Planning-while-acting: addressing the continuous dynamics of planning and action in a virtually embodied task

Davide Nuzzi[1], Paul Cisek[2], and Giovanni Pezzulo[1,*]

[1]Institute of Cognitive Sciences and Technologies, National Research Council, Rome, Italy

[2]University of Montreal, Montreal, Canada

[*]Corresponding author: giovanni.pezzulo@istc.cnr.it

March 21, 2025

Abstract

Everyday tasks, such as selecting routes when driving or preparing meals require making sequences of embodied decisions, in which planning and action processes are intertwined. In this study, we address how people make sequential embodied decisions, requiring balancing between immediate affordances and long-term utilities of alternative action plans. We designed a novel virtually embodied, game-like task in which participants controlled an avatar tasked with "crossing rivers", by jumping across rocks. The task permitted us to assess how participants balanced between immediate jumping affordances ("safe" versus "risky" jumps) and the utility (length) of the ensuing paths to the goal. Behavioral and computational analyses revealed that participants planned ahead their path to the goal rather than simply focusing on the most immediate jumping affordances. Furthermore, embodied components of the task influenced participants' decision strategies, as evident by the fact that participants' current direction of movement influenced their choice between safe and risky jumps. We also found that participants showed (pre)planning before making the first jump, but they continued deliberating during it, with movement speed decreasing at decision points and when approaching them. Finally, computational modeling indicates that farsighted participants who assigned greater weight to the utility of future jumps showed a better performance, highlighting the usefulness of planning in embodied settings. Our findings underscore the importance of studying decision-making and planning in ecologically valid, embodied settings, providing new insights into the interplay between action and cognition in real-world planning-while-acting scenarios.

Keywords: embodied decisions; planning-while-acting; affordances; affordance competition

1 Introduction

Most everyday tasks, such as selecting a path in a crowded street, preparing a table or a backpack for a trekking, require sequences of embodied decisions, in which planning and action processes unfold simultaneously. Traditional studies of decision-making and planning in psychology and (neuro)economics focused instead on static, non-embodied scenarios, such as abstract problem-solving or choices between predefined offers, in which action requirements are minimal. Therefore, it is unclear how individuals plan ahead and how they balance the immediate affordances and the longer-term utility of different courses of actions in embodied settings.

Recent work addressed some of these limitations by focusing on simple embodied decisions and asked how people select between different courses of action, by weighting immediate affordances – or opportunities for action – and utilities (Cisek and Pastor-Bernier, 2014; Gibson, 1979; Pezzulo and Cisek, 2016; Cisek and Kalaska, 2010; Yoo et al., 2021; Priorelli et al., 2024; Wispinski et al., 2020; Gallivan et al., 2018). However, most of the above studies focused on immediate actions and affordances, such as selecting between reaching, stepping, or pressing actions, or choosing between following a current target or switching to a new one, without addressing sequential decisions and planning (Cos et al., 2021; Michalski et al., 2020; Connors et al., 2018; Kobayashi and Kimura, 2022; Lepora and Pezzulo, 2015; Grießbach et al., 2022; Ulbrich and Gail, 2023; Warren, 2006; Gallivan et al., 2015; Nashed et al., 2014). Other studies investigated decision-making and affordance competition in sports, highlighting the importance of embodied strategies for successful performance (Raab, 2020; Gordon et al., 2021). While these studies advanced our understanding of simple forms of embodied decisions that exploit immediate affordances, the ways people solve sequential embodied tasks requiring the continuous construction and evaluation of novel affordances remains incompletely known. We still do not know how, in sequential embodied settings, people prefer plans whose immediate affordance for movement or long-term utility is greater; whether they preplan action sequences before initiating movements; whether the planning phase ends up at the beginning of movement or continues during it; whether embodied constraints of the task, such as the current movement or the cost of changing trajectory influence decision-making. Addressing these open questions requires designing novel experiments and computational models that address sequential decision-making in embodied settings, in which decisions and actions are intertwined, i.e., planning-while-acting studies.

Here, we address the above questions, by designing a novel experiment in which participants control an avatar tasked with "crossing rivers", by jumping across rocks. Crucially, the task requires making a series of jumping choices, in a "virtually embodied" setting, in which the movements of the avatar follow coherent spatial and physical rules. Sequential decisions are often studied in psychology and (neuro)economics in non-embodied settings, in which the movements required to make the first choice are irrelevant for the subsequent choices (Daw et al., 2011). Rather, in our virtually embodied setting, choices and movements are interdependent: each jump changes the direction and velocity of the avatar and the geometry of the problem. This allows us to probe whether and how some of the factors inherent in decisions made during real physical behavior, such as affordances, spatial relationships, and momentum, impact choice dynamics.

We created multiple sequential choices between "safe" versus "risky" jumps, corresponding to closer-and-bigger versus farther-and-smaller rocks, and between paths having different utilities (here, related to the path length to the goal). By analyzing participants' behavioral choices and formalizing their strategies using computational models, we aim to reveal how individuals balance affordance and utility in embodied settings, how they form sequential plans, and how their choice strategies are influenced by embodied aspects of the task.

2 Methods

2.1 Participants

We recruited participants via the Prolific platform, indicating native English and age between 18 an 40 years as eligibility criteria. We only included in the analysis participants who completed all levels of the game, resulting in a sample of 40 participants (19 males and 21 females), with a median age of 30 years, while 5 participants were excluded for not finishing the experiment. All participants gave informed consent online before starting the experiment. All our procedures were approved by the CNR Ethics Committee.

2.2 Experimental setup

In this study, we develop a novel 3D video game using Unity (https://unity.com/). The game requires participants to control an avatar (an anthropomorphic frog) to cross rivers, by jumping from rock to rock until a goal platftorm indicated by a red flag (Figure 1). Participants use the mouse to adjust the camera view, the 'W' key to move forward, and the spacebar to jump. If the avarar falls into the water, the trial restarts from the last rock on which it stood, after a brief animation showing the avatar being transported back. The avatar is 1.75 units tall. The game's world units are arbitrary, meaning that the relative proportions between them are what matter, rather than their absolute values. The gravity is simulated realistically, with a constant gravitational acceleration of -15 units/s^2. Collisions between the character and the platforms, as well as with the water, are handled via an invisible capsule-shaped collider attached to the avatar. A non-realistic simplification was introduced, allowing the character to move forward and rotate even while airborne. However, no additional platforming tricks, such as "coyote time", were implemented to artificially facilitate movement. The forward movement speed is 5 units per second when the 'W' key is pressed. Acceleration from rest to this speed, as well as deceleration when 'W' is released, is nearly instantaneous. The vertical speed imparted when pressing 'SPACE' to jump is calibrated so that the avatar reaches a maximum height of 1.2 units during a jump. This means the initial vertical speed at takeoff is 6 units per second, the time required to reach the peak is 0.4 seconds, and the total duration of a jump is 0.8 seconds. Keeping SPACE pressed for longer does not alter the height of the jump. If the player holds 'W' for the entire duration of the jump, the maximum horizontal distance that can be covered is 4 units .Finally, the avatar's orientation (i.e., rotation around the vertical axis) is fully controlled by mouse movements. This ensures that

the direction the avatar is facing always matches the player's camera view, meaning there is no discrepancy between where the character is looking and what the player sees.

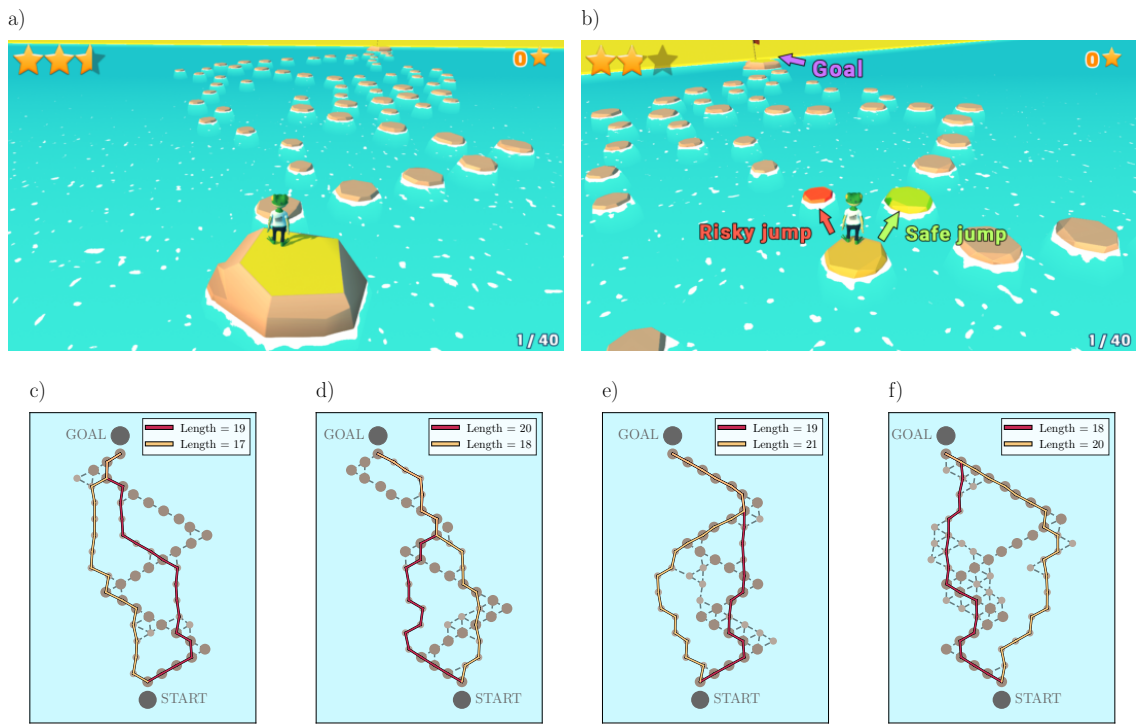

Figure 1: Experimental setup. a) A screenshot from the game showing the start of a level. The avatar (frog) is positioned on an elevated platform, providing a complete view of the entire map and the distant goal platform, marked with a red flag. **b)** A screenshot illustrating a decision point where the player must choose between two types of jumps: a "safe" jump to a larger rock or a "risky" jump to a smaller rock. The goal platform is visible in the distance. **c-f)** Examples of the four categories of maps used in the experiment, based on the properties of the first decision point. The first *risky* jump can be on the shorter path to the goal (c,d) or on the longer one (e,f). The direction of the first *risky* jump can have a smaller angle with respect to the goal direction (c,e) or a larger angle (d,f). In each map, the *safe* path is shown in red, and the *risky* paths are shown in yellow. The difference in path length (number of jumps) between the shorter and longer paths is always 2 jumps.

