

The Naïve Utility Calculus unifies spatial and statistical routes to preference

Julian Jara-Ettinger*, Felix Sun*, Laura Schulz, & Joshua B. Tenenbaum

Department of Brain & Cognitive Sciences, MIT, Cambridge, MA, 02139

* These authors made an equal contribution.

Abstract

Humans can seamlessly infer what other people like, based on what they do. Broadly, two types of accounts have been proposed to explain different aspects of this ability. A first account focuses on inferences from spatial information: agents choose and move towards things they like. A second account focuses on inferences from statistical information: uncommon choices reveal preferences more clearly compared to common choices. Here we argue that these two kinds of inferences can be explained by the assumption that agents maximize utilities. We test this idea in a task where adult participants infer an agent's preferences using a combination of spatial and statistical information. We show that our model predicts human answers with higher accuracy than a set of plausible alternative models.

Keywords: Computational modeling; Naïve Utility Calculus; Theory of mind; Social cognition.

Introduction

As humans, we understand that other people have minds, and we can infer what they know and what they want by watching their behavior. Imagine, for instance, that a man walks towards a cookie jar, opens it, peeks in, and then closes it again. Although we cannot see the inside of the man's mind or of the cookie jar, we nevertheless suspect that the man likes cookies, that he planned to eat a cookie, that he believed there were cookies in the cookie jar, and that he was wrong: the cookie jar was empty.

Our ability to infer other people's preferences, in the service of interpreting their actions and predicting their future behavior, is at the heart of this ability. A large body of work suggests that preference inferences rely on spatial information. When we watch an agent navigate, a first focus is on the path's *end state*: Agents navigate to complete goals that fulfill their desires (Woodward, 1998). A second focus is on the path's *directedness*: We expect agents to navigate efficiently, and we use this expectation to attribute goals (Gergely & Csibra, 2003). Thus, if an agent does not take the shortest path towards a goal this implies there is a constraint in the way (Csibra, Biró, Koós, & Gergely, 2003), a subgoal that the agent completed within the path (Baker, Saxe, & Tenenbaum, 2009), or that the actions themselves are the goal (Schachner & Carey, 2013).

When we infer preferences, however, we not only rely on what agents choose; we also take into account what they don't choose. Suppose that an agent can pick a fruit from a bag filled with a hundred apples and one orange. If the agent takes an apple, she doesn't necessarily like them better than oranges. But if she takes the only orange, then she probably likes them better than apples. Intuitively, the second situation reveals a stronger preference, even though the

agent could have chosen either fruit in both cases. In other words, the strength of the preference inference depends on the statistical information of the possible choices.

The ability to infer preferences using spatial and statistical information are both at work from early in life. Infants as young as three months old expect agents to navigate efficiently to some extent (Skerry, Carey, & Spelke, 2013) and show a robust expectation by their first birthday (Gergely & Csibra, 2003). Similarly, the ability to draw inferences from statistical information has its roots in infancy and it plays a role in how we learn what other people like (Kushnir, Xu, & Wellman, 2010; Wellman, Kushnir, Xu, & Brink, 2016), how we learn about the world (Gweon, Tenenbaum, & Schulz, 2010), and even how we learn the meaning of new words (Xu & Tenenbaum, 2007).

Together, these two lines of evidence suggest a dual system for inferring preferences: one that relies on spatial information, and one that relies on statistical information. But real-world situations do not break down so cleanly. Agents usually combine both spatial and statistical distributions of potential rewards in their environment, and so should our judgments about their preferences from observing their actions.

Here we propose that, rather than being supported two systems of knowledge, preference inferences from spatial and statistical information are derived from a single intuitive theory of agents: the naïve utility calculus (Jara-Ettinger, Gweon, Tenenbaum, & Schulz, 2015; Jara-Ettinger, Tenenbaum, & Schulz, 2015). Critically, our goal here is not to compare the naïve utility calculus with formal theories of decision-making, but with other theories of intuitive decision-making. Here we show how the naïve utility calculus (NUC) supports inferences from spatial and statistical information. We test our proposal by implementing and comparing a spatial inference model, a statistical inference model, and a NUC model against adult performance on a preference-inference task. We end by discussing the implications of our findings on understanding the development of commonsense psychology.

