Decision-making times reveal that people’s thinking plans adapt to the problem they are trying to solve

Joan Danielle K. Ongchoco, Joshua Knobe & Julian Jara-Ettinger

Department of Psychology, Yale University

Abstract

Much of our thinking focuses on deciding what to do in situations where the space of possible options is often too large to evaluate exhaustively. Previous work has found that people do this by learning the general value of different behaviors, and prioritizing thinking about high-value options in new situations. Is this *good action bias* always the best strategy, or can thinking about low-value options sometimes become more beneficial? Can people adapt their thinking accordingly based on the situation? And how do we know what to think about in novel events? Here, we developed a block-puzzle paradigm that enabled us to measure people’s thinking plans and compare them to a computational model of rational thought. Our results suggested that people can quickly estimate the apparent value of different options and use this to guide what to think about—in particular, the potentially high-value options (Experiment 1). But when the situation made it more beneficial to focus on potentially negative options, participants—like our model—focused on low-value options, leaving them less time to think about high-reward choices (Experiment 2). Our results suggest that thinking plans are flexible: what we think about depends on the structure of the problems we are trying to solve.

Key words: Computational modeling, Thinking

1. Introduction

We often have to make decisions involving a wide array of options: actions we can take, things we can say, or people we can interact with. Intuitively, much of our thinking involves considering different options in anticipation of a decision, as when we think about what we want to eat for dinner, or what we want to say before a meeting. But we don’t have the time and computational resources to consider every option before we have to make a choice (see Lieder & Griffiths 2020, 2017; Vul et al. 2014). So we face a complex selection problem:

9 Given that we can only think a few thoughts at a time, what should we think
10 about?

11 A growing body of computational and empirical work suggests that people
12 solve this problem through a ‘good-action bias’ (Bear et al. 2020; Icard et al.
13 2017; Mattar & Daw 2018): Over time, people learn the values associated with
14 different options or behaviors in a broad class of situations. Then, when people
15 face a specific situation and can only consider a small number of options about
16 how to behave within it, they tend to consider the options that they have
17 determined to generally have high value in that broader class of situations.
18 This process leads to more accurate representations of the high-value options
19 (e.g. Icard et al. 2018; Gelly & Silver 2011).

20 Despite the power and usefulness of a good-action bias, this strategy is only
21 useful in situations where people have readily available value representations
22 associated with different possible options. Yet, many important decisions often
23 come in the context of novel situations, where relevant past experiences can be
24 scarce, or carry little value information for the context that we’re in. In these
25 situations, what do people tend to think about?

26 *1.1. What do we think about?*

27 There are least two possibilities for what happens when people face entirely
28 new situations. A first possibility is that the good-action bias reflects a thinking
29 strategy that people use across the board. When faced with a new situation,
30 people might initially represent all options as having roughly equal value (with
31 high uncertainty). In such situations, people might show no systematicity in
32 what they choose to think about—thinking about both high and low value op-
33 tions, as the values are unknown to them—with the good-action bias becoming
34 visible once people have gotten the opportunity to begin refining their value es-
35 timates (Callaway et al. [2021b]). Alternatively, however, people might quickly

36 build an initial estimate of value based on superficial cues of the environment.
37 This would allow people to prioritize thinking about potentially valuable options
38 from the start. As we will see, whether these generalizations of the good-action
39 bias are useful depends on the structure of the decision problem, as there are
40 certain structures in which they may not be appropriate.

41 In contrast, another possibility is that the good-action bias is limited in
42 scope, and people can generate and use different thinking strategies based on
43 the context that they’re in. In other words, in the same way we can create
44 ad-hoc action plans in a novel physical task (as when we figure out which path
45 to take based on the layout of a new room), people may also be able create
46 ad-hoc “thinking plans” in a novel thinking task. Thus, people might be able
47 to execute thinking plans that are not necessarily aligned with the good-action
48 bias, depending on the problem that they’re attempting to solve. This said, gen-
49 erating novel ad-hoc thinking plans requires time and computational resources
50 (Griffiths et al. [2015]). If we were able to do this quickly, this would suggest
51 a sophisticated mechanism that generates thinking plans while working around
52 computational intractability.

53 The present study aims to explore these possibilities. How do people know
54 what to think about in novel situations? Do their strategies reflect a generalized
55 form of the good-action bias? Or can people flexibly adapt to the right strategy
56 based on the structure of the problem at hand?

57 1.2. *The present studies*

58 To answer these questions, we developed an experimental paradigm that
59 allows us to determine what people have (and have not) thought about in a novel
60 situation under time pressure. Participants saw novel incomplete block-puzzles
61 and an array of six different puzzle-piece options (Figure 1), some of which had
62 potentially high values, while others had potentially low values. Puzzle pieces

Figure 1: The two-part design of the experiments. Participants first went through a (a) thinking phase where they were allowed to freely evaluate how well each of the six pieces fit into the base (bottom piece) (b) Decision phase where participants were asked to choose between two pieces.

varied in their surface features (e.g. their size), which provided information about their apparent value. Participants could then solve the puzzle by mentally rotating the various puzzle pieces to see which pieces fit and which did not, which in turn revealed their value. Because mental rotation is a relatively slow cognitive operation ([Cooper, 1975; Shepard & Metzler, 1971]), this approach enabled us to probe whether participants had previously thought about a puzzle piece by measuring their reaction time; a slow reaction would suggest that the participant is evaluating the puzzle piece (i.e., performing a mental rotation and computing its value) only after being asked about it, while a fast reaction time would suggest that the participant had already mentally rotated the piece and computed its value before being asked.

Participants first completed a *thinking phase*, where they could freely think about different puzzle pieces (via mental rotation; Figure 1a), followed by a *decision phase*, where participants were explicitly prompted to select the best piece from a subset of pieces (Figure 1b). Therefore, participants' response

78 times could reveal whether they had already considered the value of the pieces
79 presented to them in the decision phase: participants should be faster when
80 deciding between two pieces that they had already evaluated, and slower when
81 they hadn't thought about some (or all) of the presented pieces.

82 This paradigm enabled us to test what people think about in novel situations,
83 as a function of the context that they're in. In Experiment 1, we first considered
84 a situation where observable cues contained information about the apparent
85 value of different pieces, and we tested whether these cues influenced what
86 people chose to think about first. Participants were told that only puzzle pieces
87 that fit into the main structure (shown in yellow in Figure 1a) would give them
88 points, and that the object's final size, after attaching the puzzle piece, would
89 determine the number of points that they obtain. Therefore, participants could
90 estimate the apparent value of a piece based on its size, but would need to think
91 about (and mentally rotate) it to discover its exact value.

92 Beyond what thinking strategy people will use in novel situations, we were
93 also interested in whether people flexibly adapt their thinking plans depending
94 on the situation itself. To see if people don't just default to potentially high-
95 value options, we also needed to identify situations where a better strategy
96 would be to think of potentially low-value options (therefore acting against a
97 good-action bias). To explore this possibility, we formalized our block-puzzle
98 paradigm in a computational framework and varied different task parameters to
99 explore the space of decision problems. This allowed us to identify contexts in
100 which the best strategy would be to think about potentially low-value options
101 first, which we then test in Experiment 2.

