

Imagining the good: An offline tendency to simulate good options even when no decision has to be made

Joan Danielle K. Ongchoco (joan.ongchoco@yale.edu)

Julian Jara-Ettinger (julian.jara-ettinger@yale.edu)

Joshua Knobe (joshua.knobe@yale.edu)

Abstract

Even when we are not faced with any decision, we sometimes engage in offline cognition where we simulate various possible actions we can take. In these instances, which options do we tend to simulate? Computational models have suggested that it is better to focus our limited cognitive resources towards simulating and refining our representations of options that appear, at first blush, to have higher values. Two experimental studies explore whether we use this strategy. Participants went through an ‘offline’ thinking phase, and an ‘online’ decision-making phase. Participants first freely viewed various options, which they had to simulate to determine their actual values. They were later asked to decide between good or bad options. Offline simulation produced faster online response times for the options that appeared to have higher values, indicating a pre-computation benefit for these items. These results suggest that people focus their offline cognition on the apparently good.

Keywords: Sampling; simulation; decision-making; mental rotation

Introduction

When people are trying to make decisions, they sometimes proceed by simulating possible options and asking what the outcome would be for each. Existing research has explored the various ways people use such simulations not just in making inferences about the world in general (e.g. Battaglia, Hamrick, & Tenenbaum, 2013; Callaway, Hamrick, & Griffiths, 2017; Gerstenberg, Goodman, Lagnado, & Tenenbaum, 2014), but also in the specific context of decision-making (e.g. Barron, Dolan, & Behrens, 2013; Hamrick, Smith, Griffiths, & Vul, 2015; Hamrick et al., 2016; Lieder, Griffiths, & Hsu, 2018; Wimmer & Shohamy, 2012). Because of the limited time and capacity during online processing, a central question is how to efficiently allocate computational resources. Previous work has investigated how people determine which simulations and how many to run at the moment when they have to make a decision (e.g. Callaway, Gul, Krueger, Griffiths, & Lieder, 2018; Hamrick & Griffiths, 2014; Srivastava, Miller-Trede, Schrater, & Vul, 2016; Vul, Goodman, Griffiths, & Tenenbaum, 2014).

Importantly, however, people are also capable of simulating different possible options offline, i.e., considering possible options when they are not faced with any immediate decision (see, e.g. Gershman, Markman, & Otto, 2014). For example, even when you are not out with someone on a dinner date, you may find yourself simulating various possible ways you might introduce yourself to a (perhaps hypothetical) person. This offline simulation may then prove helpful

when you later face an actual online decision-making problem.

Though our capacity for running simulations offline is not quite as limited as our capacity for running simulations online, we still cannot simulate all possible options. Thus, if you are thinking offline about how to introduce yourself on a date, you would inevitably simulate some options (e.g., talking about your background and interests) but not others (e.g., talking in detail about how loudly you snore). This raises the question—which options do people tend to simulate when thinking offline?

What should we think about offline?


One way in to this problem is to begin by asking which options it would actually be rational to simulate. Suppose our aim is to select the best action during a subsequent episode of online decision-making. Given this aim, which options would it be best to simulate offline?

Of course, one possible answer would be that it does not matter which specific options we end up simulating. Simulating different possible options for hypothetical situations might simply be helpful in a broad way, for learning the general features of good versus bad options, without having to specifically compute which option is better than another. In other words, running simulations may be a good way of discovering various heuristics about different options that we can then use later, during online decision-making.

An alternative possibility, however, is that simulating offline is not just good for learning various decision-making heuristics, but can also help us get better value estimates for specific options. When we simulate an option, we can improve our representation of the value of that option. This ‘pre-computed’ value can come in handy when we have to make decisions in the pressure of the moment, when we do not have much time to think.

The problem can then be formulated as follows. At any given point, we have a representation of the value of each option. Some options are represented with high values (i.e. as good options), others with low values (i.e. as bad options), and others as having an intermediate level of value. At first blush, all of these representations will be at least somewhat inaccurate. We may have a sense of what is good or bad, but generally need to think more about which of these is actually the best or the worst. In simulating a specific option, we can then improve our representation of its value. However, we cannot run simulations for all options. Thus, we have to

(a) Sample Puzzle



(b) Experiment Procedure




Figure 1: (a) A sample puzzle of the sort used in the experiments. The yellow blocks constitute the puzzle, the blue blocks constitute the pieces. While looking at the puzzle, readers may simulate which pieces would fit or not fit. (b) A caricatured depiction of a sample experimental procedure without the instructions. Participants first go through a series of practice trials, are given a unique puzzle, and are then asked to decide between two options.

decide which of our representations we want to have a more accurate value estimate of, while allowing others to remain inaccurate.