The Naïve Utility Calculus

A growing set of studies suggests that humans reason about agents in terms of utility maximization (Jara-Ettinger et al., 20215; Jern et al, 2011; Johnson & Rips, 2015; Lucas et al, 2014). Specifically, humans have an intuitive theory of how utilities are comprised of costs and rewards, and how, together, they guide what others do. According to this *Naïve Utility Calculus*, agents act by estimating the costs and rewards associated with each possible plan, and by selecting the plan with the highest utility (the difference between rewards and costs). That is, when people watch an agent,

they assume that her behavior yielded high utilities, and they use this assumption to infer the agent's competence (her costs) and her motivation (her rewards). To illustrate inferences through the NUC, consider an agent who chooses an apple over an orange. This implies that the utility for the apple, $U(a)$, is higher or equal than the utility for the orange, $U(o)$. By decomposing each utility into its costs ($C(a)$ and $C(o)$, respectively) and rewards ($R(a)$ and $R(o)$, respectively), the agent's choice implies that $R(a)-C(a) \geq R(o)-C(o)$. If both fruits were equally easy to get, then $C(a)=C(o)$ and, therefore, $R(a) \geq R(o)$. That is, when two options are matched for costs, agents choose what they like best. Suppose instead that the apple, the agent's choice, was more costly to get than the orange. Because $R(a)-C(a) \geq R(o)-C(o)$ and $C(a)>C(o)$, then $R(a)>R(o)$. That is, agents unambiguously reveal their preferences when they choose the more costly option. Last, if the apple was easier to get, then $R(a)-C(a) \geq R(o)-C(o)$, and $C(a)<C(o)$. Under these circumstances, $U(a)$ ($R(a)-C(a)$) may be higher than $U(o)$ ($R(o)-C(o)$) because the apple's reward ($R(a)$) was high, or because the orange's cost ($C(o)$) was high. Thus, when agents choose low cost options their preferences are not revealed.

Although its developmental origins are unclear, the NUC is at work from early childhood, supporting fundamental inferences by age five and with some aspects already at work by age two (Jara-Ettinger, et al. 2015).

Inferences from spatial information

The NUC explains why humans are sensitive to spatial information. Suppose an agent takes a sequence of actions to complete a goal. If the agent maximized utilities, then two things must be true. First, the reward must outweigh the costs. Otherwise, the plan's utility would be negative and the agent could obtain a final higher utility by not acting at all. Second, the agent must be minimizing costs: the smaller the costs the agent incurs, the higher the utility she obtains. In spatial contexts, cost minimization reduces to efficient navigation. Thus, expecting agents to maximize utilities implies that a path's *directedness* and *end state* can help reveal preferences. If, however, humans do so through a naïve utility calculus, then, as the example above reveals (see apple-orange example), humans should also be sensitive to a third feature of spatial navigation: its *cost*.

Inferences from statistical information

Inferences from statistical information ultimately rely on the assumption that rare choices reveal stronger preferences. Although intuitive, the causes underlying this assumption are unclear. The NUC, however, naturally produces this expectation. Suppose that an agent can take any object from a box. If she doesn't have a preference, then taking whichever object is easiest to get maximizes her utilities (if all objects have the same reward, the option with the lowest cost yields the highest utility). In contrast, if the agent prefers one type of object to the others, then she will have to incur a higher cost in terms of time, effort, attention, and distance to locate the object of the desired category and to retrieve it. The less common an object's category is, the

higher the cost the agent must incur to locate it and obtain it. Thus, retrieving rare objects suggests that the agent incurred a higher cost, and, if agents maximize utilities, this cost is only warranted if the reward associated with the rare object is higher than the reward associated with more common objects.