102 Altogether, these experiments allowed us to probe what people would think
103 in a range of decision-problems, varying different dimensions: a case where there
104 was an equal number of high and low-value options versus a case where there

were more low-value options; a case where there was a potentially high-value option (needle in the stack) versus a case where there was a potentially really low-value option (snake in the stack); and a case where the critical low-value option (the ‘snake’) could be more or less difficult to find.

2. Computational Framework

To formalize the problem of which options would be best to think about, we used partially-observable Markov decision processes (POMDP; Cassandra 1998; Sutton et al. 1998). POMDPs were originally developed in robotics and AI to model action-planning in complex spatial environments under partial information, but more recent work has shown that this framework can be used not only to model decisions between different physical actions but also decisions between different thought processes (see Callaway et al. 2021a; Chen et al. 2021; Griffiths et al. 2019; Lieder & Griffiths 2020). Thus, here we adapted the framework to implement a space of thinking actions rather than physical actions. Before introducing our model, we will begin by briefly introducing POMDPs in the context in which they are classically used.

2.1. Classical uses of POMDPs

To illustrate the logic of POMDPs, consider a simple situation where an agent’s goal is to take an object in a house. To achieve this, the agent must first find the object and then retrieve it. To formalize this problem, POMDPs define a state space S consisting of all possible physical states of the world. In a case like this one, the state space would include the combination of any position where the agent and the object might be at any given time point. Given this state space, the agent’s goals can be represented as a reward function R that assigns a numerical reward to combinations of states and actions. In this

130 example, the goal can be encoded as a reward function that returns a high
131 positive value in states where the agent is holding the object.

132 To obtain these rewards, the agent can take sequential actions (from a set of
133 actions A) that change the state of the world. The relationship between actions
134 and states is captured by the transition function T , where $T(s, a, s')$ represents
135 the probability that the world will change from state s to s' when the agent
136 takes action a . For instance, an agent taking the action ‘walk north’ in a state
137 should assign a high probability to the state where the agent is now one spot
138 north of where they used to be.

139 At their core, POMDPs assume that agents can have partial or incomplete
140 knowledge about the world. For instance, the agent may know their position
141 in space but not know where the object is located. To achieve this, POMDPs
142 introduce belief representations, expressed as probability distributions over the
143 state space. To model how the agent’s beliefs change as they move in space,
144 POMDPs introduce an observation function O which determines what informa-
145 tion is made available to the agent in different states, where $O(i, s, a)$ is the
146 probability that the agent receives information i when taking action a in state
147 s . For instance, a simple observation function might encode that the agent can
148 see whether the object is present or absent in any room as soon as she enters
149 it. Formally, this is achieved by defining a space of observations, and specify-
150 ing which observations are associated with each state through an observation
151 function (which can include probabilistic components).

152 Given the six-tuple defined above—a state space, an action space, an ob-
153 servation space, a reward function, a transition function, and an observation
154 function—it is possible to compute the series of actions that, given an agent’s
155 knowledge, maximize the long-term rewards that the agent obtains (requiring
156 one additional parameter, λ , that specifies how rewards are discounted over

time). Computing the exact solution to a POMDP is computationally demanding and often intractable in practice, particularly in problems with large state spaces. Nonetheless, research in the past two decades has led to the development of multiple algorithms that provide approximate solutions to POMDPs (such as by not computing the actions that would be associated with implausible belief states that the agent could have), making them a useful practical framework for determining rational action under imperfect information (Hsu et al. 2008; Ng & Jordan 2000; Kurniawati et al. 2008).

2.2. Modeling thinking through POMDPs

Although POMDPs are most typically used to model choices between physical actions, they can also be used to model choices between different *thoughts*. For simplicity, we explain this model structure in the context of Experiment 1a (a more comprehensive presentation is available in the Supplemental Materials). Here, participants were presented with a puzzle like the one shown in Figure 1a, and they learned that pieces that do not fit into the puzzle have no value, and pieces that fit into the puzzle have larger values whenever they increase the overall number of blocks in the structure. After being given time to think about whichever pieces they like (i.e., rotating them mentally to see if they fit in the puzzle during the thinking phase), participants were asked to quickly determine which of two puzzle pieces had higher value (decision phase). Although the full problem is ultimately encoded in a single POMDP, we first explain the components at use during the thinking phase, and then turn to the ones used in the decision phase.

In this context, the apparent value of a puzzle piece can be determined visually: the larger the piece, the more likely it is to be valuable. The true value of a piece, however, can only be revealed by thinking about it (i.e., mentally rotating it to test if it fits in the puzzle or not). Thus, our model represents each

184 piece in terms of its *apparent* value, determined by its size, and its *true* value,
 185 which equals the apparent value when the piece fits into the puzzle, and 0 when
 186 it does not. Given the two hypotheses about each piece (the true value either
 187 equals the apparent value or 0), we defined the state space as every possible
 188 setting over which pieces' true value matches the apparent value and which do
 189 not (i.e., in the case with six pieces, the state space consists of $2^6 = 64$ states).

190 To model the thinking phase, we gave our model the ability to execute
 191 *thinking* actions, which revealed the true value of whichever piece the agent
 192 chose to think about (through an observation). While the thinking actions
 193 do not have any causal impact on the state space, there was always a small
 194 probability that, at any given point, the state space may switch to a decision
 195 phase (capturing the idea that the participant knew that at any point they
 196 might get asked to choose a piece among a subset). To model this component,
 197 we included an additional set of states, where each state encoded a forced choice
 198 between two possible pieces selected from a uniform distribution (as in Fig 1a).
 199 To select a piece, the agent could take a 'selection' action, obtaining a reward
 200 depending on which piece they picked.

201 These specifications enabled us to use the POMDP framework to compute
 202 what sequence of thinking actions were best suited to the decision problem, with
 203 the goal of maximizing the agent's expected reward in the decision phase.

204 **3. Experiment 1**

205 In Experiment 1, we first sought to test what people think about in a novel
 206 situation where the apparent value of different choices can be estimated based
 207 on superficial observable cues. Participants saw a display like the one shown
 208 in Figure 1a. The value of each piece was given by the final size of completed
 209 puzzle, after the piece was attached (and 0 value for pieces that did not fit the

base puzzle). Thus, each piece’s size gave a superficial cue about its apparent value, but participants needed to mentally rotate each piece to test if it indeed fit the puzzle. We first confirmed that, in a context like this one, our model predicts that people should prioritize thinking about potentially high-value pieces, and we then tested this effect empirically.

3.1. Model Simulations and Results

To determine what people should think about in this situation, we implemented this puzzle in our computational framework (section 2.2). Because people can sometimes make reasoning errors, our model included a small probability that a mental rotation would lead to an incorrect conclusion (i.e., believing that a piece fits when it does not, or believing that a piece does not fit when it does). To achieve this, we included a certainty parameter, which we varied from 95% (i.e., a 5% chance of making a reasoning error) to 100% (i.e., full confidence of no reasoning error) in steps of 1%, each implemented on a separate POMDP. In these models, the probability of switching from the thinking phase to the decision phase was set to 30% and we used a future discount parameter of 0.95. These parameters were all set prior to model evaluation.

