

Uncertainty affects planning effort, but not plans

Jordan Lei (jordan.lei@nyu.edu)

Center for Neural Science, New York University
New York, NY, USA

Wei Ji Ma (weijima@nyu.edu)

Center for Neural Science and Department of Psychology, New York University
New York, NY, USA

Abstract

When people plan, they often do so in the face of uncertainty. However, little is known about how uncertainty affects planning. To study these effects, we used a reward gathering task in which the reliability of announced rewards varied from certain to completely random. We quantitatively compared several planning models. We found that participants used a suboptimal approach, failing to directly incorporate stochasticity into their planning. Instead, they “compensated” for uncertainty by decreasing their planning effort as stochasticity increased. First-move response time correspondingly decreased with increasing stochasticity. Our findings generalized to a manipulation of transition uncertainty. Together, these findings open the door to a more comprehensive and computationally grounded understanding of the role of stochasticity in planning.

Keywords: stochasticity; uncertainty; sequential decision-making; planning; behavioral modeling

Introduction

Planning is difficult. One of the challenges of planning is the fact that the future contains some level of irreducible uncertainty. This uncertainty, in turn, influences how people plan. Intuitively, in a very random environment, one might feel that it is simply “not worth” putting the effort into making a plan. If the future is uncertain, why bother? Here we attempt to understand this phenomenon.

Early cognitive studies of planning focused on tasks that had minimal uncertainty, focusing instead on difficulty derived from the task structure. In tasks like the Tower of London (Shallice, 1982) and Tower of Hanoi (Simon, 1975), the participant knows with full certainty the outcome of their actions. While this work laid a strong foundation for understanding the cognitive and neural underpinnings of prospective reasoning (Owen, Downes, Sahakian, Polkey, & Robbins, 1990), the tasks did not reflect the changing and often stochastic nature of the real world. More modern planning tasks, such as the two-step task (Daw, Gershman, Seymour, Dayan, & Dolan, 2011), introduced randomness in the transition structure of the task to probe how people make decisions in uncertain environments. It is noteworthy, however, that most studies using variants of the two-step task (so-called N-step tasks) did not vary the uncertainty of the environment across different conditions. Therefore canonical planning studies, old and new, have failed to probe the role of uncertainty in planning.

For studies addressing decision making in uncertain environments, we look towards learning and navigation, two domains where uncertainty is inextricably intertwined with decision making. Behrens, Woolrich, Walton, and Rushworth (2007) demonstrated that people adapt their learning rates according to the uncertainty of their environment, with higher learning rates in uncertain environments. While not explicitly a planning task, this work showed that people infer optimal strategies in stochastic learning environments. Wiener, Lafon, and Berthoz (2008) used a navigation task (a variant of the Traveling Salesman Problem) where participants were asked to retrieve a hidden object at one of six locations. They found that people were able to effectively learn and use probabilistic information about the object’s position to better retrieve the target object; in other words, humans adapt their planning strategies to the uncertainty of their environment. More recently, Mitra, Srivastava, and Srinivasan (2023) demonstrated that participants in a farming task chose shorter-horizon crops when faced with negative resource shocks. However, this study did not model how people plan or how planning effort is modulated by uncertainty. In sum, the field lacks an understanding of how people change their planning behavior in response to uncertainty.

There are a few ways that people could adapt their plans in the face of randomness. One hypothesis is that people *explicitly* factor randomness into their plans by estimating the expected value of outcomes given the uncertainty in their environments. Another hypothesis is that people *implicitly* factor randomness in their plans, by acting as if the world was certain but changing the *effort* they exert in response to uncertainty. Of course, it is also possible that people don’t adapt their plans to uncertainty at all.

To test these hypotheses, we introduce a planning task where the environment is subject to five graded levels of stochasticity. We model the participant’s behavior with a flexible parameter reflecting planning effort across the stochasticity conditions. We evaluate this effect across two different forms of stochasticity.