Participants start each level from a starting platform, which provides them with a full view of the map (Figure 1.a). Crossing the river requires jumping across stones of two types: large rocks, affording a *safe* jump to a larger landing area and small rocks, affording a *risky* jump to a smaller landing area. The large and small rocks colliders have an approximate radius of 0.72 and 0.48 units respectively. The game levels are generated procedurally. They include a *safe* path composed only of large stones arranged in a zig-zag pattern on a hexagonal grid from the starting platform to the goal platform. Furthermore, they include multiple *risky* paths composed of smaller rocks, which branch off the *safe* path at random points. The *risky* paths move toward the goal by selecting one of three directions on the hexagonal grid: forward left, forward, or forward right. This design creates multiple decision points, between safe or risky jumps and paths. Notably, for each problem,

there is a single shortest path, which could be either the *safe* path or a *risky* path. Regardless of which one is shorter, the difference in distance between the shortest path and the other is always two jumps. We fixed the distance between points on the hexagonal grid to 3.9 units, preventing participants from jumping to non-adjacent rocks. Furthermore, we applied a positional jitter to the rocks after the initial map generation, to create a more naturalistic arrangement. Specifically, each platform was moved in a random directions by 0.4 units. We made sure that the randomization was the same for all players and that, after the random jitter, the border-to-border distance between adjacent platforms (w.r.t. the hexagonal grid) was always less than the maximum jump distance of the avatar (4 units).

Our experiment used 40 game levels and a 2 x 2 design. The two factors reflect the decision to be made at the first decision point between a *safe* and a *risky* jump (Figure 1.b). The first factor reflects whether the the shortest path to the goal (measured as the minimum number of jumps) started with the *safe* or the *risky* jump. The second factor is the angle between the vector connecting the decision point to the chosen rock and the vector connecting the decision point to the goal, indicating how aligned the jump and the direction of the goal are. As a single jump is insufficient to indicate the path's overall direction, we considered the average direction over the first few rocks following the jump. Note that path length is typically a better proxy for performance than the angle to the goal, as fewer jumps generally require less time. However, the angle to the goal provides an immediate, visually perceivable cue, whereas estimating path length is more challenging – especially considering that paths often involve more than 20 jumps. By considering both factors, we aimed to assess how participants balance efficiency and perceptual availability. Based on these two factors, we created four groups of ten levels each, for a total of 40 levels (Figure 1.c-f). We procedurally generated multiple levels for each factor and selected manually the best 40, prioritizing those providing more useful decision points for the subsequent analyses.

Participants are informed that they have a maximum of 80 seconds to complete each level and that they are awarded from one to three stars depending on problem completion time (or zero stars if they fail to complete the problem within the time limit). The total score, i.e. the total number of collected stars, was always visible in the top-right corner of the screen. Before the experiment, participants completed three tutorial levels to familiarize with the controls and game mechanics.

2.3 Data analysis methods

We used mixed-effects models to account for both fixed effects of experimental factors and random effects associated with individual variability. We employed the *pymer4* package in Python, using the *Lmer* function for both generalized linear mixed models (GLMMs) and linear mixed models (LMMs), depending on the dependent variable.

In the case of binary outcome variables (e.g., probabilities), we used GLMMs with a logit link function, whereas in the case of continuous variables we used LMMs without a logit link function. In both cases, we included in these models random intercepts for each participant, to capture participant-level variability.

For each analysis, we compared models with or without interaction terms between factors,

using Bayesian Information Criterion (BIC) and we retained the model having the lower BIC score, indicating a better fit. Following the model fitting procedure, we conducted post-hoc analyses when we found main effects or interactions to be significant. We carried out these post-hoc pairwise comparisons using the *post_hoc* function in *pymer4*, with Tukey's Honest Significant Difference (HSD) correction applied to adjust for multiple comparisons. This procedure ensured that we identified specific differences between conditions while controlling for Type I errors.

We set the statistical significance at $p < 0.05$ for all analyses and we indicated in the boxplots all the significant differences from the post-hoc tests.

3 Results

We first considered whether participants were able to solve the game effectively, by looking at their general performance metrics (reported below as mean $\pm$ standard deviation) across the experiment. A total of 26 out of 40 participants completed all trials within the time limit, while others failed to complete one or more levels before the timer expired. On average, participants successfully completed 36.05 ± 8.46 trials, making 21.30 ± 2.97 jumps per level. The total time they spent actively playing the levels (excluding tutorials and transitions) was 31.47 ± 14.63 minutes.

Participants demonstrated a high level of accuracy in the task, with an overall jump success rate of $93\% \pm 5\%$, falling into the water 1.61 ± 2.02 times per level on average. This indicates that errors were relatively rare, and performance was generally robust across conditions. Performance varied across individuals, with participants earning an average of 76.90 ± 30.90 stars across all levels, out of the 120 possible stars if they completed each level perfectly.

3.1 Participants rely on both path length and angle to goal to decide between *safe* or *risky* jumps

We analyzed the probability of making the first *risky* jump (i.e., jumping onto the small rock at the first decision point) in relation to our two experimental factors: namely, the length of the path to the goal starting with the *risky* jump and the angle between the jump direction and the goal. We computed the probabilities per participant across all levels corresponding to the four specific combinations of these factors (Figure 2.a). Participants select the *risky* jump more often when both the path starting with it is shorter and the angle to the goal is smaller; and less often when the path is longer and the angle is larger. When either the path is shorter with a larger angle or longer with a smaller angle, participants selected the *risky* jump with intermediate probabilities.

To assess the relative influence of the two experimental factors on the probability of taking a *risky* jump, we fitted a generalized linear mixed-effects model (GLMM). We considered a model that includes an interaction term between the factors, as this resulted in a lower BIC than a model without interaction (Table 1). Both path length and angle significantly affect the probability of choosing the *risky* jump ($p < 0.001$). However, the coefficient for path length is larger, suggesting a stronger influence of path length compared to angle to the goal. We compared all pairs of conditions using Post-hoc tests (Figure 2.a and Table S1). All the differences are significant ($p < 0.001$), except

for the two intermediate conditions (shorter path length, larger angle vs longer path length, smaller angle), indicating a lack of clear dominance of either factor.

These results suggest that participants consider both path length and angle to the goal when deciding which rock to jump onto at the first decision point. Moreover, the impact of path length on the choice is greater than angle to the goal, despite the small difference (two jumps) between the lengths of *risky* and *safe* paths and the fact that the angle to the goal is more readily available perceptually.