After training each POMDP (i.e., computing the optimal policies), we obtained an expected thinking plan by running 100 simulations under each certainty parameter with a random set of puzzle pieces (set to always have two large pieces, two medium pieces, and two small pieces to match the experiment). To ensure that all simulations revealed the full thinking plan, we modified the simulations dynamics to ensure that the model would not switch to the decision phase until after the model had the chance to think of all pieces. That is, the model’s training reflected the belief that the agent might be prompted to make a decision before having had the opportunity to think about all six pieces, but the simulations were modified to stay in the thinking phase long enough for us

237 to observe the model’s full thinking pattern. We then averaged the performance
238 of all simulations over all certainty values to obtain an expected thinking plan.

239 Figure 2a shows the results from this simulation. As this simulation shows,
240 our model always uses the first two time-steps to think about the two pieces
241 that have the highest probability of being valuable. Afterwards, the model shifts
242 to thinking about intermediate-value pieces. Note that our model continues to
243 assign some probability to considering the highest-value pieces in steps 3 and
244 4. This reflects the performance of the models that believe reasoning errors are
245 likely. In these cases, our model believes it is better to double check potentially
246 high-value pieces to confirm their true value. Finally, the model only begins
247 considering the lowest-value pieces at the very end of the task.

248 *3.2. Experiment 1a*

249 Experiment 1a tested a situation where the value of pieces was correlated
250 with size: larger pieces that fit into the puzzle were more valuable than smaller
251 pieces that fit into the puzzle. However, pieces that did not fit into the puzzle
252 had no value, regardless of their size. After the thinking phase—where
253 participants could freely focus on different pieces to check if they’d fit into the
254 puzzle—participants were asked to select the best option among either two large
255 pieces or two small pieces. We predicted that, if participants prioritize thinking
256 about potentially high-value options, their response time should be significantly
257 lower when asked to select which is better of two large pieces relative to when
258 asked to select which is better of two small pieces.

259 *3.2.1. Method*

260 All methods and analyses were pre-registered ([http://aspredicted.org/
261 blind.php?x=rd2bd2](http://aspredicted.org/blind.php?x=rd2bd2)). Data and code for all experiments reported here are
262 available on:
263 https://osf.io/n8em7/?view_only=9f9bb51af6b24fdaa1eebc2a041ce910.

264 Participants. Sixty participants from the Yale and New Haven community
265 participated. Three additional participants were recruited but excluded from
266 the study based on pre-registered criteria (see Results). The sample size was
267 determined based on a power analysis run on pilot data (see Supplemental
268 Materials for details).

269 Apparatus. Stimuli were presented using custom software written in Python
270 with the PsychoPy libraries (Peirce et al. 2019) and were displayed on a mon-
271 itor with a 60Hz refresh rate. Participants completed the study on a 13-inch
272 MacBook Air with a 1440 x 900 resolution.

273 Stimuli. Each trial consisted of a yellow puzzle base and six blue puzzle
274 pieces (see Figure 1a for an example). Each participant was presented with a
275 randomly generated puzzle base, which always had three to four blocks missing
276 at the top. Each of the six puzzle pieces options consisted of the three/four
277 blocks that were missing, along with additional blocks that determined the
278 value of the piece. For example, a piece with a value of 7 points is one that
279 locks into the puzzle base after being rotated and had 7 additional blocks that
280 go beyond the puzzle’s 5x4 rectangular shape (e.g., the upper leftmost piece in
281 Figure 1a). Critically, some puzzle pieces did not have the correct shape to lock
282 into the puzzle and therefore had value 0 (e.g., the lower middle piece in Figure
283 1a, which has an apparent value of 7, but its true value is 0 because it does
284 not fit into the puzzle). The six puzzle pieces were randomly generated but
285 always consisted of three pairs of apparent values: two potentially high-value
286 options (apparent value: 7), two potentially medium-value options (apparent
287 value: 5), and two potentially low-value options (apparent value: 3). In each
288 pair, one piece would always fit (true value equals its apparent value), and the
289 other would not (true value equals 0).

290 Procedure and Design. Participants first read a brief set of task instructions
291 where they learned the logic of the task and they were told that their goal was
292 to earn as many points as possible. A ‘Total Points’ counter was visible on the
293 top-left of the screen throughout the entire experiment. Participants were then
294 shown a sample block-puzzle, and were asked questions about these different
295 options to test their understanding of the task and the point system. Partici-
296 pants obtained points for correct answers, which we used as an exclusion criteria
297 for people who did not understand the task. To ensure that participants would
298 not just look for puzzle pieces whose bottom part resembled the structure of
299 the missing section of the puzzle base, they were told that the pieces can only
300 be rotated but never simply flipped. After these instructions, participants were
301 shown a new block-puzzle. At the start of the thinking phase, participants were
302 told that six pieces would now appear above the puzzle, in two rows of 3 pieces
303 each (as in Figure 1a). Participants were told that they did not have to do
304 anything but just look and study the pieces. Once the puzzle pieces appeared,
305 participants were given five seconds to study them, with a countdown timer
306 shown above the puzzle. When the countdown timer reached 0, the task au-
307 tomatically switched to the decision phase. Here, two of the six pieces turned
308 green (as in Figure 1b) and participants were asked “Which piece is better?”
309 Critically, half of our participants were tasked with identifying which is better
310 of the two potentially high-value options, and the other half were tasked with
311 identifying which is better of the two potentially low-value options. To incen-
312 tivize participants to really think about the options during the thinking phase,
313 they were also told that they would get bonus points for responding quickly.

314 3.2.2. Results

315 Three participants were excluded because their mean performance in the
316 comprehension questions was 2 standard deviations below the grand popula-

Figure 2: (a) Model predictions about what pieces to think about as a function of time step in Experiment 1. The model could think about the potentially high-value pieces (green), the medium-value pieces (purple), or the low-value pieces (red). Our rational model suggests that the best strategy is to begin by thinking about the high-value pieces, followed by medium-value pieces, and then finally the low-value pieces. (b) Results from Experiment 1a. The bars depict mean response times in each condition. Error bars reflect 95% confidence intervals. (c) Results from Experiment 1b. (d) Results from Experiment 1c.

tion mean. These participants were replaced, until a total of 60 participants was reached. Response accuracy and response times for the single trial were recorded for each participant. Only response times where participants responded correctly were included in the analysis. Participants who chose between the potentially high-value options responded faster in the decision phase ($M=2.56s$, $SD=1.77s$) than participants who chose between the potentially low-value options ($M=5.20s$, $SD=3.82s$), as depicted in Figure 2b ($t(43.37) = 3.65$, $p < .001$, $d = 1.06$ over the logarithm of the reaction time). Including the incorrect answers did not yield any different results, $t(55.41) = 3.51$, $p < .001$, $d = 0.91$. There was no significant difference between the percentage of people who responded accurately when choosing between potentially high-value options vs. potentially low-value options (86.67% vs. 66.67%, Fisher's exact: $p = .125$).