Experiment 1: Reliability

Imagine you are planning to explore the best pizza spots in New York City with a group of friends. One of them has compiled a list of venues from the last time they visited the city. However, because the list is old, it is possible that the

quality of the restaurants has changed over time - some for worse, some for better. Here, planning is made difficult by the *reliability* of your information. You may find that some of the restaurants that you visited were better than your friend last remembered them; other restaurants might be much worse. The more unreliable the information you get, the harder it is to plan. In our first task, we explored this relationship - how unreliability affects people’s planning effort.

Task Design

We designed a minimal task to study how people adapt their planning effort in response to unreliability, if at all. We designed our task with the following criteria: (1) the task needed some component of planning or prospective thinking, (2) the task needed an environment-wide stochasticity parameter that could be varied across conditions, and (3) we must be able to infer a player’s depth of planning from their behavior.

For the basic design of the task (ignoring stochasticity), we used a variant of the game proposed in Snider, Lee, Poizner, and Gepshtein (2015). Participants are presented with a board containing several treasure chests arranged in a triangle (Fig 1A). Players begin at the topmost “start” node. For each move, the player must travel either left-and-down or right-and-down to a neighboring treasure chest. Each treasure chest has a number (between 1 and 9, inclusive) written on it corresponding to how many points it is worth. The goal is to maximize the total number of points as they move downward.

We consider five different stochasticity conditions, corresponding to varying levels of unreliability in the environment. Every treasure chest can be one of two types: a “normal” treasure chest, where the number of points received is the same as the number written on the treasure chest, or a “mystery” treasure chest, where the number of points received is a random number drawn from 1 to 9, regardless of the number written on the treasure chest (Fig 1B). In each game, each treasure chest has a probability q that it is a mystery chest. Across the five conditions, we vary q : $q = 0\%$, 25% , 50% , 75% , and 100% probability, where 0% is completely certain and 100% is completely random (i.e. maximally unreliable). Participants do not know in advance which specific treasure chests are “normal” or “mystery”. The only information they are given about the status of the treasure chests is the probability q , which is fixed for each game.

Participants

We recruited 100 participants through Prolific. The estimated duration of the task was approximately an hour. Each participant played 10 blocks of 15 games, where the stochasticity level q was held constant within each block. Participants were paid 8 USD for taking part in the study and rewarded bonuses proportional to their above-chance performance on a random sample of 5 games in the session.

All participants were given comprehension checks and opportunities to practice (minimum six practice games: 1 miniature version of the game and one full practice game for each stochasticity condition) before beginning the full experiment.

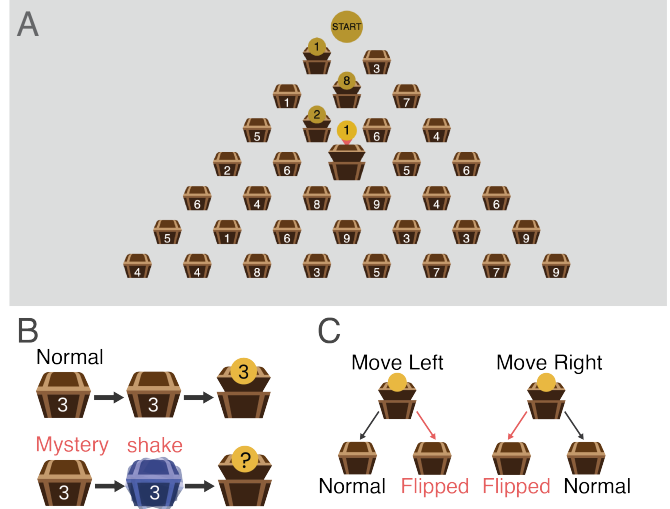


Figure 1: Experiment Setup. (A) Task display. The participant begins at the root node and moves down to the left or right at every step. (B) Experiment 1. In the reliability experiment, there is some probability q that any given treasure chest is a “mystery chest”, which gives you a random number of points between 1 and 9. (C) Experiment 2. In the transition noise experiment, there is some probability q that your move will be “flipped”.