One intriguing finding from existing research is that, other things being equal, it is generally a good idea to run offline simulations of options that we initially think have high values (e.g. Gelly & Silver, 2011; Icard, Cushman, & Knobe, 2018). In other words, if we can only improve our representation of some options, it is better to improve our value estimates of the options we initially represent to have high value than to improve our value estimates of the options we now represent to have low value.

Existing computational work has explored this point within the framework of reinforcement learning (e.g. Icard et al., 2018), but the core intuition is easy to grasp even independently of any formal framework. Suppose that there are now two different inaccuracies in your representations: (a) the option that you mistakenly represent as second-best is actually the best, and (b) the option that you mistakenly represent as second-worst is actually the worst. Now suppose that you are only able to correct one of these inaccuracies, which would you focus on?

The key point is that when you later use these representations in online decision-making, you would ideally want to choose the best option. Thus, it is important to be highly accurate about which of the good options truly is the best, but it is not nearly as important to be accurate about which of the

bad options truly is the worst. You should therefore devote your limited offline cognition to the options that you initially think to have high value, and then improve your representations from there.

A question now arises as to whether human cognition actually works in this way. When people only have a limited amount of time to devote to offline cognition, do they tend to run simulations of the options they regard as having high value, even when they do not have to (as there are infinite possibilities one can simulate offline, and there is no immediate specific decision that has to be made)?

The present studies

To address these questions, we conducted two studies in which participants had an opportunity to go through an ‘offline’ thinking phase before a subsequent ‘online’ decision-making phase. In the offline phase, participants were given an array of options to freely think about. Crucially, they had to simulate these different options to determine their actual values. In the online phase, participants were asked to decide between two options. The key question was which options would participants think about during the offline phase, when they were not told what decisions they eventually would have to make. To tap into participants’ tendencies during offline simulation, we used their response times during the online decision-making phase. We reasoned that if participants were refining their values about specific options and

computing which ones were better than others during the offline phase, then they should respond faster when choosing between those same options during the online phase.

In a novel paradigm, we used incomplete block-puzzles (that look like Tetris) with arrays of different puzzle pieces that would either fit the puzzle or not. In this design, determining whether the puzzle pieces fit would require participants to manipulate these pieces in their minds, akin to classic mental simulation and rotation studies (e.g. Cooper, 1975; Shepard & Metzler, 1971). Moreover, this block-puzzle design allowed us to also specify a ‘surface’ or apparent value (what people initially think the value of a piece to be) and an actual value (the value of the piece after being simulated) for each puzzle piece, where value could be defined by both the number of blocks the piece had, and whether the piece would actually fit the puzzle.

The idea behind this particular design was that, at first glance, one should be able to immediately ‘see’ the surface values of the different puzzle pieces, such that some pieces would clearly have higher values (as indicated by the brute number of blocks) than others. From this surface value, one can either simulate the apparently good or the apparently bad. Crucially, it is only by mentally rotating and simulating how these pieces would fit the incomplete block-puzzle that one can get a better sense of the actual values of these pieces. However, one cannot simulate all pieces during the limited window of the offline phase. This time limit allowed us to check which pieces people would simulate over others.

In Experiment 1, we looked at whether people systematically responded faster during online decision-making to some options over others as a function of offline simulation. In Experiment 2, we investigated the mechanism by which offline simulation may lead to benefits in online decision-making, as a function of developing broad heuristics about what options are good and bad in general versus actually pre-computing and refining the value of specific options. These experiments altogether explore the principles governing online and offline thinking, and suggest that these may in fact be more closely related than we previously thought: people systematically and actively imagine the good not only when there is an immediate judgement or decision to be made, but also offline, even when they do not have to.

Experiment 1

Participants were given an array of six rotated puzzle pieces per incomplete puzzle during the offline phase. Each piece had a specific value, defined by how many blocks would end up above the puzzle, once the piece fit. In general, the more blocks a piece had, the better. However, to determine the precise value of the piece, participants had to simulate the different pieces. In Figure 1a, the upper leftmost piece would fit the puzzle when rotated counter-clockwise, and would have 7 blocks above the completed puzzle. If participants selected this piece, they would get 7 points. In contrast, the lower middle piece has the same number of blocks, but would not

fit the puzzle. If participants chose this piece, they would get 0 points. Thus, we wanted to see whether people would consider the pieces that have a high surface value (like 7) during the offline thinking phase. If participants consider some pieces more than others, they might respond faster when they have to decide between these specific pieces.