329 3.2.3. Discussion

330 The results from Experiment 1a revealed that people asked to identify which
331 is better of two potentially high-value pieces were about twice as fast as peo-
332 ple asked to identify which is better of two potentially high-value pieces. This
333 suggests that participants prioritized thinking about potentially high-value op-
334 tions during the free thinking phase, producing the response time benefit in the
335 decision phase. Thus, in novel situations, people may adopt a strategy of think-
336 ing of the options that appear likely to have high values. This aligns with the
337 sequence of thinking actions from our computational framework, and suggests
338 that people may be adopting the best strategy (as determined by our model),
339 given the decision problem.

340 3.3. Experiment 1b

341 In Experiment 1a, potentially high-value options always had more blocks,
342 while potentially low-values options always had fewer blocks. Might people
343 simply have been attracted to larger pieces, independent of their value? To
344 ensure that our results were not just a matter of the brute number of blocks,
345 Experiment 1b flipped the size-value relationship: smaller puzzle pieces were
346 now more valuable than larger puzzle pieces.

347 3.3.1. Method.

348 This experiment and stimuli was identical to Experiment 1a, except as noted.
349 Seventy new participants from the Yale and New Haven community participated.
350 Two additional participants were recruited but excluded in the study based on
351 pre-registered criteria (see Results). The sample size was determined before data
352 collection began based on a power analysis run on pilot data (see Supplemental
353 Materials for details).

354 Participants first read a brief set of task instructions that explained the logic

of the task. In contrast to Experiment 1a, participants learned that the number of additional blocks in a piece that fits now reflected the number of points that would be deducted. For instance, if a piece fit into the puzzle base and had 7 additional blocks, then 7 points would be deducted. If the piece did not fit, then 10 points would be deducted. In this case, the potentially high-value pieces were the ones with fewer blocks, and the potentially low-value pieces were the ones with more blocks. Participants began with a score of 50 points and the thinking and decision phase proceeded in the same way as Experiment 1a. All methods and analyses were pre-registered (<http://aspredicted.org/blind.php?x=bw5n59>).

3.3.2. Results and Discussion

Two participants were excluded because their mean performance in the comprehension questions was 2 standard deviations below the grand population mean. These participants were replaced, until a total of 70 participants was reached. Response accuracy and response times for the single trial were recorded for each participants. Only response times where participants responded correctly were included in the analysis. Participants who chose between the potentially high-value options responded faster in the decision phase ($M=2.37s$, $SD=1.73s$) than participants who chose between the potentially low-value options ($M=3.99s$, $SD=2.23s$), as shown in in Figure 2c ($t(40.37) = 3.88$, $p < .001$, $d = 1.08$ over the logarithm of the reaction time). Including the incorrect answers did not yield any different results ($t(65.24) = 4.14$, $p < .001$, $d = 0.99$). While we found a marginal difference between the percentage of people who responded accurately when choosing between potentially high-value options vs. potentially low-value options (71.43% vs. 91.43%, Fisher's exact: $p=.062$), accuracy was nonetheless high in both conditions. And this marginally higher accuracy when choosing between low-value options may reflect participants' determination to not make a wrong choice. Thus, our results suggest that people

are not just attracted to thinking about visually larger puzzle pieces, and they instead thought about particular puzzle pieces based on the pieces' underlying values.

3.4. Experiment 1c

So far, the results from Experiments 1a-1b suggest that people prioritize thinking about potentially high-value options, even in novel situations. This strategy produced a response time benefit when people had to decide between two potentially high-value options, relative to when they had to decide between two potentially low-value options. What is the scope of this advantage? One possibility is that the reaction time benefit is specific to puzzle pieces that participants mentally rotated during the thinking phase (recalling their value during the decision phase). Alternatively, however, it is possible that participants not only learned the specific value of pieces that they thought about, but also learned a general strategy for how to quickly think about good puzzle pieces. If this is the case, then participants should show a reaction time advantage over good pieces, even when these pieces weren't available during the thinking phase.

To explore this possibility, Experiment 1c replicated Experiment 1a, with the difference that, in the decision-making phase, participants were shown puzzle pieces that were not in the thinking phase. If participants learned the specific value of pieces that they had mentally rotated, the reaction time benefit should disappear. However, if participants learned general strategies for how to quickly evaluate the true value of good puzzle pieces, then we should see a reaction time difference, even for new puzzle pieces.

3.4.1. Method

This experiment was identical to Experiment 1a, except as noted. Sixty new participants from the Yale and New Haven community participated. One

408 additional participant was recruited but excluded from the study based on pre-
409 registered criteria (see Results). The sample size was determined before data
410 collection began based on a power analysis run on pilot data (see Supplemental
411 Materials for details). During the decision phase, a new pair of high-value
412 or low-value options were generated and presented to the participants. All
413 methods and analyses were pre-registered (<http://aspredicted.org/blind.php?x=bw5n59>).
414

415 3.4.2. Results and Discussion

416 One participant was excluded because their mean performance in the com-
417 prehension questions was 2 standard deviations below the grand population
418 mean. This participant was replaced, until a total of 60 participants was reached.
419 Response accuracy and response times for the single trial were recorded for each
420 participant. Only response times where participants responded correctly were
421 included in the analysis. This time, participants who chose between potentially
422 high-value options were in fact slower ($M=4.54s$, $SD=2.51s$) than participants
423 who chose between potentially low-value options ($M=2.52s$, $SD=1.12s$), as de-
424 picted in Figure 2d ($t(41.76) = 4.10$, $p < .001$, $d = 1.21$ over the logarithmic
425 transformations of the distributions). Including the incorrect answers again did
426 not yield any different results ($t(55.15) = 2.77$, $p = .007$, $d = 0.72$). There
427 was no significant difference between the percentage of people who responded
428 accurately when choosing between potentially high-value options vs. potentially
429 low-value options (76.67% vs. 76.67%, Fisher's exact: $p=1$). To compare these
430 results with those of Experiment 1a, we ran a 2 choice types (old vs. novel) x 2
431 choice values (potentially high-value options vs. potentially low-value options)
432 ANOVA. There was no main effect of choice type, $F(1, 88) = 0.12$, $p = .728$,
433 $\eta^2 = .002$, or of choice value, $F(1, 88) = 0.33$, $p = .570$, $\eta^2 = .004$. Crucially,
434 there was a significant interaction, $F(1, 88) = 20.99$, $p < .001$, $\eta^2 = .193$. These

435 results suggest that participants in Experiments 1a-1b had a reaction time ben-
436 efit for potentially high-value pieces because they learned their specific value
437 during the thinking phase, and not because they developed a general strategy
438 for thinking about those pieces more quickly during the decision phase.

439 **4. Experiment 2**

440 Experiment 1 showed that, in novel situations, people can preferentially
441 think about options that are likely to be of high-value. This behavior was con-
442 sistent with our computational model, which predicted that a rational thinking
443 strategy should prioritize thinking about options with a higher expected value.
444 At the same time, this experiment only considered situations where thinking
445 about the good was the best strategy. We therefore do not know if people have
446 a good-action bias in all situations, or if they can revise this strategy when
447 necessary.