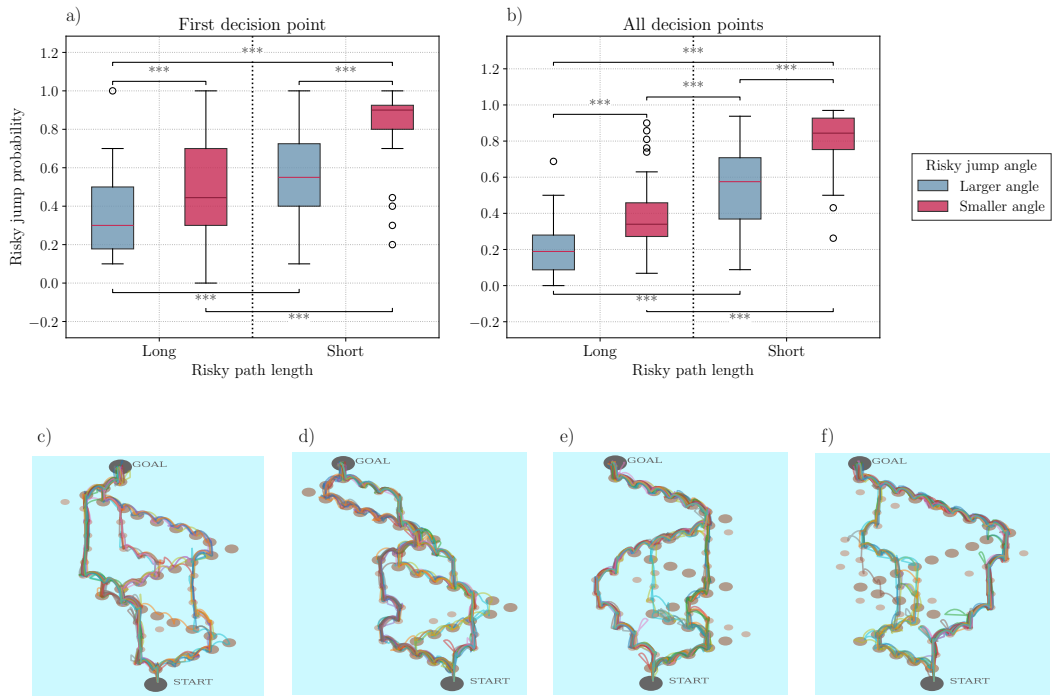

Figure 2: Both path length and angle to goal influence the probability of making a *risky* jump. a) Effect of path length and angle on the probability of making the *risky* jump, at the first decision point. Participants show a greater preference for the *risky* jump when the path is shorter and the angle is smaller, and a smaller preference when the path is longer and the angle is larger. There is no significant difference between the two intermediate conditions (shorter path with a larger angle vs. longer path with a smaller angle). **b)** Effect of path length and angle to the goal on the probability of making a *risky* jump, across all decision points. Both path length and angle significantly influence the probability of making a *risky* jump, with path length having a stronger effect than angle to the goal. **c-f)** Trajectories of all 40 participants across the same levels of Figure 1, showcasing the diversity of paths taken at each decision point. Each colored path corresponds to the trajectory of a single participant.

We extended the analysis to include all decision points, not just the first one. A decision point was defined as a condition where participants faced a forced choice between a *risky* and a *safe* jump. We excluded decisions points with more than two options or where both options where either *risky* or *safe*.

	Fixed effects					Random effects
	Estimate	Std Error	[0.025	0.095]	p-value	Variance
Intercept	-0.787	0.176	-1.132	-0.441	**< 0.001**	0.719
Risky Jump Shorter Path	1.082	0.158	0.773	1.391	**< 0.001**	
Risky Jump Smaller Angle	0.721	0.161	0.405	1.037	**< 0.001**	
Interaction	0.852	0.242	0.377	1.326	**< 0.001**	

Table 1: GLMM results for the *risky* jump probability at the first decision point. Both path length and angle to the goal significantly influence the probability of choosing the *risky* jump. The larger β coefficient for path length compared to angle to the goal suggests a stronger effect of path length. The interaction between the factors is significant, too.

For each decision point, we calculated the length of the path and the angle to the goal for both options and considered whether the *risky* jump was associated with the shorter path or smaller angle than the *safe* jump, forming the same four categories as in the analysis of the first decision point. We then computed the probability of choosing the *risky* jump for each participant across these categories (Figure 2.b). As with the first decision point, we fitted a GLMM to the data (Table 2). This time, we considered the model without the interaction term, as it resulted in a smaller BIC. The results of the GLMM confirm that both path length and angle to the goal are significant predictors of *risky* jump probability ($p < 0.001$). This indicates that participants consistently rely on both factors when making decisions across all decision points, not just at the first decision point. Post-hoc comparisons reveal significant differences across all conditions ($p < 0.001$, Table S2), with path length again exerting a stronger influence than angle, as evidenced by both the coefficient sizes and the pairwise tests.

	Fixed effects					Random effects
	Estimate	Std Error	[0.025	0.095]	p-value	Variance
Intercept	-1.610	0.133	-1.870	-1.349	**< 0.001**	0.478
Risky Jump Shorter Path	1.904	0.072	1.762	2.046	**< 0.001**	
Risky Jump Smaller Angle	1.230	0.075	1.084	1.377	**< 0.001**	

Table 2: GLMM results for *risky* jump probability, across all decision points. Both path length and angle to the goal significantly affect the probability of selecting a *risky* jump. Path length has a stronger effect than angle to the goal, as shown by the size of the coefficients and the results of pairwise post-hoc tests.

3.2 Current travel direction influences subsequent jump choices

We tested the hypothesis that participants' choice between *risky* and *safe* jumps were influenced by embodied aspects of the task – specifically, by participants' current travel direction at the moment they made the choice.

To test this hypothesis, we investigated whether the direction of travel before reaching a decision point affects jump choice probability. To analyze this, we introduced the concept of a "trajectory

angle", defined as the angle between the participant's direction of movement and the direction of the jump to be selected; see Figure 3.a for a schematic illustration of how this angle is computed.

Given that each decision point presents two possible choices – either a *risky* or a *safe* jump – we analyzed these two possibilities separately. First, we considered how often participants selected the *risky* jump, based on the trajectory angle. For ease of analysis, we categorized the trajectory angles (which are continuous variables) into a categorical variable with three levels: i) 0° to 45°; ii) 45° to 90°; iii) greater than 90°. This classification reflects the arrangement of rocks in a hexagonal grid, where typical angles between directions (ignoring jitter) are around 0°, 60°, and occasionally 120° or 180° when participants backtrack (note that the latter condition occurs rarely in our dataset). We fit a GLMM to model the probability of choosing a *risky* jump as a function of the trajectory angle (Table 3). Using the $[0°, 45°]$ class as the baseline, both other angle classes exhibit significant differences ($p < 0.001$). This indicates that the trajectory angle significantly influences participants' *risky* jump decisions. The results are further illustrated in Figure 3.b, showing that all conditions are significantly different from each other (see Table S4 for the post-hoc comparisons results).

We next made the same analysis for *safe* jumps (Figure 3.c, Table 4). Similar to *risky* jumps, the results reveal a significant effect of the trajectory angle ($p < 0.001$) and that the probability of performing the *safe* jump decreases significantly as the trajectory angle increases (see Table S5 for post-hoc comparison results). These results show that participants are more likely to select either jump option, when their current travel direction is aligned with it. This suggests that participants do not consider only the static costs and benefits of the choice alternatives but also the continuity of movement and the ease with which they can execute an action within the context of their current dynamic state, demonstrating the relevance of embodied factors like inertia and momentum for the choice. This embodied effect emerges in our task since the sequential choices are not independent from one another and the task does not break down the natural contingency between movement and subsequent perception. This feature is often absent from non-embodied, sequential decisions often addressed in (neuro)economics, which comprise series of discrete and independent choices (Dewey, 1896; Pezzulo et al., 2011; Maselli et al., 2023).

	Fixed effects					Random effects
	Estimate	Std Error	[0.025	0.095]	p-value	Variance
Intercept	-0.456	0.106	-0.664	-0.247	**< 0.001**	0.265
Angle [45°,90°]	-3.022	0.142	-3.301	-2.744	**< 0.001**	
Angle > 90°	-1.013	0.089	-1.187	-0.839	**< 0.001**	

Table 3: GLMM results for *risky* jump probabilities as a function of trajectory angle. The trajectory angle is divided into three categories: $[0°, 45°]$, $[45°, 90°]$, and greater than 90°. The $[0°, 45°]$ category serves as the baseline. The table shows significant decreases in *risky* jump probability as the trajectory angle increases. Random effects account for individual participant variability.

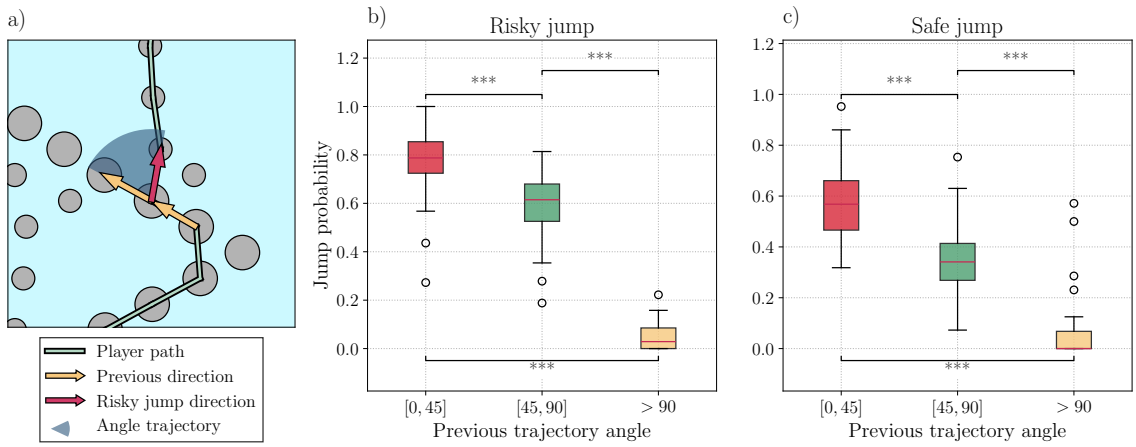

Figure 3: Influence of trajectory angle on *risky* and *safe* jump probabilities. a) Illustration of how the trajectory angle between the previous direction of movement and the jump direction is calculated. **b)** Boxplot showing the probability of performing a *risky* jump as a function of the previous trajectory angle. The trajectory angle was divided into three categories: $[0°, 45°]$, $[45°, 90°]$, and greater than $90°$. Significant differences are observed across all three categories ($p < 0.001$) indicating a sharp decrease in *risky* jump probability as the angle increases. **c)** Boxplot showing the probability of performing a *safe* jump, analyzed in the same way as the *risky* jumps. Here, too, the probability decreases significantly as the trajectory angle increases, demonstrating the same trend observed for *risky* jumps.