If a participant answered any of the comprehension check questions incorrectly, they would be required to reread the instructions and practice games for that instruction module. Three attention checks were given to each participant; failure of two attention checks terminated the experiment.

Models

In this task, every move can take one of two possibilities: left or right. So any model of behavior must take as input the board state S , the stochasticity level q , and return a probability that the participant will move left (from which we can also derive the probability of moving right). That is,

$$P(a = L|S, q) = f(S, q)$$

We assume that this decision to move left (or right) is based on a participant’s utility of moving left or right.

$$P(a = L|S, q) = f(V(S, q, a = L), V(S, q, a = R))$$

Since players cannot move back up the tree, the utility of moving left must be some value function V over the left subtree S_L , excluding the rightmost path (the opposite is true for moving right).

$$P(a = L|S, q) = f(V(S_L, q), V(S_R, q))$$

Specifically, for f , we use a Logistic function over the difference between the value of moving left and the value of moving right. Here β is an inverse temperature parameter corresponding to the decision noise of the participant.

$$P(a = L|S, q) = \text{Logistic} [\beta(V(S_L, q) - V(S_R, q))]$$

It is unreasonable to assume that participants consider the entire tree of possible paths from the root. Instead, participants may consider only a subset of the nodes in the tree, performing a kind of “filtering” to select only the nodes they wish to consider in their plans. This begs the question - how do participants choose which nodes to pay attention to? One possibility is that participants filter the board in a depth-limited way, only considering paths of length d from the root. Another possibility is that people first look at the most rewarding values in the board and ignore the rest, for example filtering for only the highest k values before planning. We therefore consider two broad classes of planning models: *depth-limited* and *value-limited* models. We propose that this filtering process is related to the effort that participants are willing to exert in planning - given a particular level of effort, participants might choose to pay attention to more or fewer nodes on the board.

Depth-limited Models. In *depth-limited* models, the participants consider the boards in a top-down manner, considering only the subset of the board up to depth d from the root. We refer to this quantity d as a participant’s *planning depth*, and it reflects how deeply a participant is willing to plan. For example, a participant with planning depth 3 will only consider the next 3 possible moves from any given position. We are interested in seeing how this planning depth d might change across different stochasticity conditions q , so we incorporate it into our model of behavior as a flexible parameter that is fitted for each condition. In all cases, $0 \leq d_q \leq 7$, where a depth of 0 corresponds with random behavior $P(a = L|S, q) = 0.5$.

$$P(a = L|S, q) = \text{Logistic} [\beta(V(S_L, q, d_q) - V(S_R, q, d_q))]$$

People can change their plans in response to uncertainty in a couple ways. On one hand, people can *explicitly* factor in stochasticity when planning, for example by calculating the expected value of every path given some stochasticity level q . In the explicit case, the utility function V is sensitive to value q . On the other hand, people can *implicitly* factor stochasticity when planning *as if* there were no stochasticity involved, but modulate their effort - here represented by planning depth d_q - in response to stochasticity. In the implicit case, the utility function V is insensitive to the value q and responds only to d_q . With this in mind, we consider several possible models:

- **Depth Optimal.** Optimal behavior in this task is when $V(S, q, d_q)$ is the maximum expected sum of any path in S , up to depth d_q , taking into account uncertainty q . This model accounts for stochasticity *explicitly*.
- **Depth MaxPath.** In this model, $V(S, q, d_q)$ corresponds to the maximum sum of any path in S , up to depth d_q .

It is worth emphasizing the difference between this model, which only *implicitly* accounts for stochasticity, and the optimal model, which *explicitly* accounts for stochasticity. While the optimal model sets V to the *expected* value,

which accounts for the uncertainty in the environment, the MaxPath model sets V to the value of the path *as if* uncertainty were zero. This distinction is subtle, but it leads to different predictions about behavior. Careful readers may note that the best path predicted by both models is the same; while this is true, the *value* these models assign to the best paths are different, which affects the models’ predictions of the probability of a leftward or rightward move. This subtle distinction is easiest to appreciate when $q = 1$ (maximum stochasticity). In this case, the optimal model will behave randomly for all depths because the value of the left and right subtrees will be exactly equal, whereas the MaxPath model will still prefer the path that has a larger sum according to the numbers written on the treasure chests, even if the numbers have no correspondence with the underlying value.

