Plans or Outcomes: How do we attribute intelligence to others?

Marta Kryven ^{*1}, Tomer D. Ullman ^{†2}, William Cowan ^{‡3} and Joshua B. Tenenbaum ^{§1}

¹Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology

²Department of Psychology, Harvard University

³Department of Computer Science, University of Waterloo

keywords: intelligence attribution, planning, social perception, theory of mind

Abstract

Humans routinely make inferences about both the contents and the workings of other minds based on observed actions. People consider what others want or know, but also how intelligent, rational, or attentive they might be. Here, we introduce a new methodology for quantitatively studying the mechanisms people use to attribute intelligence to others based on their behavior. We focus on two key judgments previously proposed in the literature: judgments based on observed outcomes (you're smart if you won the game), and judgments based on evaluating the quality of an agent's planning that led to their outcomes (you're smart if you made the right choice, even if you didn't succeed). We present a novel task, the Maze Search Task (MST), in which participants rate the intelligence of agents searching a maze for a hidden goal. We model outcome-based attributions based on the observed utility of the agent upon achieving a goal, with higher utilities indicating

^{*}(mkryven@mit.edu)

[†]tullman@fas.harvard.edu

[‡]wmcowan@cgl.uwaterloo.ca

[§]jbt@mit.edu

higher intelligence, and model planning-based attributions by measuring the proximity of the observed actions to an ideal planner, such that agents who produce closer approximations of optimal plans are seen as more intelligent. We examine human attributions of intelligence in 3 experiments that use MST, and find that participants used both outcome and planning as indicators of intelligence. However, observing outcome was not necessary, and participants still made planning-based attributions of intelligence when the outcome was not observed. We also found that the weights individuals placed on plans and on outcome correlated with an individual's ability to engage in cognitive reflection. Our results suggest that people attribute intelligence based on plans given sufficient context and cognitive resources, and rely on outcome when computational resources or context are limited.

1 Introduction

We know intelligent behavior when we see it. When we see a child building a tower of books to reach a shelf, a computer solving a challenging game-board, or a cat sneaking up on a bird, we intuitively evaluate the quality of such behaviors, and label them as intelligent or unintelligent. Humans attribute intelligence in a remarkably broad way, including to individuals (Kepka & Brickman, 1971), groups (Mackie et al., 1990), other species (Amodio et al., 2018; Godfrey-Smith, 2017; Weisman et al., 2016), machines (Turing, 1950), and other autonomous agents (Lee, 2020; Legg & Hutter, 2007). However, the mental computations underlying such attributions remain poorly understood.

Attributions of dispositions are traditionally studied along the dimensions of warmth and competence (Asch, 1946; Wojciszke et al., 1998), and the study of intelligence attribution partly stems from the literature on attributing competence (Kun & Weiner, 1973; Mackie et al., 1990; Reeder et al., 1992; Surber, 1984). The traits *intelligent* and *unintelligent* describe the high and low points of the competence dimension (Fiske et al., 2007). Yet, intelligence differs from competence in important ways.

Intelligence is commonly understood as the ability to think rationally, learn, and use one’s abilities to maximize goals, which implies the use of mental computations to direct behavior (Legg & Hutter, 2007). While the attribution of competence traits can be informed by perceptual cues (Ambady & Rosenthal, 1993; Todorov et al., 2005), including anthropomorphism (Epley et al., 2007) or prior beliefs (Weisman et al., 2016), here we focus on a generalizable mechanism of attributing intelligence to an agent’s behavior, which is likely to be expressed as rational computational principles.

In this work, we propose a formal framework for intelligence attribution, based on the Naïve Utility Calculus. The Naive Utility Calculus posits that people interpret behaviors as *approximately rational* (Dennett, 1989). Rationality in this context means maximizing rewards while minimizing costs, given one’s knowledge and observations (Jara-Ettinger et al., 2016). The rationality assumption is flexible about *how efficiently* an agent is directing its behavior – as long as the agent expresses a preference for actions with higher utility. The Naïve Utility Calculus does not require that the agents being observed actually compute and maximize exact expected utilities, only that the assumption of rationality is useful to the observer attributing mental states to those agents.

The intuitive interpretation of behavior as approximately rational comes up in various areas of cognition, including development (Gergely & Csibra, 2003), visual perception (Gao & Scholl, 2011), language processing (Grice et al., 1975), theory of mind (Baker et al., 2017), and attribution of risk aversion (Liu et al., 2017). Previous work showed how the principle of rationality can be used to quantitatively predict human inferences of hidden mental states such as intentions, beliefs, and goals (e.g. Baker & Tenenbaum, 2014; Saxe et al., 2004; Ullman et al., 2009). However, these studies have not focused on the efficiency of reasoning that motivates behavior.

Here we propose that attributions of intelligence could reflect the variability in agents’ approximations of rationality. For example, it seems more intelligent to stand on a tower of hardcovers than paperbacks to reach a shelf, even though both behaviors are intentional, goal-directed, and can be used to infer preferences or beliefs. So, while goals and intentions may be inferred from less than

optimal actions, the efficiency of the actions could carry an important signal about the agent's intelligence. For example, action efficiency can affect attributions of responsibility and blame – more efficient agents are held more responsible for their actions, and are seen as expressing a stronger intent (De Freitas & Johnson, 2018; Johnson & Rips, 2015), which may be partly due to a changed expectation of how well the agent will do in the future (Gerstenberg et al., 2018). Measuring intelligence as problem-solving efficiency has also been proposed as a method of evaluating intelligence in machines, and closely overlaps with how psychologists think about intelligence quotient in humans, typically measured by problem-solving tests (Legg & Hutter, 2007). We stress however that our goal is not to study or measure intelligence per se, but to formalize and quantitatively measure common-sense intuitions about it.

We consider two ways to link efficiency and attributed intelligence. One way is to consider the efficiency of the immediate *outcome* of a given behavior, compared to alternative possible outcomes, and expect it to generalize to the future (Legg & Hutter, 2007). The advantage of this method is that the observers do not need to understand how the outcome was achieved, as long as they can evaluate it. A second way is to consider the efficiency of someone's *planning*, and reason that more optimal planning will maximize future expected outcomes (Jara-Ettinger et al., 2016). This second method can be used to detect when a good outcome was achieved by accident, and to make attributions before an outcome is achieved. Both ways of attributing intelligence have their advantages, depending on context of the attribution.

Attributing intelligence to outcomes. Outcomes are a common way of assessing decisions (Baron & Hershey, 1988; Frank, 2016; Lerner, 1980; Martin & Cushman, 2016; Olson et al., 2008, 2013). Boyd (2017) builds a theory of cultural learning on imitating others with good outcomes, as a mechanism of gaining practical skills without having to understand the reasoning behind them. Given two identical medical decisions, people give higher ratings to the decision that resulted in a better patient outcome (Baron & Hershey, 1988; Caplan et al., 1991; Sacchi & Cherubini, 2004). Likewise, people judge identical ethical decisions as different if one resulted in harm (Gino et al., 2009). The emphasis on outcome is also a part of the law: attempted but unsuccessful murder is

treated more leniently than successful murders (though, obviously, there is disagreement on the wisdom of this) (Duff, 2017).