448 To test this possibility, we used our computational framework to search for
449 problems where the best strategy would be to think about low-value options first,
450 and tested if participants can adjust their thinking strategy accordingly. As we
451 show below, this search led to a task structure that we call “snake in the stack.”
452 This task is structurally similar to the tasks that we used in Experiment 1, with
453 the difference that we introduced a “snake”: a piece that comes with a relatively
454 larger cost for participants. In this type of situation, our model predicts that
455 the best strategy is to first find the snake (i.e., think about potentially low-value
456 pieces) and then switch to thinking about potentially good options.

457 The strategy derived by our model implies that the amount of time used
458 to think about good pieces depends on how quickly one finds the bad piece.
459 For instance, if a participant was lucky and happened to identify the bad piece
460 on their first try, this would give them enough time to also evaluate high-value

461 pieces and show the same reaction time benefit from Experiment 1. However, if
 462 a participant is unlucky and slow to find the bad piece, this would come at the
 463 cost of not having the opportunity to evaluate the high-value pieces, and the
 464 reaction time benefit from Experiment 1 should disappear.

465 To test this effect experimentally, participants in Experiment 2 were told that
 466 large pieces had higher expected values, but that one of the small pieces would
 467 lead to a very high cost. Half of the participants were then shown a puzzle where
 468 all of the small puzzle pieces were snakes (Easy-Snake condition). Participants
 469 searching for the snake in this condition would always find one on the first try
 470 (perhaps believing they got lucky). The other half of participants were shown a
 471 puzzle where none of the small pieces were snakes (Impossible-Snake condition).
 472 Participants searching for the snake in this condition would therefore fail to
 473 find it, no matter how long they spent (perhaps believing they made a rotation
 474 error). At the end of the thinking phase, all participants were prompted to
 475 select which is better of two high-value pieces. If participants simply prioritize
 476 thinking about good pieces (as in Experiment 1), our manipulation should have
 477 no effect on participants, as they would not take more or less time searching
 478 for the snake. However, if participants switch their thinking strategy by first
 479 searching for the snake before thinking about good options, then participants
 480 in the Easy-Snake condition should be faster in the decision phase relative to
 481 participants in the Impossible-Snake condition (because only participants in the
 482 Easy-Snake condition had the time to evaluate the high-value pieces).

483 *4.1. Model Simulation and Results*

484 To search for problem structures where the good-action bias no longer ap-
 485 plies, we modified the decision problem along two dimensions: (1) the number
 486 of options an agent could choose between in the decision phase (set to either 2,
 487 3, 4, 5 and 6 out of 6), and (2) the number of low-value options in the puzzle

among which the snake could be hidden (set to 2, 3, 4, or 5). For instance, the problem where number of choices = 3 and number of apparent snakes = 4 corresponds to a puzzle with two large (i.e., potentially high-value) pieces and four small (i.e., potentially low-value) pieces, one of which is believed to be the snake. At the decision phase, the model would be allowed to select one of three randomly selected pieces to win (or lose) points.

We trained POMDPs to solve all 20 problems in the 5 x 4 parameter space (see Supplemental Methods for full implementation details). We then began by testing whether any region in this space led to strategies that prioritized thinking about bad options. Figure 3 shows the results from this analysis, coding the first thinking action chosen by the model. When the agent can choose among all, or nearly all, the options during the decision phase (rightmost columns), the model begins by searching for high-value options. This makes sense: when the agent can choose from all options, knowing the option with the highest value should suffice. However, when the agent will have to select the best piece from a limited subset (e.g. two options; leftmost column), the model no longer prioritizes the potentially high-value options. Instead, it allocates time depending on the number of apparent snakes. The higher the number of apparent low-value options, the more the agent thinks about them (indicated by how the green squares turn red as we increase the number of low-value options along the y-axis in Figure 3). This again makes sense: the more low-value options, the more likely that one of them will be part of the decision phase, in which case knowing whether the piece is the snake or not is critical.

Of the space of 20 models, we chose a setting with four low-value options where a decision phase with two pieces (third row, leftmost column). We chose this setting because it was one of the closest to Experiment 1 (critically involving only two choices at the decision phase), but where our analysis suggested that

Figure 3: First actions of each model systematically varying the number of apparent low-value options (among which one was a “snake” that provides negative cost), as well as the number of choices the agent is during the decision phase.

515 people might prioritize thinking about bad options. Within this setting, the
 516 time an agent devotes to evaluating potentially high-value options depends on
 517 whether an agent finds the snake immediately or not. As we discussed above, if
 518 the agent finds the snake on the first try, then it can switch to thinking about
 519 the two high-value pieces. But if it never finds the snake, it will keep thinking
 520 about the apparent snakes until time runs out (or at least realizes or determines
 521 that it’s futile to search for the snake). To probe this intuition, we tested
 522 our model in two modified situations. In one condition, we gave the POMDP
 523 policy a puzzle where all low-value options were snakes (Easy-Snake condition).
 524 In the other condition, we gave the POMDP policy a puzzle where none of
 525 the low-value options were snakes (Impossible-Snake condition). Critically, the
 526 POMDP was always trained under the assumption that one (and only one) of
 527 the small puzzle pieces would be a snake. Testing the POMDP on these new
 528 problems therefore created situations where the POMDP would be led to believe
 529 it had found the unique snake on its first try in the Easy-Snake condition, and

530 that it had failed to find the snake in the Impossible-Snake condition. Figure
531 4a-b, show our model’s thinking strategies across both conditions. In the Easy-
532 Snake the model first searches for the snake (considering low-value options) and
533 immediately switches to evaluate high-value options on the second action, having
534 believed that it found the snake. In contrast, in the Impossible-Snake condition,
535 the model devotes more time searching for the snake, at the cost of being unable
536 to spend as much time evaluating the potentially high-value pieces. Given these
537 thinking strategies, we then asked what people will do in these situations. Will
538 participants also recognize that the best strategy is to search for the snake and
539 then switch to evaluating potentially high-value options? If they do, participants
540 in the Impossible-Snake condition should be significantly slower at identifying
541 which is better of two high-value options relative to participants in the Easy-
542 Snake condition.

543 4.2. *Method*

544 This experiment was identical to Experiment 1a, except as noted. Seventy
545 new participants participated. Three additional participants were recruited but
546 excluded from the study based on pre-registered criteria (see Results). The sam-
547 ple size was determined before data collection began based on a power analysis
548 run on pilot data (see Supplemental Materials for details).

549 Participants first read a brief set of task instructions that explained the logic
550 of the task. Here, participants were told that the six puzzle pieces consisted of
551 two potentially high-value options (where one fit and the other did not) and
552 four potentially low-value options (where three pieces fit and only one did not
553 [the snake]). Participants were told that choosing the small piece that did not
554 fit would lead to a decrease in their score of 10 points. In the Easy-Snake
555 condition, all of the low-value options did not fit—in which case any option
556 people think of should be the snake. In a Impossible-Snake condition, none of

Figure 4: (a) Model predictions about what pieces to think about as a function of time step in the Easy-Snake condition of Experiment 2. The model prioritized searching for the snake, and then switched to evaluating the two high-value options. (b) Model predictions in the Impossible-Snake condition. The model, by prioritizing searching for the snake, spends the majority of time thinking about low-value pieces, at the cost of not being able to consider the high-value pieces. (c) Results from Experiment 2. The bars depict mean response times in each condition and error bars reflect 95% confidence intervals.

the low-value options did not fit—in which case people would never actually be able to find a snake. Participants began the task with a score of 0. In the decision phase, participants were always asked to decide between the two potentially high-value options. All methods and analyses were pre-registered (<https://aspredicted.org/blind.php?x=im62yy>).