- **Depth Max.** In this model, $V(S, q, d_q)$ is the maximum value of any node in S , up to depth d_q . This model accounts for stochasticity *implicitly*.
- **Depth Sum.** In this model, $V(S, q, d_q)$ is the sum of all the values of the nodes in S , up to depth d_q . This model accounts for stochasticity *implicitly*.

Value-limited Models. In *value-limited* models, the participants apply an attentional filter to the board where only the top k values of the board are retained and the rest of the values are set to zero. Instead of considering the board in a top-down manner up to depth d_q , the model takes the entire board S and filters out the top k_q values, setting the rest to zero. Here k_q is in a flexible parameter that, like d_q , is fitted for each condition. $V(S, q, k_q)$ is maximum sum of any path in S with only the top k values represented. In this case

$$P(a = L|S, q) = \text{Logistic} [\beta(V(S_L, q, k_q) - V(S_R, q, k_q))]$$

For each of the *depth-limited* models proposed, we have a *value-limited* counterpart where instead of considering the board up to depth d_q , all but the top k_q values are set to zero and the entire board is considered. We therefore have four value-limited models: Value Optimal, Value MaxPath, Value Max, and Value Sum.

In each model, we assume that all actions are conditionally independent. Also in each model, the free parameters are the inverse temperature parameter β and the planning effort in each condition (d_q for the depth-limited models or k_q for the value-limited models). We fitted these parameters to the data of individual participants using maximum-likelihood estimation. Since all models have the same number of parameters, no correction on the log likelihood is needed.

Results

Comparing the log-likelihood of the models, we found that the Depth MaxPath model significantly outperformed all other models (95% CI, 1M bootstrap samples, Fig 2A). We wanted to see whether this model replicated the trends seen

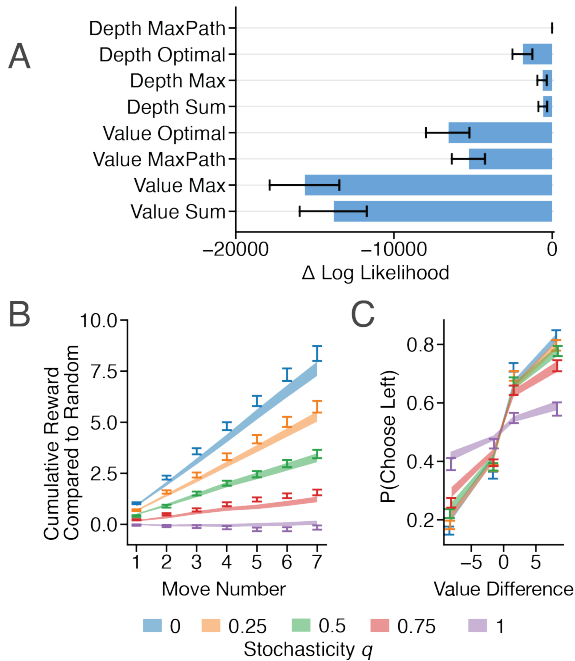


Figure 2: Experiment 1 (Reliability): Behavior and Model Comparison. In all figures, unless otherwise mentioned, error bars and shaded regions represent mean ± 1 SEM for data and model, respectively. (A) Bootstrapped differences between the log likelihood of various models and baseline (Depth MaxPath). Error bars represent 95% CIs. (B) Cumulative reward (compared to random) as a function of move number for different stochasticity levels q . (C) Proportion of left moves as a function of the difference in the value of the maximum path (with no limit on depth) between the left and right subtrees, binned by quantiles.

in human behavior. We found that participants leveraged the low-stochasticity condition to gain better rewards compared to random (Fig 2B). We also compared the participants' propensity to choose left and right as a function of the difference between the value of the maximum path of the left and right subtree, up to full depth (Fig 2C). This analysis provides a good illustration of why the Depth Optimal model fails to capture participant behavior. Recall that when $q = 1$, the optimal model predicts that participants will behave randomly, so we would expect to see a flat line at 0.5. What we see instead, however, is that participants still loosely prefer paths of higher apparent value even if this value is illusory. Our best-performing model, Depth MaxPath replicates the broad trends in these findings, across the five different conditions, suggesting that people perform depth-limited planning and implicitly account for stochasticity.