Mackie et al. (1990) tested attributions of intelligence to a group of students who took a test twice, while manipulating the outcome (pass/fail), qualifying criterion (number of problems needed to pass), and performance (number of problems solved). People who read that the students' outcome has changed attributed intelligence based on outcome to a higher extent than based on performance. This was true even if the change in outcome was due to a change in the qualifying criterion, while the number of problems solved remained the same. Notably, while people were informed about the students' outcome, they did not observe the actual behavior, making it hard to evaluate planning. For example, one could reason that failing the test resulted from poor time-management (planning), missing knowledge or skill (specific competence) or poor luck.

Outcomes carry a useful statistical cue to whether an agent has done something right, especially if they are consistent (Kepka & Brickman, 1971; Legg & Hutter, 2007). There is evidence that people disregard inconsistent outcomes when attributing ability: if prompted to consider motivation, people explain discrepant outcomes as resulting from a difference in motivation, and infer intelligence as a factor consistent between them (Kepka & Brickman, 1971). Legg & Hutter (2007) argues that machine intelligence should be evaluated by aggregating its outcomes across reasonable environments, since a universal metric of intelligence should depend only on the system's ability to adapt and maximize its goals, not on the type of hardware or implementation. The accuracy of such an evaluation should increase with observations, however this process may be difficult and time-consuming in practice.

Outcome becomes less informative when observations are sparse, and subject to stochastic effects. Frank (2016) argues that in competitive contexts people over-attribute outcomes to an agent's intrinsic qualities, such as intelligence. When a number of top contestants are equally good, performance differences between them correspond to a random process so that the person who comes first does so due to luck. Yet, common intuition has it that the winner is indeed the best (Frank, 2016).

Outcome-bias, or over-reliance of attributions on outcomes has been extensively documented in social situations where the behavior is more salient than its constraints (Heider, 1958; Gilbert & Malone, 1995).

Attributing intelligence to plans. The assumption of rationality (Dennett, 1989) suggests that agents should be seen as less intelligent when they violate this assumption. Studies of verbal reports support this view: people describe others as ‘stupid’ when observing actions that disagree with rational expectations (Aczel et al., 2015). Pantelis et al. (2016) examined intelligence attribution computationally, in a simulation of autonomous, evolving agents, competing for resources. They found that although participants’ intelligence attributions were best explained by the outcome, an agent’s evolved status also emerged as a significant, but weak predictor of attributed intelligence. Pantelis et al. (2016) emphasize that evolved agents acted in ways that were easier to interpret as goal-directed, and that behavior interpretation was connected to a concept of rationality based on the apparent fit between an agent’s actions and its environment. However, the agent models under consideration in that study did not include an analysis of an agent’s planning. This leaves open the possibility that the attributions of intelligence are sensitive to the quality of perceived planning, where the observer interprets behavior as a goal-directed planning process, and uses this inference to evaluate how well the planning maximizes the probability of achieving the agent’s goal.

An experimental framework for attributing intelligence. While most of the previous literature has focused on the importance of outcomes, we propose that both planning and outcome may be important in a framework of intelligence attribution, and take both into account in our formal modeling. To study the attribution of intelligence to behavior computationally and empirically in a way that accounts for both outcome and planning, we require an ecologically valid experimental task in which participants can observe behavior and its outcomes as well as taking part in the task themselves. Here, we introduce such an experimental design, the Maze Search Task (MST). MST is based on the classic paradigm of 2D displays with geometric agents (see e.g. Heider & Simmel, 1944). In MST agents search for a hidden exit in a maze with a known layout, which we model using rational planning with quantitatively defined utilities and outcomes. The MST can be used

in two conditions. In the MST Search Condition, participants search for a hidden exit in a maze by controlling an avatar. In the MST Attribution Condition, participants view the paths of other agents searching for the hidden exit, and rate how intelligent those agents are. The MST is a novel paradigm that supports quantitative model-based analysis of individual differences, and can help explain how humans use plans and outcomes to attribute intelligence in everyday life.

We found that attributions of intelligence in MST depended on both planning and outcome (Experiments 1, 2 and 3), with the extent to which people evaluated outcomes or plans varying between participants. Individuals consistently placed different weights on outcome and planning, in a way that depended on their cognitive reflection ability (Experiments 1 and 3). Those who scored highly on cognitive reflection attributed more intelligence based on plans and less based on outcomes. We also found that outcome was not necessary to attribute intelligence (Experiments 1 and 2). Whenever the outcome was not observed, most participants still attributed intelligence, based on planning alone. Moreover, attribution strategies were influenced by a context manipulation (Experiment 3), so that at least some individuals attributed more intelligence based on plans in MST Attribution after performing MST Search themselves. Taken together, our results suggest that people attribute intelligence based on plans given sufficient context and cognitive resources, but rely on outcome when observations or context are limited.

The remainder of this paper is organized as follows. In Section 2 we describe our computational model for assessing intelligence attributions, and the Maze Search Task (MST). Section 3 describes 3 experiments that examine people’s behavior in the MST. We conclude by discussing the possible interpretations of our results.

2 Modeling Intelligence Attribution

We model attributing intelligence to behavior by two cognitive mechanisms. First, an observer can assess *planning* by modeling an agent as implementing a rational planning procedure, and use the deviations from optimal planning to attribute a lack of intelligence, proportional to the extent of

the reward, but will appear to be unintelligent. Examples of agents with various degrees of planning efficiency are shown on Fig. 2

Figure 2: Examples of agents with various degrees of planning efficiency. The better an agent is at planning, the better it minimizes the expected effort to obtain the reward in the long run. The agent on the left is the most efficient while the agent on the right is the least efficient.

Since the location of the reward is initially unknown to the agent, an agent might be lucky and find it quickly, or it might be unlucky, and have to search the maze exhaustively. While all agents encounter a variety of lucky and unlucky outcomes, agents who are better at planning should on average take shorter paths. Seeing one trial in isolation, observers of the MST have no access to such average outcomes, but could infer how well this agent might do in the future based on the example they see. Such an inference can be based on an evaluation of planning, or it can be based on outcome. Going by planning, an agent is intelligent if it rationally plans its path, regardless of how long it took. Going by outcome, if an agent finds the exit quickly it must be intelligent, regardless of the actions it took.

