, or biological material

Policy information about studies with [human participants or human data](#). See also policy information about [sex, gender \(identity/presentation\), and sexual orientation](#) and [race, ethnicity and racism](#).

Reporting on sex and gender	Not collected.
Reporting on race, ethnicity, or other socially relevant groupings	Not collected.
Population characteristics	Participants were recruited through the crowdsourcing platform Prolific. Participants were required to live in the US and speak English as a first language.
Recruitment	Participants were recruited via Prolific, which automatically handles recruitment among the people registered for the site. We added a restriction such that people could not participate in a new games study if they had already participated in a game study with us.
Ethics oversight	MIT IRB, Tenenbaum Computational Cognitive Science Group

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☐ Life sciences ☒ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/documents/nr-reporting-summary-flat.pdf

Life sciences study design

All studies must disclose on these points even when the disclosure is negative.

Sample size	<i>Describe how sample size was determined, detailing any statistical methods used to predetermine sample size OR if no sample-size calculation was performed, describe how sample sizes were chosen and provide a rationale for why these sample sizes are sufficient.</i>
Data exclusions	<i>Describe any data exclusions. If no data were excluded from the analyses, state so OR if data were excluded, describe the exclusions and the rationale behind them, indicating whether exclusion criteria were pre-established.</i>
Replication	<i>Describe the measures taken to verify the reproducibility of the experimental findings. If all attempts at replication were successful, confirm this OR if there are any findings that were not replicated or cannot be reproduced, note this and describe why.</i>
Randomization	<i>Describe how samples/organisms/participants were allocated into experimental groups. If allocation was not random, describe how covariates were controlled OR if this is not relevant to your study, explain why.</i>
Blinding	<i>Describe whether the investigators were blinded to group allocation during data collection and/or analysis. If blinding was not possible, describe why OR explain why blinding was not relevant to your study.</i>

Behavioural & social sciences study design

All studies must disclose on these points even when the disclosure is negative.

Study description	We ran three main studies, wherein people evaluated games before play ("just think"); where people played new games ("play"); and where people watched others and predicted their play ("watch-and-predict").
Research sample	Participants were recruited through Prolific (see above). Demographic information was not collected. Prolific participants are likely more adept with online experiments than the average person (and specifically our participants are based in the US); however, running our study on Prolific allowed us to scale to levels that would have been challenging if recruiting only locally and in-person.
Sampling strategy	All studies recruited participants from Prolific. We queued a series of batches specifying a subset of games, to which participants were allocated. To ensure we had approximately equal number of participants in each batch, we queued batches in blocks. The "just think" studies were all run first. Then the play studies. Then the watch-and-predict studies. Participants were excluded from future game studies once they participated in one games study. No formal power analysis was calculated; sample size was determined by pilots for each sub-study and pre-registered for our play and watch and predict, on AsPredicted. The final participants (excluding the experiment batch errors that required collecting new data, described below) totaled 495 participants across the "just think" data, 302 for the play data, and 314 of the watch-and-predict data. Details on exclusion are described below. After quality control, this led