To study how planning effort changes with stochasticity, we fit a linear mixed-effects model of the form $y_{ij} = b_0 + b_1 x_{ij} + u_i + \epsilon_{ij}$, with $\epsilon_{ij} \sim N(0, \sigma_\epsilon^2)$ and $u_i \sim N(0, \sigma_u^2)$, where y is the planning depth, i is the participant index, and j is the index of the stochasticity condition. We found a negative effect of stochasticity level on planning depth ($p < 0.001$, Fig 3C). To validate this finding, we looked at how stochastic-

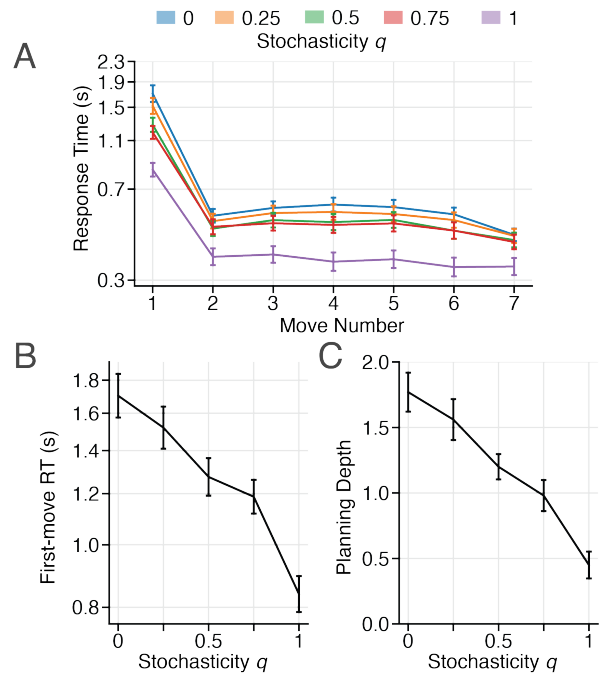


Figure 3: Experiment 1 (Reliability): Planning Depth and Response Time. (A) Response time as a function of move number, split by stochasticity level. (B) Response time of the first move as a function of stochasticity level, with logarithmic spacing on the y axis. (C) Model-fitted planning depth as a function of stochasticity level.

ity level impacted participants' response times. As response time is positively correlated with planning (Russek, Acosta-Kane, van Opheusden, Mattar, & Griffiths, 2022), we wanted to see whether stochasticity level would also decrease participants' response time. We found that participants spent most of their time thinking on the first move (Fig 3A). A linear mixed-effects model (as above, but where j is the puzzle index rather than the stochasticity condition) revealed an effect of the stochasticity level q on the log of the first move response time ($p < 0.001$, Fig 3B).

Experiment 2: Transition Noise

In experiment 1, we studied what happens when you have uncertainty in the value estimates of your states. Now that we understand how stochasticity influences planning depth when the information is unreliable, we want to show that this generalizes to other forms of stochasticity. In the real world, people may encounter uncertainty not only in their rewards but also in their actions. Here, we study uncertainty in the transition function, which we call "transition noise".

Imagine you are planning to fly home for the holidays, but you have chosen a budget airline which has no direct flights. Due to fuel constraints and budget cuts, there is some probability that any flight you take will be rerouted to a different destination, but somehow you must find your way home. Here, the uncertainty comes from the *transition noise* - i.e., you may intend to take action A but instead be forced to take

action B instead. Here we explore the relationship between transition noise and planning effort.