Method

All methods and analyses were pre-registered (<http://aspredicted.org/blind.php?x=rd2bd2>).

Data and code for all experiments reported here are available on https://osf.io/npwdq/?view_only=a808e1dd2d594b7992892bfa32fb7e8c.

Participants. Sixty subjects from the Yale University Library participated (with candy as compensation). The sample size was determined before data collection began.

Apparatus. Stimuli were presented using custom software written in Python with the PsychoPy libraries (Peirce, 2007) and were displayed on a monitor with a 60Hz refresh rate. Participants completed the study on a 13-inch MacBook Air with a 1440 x 900 resolution.

Stimuli. Puzzles were generated randomly through PsychoPy. Puzzles were made of 20 yellow blocks (0.5° black border) stacked in 4 rows of 5 blocks each. Each block was 2° in size. In each puzzle, a number—three or four—of the blocks in the top two rows would be missing. This created an incomplete section at the top of the puzzle.

The puzzle pieces were also generated randomly. Pieces comprised of the specific arrangement of blocks that were determined to be missing, along with additional blocks that made up the value the piece was assigned (e.g. a value of 7 meant that there were 7 blocks on top of the piece). Additional blocks were stacked on top of the piece randomly, for as long as they were always connected to a block in the piece. When the blocks were stacked, the piece was checked on all sides to make sure that only one side would fit in the puzzle. If the piece was not supposed to fit (i.e. have a value of 0), the bottom-most part of the piece was shifted to the left or to the right, in order to ensure that the piece would not fit the puzzle. Puzzle pieces were made out of 2° grey blocks (0.5° black border).

Procedure and design. Throughout the experiment, on the top-left of the screen, there was ‘Total Points:’ counter. All the text in this experiment was drawn in black Monaco font (0.6° in height). In a single-trial experiment, participants first went through the instructions for the task. They were told that their goal was to earn as many points as they could. They were given sample incomplete block-puzzles and arrays of possible options. They were told that they would be asked questions about these different options afterwards, and would get a number of points corresponding to the particular option they would be asked about. They were told that the value of each piece was defined by the number of blocks

that would end up above the completed puzzle once the piece fit, and that pieces would also be rotated but never flipped. They were also told that speed in responding will be important so the participants would go through a practice section first before the actual trial. After these instructions, participants would then be shown the puzzle, which subtended from 1° above to -3° below the center.

In the offline phase, participants were told that six pieces would now appear above the puzzle, in two rows of 3 pieces each. Participants were told that they did not have to do anything but just look and study the pieces. The six pieces comprised of three pairs with surface values of 3, 5, and 7. In each pair, each piece would have a different actual value: one would fit, and the other was made to not fit the puzzle. During this time, a countdown timer (6° below the center) would start from 5 and decrease per each passing second.

After the offline phase, the online decision-making phase began, where participants now had to decide which of two pieces was the better piece. The blocks of two of the pieces would turn from grey to green to indicate which pieces the participants would have to choose from, and each of these pieces was assigned either a letter j or k. Participants were simply asked, “Which piece is better?”, and indicated which piece they preferred by keying in the letter of the piece. In the practice section, participants responded to a total of six practice trials. Throughout the practice trials, if participants responded correctly, the total points counter would increase by the value of the piece they chose (if the piece fit, then this value was determined by the number of blocks above the completed puzzle; if the piece did not fit, the participants would automatically get 0 points).

After participants completed the practice section, they were told that they would be shown a different puzzle and a new set of pieces. Participants were again told that they could be asked about any of these pieces afterwards. To facilitate the pressure of having to decide in the moment, participants were now encouraged to respond as fast as possible, and were told that they would get bonus points for responding quickly. Participants first went through the offline phase, where they were again presented a new puzzle with six puzzle pieces. The countdown timer appeared again. After five seconds, participants began the online decision-making phase, responded to two pieces from the array of six options. In the Good Options condition, participants decided between the two pieces with a value of 7. In the Bad Options condition, participants decided between the two pieces with a value of 3. Participants were randomly assigned to decide either between the good options or the bad options.