	to 484 participants in the “just think” data, 284 for play, and 302 for watch-and-predict. Our pre-registrations for the play and watch data are at https://aspredicted.org/rqhy-k98m.pdf and https://aspredicted.org/2g7n-dfqm.pdf , respectively.
Data collection	Data was collected on interfaces that we developed for each experiment (see Supplement Information). Experiments were built using jsPsych. We used Pavlovia to host our server and build custom web-based experiments on. For the play experiment, where two participants played each other live, data was saved to a server using Empirica and a DigitalOceans server.
Timing	The just think experiments were run on Jan 26, Jan 29, and Jan 30 2024. The play experiments were run on Feb 28, March 3, March 4, March 5, March 7, and March 10 2025. The watch-and-predict experiments were run on March 24,25,26,28; April 8-9; and June 5, 2025.
Data exclusions	In the first “think” experiment, participants wrote in one novel game at the end; based on these open-ended responses, we excluded participants who gave a clearly non-effort response (e.g., generating an invalid win condition) or providing a clearly AI-generated response (N=11 participants). 10% of the outlier judgements that were farthest from the mean were filtered out for each game. As pre-registered and described in our Methods, participants were excluded if they gave too many “standard” judgments that were “standard” values (near 0, 50, or 100) on 80% or more of judgments for the “play” and “watch-and-predict” experiments. 18 participants were filtered out for the play data and 10 for the watch-and-predict data. While running the play experiment, we noticed a bug in a win condition for one of the games; we discarded and reran that batch. Additionally, due to a batching error in the number of matches seen by participants in the watch data between the fun and payoff prediction conditions, we reran the advantage setting (and discarded the previous data). We report the batch-corrected final number of participants in Methods.
Non-participation	57 of 555 participants returned their study in the just think experiment and 3 timed out (for unknown reasons). 88 of 543 returned their study for the watch and predict study (and three timed out). 213 of 533 participants returned their study in the play experiment; the higher rate of returns is expected by design of our synchronous study (participants need to be paired at the exact same time; as such, we at times overqueued which led participants to not be able to play and need to return their study); 4 timed out their study.
Randomization	Games were shuffled for participants in each study. Participants were quasi-randomly assigned to batches (see above).

Ecological, evolutionary & environmental sciences study design

All studies must disclose on these points even when the disclosure is negative.

Study description	<i>Briefly describe the study. For quantitative data include treatment factors and interactions, design structure (e.g. factorial, nested, hierarchical), nature and number of experimental units and replicates.</i>
Research sample	<i>Describe the research sample (e.g. a group of tagged <i>Passer domesticus</i>, all <i>Stenocereus thurberi</i> within Organ Pipe Cactus National Monument), and provide a rationale for the sample choice. When relevant, describe the organism taxa, source, sex, age range and any manipulations. State what population the sample is meant to represent when applicable. For studies involving existing datasets, describe the data and its source.</i>
Sampling strategy	<i>Note the sampling procedure. Describe the statistical methods that were used to predetermine sample size OR if no sample-size calculation was performed, describe how sample sizes were chosen and provide a rationale for why these sample sizes are sufficient.</i>
Data collection	<i>Describe the data collection procedure, including who recorded the data and how.</i>
Timing and spatial scale	<i>Indicate the start and stop dates of data collection, noting the frequency and periodicity of sampling and providing a rationale for these choices. If there is a gap between collection periods, state the dates for each sample cohort. Specify the spatial scale from which the data are taken</i>
Data exclusions	<i>If no data were excluded from the analyses, state so OR if data were excluded, describe the exclusions and the rationale behind them, indicating whether exclusion criteria were pre-established.</i>
Reproducibility	<i>Describe the measures taken to verify the reproducibility of experimental findings. For each experiment, note whether any attempts to repeat the experiment failed OR state that all attempts to repeat the experiment were successful.</i>
Randomization	<i>Describe how samples/organisms/participants were allocated into groups. If allocation was not random, describe how covariates were controlled. If this is not relevant to your study, explain why.</i>
Blinding	<i>Describe the extent of blinding used during data acquisition and analysis. If blinding was not possible, describe why OR explain why blinding was not relevant to your study.</i>
Did the study involve field work?	☐ Yes ☐ No

Field work, collection and transport

Field conditions	<i>Describe the study conditions for field work, providing relevant parameters (e.g. temperature, rainfall).</i>
------------------	--

Location	State the location of the sampling or experiment, providing relevant parameters (e.g. latitude and longitude, elevation, water depth).
Access & import/export	Describe the efforts you have made to access habitats and to collect and import/export your samples in a responsible manner and in compliance with local, national and international laws, noting any permits that were obtained (give the name of the issuing authority, the date of issue, and any identifying information).
Disturbance	Describe any disturbance caused by the study and how it was minimized.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study
☒	☐ Antibodies
☒	☐ Eukaryotic cell lines
☒	☐ Palaeontology and archaeology
☒	☐ Animals and other organisms
☒	☐ Clinical data
☒	☐ Dual use research of concern
☒	☐ Plants