Task Design

We used the same task design as in Experiment 1, with a few modifications. First, all treasure chests were “normal” (no “mystery” chests) - that is, the number written on the front of the treasure chest reflected the value you would receive in points. Instead of varying the reliability of the treasure chests, we had five different stochasticity conditions corresponding with varying levels of transition noise. In a given condition, for a given action, there is some probability q that the action will be “flipped” and opposite action will be taken instead (Fig 1C). For example, if a participant presses the button to move left, they will go right instead, and vice versa. Participants do not know in advance which specific moves will be flipped. However, the participants do know q , the probability that a given move will be flipped, q , which is kept constant for each game. Across the five conditions we vary q : $q = 0\%$, 12.5% , 25% , 37.5% and 50% probability, where 0% corresponds to actions that are completely certain and 50% corresponds to actions that are completely random (i.e. maximum transition noise).

Participants

The recruitment process for participants was the same as in Experiment 1. We recruited 100 participants for this task on Prolific. There is no overlap of participants between Experiments 1 and 2.

Models

The behavior of all models remains the same as in Experiment 1, with the exception of the optimal models (Depth Optimal and Value Optimal). As before, all models define some value function V over the left and right subtrees, respectively. The decision to move left and right is based on a strict calculation of the value of the left and right subtrees, respectively:

$$P(a = L|S, q) = \text{Logistic}[\beta(V(S_L, q) - V(S_R, q))]$$

The optimal models, however, must account for the probability q that the action is flipped. Therefore, the value of moving left also needs to account for the probability q that the participant moves *right* instead, and vice versa.

$$\begin{aligned} P(a = L|S, q) &= \text{Logistic}[\beta((1 - q)V(S_L, q) + qV(S_R, q) - \\ &\quad (1 - q)V(S_R, q) - qV(S_L, q))] \\ &= \text{Logistic}[\beta(1 - 2q)(V(S_L, q) - V(S_R, q))] \end{aligned}$$

As before, the optimal model represents the maximum possible reward you can get with any sequence of actions from the root, taking into account transition noise. The value function V can be solved with dynamic programming.

Results

We found that the Depth MaxPath model significantly outperforms all other models in log-likelihood (Fig 4A). As be-

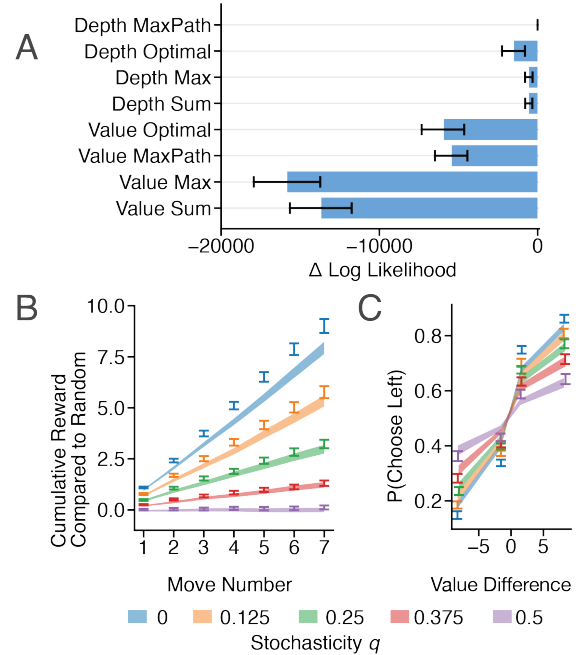


Figure 4: Experiment 2 (Transition noise): Behavior and Model Comparison. (A) Bootstrapped difference between the log likelihood of the model and baseline (Depth MaxPath). Error bars are 95 CI. (B) Cumulative reward (compared to random) as a function of move number for different stochasticity levels q . (C). Proportion of leftward choices as a function of the difference in the value of the maximum path (with no limit on depth) between the left and right subtrees, binned by quantiles.

fore, we see that participants effectively leveraged the low-stochasticity condition to gain better rewards (Fig 4B), and that the difference in the value between the left and right subtrees drove the proportion of choosing leftward moves, trends that are also well captured by the MaxPath model (Fig 4C). Here again, we can see how the optimal models fail to account for participant behavior. When uncertainty is maximal, $q = 0.5$, the optimal model predicts a flat line for the probability of moving left at 0.5 when $q = 0.5$. Instead, we see that participants still attempt to adhere to higher-valued paths.