Computational modeling

If humans infer preferences using their NUC, then a formal implementation should quantitatively predict adult preference judgments. Alternatively, if humans infer preferences through simpler ways, then simpler models should predict human inferences with equal or better accuracy. To test if participants integrate spatial and statistical information through the assumption of utility maximization we ran a preference-inference task with adult participants and we compared their performance to five computational models: our full Naïve utility calculus model and two NUC lesioned models, as well as two alternative models. These models are based on the proposals of how infants infer preferences from spatial information (spatial model) and from statistical information (statistical model). Next, to test if participants infer these preferences in a Bayesian way, we compared participants' self-reported confidence judgments with estimates from each model.

Alternative models

Spatial model The spatial model formalizes the proposal that goals directly reveal preferences. As such, it uses a limited source of spatial information: the end state. This model assumes that the distribution of choices an agent makes matches her underlying preferences. For instance, if an agent collects two red objects and one blue object, then the reward for collecting a red object is $R_{red} = 2R/3$ and the reward for collecting a blue object is $R_{blue} = R/3$, where R is a constant set to 1 (changing the value of R does not change our results as model comparison was done by z-scoring model predictions. See Results).

Statistical model The statistical model is based on proposals for how people infer preferences by relying on statistical information (Gweon, et al, 2010; Xu et al, 2007). These models were formulated in simpler domains than the one we test in our experiment so we extended them to fit our experimental design. Our model assumes that, before taking any actions, the agent goes through a decision making process to find objects that will give her high rewards. Specifically, the statistical model assumes that the agent considers one object at a time (at random) and decides whether to collect it or not based on its reward. That is, when the agent considers taking an object from category k , she selects it with a probability proportional to its reward. These assumptions imply that more common object categories are more likely to be considered and that more rewarding kinds of objects are more likely to be collected, once the agent considers them. As such, selecting an object from a rare category suggests that the agent prefers it to more common objects that the agent likely considered collecting first. In this model, the observer assumes that the agent considers each object with uniform probability (if

there are n objects, the agent considers each object with probability $1/n$.

Given the theory of how an agent chooses what to collect, we use Bayesian inference to recover the agent's preferences given her choices. Specifically, because in our experiment we use two types of objects (see Stimuli), we use Bayes' rule to estimate the relative magnitude of one reward type over the other (with 0 indicating that the first category contains all the rewards, 0.5 indicating that both categories are equally rewarding, and 1 indicating that the second category contains all the rewards), using a uniform prior.