Methods

n/a	Involved in the study
☒	☐ ChIP-seq
☒	☐ Flow cytometry
☒	☐ MRI-based neuroimaging

Antibodies

Antibodies used	Describe all antibodies used in the study; as applicable, provide supplier name, catalog number, clone name, and lot number.
Validation	Describe the validation of each primary antibody for the species and application, noting any validation statements on the manufacturer's website, relevant citations, antibody profiles in online databases, or data provided in the manuscript.

Eukaryotic cell lines

Policy information about [cell lines and Sex and Gender in Research](#)

Cell line source(s)	State the source of each cell line used and the sex of all primary cell lines and cells derived from human participants or vertebrate models.
Authentication	Describe the authentication procedures for each cell line used OR declare that none of the cell lines used were authenticated.
Mycoplasma contamination	Confirm that all cell lines tested negative for mycoplasma contamination OR describe the results of the testing for mycoplasma contamination OR declare that the cell lines were not tested for mycoplasma contamination.
Commonly misidentified lines (See ICLAC register)	Name any commonly misidentified cell lines used in the study and provide a rationale for their use.

Palaeontology and Archaeology

Specimen provenance	Provide provenance information for specimens and describe permits that were obtained for the work (including the name of the issuing authority, the date of issue, and any identifying information). Permits should encompass collection and, where applicable, export.
Specimen deposition	Indicate where the specimens have been deposited to permit free access by other researchers.
Dating methods	If new dates are provided, describe how they were obtained (e.g. collection, storage, sample pretreatment and measurement), where they were obtained (i.e. lab name), the calibration program and the protocol for quality assurance OR state that no new dates are provided.
☐ Tick this box to confirm that the raw and calibrated dates are available in the paper or in Supplementary Information.	
Ethics oversight	Identify the organization(s) that approved or provided guidance on the study protocol, OR state that no ethical approval or guidance was required and explain why not.

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Animals and other research organisms

Policy information about [studies involving animals](#); [ARRIVE guidelines](#) recommended for reporting animal research, and [Sex and Gender in Research](#)

Laboratory animals	<i>For laboratory animals, report species, strain and age OR state that the study did not involve laboratory animals.</i>
Wild animals	<i>Provide details on animals observed in or captured in the field; report species and age where possible. Describe how animals were caught and transported and what happened to captive animals after the study (if killed, explain why and describe method; if released, say where and when) OR state that the study did not involve wild animals.</i>
Reporting on sex	<i>Indicate if findings apply to only one sex; describe whether sex was considered in study design, methods used for assigning sex. Provide data disaggregated for sex where this information has been collected in the source data as appropriate; provide overall numbers in this Reporting Summary. Please state if this information has not been collected. Report sex-based analyses where performed, justify reasons for lack of sex-based analysis.</i>
Field-collected samples	<i>For laboratory work with field-collected samples, describe all relevant parameters such as housing, maintenance, temperature, photoperiod and end-of-experiment protocol OR state that the study did not involve samples collected from the field.</i>
Ethics oversight	<i>Identify the organization(s) that approved or provided guidance on the study protocol, OR state that no ethical approval or guidance was required and explain why not.</i>

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Clinical data

Policy information about [clinical studies](#)

All manuscripts should comply with the ICMJE [guidelines for publication of clinical research](#) and a completed [CONSORT checklist](#) must be included with all submissions.