3.3 Preplanning time indexes initial decision dynamics

We next examined the relations between the time participants spent preplanning before movement and their subsequent performance and strategy during the levels. We defined *preplanning time* as the time participants spent stationary on the first rock, before the beginning of movement. To assess how preplanning time influenced participants' performance during the trial, we considered the length of the path they selected. We preferred this measure compared to possible alternatives such as total level score or completion time, since they are confounded by factors such as participants' skill with the game and are therefore insufficient to isolate planning processes.

We used a linear mixed-effects model (LMM) to analyze the relationship between preplanning

	Fixed effects					Random effects
	Estimate	Std Error	[0.025	0.095]	p-value	Variance
Intercept	-1.079	0.119	-1.313	-0.846	**< 0.001**	0.318
Angle [45°,90°]	-2.369	0.164	-2.691	-2.048	**< 0.001**	
Angle > 90°	-0.559	0.101	-0.756	-0.361	**< 0.001**	

Table 4: GLMM results for *safe* jump probabilities as a function of trajectory angle. The trajectory angle is divided into three categories: $[0°, 45°]$, $[45°, 90°]$, and greater than $90°$. The $[0°, 45°]$ category serves as the baseline. The table shows significant decreases in *safe* jump probability as the trajectory angle increases. Random effects account for individual participant variability.

time and length of the selected path (Table 5) We found a significant negative effect of preplanning time on path length ($p < 0.001$), showing that participants who spent more time preplanning required fewer jumps to complete the level. A companion, correlation analysis between preplanning time and path length shows a strong negative correlation between the two variables ($R^2 = 0.82$, Figure 4). Quantitatively, each additional second of preplanning reduces path length by approximately 0.274 jumps.

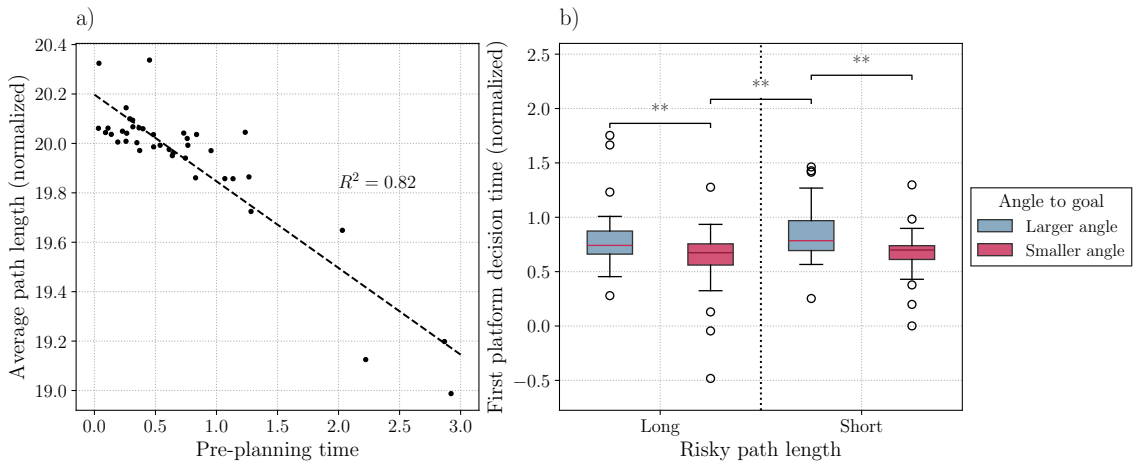

Figure 4: Preplanning time and its effects on path length and decision-making. a) Negative correlation between preplanning time and average path length ($R^2 = 0.82$), indicating that participants who spent more time planning required fewer jumps to complete the level. **b)** Boxplot showing decision time as a function of *risky* path length (long/short) and angle to the goal (larger/smaller). Significant post-hoc contrasts reveal that participants pre-plan faster when the *risky* jump has a smaller angle, with no significant difference in pre-planning time based on path length. Time values have been normalized for visualization purposes by removing the per-participant random intercept from the LMM model, accounting for group variability.

	Fixed effects					Random effects
	Estimate	Std Error	[0.025	0.095]	p-value	Variance
Intercept	20.143	0.131	19.886	20.401	**< 0.001**	0.544
Pre-planning time	-0.274	0.047	-0.366	-0.183	**< 0.001**	

Table 5: Results of the LMM examining the effect of pre-planning time on the length of the selected path. The pre-planning time significantly affects ($p < 0.001$) the length of the path chosen by each participant.

Next, we analyzed how preplanning time was influenced by the path length of the *risky* jump and the angle to the goal at the first decision point, using a generalized linear mixed-effects model (Figure 4 and Table 6). We found no significant difference in preplanning time was observed between levels where the *risky* jump resulted in a longer or shorter path. However, we found that participants spent significantly less time preplanning when the *risky* jump had a smaller angle to the goal ($p = 0.001$). Post-hoc comparisons show participants preplanned faster when the *risky*

jump had a smaller compared to a larger angle to the goal (Tables S3). These findings suggest that participants considered trials in which *risky* jumps had smaller angles to the goal more intuitive and less uncertain.

The two analyses of this section highlight distinctive roles of preplanning for selecting the first jump versus the overall path. Despite the fact that preplanning time is not affected by the length of the *risky* path evaluated at the first decision point, it strongly correlates with participants' ability to select an efficient (shorter) solution to the overall problem, suggesting that during preplanning participants might focus on a global evaluation of the problem rather than only on the first decision.

	Fixed effects					Random effects
	Estimate	Std Error	[0.025	0.095]	p-value	Variance
Intercept	0.792	0.116	0.564	1.019	**< 0.001**	0.459
Risky Jump Shorter Path	0.061	0.052	-0.040	0.163	0.237	
Risky Jump Smaller Angle	-0.171	0.052	-0.272	-0.070	**0.001**	

Table 6: Results of the LMM for pre-planning time as a function of *risky* path length and angle to the goal. There is no significant effect of the *risky* path length on the pre-planning time, whereas the angle factor has a significant impact ($p = 0.001$).

3.4 Planning time during navigation increases at decision points and at platforms preceding them

We next tested the hypothesis that in a continuous task such as crossing the river, planning processes might not only occur at the beginning of a level (i.e., preplanning), but rather continue during the level. To investigate this, we defined *planning time* as the time participants remain stationary on a platform during the solution of a level, without movement or rotation. We then tested whether planning time increases at decision points and/or at platforms immediately preceding decision points, compared to the other platforms. We categorized platforms into three types: i) *Decision points*, where participants must choose between different risky or safe jumps (as previously described); ii) *Pre-decision points*, platforms immediately preceding a decision point, which are not themselves decision points; iii) *Non-decision points*, all other platforms that are neither decision points nor pre-decision points; see Figure 5.a for a schematic illustration of these platform types in an example level.

We fit a Linear Mixed Model (LMM) to predict planning time based on platform type, considered as a categorical variable with three classes (Table 7 and Figure 5.b). The results show greater planning time for both pre-decision points and decision points ($p < 0.001$) compared to non-decision points, which was taken as the baseline. The coefficient estimates indicate that the increase in planning time is smaller for pre-decision points ($\beta = 0.049$) than for decision points ($\beta = 0.114$). See also Table S6 for post-hoc comparisons. These results suggests that planning is a continuous process during crossing the river: it occurs not only before the movement but continues during the task. Furthermore, it is distributed across platforms rather than only occurring at the decision points.

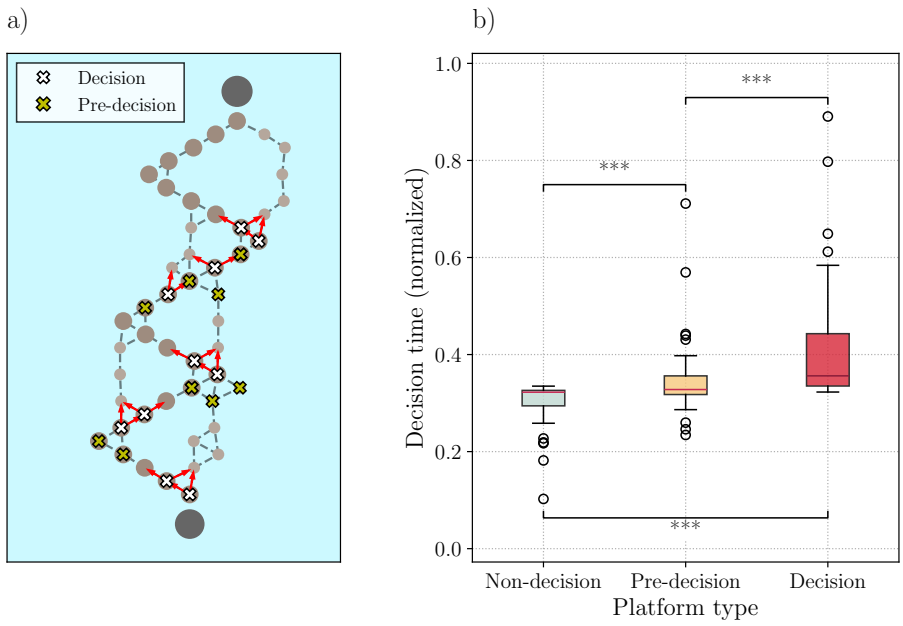

Figure 5: Planning time distribution across different platform types. a) Example level showing the distribution of three platform types: decision points, pre-decision points, and non-decision points. Decision points are platforms where participants make choices between risky and safe jumps, pre-decision points are the platforms immediately preceding a decision point, and non-decision points are all other platforms. **b)** Boxplot illustrating the distribution of planning time across the three platform types. Significant differences between platform types are indicated by lines, showing that decision points have the longest planning time, followed by pre-decision points, while non-decision points exhibit the shortest times. Post-hoc comparisons show significant differences ($p < 0.001$) across all three platform types.