4.3. Results and Discussion

Three participants were excluded because their mean performance in the comprehension questions was 2 standard deviations below the grand population mean. These participants were replaced, until a total of 70 participants was reached. Response accuracy and response times for the single trial were recorded for each observer. Only response times where participants responded correctly were included in our analysis. Participants in the Easy-Snake condition re-

sponded faster to the potentially high-value options ($M=2.11s$, $SD=1.67s$) than participants in the Impossible-Snake condition ($M=2.89s$, $SD=1.33s$; $t(46.36)=2.60$, $p=.013$, $d=0.70$ over the logarithm of the distributions; see Figure 4c). Including the incorrect answers did not yield any different results, $t(57.25)=2.85$, $p=.006$, $d=0.68$). There was no significant difference in accuracy across the Easy-Snake and Impossible-Snake conditions (62.86% vs. 80.00%, Fisher’s exact: $p=.185$). This pattern of results suggests that people in the Impossible-Snake condition might have fixated more on the apparent snakes, giving them less time to think about the potentially high-value options. Implicit in this is that people must have started thinking about the potentially low-value options first—ultimately suggesting that they flexibly switched their thinking strategy to think about the worst options instead of the best options this time.

5. General Discussion

People often think about possible actions they can perform even before they are faced with any actual decision. A question arises about which specific actions people tend to consider when they are engaged in this type of cognition. To make progress on this question, we used computational methods to determine which actions would be best to consider in various different settings and conducted a series of experiments to determine which actions people actually do tend to consider.

Computationally, we formalized this class of problems as a partially observable Markov decision process (POMDP), with thinking itself treated as a type of action. We then found the best thinking plan for a number of different specific decision problems. The results indicated that there are certain decision problems for which the best strategy is to think about the potentially high-value options and others for which the best strategy is to think about the potentially

595 low-value options.

596 More specifically, the results revealed an interaction between (a) the number
597 of options the agent is able to consider at decision time and (b) the proportion
598 of options that are low-value. When the number of options the agent is able to
599 consider at decision time is high, it is always best to consider the potentially
600 high-value options. By contrast, when the number of options the agent is able
601 to consider at decision time is low, the best strategy depends on the proportion
602 of options that are low-value (especially if some of these options are extremely
603 bad, as in the “snake” in the stack). When that proportion is small, the best
604 strategy is to consider the potentially high-value options, whereas when that
605 proportion is large, the best strategy is to consider the potentially low-value
606 options.

607 We then conducted a series of experiments that provided evidence about
608 which options people actually tend to consider. To do so, we developed an
609 experimental paradigm in which participants went through a “thinking phase”
610 followed by a “decision phase.” In the thinking phase, participants were given a
611 number of different options and had the opportunity to think about whichever
612 options they wanted to. Then, in the decision phase, participants were con-
613 fronted with just two of these options and asked to choose between them. Reac-
614 tion times in the decision phase thereby provided evidence about which options
615 participants were considering in the thinking phase. Specifically, the shorter a
616 participant’s reaction time in response to a pair of options in the decision phase,
617 the more reason we have to conclude that the participant already considered
618 those options in the thinking phase.

619 Using this method, Experiments 1a-c looked at decision problems for which
620 the formal model indicated that the best strategy would be to think about the
621 high-value options. In those experiments, the results indicated that participants

622 actually did tend to think about the high-value options. Experiment 2 then
623 turned to decision problems for which the formal model indicated that the best
624 strategy would be to think about the low-value options, and in that experiment,
625 the results indicated that people did indeed tend to think about the low-value
626 options. Taken together, the studies therefore suggest that people are able to
627 respond flexibly, thinking about either high-value options or low-value options
628 depending on which strategy is best for the specific decision problem they face.

629 These results leave us with a puzzle. On one hand, we have evidence that
630 people show a remarkable level of sophistication and flexibility in determining
631 which actions to consider. But, on the other, there is strong reason to suspect
632 that the actual cognitive process people use to make this determination involves
633 relatively little computation. For example, in our studies, participants could de-
634 vote computation to thinking about which of the six options they should think
635 about, but to the extent that they do this, they will have less computation avail-
636 able to actually think about the options themselves. Thus, it would never make
637 sense to devote more computation to this task than would have been required
638 just to think about all six options. This problem becomes even more impor-
639 tant when considering how POMDPs are associated with high computational
640 demands, suggesting that people are not using that same approach when they
641 face this problem. So then, how do people actually solve the problem they faced
642 in these experiments?

643 *5.1. Exploring the underlying process*

644 Future work should continue to explore people’s ability to show this sort of
645 flexibility in determining which options to think about. Such work will require
646 a mixture of computational and empirical research.

647 At a computational level, we face further questions about the conditions
648 under which it is best to think about potentially good options vs. potentially

649 bad options. In the present paper, we looked at two specific dimensions along
650 which decision problems could vary, but future research could continue this
651 investigation by looking at other dimensions. Such research would presumably
652 uncover other possible decision problems for which the best strategy is to think
653 about potentially high-value options as well as other possible decision problems
654 for which the best strategy is to think about potentially low-value options.

655 Work in this area might eventually lead to the development of more general
656 theories that specify the conditions under which the best strategy is to think
657 about potentially high-value vs. potentially low-value options. Such theories
658 would not be limited just to one particular setting but would provide more
659 general insights about when each strategy is best. For example, it might turn out
660 that all possible decision problems that have certain features will be problems
661 for which the best strategy is to think about potentially low-value options.

662 To the extent that we are able to develop such an account, we open up the
663 possibility of a new explanation of the effects observed in the present studies. It
664 might be that people are not determining which options to consider by using a
665 process that is even remotely like solving a POMDP. Instead, it might be that
666 people are simply checking for certain features that serve as reliable cues to
667 whether it is better to think of potentially high-value or potentially low-value
668 options. Future empirical studies could directly test specific hypotheses of this
669 form.

670 If we do find that people are responding to certain characteristics that are
671 reliable indicators of which options are best to consider, we would face a further
672 question as to how people come to be able to identify these in the first place.
673 One possible answer would be that people never need to learn them. Instead,
674 the use of these features could simply be built into people's decision-making
675 mechanisms. A second possibility would be that people are actually learning the

676 use of these features over the course of numerous episodes of decision-making.

677 If people are indeed learning to use the relevant features, a question would
678 arise as to how this learning takes place. One possible answer would be that peo-
679 ple are making use of familiar mechanisms of model-free learning (e.g., Gläscher
680 et al. 2010). On this view, people would have to be capable of model-free learn-
681 ing at an extremely abstract level. For example, over the course of numerous
682 episodes of playing chess, people would have to be learning not only about the
683 game of chess in particular but also about very abstract patterns regarding
684 which options tend to be most worth thinking about. The development of al-
685 gorithms for such learning is an important problem for future research—which
686 some researchers have in fact already begun to address (Jain et al. 2019).