Clinical trial registration	<i>Provide the trial registration number from ClinicalTrials.gov or an equivalent agency.</i>
Study protocol	<i>Note where the full trial protocol can be accessed OR if not available, explain why.</i>
Data collection	<i>Describe the settings and locales of data collection, noting the time periods of recruitment and data collection.</i>
Outcomes	<i>Describe how you pre-defined primary and secondary outcome measures and how you assessed these measures.</i>

Dual use research of concern

Policy information about [dual use research of concern](#)

Hazards

Could the accidental, deliberate or reckless misuse of agents or technologies generated in the work, or the application of information presented in the manuscript, pose a threat to:

No	Yes	
☒	☐	Public health
☒	☐	National security
☒	☐	Crops and/or livestock
☒	☐	Ecosystems
☒	☐	Any other significant area

Experiments of concern

Does the work involve any of these experiments of concern:

No	Yes
☒	☐ Demonstrate how to render a vaccine ineffective
☒	☐ Confer resistance to therapeutically useful antibiotics or antiviral agents
☒	☐ Enhance the virulence of a pathogen or render a nonpathogen virulent
☒	☐ Increase transmissibility of a pathogen
☒	☐ Alter the host range of a pathogen
☒	☐ Enable evasion of diagnostic/detection modalities
☒	☐ Enable the weaponization of a biological agent or toxin
☒	☐ Any other potentially harmful combination of experiments and agents

Plants

Seed stocks	Report on the source of all seed stocks or other plant material used. If applicable, state the seed stock centre and catalogue number. If plant specimens were collected from the field, describe the collection location, date and sampling procedures.
Novel plant genotypes	Describe the methods by which all novel plant genotypes were produced. This includes those generated by transgenic approaches, gene editing, chemical/radiation-based mutagenesis and hybridization. For transgenic lines, describe the transformation method, the number of independent lines analyzed and the generation upon which experiments were performed. For gene-edited lines, describe the editor used, the endogenous sequence targeted for editing, the targeting guide RNA sequence (if applicable) and how the editor was applied.
Authentication	Describe any authentication procedures for each seed stock used or novel genotype generated. Describe any experiments used to assess the effect of a mutation and, where applicable, how potential secondary effects (e.g. second site T-DNA insertions, mosaicism, off-target gene editing) were examined.

ChIP-seq

Data deposition

- ☐ Confirm that both raw and final processed data have been deposited in a public database such as [GEO](#).
- ☐ Confirm that you have deposited or provided access to graph files (e.g. BED files) for the called peaks.

Data access links <i>May remain private before publication.</i>	For "Initial submission" or "Revised version" documents, provide reviewer access links. For your "Final submission" document, provide a link to the deposited data.
Files in database submission	Provide a list of all files available in the database submission.
Genome browser session (e.g. UCSC)	Provide a link to an anonymized genome browser session for "Initial submission" and "Revised version" documents only, to enable peer review. Write "no longer applicable" for "Final submission" documents.

Methodology

Replicates	Describe the experimental replicates, specifying number, type and replicate agreement.
Sequencing depth	Describe the sequencing depth for each experiment, providing the total number of reads, uniquely mapped reads, length of reads and whether they were paired- or single-end.
Antibodies	Describe the antibodies used for the ChIP-seq experiments; as applicable, provide supplier name, catalog number, clone name, and lot number.
Peak calling parameters	Specify the command line program and parameters used for read mapping and peak calling, including the ChIP, control and index files used.
Data quality	Describe the methods used to ensure data quality in full detail, including how many peaks are at FDR 5% and above 5-fold enrichment.
Software	Describe the software used to collect and analyze the ChIP-seq data. For custom code that has been deposited into a community repository, provide accession details.

Flow Cytometry

Plots

Confirm that:

- ☐ The axis labels state the marker and fluorochrome used (e.g. CD4-FITC).
- ☐ The axis scales are clearly visible. Include numbers along axes only for bottom left plot of group (a 'group' is an analysis of identical markers).
- ☐ All plots are contour plots with outliers or pseudocolor plots.
- ☐ A numerical value for number of cells or percentage (with statistics) is provided.