2.2 The Ideal Planner

We first define formally an ideal rational planner, a model of an agent who maximizes rewards while minimizing costs, as described by the ‘Bayesian Theory of Mind’ (BToM) formalizing naïve utility calculus. Given an agent’s observed actions, a BToM observer views agents as implementing a planning procedure and forms a posterior probability distribution over hidden variables (e.g. goals, beliefs, constraints) that drive it. This approach has been successful in modeling a variety

of reasoning in intuitive psychology (e.g. Baker et al., 2011, 2017; Baker & Tenenbaum, 2014; Jara-Ettinger et al., 2016; Liu et al., 2017; Ullman et al., 2009). Building on BToM methodology, our ideal planner calculates the expected utilities of possible search trajectories as minimizing the expected number of steps to reach the exit. We model attributed intelligence as an evaluation of the optimality of this planning mechanism.

In the general case, the MST Search problem can be described by a Partially Observable Markov Decision Processes (POMDP). To formally describe an ideal planner, consider a maze world described by discrete time, $0 \leq t \leq T$, and a grid of cells, $W = \{w(i, j) \mid 0 \leq i \leq \text{width}, 0 \leq j \leq \text{height}, w(i, j) \in \{\text{wall}, \text{empty}, \text{goal}\}\}$. A single cell in the world contains the exit: $\exists(i_g, j_g) \rightarrow w(i_g, j_g) = \text{exit}$. The planner knows its location at time t , L_t , and acts based on its beliefs about the world at that time.

The beliefs about the location of the exit at a given time are a set of probabilities X_t , that include all possible locations of the exit. Here X_0 encodes the set of the agent’s initial beliefs. Since the location of the walls is initially known to the agent, X_0 encodes every wall cell as 0: $\forall i, j, w(i, j) = \text{wall} \rightarrow X_0(i, j) = 0$. Non-wall grid cells are initialized with positive probabilities, assuming that every cell is equally likely to contain the exit $\forall i, j, w(i, j) \neq \text{wall} \rightarrow X_0(i, j) = 1/n$, where n is the number of non-wall cells in the world. The agent moves one grid cell at a time in four cardinal directions $\{N, S, W, E\}$. Each legal action results in moving to a new location, and possibly receiving new observations.

The agent has a 360 degree vision, and can see any cells not occluded by walls. So, when the agent enters an unobserved room, the entire room becomes visible at once. The algorithm that determines cell visibility uses ray-casting: a cell is visible, if the rays cast from the agent’s location to any part of the cell do not intersect walls.

In the general case POMDP, the observations are a set of probabilities $O_t = P(W|X_t, L_t)$, such that for every visible cell (i_v, j_v) $0 \leq O_t(i_v, j_v) \leq 1$, and for every invisible cell (i_i, j_i) $O_t(i_i, j_i) = 0$. For simplicity we model deterministic observations, setting $\forall i_v, j_v \rightarrow O_t(i_v, j_v) = 1$ for visible cells.

Every time new observations are made the agent narrows the distribution over its beliefs using standard Bayesian updating, $X_{t+1} \propto X_t O_t$. The agent makes a move after updating observations, such that optimizes the value of the agent's state.

Thus, MST is a special case of a POMDP with deterministic observations, such that each observation reveals the contents of a room. To define a planning state space in MST we align states with observations. Formally, a maze-world W can be represented as a tree, where each node N_i is a state defined by the location (x, y) of the cell from which an observation can be made, and the total visible area of the maze. The root node indicates the starting position, and the adjacent nodes indicate the subsequent states that the agents can move to (see Supplement for an illustration). The value of a node corresponds to the expected number of steps to reach the exit assuming the given node is chosen. Suppose that the agent can reach a node N_i after s_i steps, which contains the exit with the probability $P(exit|N_i)$. If the exit is found, then it would take another e_i steps to reach it. Then, the value of N_i is given by:

$$Q(N_i) = P(exit|N_i)(s_i + e_i) + (1 - P(exit|N_i)) \min_{c_j \in C(N_i)} Q(c_j) \quad (1)$$

Here $C(N_i)$ is the set of all future nodes that follow from N_i . So, the value of choosing the state N_i is equal to a sum of an immediate value of this state, and an expectation of future value of the future state accessible from it, assuming the future states will be chosen optimally. The optimal agent will always choose the state with the smallest *absolute* value (the fewest expected number of steps to reach the exit), but suboptimal agents can deviate from the optimal path.

2.3 Attributing Intelligence to Planning

To attribute intelligence to *planning* the observer needs to measure the extent of deviation from the optimal plan. Given the agent's location in the maze and the revealed area, the observer can generate the optimal plan and compare it to the agent's actions. When seeing a suboptimal action, the observer evaluates by how much this action lengthens the expected path compared to the the

optimal alternative. Agents better at approximating the optimal path are more likely to choose optimally in the future, and are seen as more intelligent.

Formally, we measure the goodness of the agent’s planning by inferring the magnitude of *softmax noise* in the decision policy needed to generate the agent’s path. We assume that the agent is choosing between k state nodes with values $Q_1, \dots, Q_k$ and that the probability of choosing value Q_i is proportional to the relative magnitude of the value: $Pr(N_i|\tau_j) = \frac{\exp(-N_i/\tau_j)}{\sum_{l=1}^k \exp(-N_l/\tau_j)}$. Here τ_j is a parameter controlling *softmax noise* and j indexes hypothetical values of τ . This method of mapping expected utilities to the probabilities of making a choice is widely used to model expressed preferences, known as Luce’s choice in economics (Luce, 1959), as Boltzmann rationality in reinforcement learning (Daw et al., 2006), and as divisive normalization in neuroscience (Louie et al., 2011). The negative sign in front of the values ensures that shorter paths result in higher probabilities. As $\tau \rightarrow 0$ the agent will always choose the shortest expected path. As tau increases the agent will behave in a more noisy way, and as $\tau \rightarrow \infty$ agents will chose actions at random (see Supplement).

Notably, we are not committing to softmax as being necessarily the underlying policy, but rather use it to describe a range of policies that approximate optimal planning to a different degree. This means that the exact choice of planner the observer uses for inference is not critical, as long that it approximates optimal planning. In particular, the softmax policy is equivalent to an alternative stochastic choice models that add randomness to the value itself, but always choose the best value deterministically (see Supplement).

To infer the τ_j parameter of an agent we place a prior over tau, which we chose as a logarithmically spaced sequence of 50 numbers $\hat{\tau} = (0.05, 0.06, 0.07 \dots 50)$. Here assuming rationality means that the prior probability of $\tau_j \in \hat{\tau}$ can be defined as inversely proportional its magnitude $Pr(\tau_j) = 1/\tau_j$, normalized so that all prior probabilities sum up to 1. The posterior probabilities of τ_j after observing m node choices are given by the normalized products of probabilities of each choice: $Pr(\tau_j|N_1..N_m) = Pr(N_m|\tau_j) \dots Pr(N_1|\tau_j) Pr(\tau_j)$. The expected decision noise that is likely to have

produced the path is given by $\tau = \sum_{i=1}^{50} \tau_j Pr(\tau_j | N_1..N_m)$. Qualitatively, we categorize agents as *optimal* if their path could be produced by always choosing the best value ($\tau \rightarrow 0$) and as *suboptimal* otherwise.