Results and discussion

Three participants were excluded because their mean performance in the practice section was 2 standard deviations below the grand population mean ($M=29.08$ out of 34 total points that could be earned in the practice section; the cut-off was at 18.63). These subjects were replaced, until a




Figure 2: Results from Experiments 1 and 2. The left bar graph presents the mean response times per decision type in Experiment 1, while the right bar graph presents the mean response times per decision type in Experiment 2. The vertical axis represents the mean response time in seconds. The horizontal axis represents the key comparisons in both experiments. The error bars reflect 95% confidence intervals within experiments.

total of 60 participants was reached (30 decided between the good options, and 30 decided between the bad options). Response accuracy and response times for the single trial were recorded for each observer. Only response times where participants responded correctly were included in the analysis.

Initial inspection of the left bar graph in Figure 2 shows a lower mean response time for good options than for the bad options. This initial impression was confirmed with statistical tests. Mean reaction time for good options ($M=2.56s$, $MSD=1.77s$) was significantly faster than for bad options ($M=5.20s$, $SD=3.82s$), $t(25.31)=2.87$, $p=.008$, $d=0.93$. Because the response time distributions violated the normality assumption, we conducted a t-test on the log transformations of the distributions (which now meet the normality assumption), $t(43.37)=3.65$, $p<.001$, $d=1.06$. (We also note that 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 good options (86.67%) vs. bad options (66.67%) (Fisher’s exact, $p=.125$).

These results suggest a pre-computation ‘imagination’ benefit for the good options. In other words, it appears that when given the opportunity to freely think about an array of various options offline, people tend to simulate the options they initially think have a higher value rather than those they initially think to have lower values, even when they do not know what specific decisions they will have to make later on. In simulating the good options, people can determine offline what the actual values of these good options are, such that when it comes to having to make a decision, they re-

spond faster to the good options they had already simulated, computed, and compared beforehand.

Experiment 2

The results from the initial experiment were promising because they suggested that when given an opportunity to think, participants think about and simulate the good options more than the bad options, resulting in a response time benefit at the time of decision-making. But is this a benefit from simulating specific options and refining our representations of their values, and not just from being exposed to and learning what the good or bad options are (for instance, participants could simply have been learning throughout the practice section that the bigger pieces, regardless of their specific configurations, have higher values)? To explore the mechanism underlying the response time benefit observed in Experiment 1, we used the same block-puzzle design. We asked participants to again look at a puzzle and an array of six options. This time, during decision-making, unbeknownst to the participants, instead of presenting them with pieces that were originally in the array of six options, we presented them with a novel pair of good or bad options. Thus, none of their offline thinking strategies should have changed, since they were not told that they would be shown novel pieces. If offline simulation were simply a way of discovering heuristics about which options are good or bad, then participants should still respond faster to the good options than the bad options. However, if offline simulation involves the pre-computation of the values of specific pieces, then the pre-computation benefit observed in Experiment 1 should disappear when participants are presented with a novel pair.

Method

This experiment was identical to Experiment 1, except as noted. Sixty new participants participated, with this sample size chosen to match Experiment 1. During the decision-making phase, a new pair of pieces with the value of either 7 (i.e. Novel Good Options condition) or 3 (i.e. Novel Bad Options condition) were generated and presented to the participants. All methods and analyses were pre-registered (<http://aspredicted.org/blind.php?x=bw5n59>).

Results and discussion

One participant was excluded because their mean performance in the practice section was 2 standard deviations below the grand population mean ($M=29.08$ out of 34 total points that could be earned in the practice section; the cut-off was at 20.65). This subject was replaced, until a total of 60 participants was reached (30 decided between the good options, and 30 decided between the bad options). Response accuracy and response times for the single trial were recorded for each observer. Only response times where participants responded correctly were included in the analysis.

Initial inspection of the right bar graph in Figure 2 shows a lower mean response time for bad options than for the good options. Mean reaction time for good options ($M=4.54s$,

$SD=2.51s$) was significantly slower than for the bad options ($M=2.52s$, $SD=1.12s$), $t(30.52)=3.53$, $p=.001$, $d=1.04$. Again, because the response time distributions violated the normality assumption, we conducted a t-test on the log transformations of the distributions, $t(44.51)=3.45$, $p=.001$, $d=0.99$. (We also note that including the incorrect answers again did not yield any different results, $t(58.22)=2.36$, $p=.021$, $d=0.60$). There was no significant difference between the percentage of people who responded accurately when choosing between good options (76.67%) vs. bad options (76.67%) (Fisher's exact, $p=1$).