Methodology

Sample preparation	<i>Describe the sample preparation, detailing the biological source of the cells and any tissue processing steps used.</i>
Instrument	<i>Identify the instrument used for data collection, specifying make and model number.</i>
Software	<i>Describe the software used to collect and analyze the flow cytometry data. For custom code that has been deposited into a community repository, provide accession details.</i>
Cell population abundance	<i>Describe the abundance of the relevant cell populations within post-sort fractions, providing details on the purity of the samples and how it was determined.</i>
Gating strategy	<i>Describe the gating strategy used for all relevant experiments, specifying the preliminary FSC/SSC gates of the starting cell population, indicating where boundaries between "positive" and "negative" staining cell populations are defined.</i>

- ☐ Tick this box to confirm that a figure exemplifying the gating strategy is provided in the Supplementary Information.

Magnetic resonance imaging

Experimental design

Design type	<i>Indicate task or resting state; event-related or block design.</i>
Design specifications	<i>Specify the number of blocks, trials or experimental units per session and/or subject, and specify the length of each trial or block (if trials are blocked) and interval between trials.</i>
Behavioral performance measures	<i>State number and/or type of variables recorded (e.g. correct button press, response time) and what statistics were used to establish that the subjects were performing the task as expected (e.g. mean, range, and/or standard deviation across subjects).</i>

Acquisition

Imaging type(s)	<i>Specify: functional, structural, diffusion, perfusion.</i>
Field strength	<i>Specify in Tesla</i>
Sequence & imaging parameters	<i>Specify the pulse sequence type (gradient echo, spin echo, etc.), imaging type (EPI, spiral, etc.), field of view, matrix size, slice thickness, orientation and TE/TR/flip angle.</i>
Area of acquisition	<i>State whether a whole brain scan was used OR define the area of acquisition, describing how the region was determined.</i>
Diffusion MRI	☐ Used ☐ Not used

Preprocessing

Preprocessing software	<i>Provide detail on software version and revision number and on specific parameters (model/functions, brain extraction, segmentation, smoothing kernel size, etc.).</i>
Normalization	<i>If data were normalized/standardized, describe the approach(es): specify linear or non-linear and define image types used for transformation OR indicate that data were not normalized and explain rationale for lack of normalization.</i>
Normalization template	<i>Describe the template used for normalization/transformation, specifying subject space or group standardized space (e.g. original Talairach, MNI305, ICBM152) OR indicate that the data were not normalized.</i>
Noise and artifact removal	<i>Describe your procedure(s) for artifact and structured noise removal, specifying motion parameters, tissue signals and physiological signals (heart rate, respiration).</i>

Volume censoring

Define your software and/or method and criteria for volume censoring, and state the extent of such censoring.

Statistical modeling & inference

Model type and settings

Specify type (mass univariate, multivariate, RSA, predictive, etc.) and describe essential details of the model at the first and second levels (e.g. fixed, random or mixed effects; drift or auto-correlation).

Effect(s) tested

Define precise effect in terms of the task or stimulus conditions instead of psychological concepts and indicate whether ANOVA or factorial designs were used.

Specify type of analysis: ☐ Whole brain ☐ ROI-based ☐ Both

Statistic type for inference

Specify voxel-wise or cluster-wise and report all relevant parameters for cluster-wise methods.

(See [Eklund et al. 2016](#))

Correction

Describe the type of correction and how it is obtained for multiple comparisons (e.g. FWE, FDR, permutation or Monte Carlo).

Models & analysis

n/a | Involved in the study

☒ ☐ Functional and/or effective connectivity

☒ ☐ Graph analysis

☒ ☐ Multivariate modeling or predictive analysis

Functional and/or effective connectivity

Report the measures of dependence used and the model details (e.g. Pearson correlation, partial correlation, mutual information).

Graph analysis

Report the dependent variable and connectivity measure, specifying weighted graph or binarized graph, subject- or group-level, and the global and/or node summaries used (e.g. clustering coefficient, efficiency, etc.).

Multivariate modeling and predictive analysis

Specify independent variables, features extraction and dimension reduction, model, training and evaluation metrics.