We further transform the agent’s expected decision noise into a **metric of planning** by a mapping: $planning = zscore(\ln(1/\tau))$. This transformation ensures that outcome and planning are measured on the same scale. Here 0 corresponds to the mean planning value, and positive values indicate above average approximations of the optimal path. Observers who rely on planning should attribute higher intelligence to agents for whom the measured planning is higher.

2.4 Attributing Intelligence to Outcome

Another way of attributing intelligence is by evaluating the agent’s *outcome*, defined as the agent’s actual utility upon completing the task. Utility means the total reward received by the agent minus the costs of actions. The total reward in all trials in the MST is always the same: the agent always finds the exit. However, agents’ trajectories can have different costs, measured as the number of steps to reach the exit. We model outcome-based attributions of intelligence as inversely proportional to the incurred steps: the fewer steps the agent takes to find the exit, the more intelligent it appears to be. Formally we define $outcome = zscore(stepsDirect/steps)$. Here *stepsDirect* is the shortest possible number of steps going directly to the goal by which the exit could be reached if its position was initially known to the agent. So for the shortest possible path $stepsDirect/steps = 1$, and $stepsDirect/steps < 1$ for all other paths.

As a qualitative way of measuring outcome, we also consider whether the agent is lucky. We categorize outcomes as *lucky* if the agent found the exit in the first room it revealed (meaning, the first observed cluster of black cells that could contain the exit), *unlucky* if it had to reveal the entire maze before the exit was found, and *fair* if the exit is found in-between. When attributing intelligence to outcome, the *lucky* agents should be seen as more intelligent compared to the *unlucky*

or *fair* agents.

Outcomes differ depending on the exit location, and inefficient planning occasionally produces lucky outcomes. So, the agent’s outcome can be thought of as a random variable, such that evaluating planning measures the expectation of outcome. With these two ways of attributing intelligence formally defined, we turn to experiments that assess participant’s attributions empirically.

3 Experiments

We tested intelligence attribution with three experiments. In the first experiment we measure attributing intelligence to planning and to outcome using the MST Attribution Condition, in which participants viewed and rated different agents performing MST. In the second experiment we looked for a relationship between how participants themselves plan and how they evaluate others. To measure how participants plan, they first completed an MST Search Condition in which they controlled the agent themselves, and then completed an MST Attribution Condition as in Experiment 1. The third experiment controlled for the fact that optimal planners on average take shorter paths compared to sub-optimal planners, and replicated the findings of Experiments 1 and 2 while more cleanly dissociating between planning and outcome. All experiments received ethics clearance from a University of Waterloo Research Ethics Committee, and from the MIT Ethics Review Board. The full experimental procedure and the set of stimuli used can be downloaded from <https://marta-kryven.github.io/attribution.html>

3.1 Experiment 1

In the first experiment we measured attributing intelligence based on outcome and based on planning in the MST Attribution condition with different agents. To recap, we hypothesized that ratings can depend on **outcome** and on **planning**, which can be measured qualitatively (is it good or bad?) and quantitatively (how good is it?). We expected that observers attributing intelligence based on **outcome** should rate *lucky* agents as more intelligent than *unlucky* agents, and should attribute intelligence in proportion to the metric of outcome defined in Section 2.4. In contrast, those who

rely on **planning** should rate *optimal* agents as more intelligent than *suboptimal* agents, and attribute intelligence in proportion to the metric of planning defined in Section 2.3. Here an optimal-unlucky trajectory means that the plan was optimal, but the exit was in the last place searched, and a suboptimal-lucky trajectory means that the exit happened to be in the first place the agent looked, despite a suboptimal decision. Examples of various agents are shown on Fig. 3.

Figure 3: In this example the optimal choice would be to first visit the room on the left, given that it is closer to the agent, and both rooms are of equal size and shape. The four examples illustrate different trajectories that agents could take, assuming that the exit is found at the location indicated by the arrow. The four agent trajectories from left to right are: *optimal-lucky*, *suboptimal-lucky*, *optimal-unlucky*, and *suboptimal-unlucky*. The *optimal-lucky* agent decides to observe the closer room on the left, which is the optimal choice, and also happens to be lucky. The *suboptimal-unlucky* agent makes a suboptimal decision to go to the further room on the right. This choice is unlucky, since the exit happens to be in the other room.

The agents' trajectories were generated by replaying actions of participants previously recruited on Amazon Mechanical Turk to complete MST Search. The trajectories were generated by 60 participants, out of which we sampled representative optimal and suboptimal paths. Our goal in sampling trajectories was to select environments that did not repeat, as well as include four examples of each type of agent (see Stimuli). A detailed description of the instructions given to participants in MST Search is provided in Experiment 2. Several of the participants recruited to generate the trajectories moved seemingly at random until they stumbled upon the exit. Whatever the cause for this (e.g. distraction, technical difficulties, misunderstanding instructions), we decided to include such trajectories in MST Attribution, and we refer to them as *pseudo-random*. We expected pseudo-random trajectories to be rated as unintelligent given either attribution strategy, since they present with both poor planning and poor outcome.

Since planning-based attributions of intelligence are mentally costly compared to outcome-based attributions, and may require overriding outcome-based intuitions, we measured participants' cognitive reflection by administering a Cognitive Reflection Test (CRT) (Frederick, 2005). The CRT scores correlate with numerical ability, inhibition of incorrect impulsive responses, intelligence quotient, and rational thinking (Toplak et al., 2011; Welsh et al., 2013). Since the original CRT was used extensively in web-based studies, participants are likely to be familiar with the test in its original form (Chandler et al., 2014). However, providing analogous problems in a novel context was shown to address this concern (Chandler et al., 2014). In our version of CRT participants answered the following questions:

- A slime mold doubles in size every 2 hours. One gram of slime mold can fill a container in 8 hours. How long does it take to fill half of the same container?
- It takes 10 people 10 hours to knit 10 scarfs. How long does it take 100 people to knit 100 scarfs?
- A coffee and a sandwich cost \$12. A sandwich costs \$10 more than a coffee. How much does a coffee cost?

3.2 Participants

Sixty participants were recruited online via Amazon Mechanical Turk, restricted to US participants. The number of participants was decided based on a pilot study, to ensure enough participants for an analysis of different strategies. The exclusion procedure was decided on in advance and included a check for repeated responses from the same IP address, answering an instruction quiz correctly (see below), and coherent answers to verbal questions.¹ As a result, 4 participants were discarded for failing the instruction quiz, and two for failing verbal responses. The analysis thus included 54 participants (23 females, 31 males, median age 34).