To compare these results with those of Experiment 1, we ran a 2 (offline vs. no offline phase) $\times$ 2 (good options vs. bad options) ANOVA. There was no main effect of offline thinking, $F(1, 88)=0.12$, $p=.728$, $\eta^2=.002$, or of decision type, $F(1, 88)=0.33$, $p=.570$, $\eta^2=.004$. Crucially, there was a significant interaction, $F(1, 88)=20.99$, $p<.001$, $\eta^2=.193$.

In short, these results show a reversal of the pattern observed in Experiment 1. Since the task is constructed in such a way that the higher value pieces contain more blocks, one might expect at baseline that participants would show longer reaction times for the higher value pieces. In Experiment 1, where participants had an opportunity to engage in offline simulation, we instead found shorter reaction times for the higher value pieces. By contrast, in the present study, we find the expected baseline result: when participants do not have an opportunity to engage in offline simulation, they show longer reaction times for the higher value pieces (perhaps because they had more blocks in general).

General Discussion

There are many instances when we imagine different options without having to immediately make a decision, as when we daydream about which restaurant to go to for dinner or what to say when we are on a date or in an important meeting. In these instances of offline simulation, what do we tend to think about, and why? The present experiments explored this question in terms of the mental simulation of visual stimuli, and asked whether people tend to simulate the apparently good options over the apparently bad options.

The key takeaway from these experiments is simple to summarize: people choosing between two good options responded faster at the point of decision-making than people choosing between two bad options, suggesting that people were thinking more about the good options during the offline thinking phase, when they did not actually have to (and we note, interestingly, even when the good options were more difficult to think about and took longer to process at baseline). Moreover, this does not seem to be just a matter of general practice and exposure to deciding between good versus bad options. When presented a novel pair of good or bad options, participants no longer show this pre-computation benefit, and in fact, perform in the opposite way (responding slower to good options than the bad options). This suggests that thinking in the general does not suffice to produce the

benefit at decision-making. Rather it is thinking and mentally simulating specific possible options offline that proves particularly adaptive when eventually having to choose between these same options.

This result adds to the existing body of work that has explored what people should think about given limited computational resources (e.g. Callaway et al., 2018; Srivastava et al., 2016; Vul et al., 2014). In online decision-making, it generally makes sense to be actively sampling the options with the highest values in order to make the best decision. Our results demonstrate that simulating the best possible options also occurs in offline cognition, when people are allowed to freely think about any option, and do not have to make any decision at all. The tendency to imagine the good options may reflect a more general principle of cognition that is at play while running both online and offline simulations.

This tendency might be interestingly related to recent work on mind-wandering and on memory replay. Research on mind-wandering finds that people tend to spend a good amount of their waking hours just thinking offline (e.g. Mason et al., 2007). Such research indicates that people's minds are in general more likely to wander to pleasant topics than unpleasant topics (e.g. Killingsworth & Gilbert, 2010). A separate strand of literature, mostly focused on nonhuman animals, has explored the 'replay' of memories. Intriguingly, this literature indicates a similar tendency: animals tend to replay particular memories in proportion to potential gain, and that this process may support future decision-making (see Mattar & Daw, 2018). Future work should explore the potential connection between these two strands of research and the patterns of offline simulation observed here.

These results are also relevant to previous work on people's judgements in moral situations. Existing research suggests that moral judgments can impact people's intuitions about causation, intentional action, and a variety of other apparently non-moral issues (Knobe, 2010). One hypothesis about these effects is that they are explained by a tendency to simulate counterfactuals in which agents perform actions that are morally good, and not to simulate counterfactuals in which agents perform actions that are morally bad (e.g. Icard, Kominsky, & Knobe, 2017; Phillips, Luguri, & Knobe, 2015). Future research could ask whether this tendency is best understood as just another manifestation of the same basic pattern observed in the present studies.

The principal contribution of the present studies is its suggestion that our cognition is particularly attuned to the best possible options, regardless of whether there is an immediate decision that has to be made. One possibility is that our minds are simply wired to default to simulating the good possibilities during offline cognition and that people will therefore show this tendency even when they do not want to be thinking of the good (as when they do not want to get their hopes up), or even when it may not even be beneficial to the task to be thinking of the good (as when they need to be looking out for potential worst-case scenarios).