3.5 Computational model

We sought to formalize the sequential decision process that participants face in the experiment, using a simple sequential decision model based upon the notion of "expected value" (EV) (Von Neumann and Morgenstern, 1947). The model assigns an EV to each jump, calculated as the product of two components: utility (U), corresponding to the inverse of the length of the path after the jump, and affordance (A), corresponding to the jumpability" or "landability" (risky or safe) of the destination platform. In this context, the affordance takes the place of outcome probability in the classical economic setting Gordon et al. (2021). Formally, the EV can be expressed as:

$$\mathrm{EV}(s, a) = \mathrm{U}(s, a) \times \mathrm{A}(s, a)^{-1} \tag{1}$$

where s is the current state of the avatar (the platform it occupies) and a is the action (the choice of jumping to one of the neighboring platforms). This formulation emphasizes that the EV of each jumping action depends on both the utility and affordance of the landing platform. Note that the affordance (A) term is raised to minus one, to compensate for the fact that the utility term only assumes negative values.

	Fixed effects					Random effects
	Estimate	Std Error	[0.025	0.095]	p-value	Variance
Intercept	0.304	0.074	0.159	0.448	**< 0.001**	0.216
Pre-decision	0.049	0.010	0.029	0.069	**< 0.001**	
Decision	0.114	0.008	0.098	0.131	**< 0.001**	

Table 7: LMM results for planning time as a function of platform type. The table shows that the time spent planning on pre-decision and decision points is significantly different ($p < 0.001$) from the time spent on non-decision platforms, which serve as the baseline.

To extend the model to sequential jumps, We incorporated a forward-looking aspect, taking into account future jumping choices. For this, we included a discount factor γ, which controls how much weight is given to future jumping choices. The revised formula for expected value becomes:

$$\mathrm{EV}(s,a) = \mathrm{U}(s,a)\mathrm{A}(s,a)^{-1} + \gamma \cdot \max_{a'}\mathrm{EV}(s',a') \tag{2}$$

where $s' = f(s,a)$ represents the new state reached after action a is taken, and a' denotes the available actions from that new state. The discount factor γ determines how far-sighted the agent is. With $\gamma = 0$, the agent only considers the immediate outcome, whereas with $\gamma = 1$ it gives equal consideration to future decisions. Intermediate values of γ represent varying degrees of far-sightedness.

In this model, utility $U(s,a)$ is defined as a function of path length. If the state s is the goal platform, then $U(s,a) = 0$, as no further jumps are needed. For all other states, $U(s,a) = -1$. This negative value ensures that, when affordances are ignored (i.e., when $A(s,a) = 1$), the expected value $\mathrm{EV}(s,a)$ with $\gamma = 1$ simply corresponds to the number of jumps required to reach the goal, via the shortest path.

Affordance $A(s,a)$ indexes jumpability and has larger values for easier jumps and smaller values for more challenging jumps. While previous sections explored various types of affordances, here for simplicity we only consider platform distance. While on average greater distance implies a small platform and a *risky* jump, there might be small differences in the distances due to the jitter of rocks – hence considering distance permits a more fine-grained evaluation compared to only considering *risky* and *safe* jumps.

Formally, affordance is defined as:

$$A(s,a) = \exp\left(-\alpha \cdot d(s,s')\right) \tag{3}$$

where $d(s,s')$ is the distance between the current platform and the destination platform after jump a and α is a coefficient that controls the weight of this factor. To ensure that the model focuses on relative rather than absolute distances, we normalize the affordance values for each platform, such that the maximum value is set to 1 and the other values are scaled accordingly. This scaling allows the model to treat all platforms equally, with only the relative difficulty of jumps influencing the affordance.

Solving this system of equations iteratively, we can compute the EV for each jump, at every decision point. Figure 6 shows landscapes of EVs of an example level, normalized in such a way that greater values (closer to one) correspond to better choices. The figure also illustrates the impact of different values of the model parameters α and γ. Each row corresponds to a different value of α, representing the weight of the affordance term, while each column shows a different value of γ, representing the degree of farsightedness. For selected decision points in each level, the figure shows the EV of the two available choices (*risky* and *safe* jumps) as a color gradient on the destination platform, illustrating how the model balances affordance and utility under different conditions.

To fit the model to participants' behavior, we convert the expected values into a policy using a softmax function:

$$\pi(a|s) = \frac{\exp\left(\beta \cdot \mathrm{EV}(s,a)\right)}{\sum_{a'} \exp\left(\beta \cdot \mathrm{EV}(s,a')\right)} \tag{4}$$

We set $\beta = 3$, strongly biasing the policy towards the action with the highest EV. For each participant, we measure the negative log-likelihood of their chosen path under the model, using the following equation:

$$-\mathcal{L} = \sum_{(s,a)\in \mathrm{path}} \log \pi(a|s) \tag{5}$$

We sum this quantity for each of the 40 levels to get the total log-likelihood. We then minimize the total log-likelihood with respect to the two free parameters, γ and α, for each participant.

The fitted values of γ range from 0.798 to 0.981, indicating that all participants engage in some degree of forward planning while solving the levels. The relatively narrow range of γ values can be explained by considering that small variations in this parameter result in large differences in planning behavior, due to its exponential effect on future choices as the number of steps increases.

The fitted affordance coefficient α reveals two groups of participants. For 15 participants, the values of α are equal to zero, suggesting that this group does not prioritize affordance (jump distance) when evaluating the EV of their choices. In other words, these participants adopt a rather uniform strategy, treating all jumps similarly, regardless of distance. By contrast, the α of the other participants range between 0 and 0.5, implying that they consider the affordance (distance between platforms) when making jumping choices.

We explored whether the two parameters, farsightedness γ and affordance coefficient α, correlate with participant performance in the experiment (Figure 7). Our results reveal a significant positive correlation between γ and participant performance, with more farsighted participants (those having higher γ values) achieving better scores in the task (Pearson $r = 0.53$, $p < 0.001$; Figure 7.a). Furthermore, our results reveal a negative correlation between α and the score, with participants giving more weight to affordances (those having higher α values) achieving lower scores (Pearson $r = -0.44$, $p = 0.005$; Figure 7.b). These results suggest that high-performing participants might pay little attention to the distance between platforms.

Furthermore, we examined the number of errors made by each participant throughout the experiment, where an error was defined as falling into the water during a jump. Our results reveal a negative correlation between the farsightedness parameter γ and the number of errors, suggest-