687 5.2. *Flexibility and inflexibility*

688 Our results are also striking in light of existing research that suggests that
689 the mechanism people use to figure out which options to consider are highly
690 inflexible. Specifically, a number of studies suggest that people show an inflexible
691 tendency to think of options that are generally good and statistically frequent
692 (Bear et al. 2020; Morris et al. 2021; Phillips et al. 2019). For many decision
693 problems, it will be helpful to think about options that are statistically frequent
694 and generally good, but these studies seem to suggest that people tend to think
695 about those options even when the decision problem is structured in such a way
696 that it is obviously not helpful to think about those options.

697 For example, in one study, participants were asked to name the food they
698 would *least* want for dinner (Morris et al. 2021). Clearly, in answering this
699 question, it is helpful to think about foods that are generally bad, but the
700 results indicated that participants actually showed a tendency to think first of
701 foods that are generally good. Results like this one seem to suggest that there is
702 an inflexible cognitive mechanism that generates options for people to consider

703 when planning.

704 One way in which the present studies depart from previous work, however,
705 is in the type of situation and options that participants were presented with. In
706 previous studies, people were typically asked about options and situations for
707 which people must already have had model-free values (e.g., how many hours
708 of TV to watch a day). In such cases, it may be that the pre-existing value
709 assignment gives rise to the inflexible mode of thought that previous studies
710 have observed. In contrast, the present studies look at the options people think
711 about in novel situations, where they do not already have model-free values with
712 the various options. The absence of model-free values may have allowed for the
713 flexibility observed in our experiments.

714 Another important distinction between the present studies and previous
715 work is that between the conditions under which an option “comes to mind” and
716 the conditions under which people actually simulate forward what would hap-
717 pen if they chose an option. Within existing research on which options come to
718 mind, there is some evidence that the options that come to mind are determined
719 by an inflexible mechanism. Yet, despite this, it might be that people make use
720 of a more sophisticated and flexible mechanism to determine which simulations
721 to run. Thus, in the present studies, it might turn out that the options with
722 potentially high value are the first that come to mind in all conditions, but
723 then it might be that people make use of a different psychological mechanism to
724 determine which simulations to run, and that this mechanism sometimes selects
725 the options with potentially low value.

726 To the extent that this latter answer turns out to be correct, a further
727 question arises as to how far the flexibility extends. The present studies explore
728 people’s thinking in a context in which they are aiming to achieve a specific
729 goal, but much of our thinking is not goal-directed—as in mind-wandering, in

730 which people are not trying to address any particular decision problem (e.g.
731 Irving & Thompson 2018; Mooneyham & Schooler 2013). We therefore face a
732 question about which options people tend to simulate during mind wandering.
733 When people’s minds are wandering, do they also show the sort of flexibility
734 observed in the present studies? For example, does the degree to which they
735 think about potentially high-value vs. potentially low-value options depend in
736 part on the number of options they expect to be choosing between at decision
737 time? Regardless of what the answer to this question turns out to be, such
738 research promises to give us a real insight into the scope or boundary conditions
739 of the phenomena we have been exploring here.

740 5.3. *Conclusion*

741 When people face a choice from a set of options, they often face a difficult
742 problem. They find themselves faced with so many different options that it
743 would not be possible to consider them all, and they therefore need to have
744 some way of picking out certain specific options that are especially worthy of
745 consideration.

746 In a series of studies, we looked at the options people tend to consider and
747 found that people appear to be selecting options using remarkably sophisticated
748 criteria. People don’t simply show a general tendency to consider the potentially
749 high-value options. Instead, they consider the high-value options on certain
750 decision problems and the low-value options on others. A key task now will
751 be to explain how people are able to show this flexibility and how to reconcile
752 the flexibility they show on the sort of problems in this current study with
753 inflexibility they show on other, seemingly related problems.

754 References

- 755 Bear, A., Bensinger, S., Jara-Ettinger, J., Knobe, J., & Cushman, F. (2020).
756 What comes to mind? *Cognition*, *194*, 104057.
- 757 Callaway, F., van Opheusden, B., Gul, S., Das, P., Krueger, P., Lieder, F., &
758 Griffiths, T. (2021a). Rational use of cognitive resources in human planning.
759 *Nature Human Behavior*, .
- 760 Callaway, F., Rangel, A., & Griffiths, T. L. (2021b). Fixation patterns in simple
761 choice reflect optimal information sampling. *PLoS computational biology*, *17*,
762 e1008863.
- 763 Cassandra, A. R. (1998). A survey of pomdp applications. In *Working notes*
764 *of AAAI 1998 fall symposium on planning with partially observable Markov*
765 *decision processes*.
- 766 Chen, H., Chang, H. J., & Howes, A. (2021). Apparently irrational choice as
767 optimal sequential decision making. In *Proceedings of the AAAI Conference*
768 *on Artificial Intelligence* (pp. 792–800). volume 35.
- 769 Cooper, L. A. (1975). Mental rotation of random two-dimensional shapes. *Cog-*
770 *nitive Psychology*, *7*, 20–43.
- 771 Gelly, S., & Silver, D. (2011). Monte-carlo tree search and rapid action value
772 estimation in computer go. *Artificial Intelligence*, *175*, 1856–1875.
- 773 Gläscher, J., Daw, N., Dayan, P., & O’Doherty, J. P. (2010). States versus
774 rewards: dissociable neural prediction error signals underlying model-based
775 and model-free reinforcement learning. *Neuron*, *66*, 585–595.
- 776 Griffiths, T. L., Callaway, F., Chang, M. B., Grant, E., Krueger, P. M., &
777 Lieder, F. (2019). Doing more with less: meta-reasoning and meta-learning
778 in humans and machines. *Current Opinion in Behavioral Sciences*, *29*, 24–30.

- 779 Griffiths, T. L., Lieder, F., & Goodman, N. D. (2015). Rational use of cognitive
780 resources: Levels of analysis between the computational and the algorithmic.
781 *Topics in Cognitive Science*, 7, 217–229.
- 782 Hsu, D., Lee, W. S., & Rong, N. (2008). A point-based pomdp planner for
783 target tracking. In *2008 IEEE International Conference on Robotics and*
784 *Automation* (pp. 2644–2650).
- 785 Icard, T., Cushman, F., & Knobe, J. (2018). On the instrumental value of hy-
786 pothetical and counterfactual thought. In *Proceedings of the Annual Meeting*
787 *of the Cognitive Science Society* (pp. 517–522).
- 788 Icard, T. F., Kominsky, J. F., & Knobe, J. (2017). Normality and actual causal
789 strength. *Cognition*, 161, 80–93.
- 790 Irving, Z. C., & Thompson, E. (2018). The philosophy of mind-wandering. *The*
791 *Oxford Handbook of Spontaneous Thought*, (pp. 87–96).
- 792 Jain, Y. R., Callaway, F., & Lieder, F. (2019). Measuring how people learn
793 how to plan. In *Proceedings of the Annual Meeting of the Cognitive Science*
794 *Society* (pp. 1956–1962).
- 795 Kurniawati, H., Hsu, D., & Lee, W. S. (2008). Sarsop: Efficient point-based
796 pomdp planning by approximating optimally reachable belief spaces. In
797 *Robotics: Science and Systems*.
- 798 Lieder, F., & Griffiths, T. L. (2017). Strategy selection as rational metareason-
799 ing. *Psychological Review*, 124, 762–794.
- 800 Lieder, F., & Griffiths, T. L. (2020). Resource-rational analysis: understand-
801 ing human cognition as the optimal use of limited computational resources.
802 *Behavioral and Brain Sciences*, 43.