¹For example, we excluded participants who respond to *How did you make your decisions?* by typing unrelated responses such as *Have a nice day!*

3.3 Method

The experiment was presented on a computer screen in a web browser using our own JavaScript interface developed in the lab. Participants first read a consent page, on which they provided their age and gender, and read a short description of the experiment. Following this they read the instructions, which included an annotated image to disambiguate the meaning of terms like ‘wall’ and the reference to the agent character.

You will see videos of other people playing a Maze Game, and evaluate their intelligence. Explanation of the Maze Game: People are trying to find an exit (a red circle). People know the layout of the maze (where the walls and rooms are). People know the exit is equally likely to be behind any of the black squares. People want to find the exit in as few moves as possible.

YOUR TASK: After you watch each video, please evaluate the intelligence of the person in the video. You will use a scale from 1 (less intelligent) to 5 (more intelligent). Each video shows a different person. Walls and floor tiles change from trial to trial. Some videos show only part of a person’s solution. It is important to view each video at least once.

After reading the instructions, participants viewed 3 familiarization examples. Participants then completed a quiz checking their understanding of the instructions. The quiz contained two multiple choice questions (*The person’s task is to ..* and *My task is to ..*). Following the quiz, participants viewed 32 stimuli movies in a randomized order. After each movie finished playing, participants selected a rating from a Likert scale between 1 (least intelligent) to 5 (most intelligent). At the end of the survey participants were asked: *How did you make your decisions?*. Lastly, participants completed the CRT. The complete set of all instructions is given in the Supplementary Materials.

3.4 Stimuli

The stimuli were 32 videos showing 8 types of agent trajectories: *optimal-lucky*, *suboptimal-lucky*, *optimal-fair*, *optimal-unlucky*, *suboptimal-unlucky*, *pseudo-random*, *optimal-incomplete*, and *suboptimal-*

incomplete. Each of these types was shown four times, using a different maze layout. The maze layouts were similar to the examples shown in Figure 8, ranging in size between 5x6 to 6x8, and containing between 3 and 5 rooms. In terms of maze difficulty, our goal was to choose a range of difficulty that is non-trivial, but such that subjects can approximate optimal planning reasonably well. Examples of optimal and suboptimal trajectories with lucky and unlucky outcomes are shown on Fig. 3. The 8 *incomplete* trajectories stopped after an agent chose one of the rooms, but before the black cells were revealed. So, on incomplete trials observers saw one decision made by the agent, but did not see the outcome. In all incomplete trials agents took exactly 9 steps.

For all trajectories we inferred the **noise parameter** τ and calculated the measure of **planning**, as described in Section 2.3. We also measured the total number of **steps** in the path and calculated the measure of **outcome**, as described in Section 2.4. Since the outcome was not visible on incomplete trajectories, we assigned to them the mean outcome (which is equal to 0).

3.5 Results

Figure 4: Intelligence ratings for each type of agent in Experiment 1, z-scored across all trial types and across the entire sample of participants. Each triangle represents a participant’s mean rating of a given type of agent. The green dot is the mean.

The distribution of ratings for each type of agent are shown on Fig. 4. Each triangle represents a participant’s mean rating of agents of a given type. We performed a Mixed Effects Linear Model (MELM) analysis to assess the effect of *planning* and *outcome* on ratings. All trials were included

Figure 5: Mean rating predictions from each model compared to mean human ratings in Experiment 1. Error bars indicate 95% confidence intervals. Agent labels are: *OL* - *optimal-lucky*, *OF* - *optimal-fair*, *OUL* - *optimal-unlucky*, *PS-R* - *pseudo-random*, *SL* - *suboptimal-lucky*, *SUL* - *suboptimal-unlucky*, *OPT-part* - *optimal-part*, *SUB-part* - *suboptimal-part*. The combined (Both) model is the closest to human attributions.

in this analysis, but outcome of incomplete trials was set to zero (the mean outcome), to reflect a neutral expected outcome. To test whether both outcome and planning contribute to ratings, first we constructed a null model including only a random effect of participant. We found that this model was significantly improved by modeling random slopes for outcome per participant $\chi^2(1) = 714.49, p < .001$. The null model was also significantly improved by modeling random slopes for planning per participant, $\chi^2(1) = 472.71, p < .001$. Adding random slopes for planning per participant to the outcome-only model resulted in further significant improvement, $\chi^2(1) = 266.92, p < .001$, as did adding random slopes for outcome to the planning-only model $\chi^2(1) = 508.69, p < .001$. Together, these results establish that both outcome and planning were used to make the attributions². Mean rating predictions from each model compared to mean human ratings are shown on Fig. 5. Random slope coefficients for planning and outcome fitted to individuals are shown on Fig. 6.

Next we assess the effect of planning on ratings of incomplete trials. On incomplete trials participants could not evaluate the outcome of the trial, but they could evaluate the agent's planning based on the decision they saw. To test the effect of planning on ratings of incomplete trials we constructed a MELM null model with only a random intercept for participant. Adding a random

²To ensure that there is no effect of the stimuli that would manifest in the analysis, even if the subjects rated the stimuli at random, we re-analyzed the data with ratings permuted between individuals. Please see Supplementary Materials for details.

Figure 6: The MELM random slope coefficients for planning and outcome fitted to individuals in Experiment 1. Error bars indicate bootstrapped 95% confidence intervals. The red line visually separates participants who weight planning more than outcome, from those who weight outcome more than planning. Both weights are different from zero for all participants.

slope for planning to this null model significantly improved the model, $\chi^2(1) = 16.14, p < .0001$, meaning that in the absence of outcome planning alone had an effect on attributed intelligence. The correlation between the weights placed on planning on incomplete trials and on all trials was significant ($r = .65, p < .0001$), meaning that participants' planning-based attributions on all trials can be predicted from how they evaluated incomplete trials.

Next, we examined whether participants' intelligence attributions to outcome or to planning depended on their scores on a CRT. We fitted a linear regression model using the MELM planning weight as the independent variable, and CRT score as the dependent variable (see Fig. 7). The regression was significant ($F(1, 52) = 11.5, p = .001$), $planningWeight = 0.1CRT$, adjusted $R^2 = 0.17$, meaning that CRT scores predicted participants' tendency to attribute intelligence to planning. A linear regression model using the MELM outcome weight as the independent variable, and CRT score as the dependent variable was also significant ($F(1, 52) = 4.5, p = .04$), $outcomeWeight = 0.7 - 0.07CRT$ adjusted $R^2 = 0.06$, meaning that CRT scores predicted participants' tendency to attribute intelligence to outcome. Together these results show that individuals with higher CRT scores attributed more intelligence based on planning and less based on out-

Figure 7: Experiment 1. Participants’ planning and outcome weights, as predicted from their CRT scores

come.