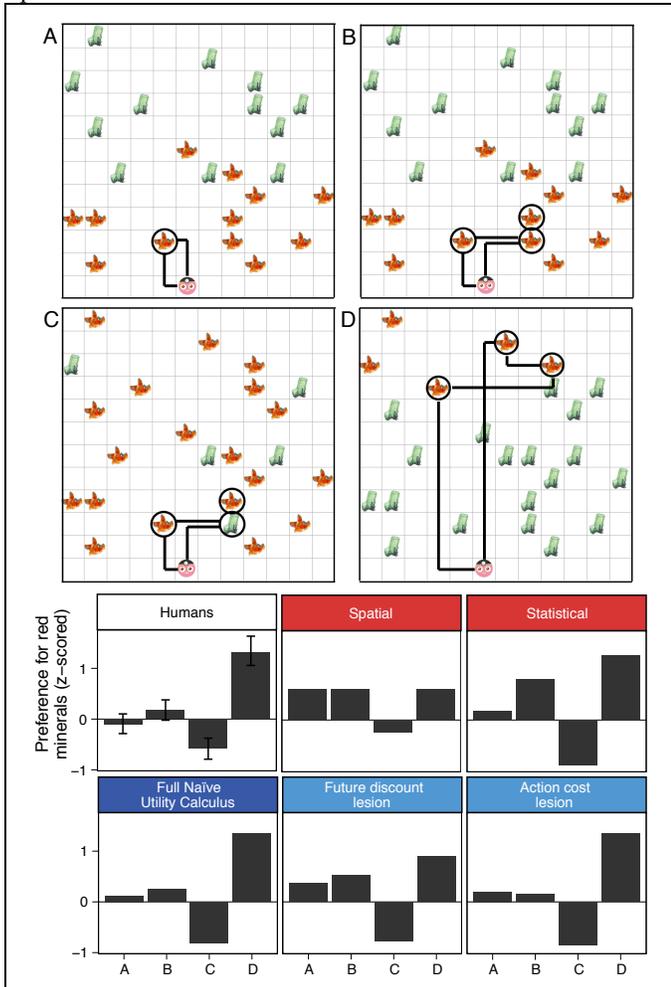


Figure 1. Stimuli examples along with participant judgments and model predictions. Because judgments are z-scored, positive and negative values are relative to the average preference inference and do not correspond to preferring red or green minerals, respectively. The negative-valued prediction in this plot indicates a weaker preference for red minerals compared to the average inference.

Naïve Utility Calculus models

The last three models are implementations of the naïve utility calculus (NUC), but they integrate costs in different ways, enabling us to understand how humans may reason about costs, rewards, and utility maximization. All models are formulated as generative models that predict agent choices given their preferences, and the inference from

choices to preferences is done through Bayes' rule with a uniform prior over the possible distribution of rewards over the object categories.

Full Naïve Utility Calculus The full NUC model assumes that agents maximize utilities. Costs are function of the number of actions the agent takes (set to be a constant cost per action = 0.01; our conclusions are robust to parameter changes) and rewards are exponentially discounted over time. Intuitively, the future discount corresponds to the assumption that the longer an agent takes to reach a reward, the less likely the reward will still be there. Thus, this model relies on spatial information in three ways: first, it expects agents to navigate efficiently because smaller sequences of actions incur fewer costs (minimizing costs), and because collecting objects faster results in higher rewards (maximizing exponentially discounted rewards); second it assumes that the agent's goals have sources of rewards; and last, and in contrast to the alternative models, it assumes that longer distances reveal stronger preferences. More formally, the cost of actions and rewards in objects are integrated into a utility function ($U=R-C$) and the utility-maximizing actions are derived through a Markov Decision Process. Further details about the computational implementations of the naïve utility calculus can be found in Jara-Ettinger et al (2015).

Future-discount lesion The future-discount lesion is identical to the NUC model but rewards aren't discounted over time. Thus, this model integrates statistical information in a full manner, and spatial information in a simplified manner. The model expects agents to navigate efficiently only because lower costs lead to higher utilities, but not because longer distances increases the chance of losing the target reward.

Action cost lesion Conversely, the action cost lesion model is identical to the NUC but it ignores action costs. Nevertheless, the model assumes that the agent's rewards are discounted over time. This model therefore integrates spatial information through the expectation that agents act efficiently because the longer it takes them to reach a reward, the less likely it will still be there when they arrive.

Experiment

To test our models, we designed a simple task where participants watched a miner collect minerals in mines with variable distributions of minerals.

Stimuli

Figure 1 shows examples of the stimuli. Each stimulus consisted of an animated display of an agent (the miner) entering a mine (a 12x12 grid world) and collecting green and/or red minerals. Each map contained 24 minerals in the same locations (which were chosen at random and kept constant across stimuli), but the proportion and the distribution of these minerals varied. The proportion varied according to three levels: more green than red (20 green and 4 red), more red than green (4 green and 20 red), or an equal number of each (12 of each). The distributions of these minerals varied according to three levels: red minerals closer, green minerals closer, or all minerals intermixed.

This generated a total of nine different maps. By varying the proportion of the objects, we can test how statistical information influences preference inferences; by varying the location of the objects we can test how spatial information influences preference inferences.