A linear mixed-effects model revealed a negative effect of transition noise level on planning depth ($p < 0.001$, Fig 5C) and a negative effect of transition noise level on log first move response times ($p < 0.001$, Fig 5B). This indicates that participants adapted to stochasticity by reducing their planning depth.

Discussion

People often plan for the future under uncertainty. It is therefore a natural question to ask how this uncertainty shapes how people plan, if at all. Here we designed a multi-step planning task and investigated how two different forms of stochasticity, reliability and transition noise, affected participants’ planning effort. In Experiment 1, we varied the reliability by changing the probability that a treasure chest was a “mystery” chest,

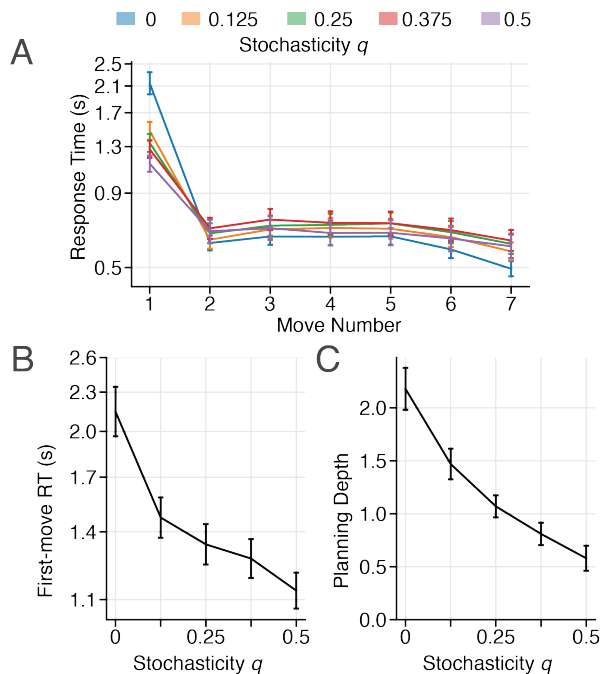


Figure 5: Experiment 2 (Transition Noise): Planning Depth and Response Time Analysis. (A) Response time as a function of move number, split by stochasticity level. (B) Response time of the first move as a function of stochasticity level, with logarithmic spacing on the y axis. (C) Model-fitted planning depth as a function of stochasticity level.

whereas in Experiment 2, stochasticity was modulated by the probability that a move would be flipped.

In both tasks, participants attained greater cumulative rewards in lower stochasticity conditions, demonstrating that they were able to successfully learn the task and adapt to different stochasticity levels. Their decision was also driven by the differences in the relative values of each subtree, demonstrating that they had learned the structure of the task.

For both experiments, the Depth MaxPath model outperformed all other models at explaining behavior. The fact that the depth-limited model outperformed all other value-limited models suggests that people largely plan from the top down in this task, rather than filtering for the highest values in the tree. Notably, the algorithm that they chose was not optimal; the MaxPath model does not take into account the stochasticity when calculating the value of each path. This suggests that people did not explicitly consider stochasticity in their approach, but instead compensated for stochasticity by modulating the effort they applied under different conditions.