3.6 Discussion

In Experiment 1 we found that both planning and outcome contributed to attributed intelligence. On average, agents with good outcomes and agents with optimal planning were rated highly, which may explain why some studies report outcome-based (Frank, 2016; Olson et al., 2008) and others find rationality-based (Pantelis et al., 2016) attributions of intelligence. Participants consistently rated the *pseudo-random* agents as the least intelligent. This finding confirms the earlier observation of Pantelis et al. (2016), that attributing intelligent behavior relies on interpreting an agent’s actions as intentional and reasonable in the first place. We also found that observing the outcome is not necessary to attribute intelligence. Even if the outcome was not observed, at least some participants, still attributed intelligence based on planning.

In addition to these average trends, we found significant differences in how strongly individuals weighted outcome and planning when attributing intelligence (see Fig. 6), with most participants placing higher weight on outcome. To characterize this phenomenon further, we observe that the

weight placed on planning was correlated with one's CRT score, and the weight placed on outcome was anti-correlated with CRT. CRT measures both one's inhibition of an impulsive response and numerical reasoning ability (Welsh et al., 2013). So, attributing intelligence based on planning may require over-riding an initial outcome-based heuristic, and it may also require numerical ability to evaluate probabilities and costs to construct optimal plans.

In the next experiment we further investigate these individual differences in two ways. First, we test to what extent people's own planning ability contributes to attributing intelligence based on planning. To measure planning, we introduce the MST Search task, in which participants control the agent themselves. If planning-based attribution in MST Attribution are limited by planning ability, then we should expect to see a positive correlation between how well people plan in MST Search, and how much they attribute intelligence based on planning, provided people plan in a way that is similar to how they expect others to plan.

It is also possible that participants evaluate outcome in preference to planning because they do not sufficiently consider how others plan. Outcome-bias occurs when behavior is more salient than its constraints (Gilbert & Malone, 1995; Heider, 1958), and in MST Attribution the length of path could be more salient than the agent's planning. Thus, we also hypothesized that engaging in MST Search could increase the salience of the agent's constraints and reduce outcome-bias in MST Attribution. This could happen by priming people to take a seeker's perspective, and increasing the understanding of the task as depending on planning. If outcome-bias is reduced, then in Experiment 2 compared to Experiment 1 we expected to see an increase in the attribution weights placed on planning, a decrease of weights placed on outcome, or both. This would indicate that intelligence attribution is flexible, and the extent to which people evaluate planning or outcome can depend on context.

4 Experiment 2

Individuals who find it easier to generate better plans may find it easier to compare the agent’s path to the ideal path. In the second experiment, participants first completed MST Search Condition to measure their own planning ability. Next, they completed the MST Attribution Condition with the same set of attribution trajectories as in Experiment 1. If one’s tendency to attribute intelligence based on planning depends on one’s planning ability, then the participants’ planning weights in MST Attribution should be predicted by their own planning in MST Search. Engaging in MST Search could additionally reduce outcome-bias, and increase the weights they place on planning in MST Attribution.

4.1 Participants

One hundred twenty participants were recruited via Amazon Mechanical Turk, restricted to US participants and further restricted to participants who had not taken part in Experiment 1. The number of participants was decided based on variance in search strategies in a pilot study of MST Search. From this initial pool, 4 were discarded for failing the instruction quiz, and two for failing verbal responses. The exclusion procedure was identical to the procedure in Experiment 1. The analysis thus included 114 participants (48 females, 66 males, median age 31).

4.2 Method

The experiment was presented on a computer screen in a web browser using a JavaScript interface developed in the lab. Participants first read a consent page, on which they provided their age and gender, and read a short description of the experiment. Participants next completed a MST Search, in which they searched for the exit in a series of mazes, similar to the task of the agent in Experiment 1, with the prompt to complete the task in as few steps as possible. Next, participants completed MST Attribution, identical to the procedure described in Experiment 1. Experiment 2 did not include a CRT.

4.2.1 MST Search

Participants were presented with the following prompt:

PLEASE READ THE FOLLOWING INSTRUCTIONS CAREFULLY. Your task is to exit the maze by reaching the red square in as few steps as possible. You can move one square at a time, by clicking on the white squares near your character. Blue squares are walls. You cannot see through the walls, so the squares you cannot see yet are black. The exit is equally likely to be hidden behind any of the black squares. At the end of the experiment we will add all steps you took and show you how you did compared to previous results. You have an opportunity to earn a bonus for completing the mazes in fewer steps.

The instructions included an annotated image to disambiguate the meaning of terms like ‘wall’ and the reference to the agent character. After reading the instructions, participants completed three practice trials, and answered a comprehension quiz. The quiz contained two multiple choice questions (*My task is to ..* and *My bonus will be bigger if I ..*). Participants who answered the comprehension quiz incorrectly proceeded with the experiment, but their responses were discarded.

On each trial participants looked for an ‘exit’ (marked when visible as a bright red circle) in a series of mazes, by controlling an agent using a mouse. The agent could move one grid square at a time: N, W, S, or E, and had a 360 degree view of the maze limited by walls. A maze started out dark (except for walls), but was uncovered as a participant opened new areas. Participants initially saw the layout of the rooms and location of the barrier walls, but did not know where the exit was. Participants received a performance-based bonus of up to \$1 for finishing all mazes with least steps. To receive the maximal bonus, participants needed to achieve a step cost within 5% of the optimal solution. The bottom 10% of participants, determined offline after we analyzed the responses, received no bonus. At the end of MST Search, participants answered a question about how they made their decisions. The decision times (times to make a move) and the path taken were recorded for each trial.

4.2.2 MST Attribution

Participants rated the intelligence of maze-solving agents, in a procedure identical to Experiment 1. As before, at the end of the task participants were asked how they made their decisions.

4.3 Stimuli

4.3.1 MST Search

There were 12 stimuli, which are shown on Fig. 8.

Figure 8: Stimuli used in MST Search in Experiment 2. Black cells indicate unseen areas. White cells are empty cells already revealed. Grey cells are walls, and are always visible. The agent can move through empty areas, but not through walls. The starting location was always at the top-left corner. The exit was hidden behind one of the black squares. Solid lines indicate the averaged path taken by the majority of participants. Here the average paths are calculated by taking a step in each cell that is taken by the majority of the participants. The average paths were identical to the optimal solution in all mazes but two – in these two case the optimal solution was the second most common path. Dashed blue lines indicate the second most common path.

Figure 9: Individual’s planning values in MST Search plotted against MST attribution weights placed on planning (left) and on outcome (right) in Experiment 2. The better participants planned in the MST Search, the more they weighted planning, and the less they weighted outcome in MST Attribution.

4.3.2 MST Attribution

MST Attribution consisted of the 32 trials identical to Experiment 1.