The miner’s paths were obtained by computing the shortest path an agent would need to take to collect all minerals of one kind, or to collect the closest minerals (which could be a combination of red and green minerals). These paths were generated in accordance to three conditions. In the first condition, the miner collected one mineral and exited the mine. In the second condition, the miner collected three minerals in a single trip and then exited the mine. And in the last condition the miner collected three minerals, but had to return to the mine’s exit after collecting each object. Thus, the first and second conditions test how the amount of data an observer receives influences observers’ inferences, and the second and third conditions together test how the costs of collecting the minerals influence observers’ inferences. The combination of the two agent types (strong preference or no preference) with the nine maps produced a total of 18 test paths per condition.

Participants

90 U.S. residents (as determined by their IP address) were recruited and tested through Amazon’s Mechanical Turk platform (Mean age = 33 years. Range = 20 - 59 years).

Procedure

Participants were randomly assigned to the one mineral condition, to the three minerals in one trip condition, or to the three minerals in three trips condition (N = 30 participants per condition). Thus, each participant only completed one-third of the trials. Participants first completed a brief tutorial that explained the task. Next, participants completed a questionnaire with three questions to ensure they understood the task. Participants who responded all questions correctly were given access to the experiment, and participants who made at least one error were redirected to the beginning of the tutorial.

In the test stage, participants saw an animated display of the miner collecting the minerals and had to respond four questions. The first two questions were multiple choice control questions asking about the proportion and distribution of the minerals. Participants who answered these questions incorrectly were asked to re-examine the stimulus. The third question asked participants to rate the miner’s preference using a slider that ranged from “Red is much more valuable” (coded as a 0) to “Green is much more valuable” (coded as a 1). The last question asked participants to rate their confidence in the preference judgment using a slider that ranged from “Not at all” (coded as a 0) to “Extremely confident” (coded as a 1).

Results

Figure 2 shows the results from the experiment. As expected, the formalizations of spatial and statistical accounts matched the qualitative pattern of participant

judgments: they predicted strong and weak preferences accurately. However, as Figure 2 shows, the NUC model captured human judgments with higher precision. To evaluate model performance more precisely we computed each model’s correlation with average human judgments (z-scored within each participant and averaged; see Table 1).

Model	Correlation (95% CI)
Spatial	.84 (0.79,0.92)
Statistical	.81 (0.74,0.90)
Naïve Utility Calculus	.97 (0.96,0.98)
Future-discount lesion	.93 (0.90,0.96)
Action cost lesion	.96 (0.92,0.97)

Table 1. Model correlations with participant responses along with 95% bootstrapped confidence intervals.

Comparison with alternative models

Overall, the NUC model had the highest correlation ($r=0.97$) between its predictions and participant responses. To evaluate this correlation we bootstrapped the correlation difference between the NUC and the alternative models. The NUC reliably outperformed the spatial model (correlation difference=0.12; 95% CI=(0.03,0.18)) and the statistical model (correlation difference=0.16; 95% CI=(0.05,0.22)).

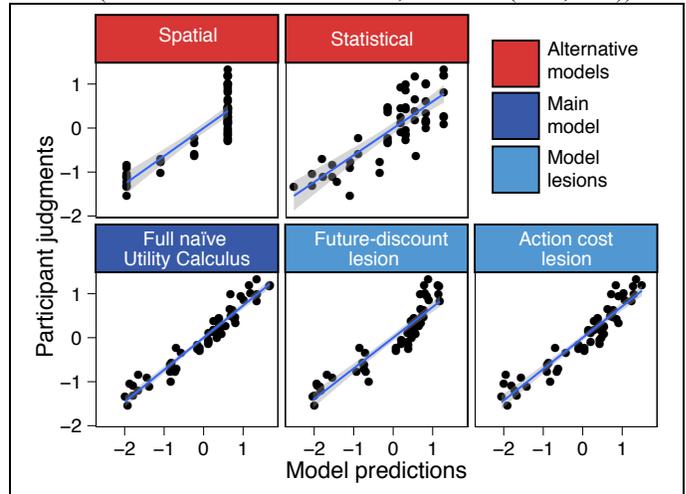


Figure 2. Experiment results. In each plot, each black dot represents a stimulus. The x-axis shows the model’s prediction (z-scored) and the y-axis shows average participant judgments (z-scored within each participant and averaged).