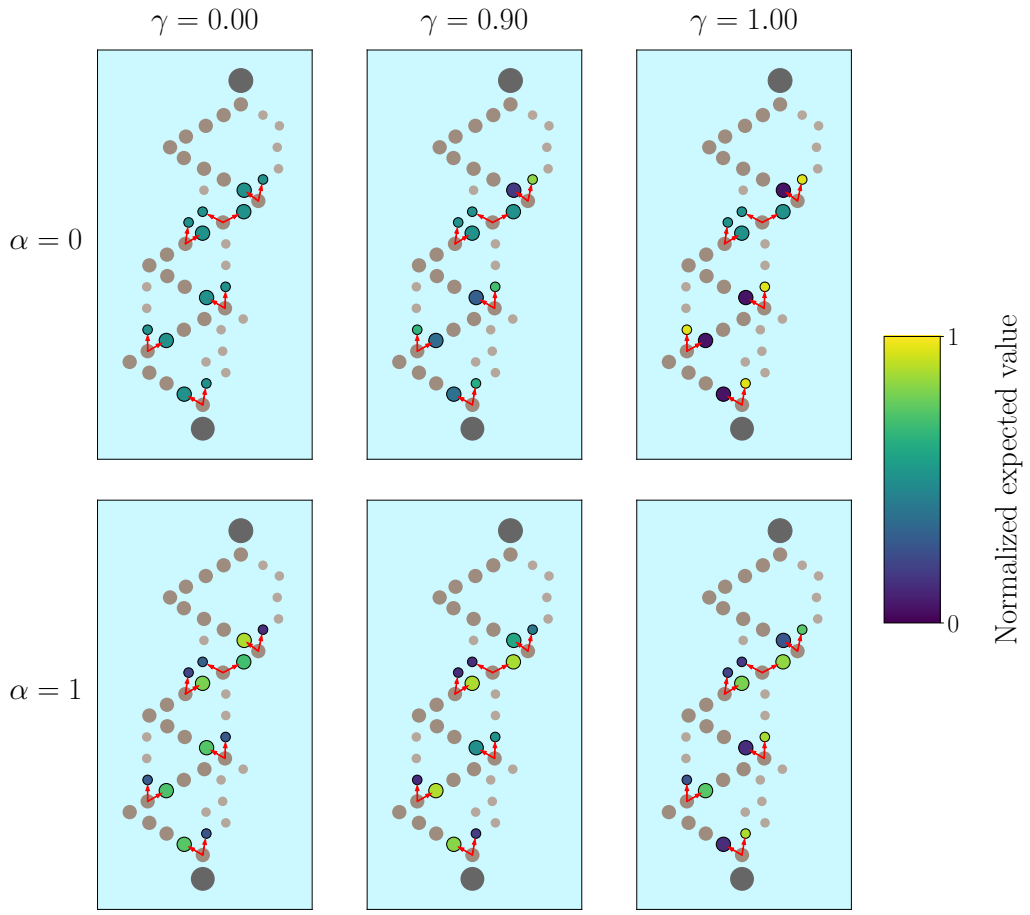

Figure 6: Normalized Expected Value (EV) landscapes for an example problem. The figure shows how normalized EVs change as a function of the parameters α (rows) and γ (columns) parameters, reflecting affordances and farsightedness, respectively. The color of each platform represents the normalized EV for the corresponding action, with a gradient from 0 (dark blue, representing less favorable actions) to 1 (yellow, indicating more favorable actions). In the top row ($\alpha = 0$), where affordances do not vary with distance, all jumps have the same affordance. For $\gamma = 0$, the agent evaluates all jumps with the same EV, since it does not consider future jumps. As γ increases (moving from left to right), the EV increases for platforms along the shortest path toward the goal, reflecting the influence of planning. In the bottom row ($\alpha = 1$), where the weight of distance-based affordances increases, the model starts to favor safer jumps to closer platforms. This tendency is evident even for $\gamma = 0.9$, where jumping decisions are also influenced by future jumps. For $\gamma = 1$, the agent shows a strong preference for the shortest path, while still slightly favoring safe jumps.

ing that more farsighted participants tended to make fewer errors (Figure 7.c). Furthermore, the results reveal a positive correlation between affordance coefficients α and number of errors, potentially reflecting the fact that participants who made more (fewer) errors tended to assign more (less) importance to the relative difficulty of making safe or risky jumps (Figure 7.c).

To further investigate the role of affordances in jump performance, we examined the relation-

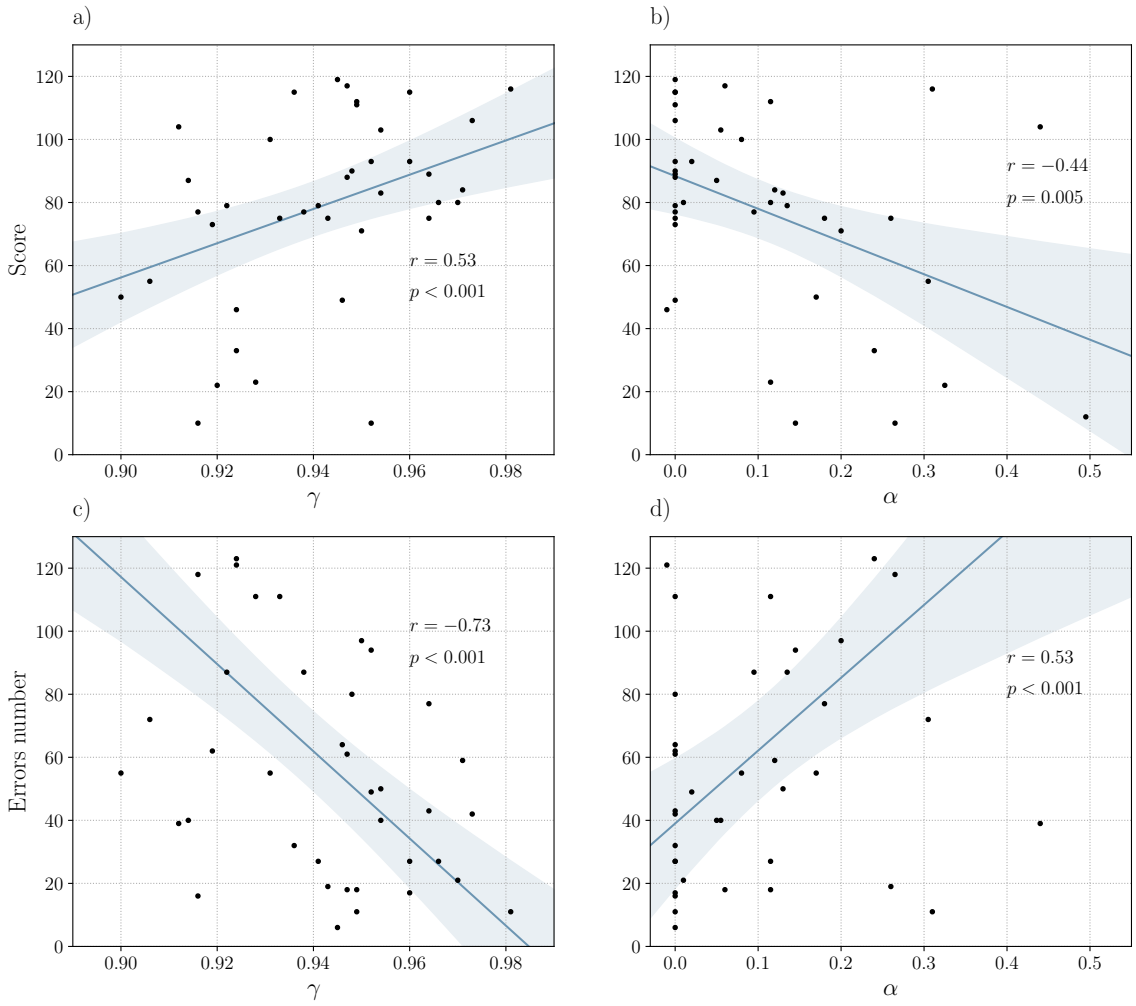

Figure 7: Correlations between participants' performance and model parameters. a) Positive correlation between the farsightedness parameter γ and the participants' scores. Higher values of γ are associated with better performance, with Pearson $r = 0.53$ and $p < 0.001$. **b)** Negative correlation between the affordance coefficient α and performance, indicating that participants who weighed affordances more heavily tended to score lower, with Pearson $r = -0.44$ and $p = 0.005$. The shaded area represents the 95% confidence interval. **c)** Negative correlation between the farsightedness parameter γ and the number of errors made by participants, with higher values of γ associated with fewer errors. Pearson $r = -0.73$ and $p < 0.001$. **d)** Positive correlation between the affordance coefficient α and the number of errors. Participants who considered affordances more tended to make more errors. Pearson $r = 0.53$ and $p < 0.001$.

ship between α and jump success probability. For consistency with the selection criteria used in other analyses and our definition of decision points, we only focused on jumps starting from large platforms. Furthermore, we only included in the analyses those failed jumps where it was clear which rock the participant was aiming to reach; instead, we discarded ambiguous jumps in which participants landed in the water near their starting rock (Figure 8).

Our analysis reveals that success rates were significantly smaller with risky compared to safe jumps, highlighting the increased difficulty of the former (Figure 8.a). Furthermore, our analysis reveals a negative correlation between α and the change in success probability from safe to risky jumps (Figure 8.b). Specifically, participants with lower α values show minimal changes in success probability, suggesting that affordance-related factors may not be relevant to them, as all jumps, whether safe or risky, are equally easy for them. In contrast, higher α participants are more sensitive to the relative difficulty posed by risky jumps.

Finally, we compared the log-likelihoods between our EV model, a random agent, and an agent that always follows the shortest path (Table 8). We found the log-likelihood of our EV model to be lower than both the random agent and the shortest path agent. The shortest path agent performs better than the random agent, but exhibits a much higher log-likelihood than our EV model, along with a much larger standard deviation. This comparison suggests that participants make decisions that take into account both future planning and affordances, as captured by the fitted parameters of our model, rather than exclusively following the shortest path.

	Negative log-likelihood	
	Mean	Std
Our EV model	152.46	31.11
Random agent	234.77	23.74
Shortest path agent	209.95	83.11

Table 8: Comparison of Negative Log-Likelihood Across Models. The negative log-likelihood for the three models: our EV model, a random agent, and a shortest path agent. Lower values indicate better model fit to participant behavior.