Our key finding was that as stochasticity increased, planning depth monotonically decreased. This effect was statistically significant in both experiments, demonstrating that people do in fact consider stochasticity when planning, albeit implicitly. We validated our model predictions by analyzing the first-move response times, which showed a similar negative relationship between response time and stochasticity level. For an intuitive explanation for this phenomenon, we

can turn to resource rationality (Lieder & Griffiths, 2020), which argues that participants take into account the computational costs of cognition when making decisions. Recent work suggests that people plan in a resource-rational way (Callaway et al., 2018), adapting how deeply they plan by comparing the cognitive costs of planning and the potential benefits (Sezener, Dezfouli, & Keramati, 2019). Because the expected reward diminishes with stochasticity while the cost of computation remains the same, it makes sense that participants would try to conserve computational cost on higher stochasticity conditions. We have identified one mechanism by which participants can conserve computational load: by reducing the depth at which they are willing to plan. In this specific task, participants used a suboptimal but easily computable algorithm (MaxPath) to plan but used resource-rational principles to scale their cognitive computation (depth-limiting). In other words, people used effort to *implicitly* account for stochasticity rather than *explicitly* formulating the optimal solution. This finding raises the possibility that people might use cognitive resource allocation (e.g. effort, attention, memory) as a surrogate for optimal or near-optimal behavior.

While we have identified a connection between our work and the framework of resource rationality, it is worth noting that our models are not explicitly resource-rational. A resource-rational model would quantify the computational costs associated in planning under different stochasticity levels and *derive* the dependence of planning depth on stochasticity. Additionally, there are other forms of stochastic environments that we have not considered here, for example a “volatile” environment in which the value of each treasure chest is reliable but has a probability of changing after each decision. Another possible extension would be to investigate the costs that participants are willing to incur to reduce uncertainty. Finally, in multi-player games, people may view their opponents or collaborators as additional sources of uncertainty and may adapt their plans accordingly. Taken together, our work could provide the foundation for a deeper understanding of how people allocate cognitive resources to adapt to uncertainty in planning.

References

- Behrens, T. E. J., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, *10*(9), 1214–1221. doi: 10.1038/nn1954
- Callaway, F., Lieder, F., Das, P., Gul, S., Krueger, P. M., & Griffiths, T. L. (2018). A resource-rational analysis of human planning. In *Proceedings of the 40th Annual Meeting of the Cognitive Science Society, CogSci 2018* (pp. 178–183). The Cognitive Science Society.
- Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P., & Dolan, R. J. (2011). Model-based influences on humans' choices and striatal prediction errors. *Neuron*, *69*(6), 1204–1215.
- Lieder, F., & Griffiths, T. L. (2020). Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, *43*, e1. doi: 10.1017/S0140525X1900061X
- Mitra, A., Srivastava, N., & Srinivasan, N. (2023). Unpredictability shortens planning horizons. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *45*(45).
- Owen, A. M., Downes, J. J., Sahakian, B. J., Polkey, C. E., & Robbins, T. W. (1990). Planning and spatial working memory following frontal lobe lesions in man. *Neuropsychologia*, *28*(10), 1021–1034. doi: 10.1016/0028-3932(90)90137-D
- Russek, E., Acosta-Kane, D., van Opheusden, B., Mattar, M. G., & Griffiths, T. (2022). Time spent thinking in online chess reflects the value of computation. doi: 10.31234/osf.io/8j9zx
- Sezener, C. E., Dezfouli, A., & Keramati, M. (2019). Optimizing the depth and the direction of prospective planning using information values. *PLOS Computational Biology*, *15*(3), e1006827. doi: 10.1371/journal.pcbi.1006827
- Shallice, T. (1982). Specific Impairments of Planning. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, *298*, 199–209. doi: 10.1098/rstb.1982.0082
- Simon, H. A. (1975). The functional equivalence of problem solving skills. *Cognitive Psychology*, *7*(2), 268–288. doi: 10.1016/0010-0285(75)90012-2
- Snider, J., Lee, D., Poizner, H., & Gepshtein, S. (2015). Prospective Optimization with Limited Resources. *PLOS Computational Biology*, *11*(9), e1004501. doi: 10.1371/journal.pcbi.1004501
- Wiener, J. M., Lafon, M., & Berthoz, A. (2008). Path planning under spatial uncertainty. *Memory & Cognition*, *36*(3), 495–504. doi: 10.3758/MC.36.3.495