Figure 1 shows four example trials that reveal how the NUC outperforms the alternative models. The spatial model fails to capture differences between trials A, B, and C, as it is not sensitive to the amount of evidence. The statistical model roughly captures human responses, but it attributes a stronger preference to the miner in trial B, as it neglects the spatial distribution. In contrast, the NUC models show sensitivity to the amount of data, the spatial information, and the statistical information.

Comparison with model lesions

Both model lesions had a lower correlation with participant judgments compared to the full NUC model (see Table 1). Removing the future-discount parameter led to a significant decrease in the model’s correlation with human

judgments (correlation difference = .042; 95% CI = (0.004,0.072)). This suggests that participants are sensitive to an exponential discounting of the mineral rewards over the length of the miner’s trajectory. Similarly, removing the cost of travelling decreased the model’s correlation with human judgments (difference = .022; 95% CI = (-0.007,0.047)). However, 13% of the mass of the 95% confidence interval was on the negative region. This suggests that integrating a linear cost over the future-discount may better fit human judgments, but the results are inconclusive.

Confidence judgments

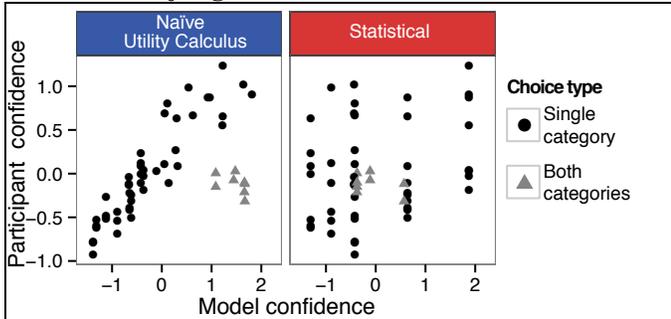


Figure 3. Confidence judgments. The models’ confidence ratings were obtained by computing the standard deviation of the posterior distribution of each stimulus, multiplying them by -1 (so as to match the qualitative order in participant judgments), and then z-scoring the values. Participant confidence judgments were z-scored within participant and averaged.

Our evidence so far suggests that humans infer preferences through the assumption of utility-maximization. Nevertheless, this inference is not necessarily Bayesian. Participants may, for instance, approximate the responses from a Bayesian models through simpler heuristics. To explore this possibility, we asked participants to report confidence judgments on each trial (see Methods section) and we compared them with a rough measure of each model’s uncertainty: the posterior distribution’s standard deviation. If participants are inferring preferences in a probabilistic manner, then the NUC’s standard deviation should correlate with participant confidence judgments. However, if participants infer preferences through some heuristics that approximate Bayesian inference, then their confidence should not necessarily be related to the one in our model. Moreover, the statistical model, being Bayesian, also produces confidence judgments (the spatial model generates a single inferred estimate with full confidence), enabling us to further test its validity.

Figure 3 shows each model’s negative standard deviation along with participants’ confidence judgments. Although the alternative models all captured preference inferences in a coarse way (see Figure 2), their measures of confidence did not resemble participant’s confidence judgments (see Figure 3). In contrast, the NUC model and its lesions predicted with far higher accuracy participants’ confidence judgments. Table 2 shows the correlations and confidence intervals. Although the NUC’s correlations were reliably greater than 0, Figure 3 reveals that it failed to capture the variation in a

small set of stimuli (the results were qualitatively identical for the NUC model lesions). Post-hoc inspection of these outliers revealed that they were all cases where the miner had selected a combination of red and green minerals (because of the way we generated the stimuli, the miner only took a combination of red and green minerals whenever these were the closest and the agent had no preference; see Stimuli section). Consistent with this, we found that when we decomposed the stimuli into trials where the agent collected only one type of mineral (Single category), the NUC model and its lesions showed high correlations and performed roughly as well. In contrast, in the stimuli where the agent collected various kinds of minerals (Both categories), none of the models predicted human confidence judgments (see Table 2). Nevertheless, it is important to note that this subset of stimuli consists of seven data points, making it difficult to draw conclusions from the correlations.