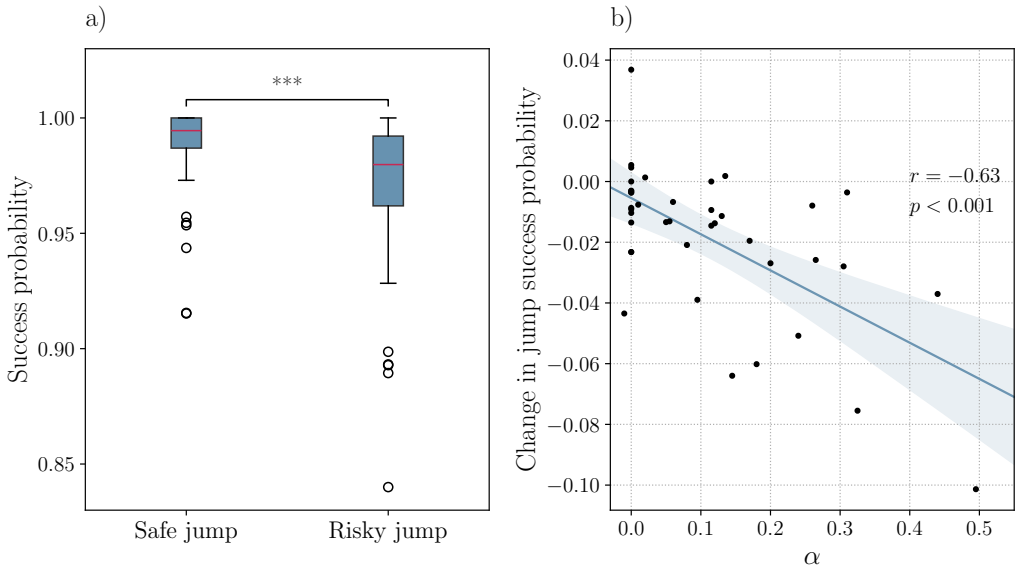

Figure 8: Success probability and its relationship with jump type (safe or risky) and the affordance parameter. a) Success probability of safe versus risky jumps, starting from large platforms. Success rates are significantly lower for risky jumps, indicating their increased difficulty. **b)** Correlation between the affordance parameter α and the change in success probability when making risky jumps. Participants with lower α values show minimal changes in success probability from safe to risky jumps, suggesting a lower sensitivity to affordances, while higher α participants exhibit a stronger differentiation between the two types of jumps.

4 Discussion

Every day, we face a myriad of embodied decisions, such as when deciding what path to take in a busy road or which objects to use and in which order to set up a table for breakfast or prepare a backpack for a trekking. These and other tasks present sequences of choices between immediate and more distal affordances and often require simultaneously making sequential decisions and actions (Cisek and Pastor-Bernier, 2014; Gibson, 1979; Pezzulo and Cisek, 2016; Cisek and Kalaska, 2010; Yoo et al., 2021; Priorelli et al., 2024; Wispinski et al., 2020; Gallivan et al., 2018; Nashed et al., 2014). Despite their importance, embodied decisions remain poorly investigated. Most studies of decision-making and planning in human psychology and (neuro)economics focus on non-embodied settings and the few studies of embodied decisions tend to focus on single embodied decisions, not sequential decisions (Cisek and Pastor-Bernier, 2014; Gibson, 1979; Pezzulo and Cisek, 2016; Cisek and Kalaska, 2010; Yoo et al., 2021; Priorelli et al., 2024; Gallivan et al., 2015). Hence, we still do not know whether, when making sequences of embodied decisions, people plan ahead by considering how different choices will lead to new affordances, whether they preplan an entire task before starting or continue planning even while acting, and whether embodied aspects of action execution influence the decisions that are made.

We addressed the above questions by designing a game-like experiment in which participants controlled an avatar required to "cross rivers", by jumping on a series of rocks of different sizes

forming different paths to the goal destination. Crucially, the task presented participants with a series of choices between "safe" and "risky" jumps (corresponding to bigger-and-closer versus smaller-and-farther rocks) and between longer and shorter paths to the goal. This setup permitted us to study how people balance affordance and utility in sequential embodied settings. To get more insights about participants' strategies, we also developed a computational model implementing jumping decisions as a function of current affordances ("safe" and "risky" jumps) and the utility (length) of the paths that they open up. By fitting the models' parameters (the weight of the affordance term and the degree of forward planning) to participants' behavior, we aimed to formalize their decision strategies.

Our behavioral and computational analyses highlight five main findings. First, participants plan ahead their path to the goal rather than simply focusing on the most immediate jumping affordances. At decision points between safe and risky jumps, they balance jumping affordances and their long-term utility, indexed by the length of the paths that they open up and the angle that they form with the goal. Interestingly, despite the angle to goal being a more perceptually available proxy for path utility (albeit potentially misleading), participants take into more consideration the path length, which is a veridical index of utility. The fact that participants show effective planning strategies during "crossing the river" illustrates the effectiveness of embodied tasks to study higher cognitive processes (Maselli et al., 2023).

Second, embodied components of the task influence participants' decision strategies. This is demonstrated by the fact that participants' current direction of movement influences their choices between risky and safe jumps. This is possible since, despite being merely virtually embodied, our game requires detailed sensorimotor control of the avatar and the consideration of environmental affordances. Furthermore, the movements of the avatar follow coherent physical and spatial rules and the ease of executing a jump while continuing along the same trajectory (with some inertia and without the motor cost required to change direction) may impact the participant's decisions. This finding illustrates that in (virtually) embodied settings, action dynamics can influence decisions (Lepora and Pezzulo, 2015; Priorelli et al., 2024). From a methodological point of view, studying sequential choices in embodied settings reveals subtle aspects of decision processes that might be relevant to understand everyday cognition, but might be missed when only focusing on non-embodied settings, as usually done in psychology and (neuro)economics.

Third, participants preplan part of the path at the beginning of the task, before making the first jump. This is evidenced by the fact that preplanning time is influenced by the perceived angle between the first risky jump and the goal and that longer preplanning time can lead to improved solutions (shorter paths). This finding parallels studies of preplanning during more abstract problem solving, such as the fact that humans think for different durations in different contexts (Jensen et al., 2024), highlighting that participants are sensitive to the sequential planning nature of "crossing the river".

Fourth, planning is not completed before movement but continues during the task. This is shown by the fact that participants' deliberation time (indexed by slowing down during the task, not related to motor requirements) increases at decision points and even before them. This finding illustrates that in embodied settings, participants can simultaneously plan and act – and that a strict

separation between the two processes does not allow accounting for subtle aspects of naturalistic behavior, such as the fact that movement kinematics index decision variables such as decision uncertainty and confidence (Song and Nakayama, 2009; Barca and Pezzulo, 2012; Dotan et al., 2018) and that participants can change their mind along the way (Spivey, 2008).

Finally, computational modeling shows that farsighted participants (which assigned greater weight to the utility of future jumps) show better task performance, above and beyond their jumping skill (which varies across participants independent of decision strategy). This finding highlights the possibility to adopt computational models that formalize the usual trade-offs that characterize value-based decisions, such as between the probabilities of the different outcomes (here, of successful jumps) and their utilities (here, path length) but in embodied settings (Maselli et al., 2023). Besides providing a way to quantify participants' strategies, the variables inferred by the model – namely, affordance and utility landscapes at decision points – could be potentially used a predictors of neural representations in future studies that address the neuronal underpinnings of embodied decision and planning processes Thura et al. (2022); Cisek and Kalaska (2005); Klaes et al. (2011); Siegel et al. (2015).

Taken together, our results show that when solving an embodied task (even a "virtually embodied" one) – crossing the river by continuously selecting and executing "risky" or "safe" jumps – participants adopt effective planning strategies that balance immediate jump affordances and long-term path utility. Notably, participants' decision strategies are influenced by embodied characteristics of the task such as their current travel direction and reciprocally, their movement speed is influenced by the necessity to continue planning along the way. These findings emphasize the interplay between decision and action in embodied tasks, shedding light on the cognitive mechanisms underlying real-world planning-while-acting scenarios. Furthermore, they underscore the importance of studying higher level processes such as planning in settings that reflect the ecological conditions of our everyday decisions and actions.

Acknowledgements

This research received funding from the European Research Council under the Grant Agreement No. 820213 (ThinkAhead), the Italian National Recovery and Resilience Plan (NRRP), M4C2, funded by the European Union – NextGenerationEU (Project IR0000011, CUP B51E22000150006, "EBRAINS-Italy"; Project PE0000013, "FAIR"; Project PE0000006, "MNESYS"), and the Ministry of University and Research, PRIN PNRR P20224FESY and PRIN 20229Z7M8N. PC was supported by a grant from the Natural Sciences and Engineering Research Council of Canada (RGPIN-2022-05345). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript. The GEFORCE Quadro RTX6000 and Titan GPU cards used for this research were donated by the NVIDIA Corporation.

References

Barca, L. and Pezzulo, G. (2012). Unfolding visual lexical decision in time. *PloS one*, 7(4):e35932.

Cisek, P. and Kalaska, J. F. (2005). Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. *Neuron*, 45(5):801–814.

Cisek, P. and Kalaska, J. F. (2010). Neural mechanisms for interacting with a world full of action choices. *Annual review of neuroscience*, 33(1):269–298.

Cisek, P. and Pastor-Bernier, A. (2014). On the challenges and mechanisms of embodied decisions. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1655):20130479.

Connors, B. L., Moore, C.-L., Rende, R., and Colton, T. J. (2018). Movement pattern analysis (mpa): decoding individual differences in embodied decision making. *The SAGE Handbook of Personality and Individual Differences*, pages 257–277.

Cos, I., Pezzulo, G., and Cisek, P. (2021). Changes of mind after movement onset depend on the state of the motor system. *Eneuro*, 8(6).

Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., and Dolan, R. J. (2011). Model-based influences on humans' choices and striatal prediction errors. *Neuron*, 69(6):1204–1215.

Dewey, J. (1896). The reflex arc concept in psychology. *Psychological review*, 3(4):357.

Dotan, D., Meyniel, F., and Dehaene, S. (2018). On-line confidence monitoring during decision making. *Cognition*, 171:112–121.

Gallivan, J. P., Barton, K. S., Chapman, C. S., Wolpert, D. M., and Randall Flanagan, J. (2015). Action plan co-optimization reveals the parallel encoding of competing reach movements. *Nature communications*, 6(1):7428.