4.4 Results

4.4.1 MST Search

We measure a participant’s planning ability in the same way as we measure an agent’s planning in Experiment 1. We infer the decision noise parameter τ required to generate the participant’s path and calculate planning values as $z - \text{score}(\ln(1/\tau))$ (see Section 2.3). We found a significant correlation between participants’ planning in MST Search and the weights they placed on planning in MST Attribution, $r = .45, p < .001$ (see Fig. 9 and MST Attribution results below), suggesting that both processes use a common mental planning mechanism. We also found a significant anti-correlation between an individual’s planning in MST Search and the weights placed on outcome ($r = -0.32, p > .001$). This suggests that participants who were better at planning in MST Search

also placed less weight on outcome in MST Attribution. To further characterize behavior in general, we calculated the average paths, by taking a step in each cell that is taken by the majority of the participants. The average paths, and the second most common paths, are shown on Fig. 8. The average paths were identical to the optimal trajectory in all mazes but two – in these two case the optimal solution was the second most common path.

4.4.2 MST Attribution

The distribution of z-scored ratings for each type of agent are shown on Fig. 10. Each triangle represents a participant’s mean rating of four agents of a given type.

Figure 10: Intelligence ratings for each type of agent in Experiment 2, following the same conventions as in Experiment 1

We replicated the results of the MELM analysis from Experiment 1 to assess the effect of planning and outcome on ratings of complete trials. We constructed a null model including only a random effect of participant, and found that it was significantly improved by modeling random slopes for outcome per participant $\chi^2(1) = 1433.9, p < .001$. The null model was also significantly improved by modeling random slopes for planning per participant, $\chi^2(1) = 1437.2, p < .001$. Adding random slopes for planning per participant to the outcome-only model resulted in further significant improvement, $\chi^2(1) = 850.4, p < .001$, as did adding random slopes for outcome to the planning-only model $\chi^2(1) = 847.06, p < .001$. Mean rating predictions from each model compared to mean human ratings are shown on Fig. 11. The outcome and planning weights fitted to individuals are

Figure 11: Mean rating predictions from each model compared to mean human ratings in Experiment 2. Error bars indicate 95% confidence intervals. Agent labels are: *OL* - *optimal-lucky*, *OF* - *optimal-fair*, *OUL* - *optimal-unlucky*, *PS-R* - *pseudo-random*, *SL* - *suboptimal-lucky*, *SUL* - *suboptimal-unlucky*, *OPT-part* - *optimal-part*, *SUB-part* - *suboptimal-part*. The combined (Both) model is the closest to human attributions.

shown on Fig. 12.

We also replicated the effect of planning on ratings of incomplete trials. We constructed a MELM null model with only a random intercept for participant, and found that it was significantly improved by adding a random slope for planning, $\chi^2(1) = 41.211, p < .001$, meaning that on trials where outcome was not shown planning alone still had an effect on attributed intelligence. The correlation between the weights placed on planning on incomplete and on all trials was significant ($r = .73, p < .0001$), meaning that participants' planning-based attributions on all trials can be predicted from how they evaluated incomplete trials.

4.4.3 Task Order Effects

We compared the planning and outcome MELM random slope coefficients of participants in Experiment 1 and 2 (Fig. 13) using an independent sample t-test $t(166) = -2.68, p = .008$, and found that in Experiment 2 participants' planning weights were higher than in Experiment 1. There was no significant difference between the weights placed on outcome in Experiment 1 and 2. An independent sample sample t-test comparing outcome coefficients was not significant, $t(166) = 1.9, p = .06$.

Figure 12: The MELM random slope coefficients for outcome and planning fitted to individuals in Experiment 2. Error bars indicate bootstrapped 95% confidence intervals. The red line visually separates participants who weight planning more than outcome, and who weight outcome more than planning. Both weights are different from zero for all participants.

4.5 Discussion

In the second Experiment we replicated the findings of Experiment 1. We additionally found a significant relationship between one’s performance in MST Search and attributing intelligence based on planning in MST Attribution. Specifically, we found that participants’ proximity to the optimal planner in MST Search predicted planning-based attributions of intelligence in MST Attribution. This suggests that intelligence attribution depends on cognitive abilities related to planning, as measured by MST Search and CRT.

MST Search and MST Attribution both require approximating rational planning. The finding that the quality of planning approximations during search and during evaluation are correlated, suggests that people may use a similar cognitive mechanism to plan for themselves, and to evaluate others. If the two mechanisms were unrelated, we would have seen no such correlation. For example, one could do the task by guessing, but expect others to follow the optimal plan. One could also do the task rationally, and yet evaluate others based on outcome. Our results exclude this possibility, and

Figure 13: MELM random slope coefficients of individual participants in Experiment 1 and 2. Participants in Experiment 2 (who completed the MST Search) were fitted with higher planning coefficients.

suggest that there is a dependence between how well people plan, and how strongly they expect good planning from others.

We also found that in the second experiment outcome-bias was reduced through an increase in the weights placed on planning in MST Attribution after completing MST search. At the same time, we have observed a marginal difference between outcome weights placed in Experiments 1 and 2 ($p = .06$), meaning that the weights placed on outcome may have been reduced as well. This suggests that, at least for some participants, the choice of attribution strategy is flexible, and can be influenced by context. In Experiment 3, we further verify this result by presenting MST Attribution and MST Search in a counter-balanced order.

We note that the reduction in outcome bias in MST Attribution could happen by training the participants planning skills, in addition to increasing the salience of task constraints. If so, then MST Search performance could be affected by task order in Experiment 3, assuming MST Attribution and Search tasks train one another.

Notably, in Experiments 1 and 2 the effect of optimal planning is consistent with real-world expectations: optimal planning minimizes unlucky outcomes, so that when optimal agents are unlucky, their outcome is on average better, compared to suboptimal-unlucky agents. In Experiment 3 we

designed a new set of stimuli that equalized the number of steps across optimal and suboptimal agents, to preempt the concern that outcome depends on planning, and potentially increase the discriminating strength of the design. We achieved this by manipulating the starting location of each agent: suboptimal agents tended to start closer to the goal, so that on average they took the same number of steps as optimal agents and differed only in the quality of their planning. As a result of this manipulation, the optimal agents in Experiment 3 "play with a handicap", which means that the aggregate statistics of the stimuli may differ from what subjects expect in real life. By replicating the results of Experiments 1,2 in Experiment 3 we show that our findings generalize across the environments, and are independent of the aggregate outcomes of optimal and suboptimal agents.

5 Experiment 3

In the third experiment our goal was to replicate the results of Experiments 1 and 2, controlling for the total number of steps taken by different types of agents. We compared four types of agents: optimal-lucky, suboptimal-lucky, optimal-unlucky and suboptimal-unlucky, and varied the agent's starting location so that optimal and suboptimal agents always took an equal number of steps. The lucky agents took between 7 and 11 steps. The unlucky agents took between 14 and 23 steps. In Experiment 3 we also use a different set of MST Search stimuli to ensure the generalizability of our MST Search results. As a result of this manipulation, we expected to see a more interpretable difference between planning-based and outcome-based attributions. We also presented MST Search and MST Attribution in a counterbalanced order, to verify the task order effects observed between Experiments 1 and 2.