803 Mattar, M. G., & Daw, N. D. (2018). Prioritized memory access explains plan-
804 ning and hippocampal replay. *Nature neuroscience*, *21*, 1609–1617.

805 Mooneyham, B. W., & Schooler, J. W. (2013). The costs and benefits of mind-
806 wandering: a review. *Canadian Journal of Experimental Psychology/Revue*
807 *canadienne de psychologie expérimentale*, *67*, 11–18.

808 Morris, J., Adam end Phillips, Huang, K., & Cushman, F. (2021). Generating
809 options and choosing between them rely on distinct forms of value represen-
810 tation. *Psychological Science*, *32*, 1731–1746.

811 Ng, A. Y., & Jordan, M. I. (2000). Pegasus: A policy search method for large
812 mdps and pomdps. In *Uncertainty in Artificial Intelligence Proceedings* (pp.
813 406–415).

814 Peirce, J., Gray, J. R., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H.,
815 Kastman, E., & Lindeløv, J. K. (2019). Psychopy2: Experiments in behavior
816 made easy. *Behavior Research Methods*, *51*, 195–203.

817 Phillips, J., Morris, A., & Cushman, F. (2019). How we know what not to think.
818 *Trends in Cognitive Sciences*, *23*, 1026–1040.

819 Shepard, R. N., & Metzler, J. (1971). Mental rotation of three-dimensional
820 objects. *Science*, *171*, 701–703.

821 Sutton, R. S., Barto, A. G. et al. (1998). *Introduction to reinforcement learning*
822 volume 135. MIT Press Cambridge.

823 Vul, E., Goodman, N., Griffiths, T. L., & Tenenbaum, J. B. (2014). One and
824 done? optimal decisions from very few samples. *Cognitive science*, *38*, 599–
825 637.

826 6. Supplementary Information

827 6.1. POMDP Implementation Details for Experiment 1

828 To implement the experiment design formally, all pieces were designed so
829 that they could only take one of four possible values: 0, 3, 5, and 7 (determined
830 by the number of additional blocks the puzzle piece added after fitting into
831 the block). In any given context, each piece could take on two possible states,
832 capturing the possibility that the piece’s true value matched the observed one
833 or that it did not. For instance, the largest piece could have state ‘appears to
834 have value 7 but its true value is 0’, or ‘accurately appears to have value 7’.
835 These representations produced a set of 64 basic states (2^6) for any given trial.
836 We then supplemented each of these basic states with a corresponding decision
837 state, which identified a subset of actions that the agent can choose from. For
838 instance, if the decision phase asked participants to choose one of two pieces,
839 this would result in an additional 15 decision states (as there were 15 ways to
840 select two objects from a set of six options) associated with each of the 64 basic
841 states. Thus, our task state space consisted of 1024 states: 64 basic thinking
842 states, and 960 (64×15) possible decision states.

843 We next defined the state of actions as consisting of six thinking actions and
844 six selection options (one associated with each piece). To model the result of
845 these actions, we assumed that a thinking action reveals the true value of the
846 selected piece. We defined the observation space as consisting of four possible
847 observations—0, 3, 5, and 7. To account for human reasoning errors, we varied
848 the probability of getting a correct observation (i.e., believing that a piece fits
849 when it does not, or believing that a piece does not fit when it does). To
850 achieve this, we included a certainty parameter, which we varied from 95%
851 (i.e., a 5% chance of making a reasoning error) to 100% (i.e., full confidence of
852 no reasoning error) in steps of 1%, each implemented on a separate POMDP.

853 In these models, the probability of switching from the thinking phase to the
854 decision phase was set to 30%. By contrast, we assumed that a selection always
855 produces no observation (formally captured by adding an additional 'empty'
856 observation). Thus, choosing to think about a piece provides information about
857 the world, but choosing one piece does not.

858 To model the state-space change, we assumed that during the thinking phase,
859 a thinking action had a 70% probability of leaving the world unchanged, and a
860 30% probability that the world would switch to the decision-making phase, with
861 a discount parameter of 0.95. Critically, this probability is independent of which
862 thinking action the agent takes. Thus, this probability does not represent a
863 causal relation between thinking and state change. It instead expresses the idea
864 that thinking requires time and that there is always a possibility that time will
865 run out and the agent will have to make a choice. In contrast to thinking actions,
866 selection actions always terminated the task. Whenever selection actions were
867 chosen during the thinking phase, the task was terminated with no reward.
868 Similarly, selecting a piece that was not available for choice during the final
869 phase also terminated the task with no reward, but selecting one of the pieces
870 offered yielded the reward associated with the piece.

871 Finally, we assumed that agents' initial beliefs only assigned a positive prob-
872 ability to states where the observable values matched the sizes of the pieces, but
873 leaving the true values unknown.

874 With these specifications, it is now possible to use POMDPs to compute
875 how participants should prioritize different thinking actions as a function of
876 different task variables, including the costs and reward of different actions, as
877 well as the number of pieces that the agent will ultimately be able to select
878 from. To estimate this solution, we used the SARSOP algorithm (Successive
879 Approximations of the Reachable Space under Optimal Policies; Kurniawati

et al. 2008).

6.2. POMDP Implementation Details for Experiment 2

We explored 20 different decision problems: 5 different choice parameters x 4 different snake parameters. In the choice parameter, we tested what strategy the agent would use when given 2, 3, 4, 5, or 6 choices in the decision phase. In the snake parameter, we tested what strategy the agent would use when there were 2, 3, 4, or 5 possible snakes (of the 6 puzzle pieces). To account for the mismatch of beliefs from the possible states, we introduced ‘noise’ to the observations, such that the probability of getting the correct observation was set to 97%, and the 3% distributed among other possible observations. In the actual simulation the policy, we then ‘hacked’ the simulation, such that the agent always obtained observations of the piece not fitting in the ‘easy snake’ condition, and observations the piece fitting in the ‘impossible snake’ condition. The agent received a reward of -500 if it selected the piece was the designated ‘snake’ of the stack.

6.3. Power Analyses

A pilot experiment on 16 participants was initially run for Experiment 1a, which suggested a sample size of $n=27$ per group to obtain a power of 98%. We rounded this up to $n=30$ per group (for a total of $n=60$), and this served as the sample size for Experiments 1a and 1c. For Experiment 1b, pilot experiments suggested a sample size of $n=35$ per group to obtain a power of 98% (for a total of $n=70$). This was then matched to be the sample size for Experiment 1b and 2.