Gallivan, J. P., Chapman, C. S., Wolpert, D. M., and Flanagan, J. R. (2018). Decision-making in sensorimotor control. *Nature Reviews Neuroscience*, 19(9):519–534.

Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Houghton Mifflin.

Gordon, J., Maselli, A., Lancia, G. L., Thiery, T., Cisek, P., and Pezzulo, G. (2021). The road towards understanding embodied decisions. *Neuroscience & Biobehavioral Reviews*, 131:722–736.

Grießbach, E., Raßbach, P., Herbort, O., and Cañal-Bruland, R. (2022). Embodied decisions during walking. *Journal of Neurophysiology*, 128(5):1207–1223.

Jensen, K. T., Hennequin, G., and Mattar, M. G. (2024). A recurrent network model of planning explains hippocampal replay and human behavior. *Nature Neuroscience*, pages 1–9.

Klaes, C., Westendorff, S., Chakrabarti, S., and Gail, A. (2011). Choosing goals, not rules: deciding among rule-based action plans. *Neuron*, 70(3):536–548.

Kobayashi, A. and Kimura, T. (2022). Compensative movement ameliorates reduced efficacy of rapidly-embodied decisions in humans. *Communications biology*, 5(1):294.

Lepora, N. F. and Pezzulo, G. (2015). Embodied choice: how action influences perceptual decision making. *PLoS computational biology*, 11(4):e1004110.

Maselli, A., Gordon, J., Eluchans, M., Lancia, G. L., Thiery, T., Moretti, R., Cisek, P., and Pezzulo, G. (2023). Beyond simple laboratory studies: developing sophisticated models to study rich behavior. *Physics of Life Reviews*.

Michalski, J., Green, A. M., and Cisek, P. (2020). Reaching decisions during ongoing movements. *Journal of Neurophysiology*, 123(3):1090–1102.

Nashed, J. Y., Crevecoeur, F., and Scott, S. H. (2014). Rapid online selection between multiple motor plans. *Journal of Neuroscience*, 34(5):1769–1780.

Pezzulo, G., Barsalou, L. W., Cangelosi, A., Fischer, M. H., McRae, K., and Spivey, M. J. (2011). The mechanics of embodiment: A dialog on embodiment and computational modeling. *Frontiers in psychology*, 2:5.

Pezzulo, G. and Cisek, P. (2016). Navigating the affordance landscape: feedback control as a process model of behavior and cognition. *Trends in cognitive sciences*, 20(6):414–424.

Priorelli, M., Stoianov, I. P., and Pezzulo, G. (2024). Embodied decisions as active inference. *bioRxiv*, pages 2024–05.

Raab, M. (2020). *Judgment, decision-making, and embodied choices*. Academic Press.

Siegel, M., Buschman, T. J., and Miller, E. K. (2015). Cortical information flow during flexible sensorimotor decisions. *Science*, 348(6241):1352–1355.

Song, J.-H. and Nakayama, K. (2009). Hidden cognitive states revealed in choice reaching tasks. *Trends in cognitive sciences*, 13(8):360–366.

Spivey, M. (2008). *The continuity of mind*. Oxford University Press.

Thura, D., Cabana, J.-F., Feghaly, A., and Cisek, P. (2022). Integrated neural dynamics of sensorimotor decisions and actions. *PLoS biology*, 20(12):e3001861.

Ulbrich, P. and Gail, A. (2023). Deciding while acting—mid-movement decisions are more strongly affected by action probability than reward amount. *Eneuro*, 10(4).

Von Neumann, J. and Morgenstern, O. (1947). *Theory of games and economic behavior, 2nd rev.* Princeton university press.

Warren, W. H. (2006). The dynamics of locomotion: Perception and action in complex environments. *Experimental Brain Research*, 173:623–632.

Wispinski, N. J., Gallivan, J. P., and Chapman, C. S. (2020). Models, movements, and minds: bridging the gap between decision making and action. *Annals of the New York Academy of Sciences*, 1464(1):30–51.

Yoo, S. B. M., Hayden, B. Y., and Pearson, J. M. (2021). Continuous decisions. *Philosophical Transactions of the Royal Society B*, 376(1819):20190664.

A Supplementary materials

Contrast	Estimate	Std Error	[0.025	0.095]	p-value
Larger / Longer - Smaller / Longer	-0.721	0.161	-1.135	-0.307	**< 0.001**
Larger / Longer - Larger / Shorter	-1.082	0.158	-1.487	-0.677	**< 0.001**
Larger / Longer - Smaller / Shorter	-2.655	0.190	-3.142	-2.168	**< 0.001**
Smaller / Longer - Larger / Shorter	-0.361	0.157	-0.765	0.044	0.100
Smaller / Longer - Smaller / Shorter	-1.934	0.188	-2.416	-1.452	**< 0.001**
Larger / Shorter - Smaller / Shorter	-1.573	0.182	-2.042	-1.104	**< 0.001**

Table S1: Pairwise post-hoc comparisons for the probability of making a *risky* jump at the first decision point. The contrasts between different combinations of path length and angle are shown. For brevity, larger/smaller refer to the angle to the goal, and longer/shorter refer to the path length. The table provides the estimated effects, standard errors, confidence intervals, and p-values for each contrast. Significant differences between conditions are highlighted ($p < 0.001$).

Contrast	Estimate	Std Error	[0.025	0.095]	p-value
Larger / Longer - Smaller / Longer	-1.014	0.116	-1.313	-0.716	**< 0.001**
Larger / Longer - Larger / Shorter	-1.665	0.123	-1.980	-1.350	**< 0.001**
Larger / Longer - Smaller / Shorter	-3.037	0.113	-3.327	-2.746	**< 0.001**
Smaller / Longer - Larger / Shorter	-0.651	0.101	-0.910	-0.391	**< 0.001**
Smaller / Longer - Smaller / Shorter	-2.022	0.088	-2.250	-1.795	**< 0.001**
Larger / Shorter - Smaller / Shorter	-1.372	0.096	-1.617	-1.126	**< 0.001**

Table S2: Pairwise post-hoc comparisons for the probability of making a *risky* jump across all decision points. The contrasts represent combinations of angle to the goal (larger/smaller) and path length (longer/shorter). The table shows the estimated effects, standard errors, confidence intervals, and p-values for each pairwise comparison. Significant results are marked ($p < 0.001$).

Contrast	Estimate	Std Error	[0.025	0.095]	p-value
Larger / Longer - Smaller / Longer	0.171	0.052	0.038	0.304	**0.005**
Larger / Longer - Larger / Shorter	-0.061	0.052	-0.194	0.072	0.638
Larger / Longer - Smaller / Shorter	0.110	0.073	-0.079	0.298	0.438
Smaller / Longer - Larger / Shorter	-0.232	0.073	-0.420	-0.044	**0.008**
Smaller / Longer - Smaller / Shorter	-0.061	0.052	-0.194	0.072	0.638
Larger / Shorter - Smaller / Shorter	0.171	0.052	0.038	0.304	**0.005**

Table S3: Pairwise post-hoc comparisons for the pre-planning time. The contrasts represent combinations of angle to the goal (larger/smaller) and path length (longer/shorter). The table shows the estimated effects, standard errors, confidence intervals, and p-values for each pairwise comparison. Significant results are marked ($p < 0.001$).

Contrast	Estimate	Std Error	[0.025	0.095]	p-value
$[0°, 45°]$ - $[45°, 90°]$	0.896	0.074	0.722	1.070	$<$ **0.001**
$[0°, 45°]$ - Greater than 90°	4.274	0.201	3.803	4.745	$<$ **0.001**
$[45°, 90°]$ - Greater than 90°	3.378	0.196	2.919	3.837	$<$ **0.001**

Table S4: Post-hoc comparison of trajectory angle categories for *risky* jump probabilities. results of the post-hoc pairwise comparisons between the three trajectory angle categories ($[0°, 45°]$, $[45°, 90°]$, and greater than 90°) for *risky* jump probabilities. Estimates, standard errors, confidence intervals, and p-values are reported for each contrast. All comparisons show significant differences ($p < 0.001$) between angle categories.

Contrast	Estimate	Std Error	[0.025	0.095]	p-value
$[0°, 45°]$ - $[45°, 90°]$	0.991	0.065	0.838	1.145	$<$ **0.001**
$[0°, 45°]$ - Greater than 90°	3.351	0.232	2.807	3.895	$<$ **0.001**
$[45°, 90°]$ - Greater than 90°	2.360	0.231	1.819	2.900	$<$ **0.001**

Table S5: Post-hoc comparison of trajectory angle categories for *safe* jump probabilities. results of the post-hoc pairwise comparisons between the three trajectory angle categories ($[0°, 45°]$, $[45°, 90°]$, and greater than 90°) for *safe* jump probabilities. Estimates, standard errors, confidence intervals, and p-values are reported for each contrast. All comparisons show significant differences ($p < 0.001$) between angle categories.

Contrast	Estimate	Std Error	[0.025	0.095]	p-value
Nondecision - Predecision	-0.049	0.010	-0.073	-0.025	$<$ **0.001**
Nondecision - Decision	-0.114	0.008	-0.134	-0.095	$<$ **0.001**
Predecision - Decision	-0.065	0.012	-0.094	-0.037	$<$ **0.001**

Table S6