But another possibility is that our offline tendencies are more flexible depending on the context. Here we explored cases where people can choose which option they want, making it rational to identify the best possible ones. Yet, critical life events are often out of our control, and we can do nothing but prepare for what may come. In cases like these, we may hope for the best and prepare for the worst, making it rational to switch our offline tendencies to focus on bad outcomes to decide what to do in response. We are curious about whether people will show this same tendency when they have less control over which options they end up with, or when there is greater uncertainty about the bad outcomes. Future work can explore the boundaries of this offline tendency to imagine the good.

References

- Barron, H. C., Dolan, R. J., & Behrens, T. E. (2013). Online evaluation of novel choices by simultaneous representation of multiple memories. *Nature neuroscience*, 16, 1492–1498.
- Battaglia, P. W., Hamrick, J. B., & Tenenbaum, J. B. (2013). Simulation as an engine of physical scene understanding. *Proceedings of the National Academy of Sciences*, 110, 18327–18332.
- Callaway, F., Gul, S., Krueger, P. M., Griffiths, T. L., & Lieder, F. (2018). Learning to select computations. In *NIPS workshop on cognitively inspired Artificial Intelligence*.
- Callaway, F., Hamrick, J. B., & Griffiths, T. L. (2017). Discovering simple heuristics from mental simulation. In *Proceedings of the Annual Conference of the Cognitive Science Society* (Vol. 39).
- Cooper, L. A. (1975). Mental rotation of random two-dimensional shapes. *Cognitive psychology*, 7, 20–43.
- Gelly, S., & Silver, D. (2011). Monte-Carlo tree search and rapid action value estimation in computer Go. *Artificial Intelligence*, 175, 1856–1875.
- Gershman, S. J., Markman, A. B., & Otto, A. R. (2014). Retrospective revaluation in sequential decision making: A tale of two systems. *Journal of Experimental Psychology: General*, 143, 182–194.
- Gerstenberg, T., Goodman, N., Lagnado, D., & Tenenbaum, J. (2014). From counterfactual simulation to causal judgment. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 36).
- Hamrick, J. B., & Griffiths, T. L. (2014). What to simulate? Inferring the right direction for mental rotation. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 36).
- Hamrick, J. B., Pascanu, R., Vinyals, O., Ballard, A., Heess, N., & Battaglia, P. (2016). Imagination-based decision making with physical models in deep neural networks. In *NIPS 2016 workshop on intuitive physics*.
- Hamrick, J. B., Smith, K. A., Griffiths, T. L., & Vul, E. (2015). Think again? The amount of mental simula-

- tion tracks uncertainty in the outcome. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 37).
- Icard, T. F., Cushman, F., & Knobe, J. (2018). On the instrumental value of hypothetical and counterfactual thought. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 40).
- Icard, T. F., Kominsky, J. F., & Knobe, J. (2017). Normality and actual causal strength. *Cognition*, 161, 80–93.
- Killingsworth, M. A., & Gilbert, D. T. (2010). A wandering mind is an unhappy mind. *Science*, 330, 932–932.
- Knobe, J. (2010). Person as scientist, person as moralist. *Behavioral and Brain Sciences*, 33, 315–329.
- Lieder, F., Griffiths, T. L., & Hsu, M. (2018). Overrepresentation of extreme events in decision making reflects rational use of cognitive resources. *Psychological review*, 125, 1–32.
- Mason, M. F., Norton, M. I., Van Horn, J. D., Wegner, D. M., Grafton, S. T., & Macrae, C. N. (2007). Wandering minds: the default network and stimulus-independent thought. *Science*, 315, 393–395.
- Mattar, M. G., & Daw, N. D. (2018). Prioritized memory access explains planning and hippocampal replay. *Nature Neuroscience*, 21, 1609–1617.
- Peirce, J. W. (2007). PsychoPyPsychophysics software in Python. *Journal of neuroscience methods*, 162, 8–13.
- Phillips, J., Luguri, J. B., & Knobe, J. (2015). Unifying moralities influence on non-moral judgments: The relevance of alternative possibilities. *Cognition*, 145, 30–42.
- Shepard, R. N., & Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science*, 171, 701–703.
- Srivastava, N., Miller-Trede, J., Schrater, P. R., & Vul, E. (2016). Modeling sampling duration in decisions from experience. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 38).
- Vul, E., Goodman, N., Griffiths, T. L., & Tenenbaum, J. B. (2014). One and done? Optimal decisions from very few samples. *Cognitive science*, 38, 599–637.
- Wimmer, G. E., & Shohamy, D. (2012). Preference by association: how memory mechanisms in the hippocampus bias decisions. *Science*, 338, 270–273.