Model	Correlation (95% CI)	Single category correlation	Both categories correlation
Statistical	0.28 (0.04,0.56)	0.29 (0.05,0.57)	-0.43 (-1,0.01)
Naive Utility Calculus	0.65 (0.49,0.83)	0.91 (0.88,0.95)	-0.45 (-1,-0.04)
Future-discount lesion	0.33 (0.12,0.51)	0.84 (0.79,0.89)	-0.45 (-1,-0.04)
Action cost lesion	0.68 (0.53,0.86)	0.91 (0.88,0.96)	-0.32 (-0.98,0.18)

Table 2. Correlation between the standard deviation of the model’s posterior distribution and participant confidence judgments, along with 95% bootstrapped confidence intervals. The first column shows the overall correlations, and the last two columns show the correlations after splitting the stimuli into the group where the miner only collected one type of mineral (single category) and when the miner collected a combination of red and green minerals (both categories). The spatial model is not presented as it only produces a point estimate rather than a probability distribution.

Discussion

Here we reviewed evidence that, from early in life, humans can infer preferences using statistical and spatial information, and we proposed that these two types of inferences are driven by the naive utility calculus (NUC) – our intuitive theory of how agents select their goals by estimating and maximizing utilities. We tested our proposal by implementing a formal model of the naive utility calculus and comparing it to other accounts that rely on spatial and statistical information separately. Our results show that adults were both sensitive to the spatial and statistical information of an agent’s behavior, and that this variation was best captured by the NUC model.

Critically, all accounts fit participant judgments qualitatively. Thus, implementing formal computational models was critical for generating precise predictions and assessing whether they explained variation in human judgments in a fine-grained manner. Our results show that

the NUC model significantly outperformed the alternative models at a detailed level.

In order to better understand the NUC's performance we implemented two model lesions. In one model lesion we removed the future discount parameter (future-discount lesion) and in the second model lesion we removed the cost for traveling (action cost lesion). Critically, both model lesions were still sensitive to the statistical information, and they both expected the agent to navigate efficiently. The NUC correlated with human responses better than both of the model lesions, but this difference was only reliable when comparing the NUC model with the cost sensitive lesion and not when comparing it with the action cost lesion. Our results suggest that a non-linear reward discount is critical for how humans reason about efficiency. However, once a model integrates a future-discount parameter, adding a cost of traveling only produces a modest improvement.

Although the alternative models roughly predicted human responses, a comparison of the models' posterior standard deviation (a measure of the model's uncertainty) against participant confidence judgments revealed strong discrepancies. In contrast, the NUC and its lesions predicted our participant's confidence judgments for a large set of stimuli (see Figure 3 and Table 2, columns 1 and 2). Nevertheless, all models failed to capture human confidence judgments in the trials where the miner collected a combination of red and green minerals closest to the mine's entrance (see last column of Table 2). In these situations, the NUC models were confident that the miner liked both minerals roughly as much, and that she was therefore collecting the closest ones. Participants made similar judgments, but they were less confident. One possible explanation for this discrepancy is that our model assumes that the cost for traveling is fixed and observable, whereas participants may not. Instead, participants may be uncertain about how exhausting it is to travel the mine, and this may lead to a confound in the miner's behavior: she might be taking the closest minerals because she likes all minerals just as much, or because she finds traveling deep into the mine to be very costly. A richer version of the NUC that integrates uncertainty over the costs and rewards is needed to evaluate this possibility.

Altogether, our results show that the NUC explains why and how humans rely on spatial and statistical information when inferring preferences. Empirical results show that the ability to infer preferences from spatial information and from statistical information arises in early childhood (Gweon et al., 2010; Gergely & Csibra, 2003). However, these sources of information have been studied separately, and different accounts have been proposed to explain how we draw these inferences. Our finding that inferences from statistical and spatial inferences are unified in adults raises two hypotheses about the development of this reasoning. A first possibility is that the NUC is already at work in infancy. If so, infants may use it to solve tasks involving spatial information (e.g., Gergely & Csibra, 2003), and tasks involving statistical information (e.g., Kushnir, et al, 2010).

A second possibility, however, is that the NUC emerges later in life. Under this account, infants must rely on simpler expectations about agents to reason about spatial and statistical information (perhaps driven by two separate systems of understanding; a spatial one and a statistical one). If this is true, then the proposed explanations for how infants use spatial and statistical information (formalized in the spatial and the statistical models), may be correct and serve as the bedrock for a richer unifying intuitive theory: the naïve utility calculus.

Acknowledgments

This material is based upon work supported by the Center for Brains, Minds, and Machines (CBMM), funded by NSF-STC award CCF-1231216, and by the Simons Center for the Social Brain.

References

- Baker, C. L., Saxe, R., & Tenenbaum, J. B. (2009). Action understanding as inverse planning. *Cognition*, *113*(3), 329-349.
- Csibra, G., Biró, S., Koós, O., & Gergely, G. (2003). One-year-old infants use teleological representations of actions productively. *Cognitive Science*, *27*(1), 111-133.
- Gergely, G., & Csibra, G. (2003). Teleological reasoning in infancy: The naïve theory of rational action. *TiCS*, *7*.
- Gweon, H., Tenenbaum, J. B., & Schulz, L. E. (2010). Infants consider both the sample and the sampling process in inductive generalization. *PNAS*, *107*(20).
- Jara-Ettinger, J., Schulz, L. E. & Tenenbaum, J. B. (2015). The naïve utility calculus: Joint inferences about the costs and rewards of actions. *Proceedings of the 37th Annual Conference of the Cognitive Science Society*.
- Jara-Ettinger, J., Gweon, H., Tenenbaum, J. B., & Schulz, L. E. (2015). Children's understanding of the costs and rewards underlying rational action. *Cognition*, *140*, 14-23.
- Jara-Ettinger, J., Tenenbaum, J. B., & Schulz, L. E. (2015). Not So Innocent Toddlers' Inferences About Costs and Culpability. *Psychological science*, *26*(5), 633-640.
- Jern, A., Lucas, C. G., & Kemp, C. (2011). Evaluating the inverse decision-making approach to preference learning. In *NIPS*.
- Johnson, S. G., & Rips, L. J. (2015). Do the right thing: The assumption of optimality in lay decision theory and causal judgment. *Cognitive psychology*, *77*, 42-76.
- Kushnir, T., Xu, F., & Wellman, H. M. (2010). Young children use statistical sampling to infer the preferences of other people. *Psychological Science*, *21*(8), 1134-1140.
- Lucas, Christopher G., Thomas L. Griffiths, Fei Xu, Christine Fawcett, Alison Gopnik, Tamar Kushnir, Lori Markson, and Jane Hu. "The child as econometrician: A rational model of preference understanding in children." *PloS one*.
- Schachner, A., & Carey, S. (2013). Reasoning about 'irrational' actions: When intentional movements cannot be explained, the movements themselves are seen as the goal. *Cognition*.
- Skerry, A. E., Carey, S. E., & Spelke, E. S. (2013). First-person action experience reveals sensitivity to action efficiency in pre-reaching infants. *PNAS*, *110*(46).
- Xu, F., & Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological review*, *114*(2), 245.
- Wellman, H. M., Kushnir, T., Xu, F., & Brink, K. (2016). Infants use statistical sampling to understand the psychological world. *Infancy*.
- Woodward, A. L. (1998). Infants selectively encode the goal object of an actor's reach. *Cognition*, *69*(1), 1-34.