5.1 Participants

One hundred and thirty six participants were recruited online via Amazon Mechanical Turk³, restricted to US participants and to participants who had not taken the previous experiments in this

³We recruited participants until we had 120 who passed the exclusion criterion.

series. Of our initial pool, 14 participants were discarded for failing the instruction quiz, and two failed the verbal responses. The exclusion procedure was identical to the procedure in Experiments 1 and 2. The analysis thus included 120 participants (47 females, 73 males, median age 34, $SD = 9.7$).

5.2 Method

The experiment was presented on a computer screen in a web browser using a JavaScript interface developed in our lab. Participants first read a consent page, on which they provided their age and gender and read a short description of the experiment. Following the consent participants completed MST Search and MST Attribution (detailed below). The task order was counterbalanced between participants. At the end of the experiment participants answered 3 CRT questions, identical to those in Experiment 1.

5.2.1 MST Search

Participants searched for an exit from a maze, as in Experiment 2. After reading the instructions (similar to Experiment 2, except that the subjects did not receive a bonus), participants completed three practice trials and answered a comprehension quiz. As before, participants who answered the comprehension quiz incorrectly proceeded with the experiment, but their responses were discarded. At the end of the task participants answered a question: *How did you make your decisions?*

5.2.2 MST Attribution

This task was similar to the MST Attribution used in Experiments 1 and 2. As before, after reading instructions (same as Experiments 1 and 2), participants answered a comprehension quiz. Participants who answered the comprehension quiz incorrectly proceeded with the experiment, but their responses were discarded. Participants then viewed and rated three familiarization examples followed by 24 more trials. At the end of the task participants answered a question: *How did you make your decisions?*

Figure 14: Individual planning in MST Search plotted against weights placed on planning and on outcome in MST attribution in Experiment 3. Error bars indicate 95% confidence intervals.

5.3 Stimuli

5.3.1 MST Search

The search stimuli included 12 mazes of varying difficulty, as in Experiment 2.

5.3.2 MST Attribution

Stimuli were 24 videos of agents searching for the exit in a maze. There were four types of agents: optimal-lucky, optimal-unlucky, sub-optimal-lucky and sub-optimal-unlucky, which were shown in 6 different mazes similar to stimuli used in Experiments 1 and 2 with controlling for the length of path.

5.4 Results

Of the 120 participants, 60 completed the MST Attribution first and 60 completed MST Search first.

5.4.1 MST Search

As an Experiment 2, we assessed the quality of participants' planning as described in Section 2.3. We replicated the results of Experiment 2, and found a significant correlation of individual's quality of planning in MST Search with the weight that they place on planning in MST Attribution (see Fig. 14 and MST Attribution results below), $r = .24, p = .008$. However, unlike in Experiment 2, no significant correlation was observed between individuals' quality of planning in MST Search and outcome weights in MST Attribution ($p = .3$). The correlations between performance in MST Search and outcome weights in MST Attribution were not significant both for subjects who did MST Search first ($p = .35$) and second ($p = .37$). This confirms the finding of Experiment 2, that participants' planning quality predicts their tendency to attribute intelligence based on planning.

5.4.2 MST Attribution

The distribution of ratings for each type of agent are shown on Fig. 15. Each triangle represents a participant's mean rating of four agents of a given type.

Figure 15: Intelligence ratings for each type of agent in Experiment 3, following the same conventions as in Experiments 1 and 2.

We replicated the results of the MELM analysis from Experiment 1 and 2 to assess the effect of planning and outcome on ratings of complete trials. We constructed a null model including only a random effect of participant, and found that it was significantly improved by modeling random

Figure 16: Mean rating predictions from each model compared to mean human ratings in Experiment 3. Error bars indicate 95% confidence intervals. Agent labels are: *OL* - *optimal-lucky*, *OUL* - *optimal-unlucky*, *SL* - *suboptimal-lucky*, *SUL* - *suboptimal-unlucky*. The combined (Both) model is the closest to human attributions.

slopes for outcome per participant $\chi^2(1) = 869.24, p < .001$. The null model was also significantly improved by modeling random slopes for planning per participant, $\chi^2(1) = 489.84, p < .001$. Adding random slopes for planning per participant to the outcome-only model resulted in further significant improvement, $\chi^2(1) = 659.65, p < .001$, as did adding random slopes for outcome to the planning-only model $\chi^2(1) = 1039, p < .001$. Mean rating predictions from each model compared to mean human ratings are shown on Fig. 16. The random slope coefficients for outcome and planning fitted to individuals by the two-factor model are shown on Fig. 17

We also replicated the finding of Experiment 1 that participants with higher CRT scores placed more weight on planning when attributing intelligence. A linear regression model using the MELM planning weight as the independent variable and CRT as the dependent variable was significant ($F(1, 118) = 21.57, p < .001$), $planningWeight = -0.02 + 0.1CRT$, adjusted $R^2 = 0.15$ (see Fig. 18). A linear regression model using the MELM outcome weight as the independent variable and CRT as the dependent variable was also significant ($F(1, 118) = 11.11, p = .001$), $outcomeWeight = 0.6 - 0.07CRT$, adjusted $R^2 = 0.08$. Together these results show that participants with higher CRT scores placed higher weights on planning and lower weights on outcome when attributing intelligence.

Figure 17: The MELM random slope coefficients for outcome and planning fitted to individuals in Experiment 3. Error bars indicate bootstrapped 95% confidence intervals. The red line visually separates participants who weight planning more than outcome, and who weight outcome more than planning. Both weights are different from zero for all participants.

5.4.3 Task Order Effects

We tested our earlier finding that participants who complete MST Search prior to MST Attribution place higher weight on planning in MST Attribution. To determine which kind of t-test is appropriate, we first conducted an F-test comparing variances of planning weights in the two groups of participants. The variances were equal ($F(1, 59) = 0.88, p = .6$, ratio of variances 0.88), and so we conducted an independent-samples t-test with homogeneous variance to compare the mean planning weights of the two samples. We found that participants who did MST Search first were fitted with higher MELM planning weights compared to participants who did MST Search second, $t(118) = 2.48, p = .015$. The variances of outcome weights for the two groups of participants were unequal ($F(1, 59) = 2.5, p < .001$, ratio of variances 2.5), and so we conducted an independent-samples t-test with heterogeneous variances to compare the mean outcome weights of the two samples. We found no significant difference in MELM outcome weights of participants who did MST Search first and those who did MST Search second, $t(118) = -1.84, p = .07$. The distri-

Figure 18: Participants’ outcome and planning weights plotted against CRT scores in Experiment 3. CRT scores are correlated with planning, and anti-correlated with outcome.

bution of planning and outcome weights for the two groups of participants are shown on Fig. 19. This result confirms our earlier finding that participants place higher weight on planning in MST Attribution after completing MST Search.

We also tested whether participants performance in MST Search depends on task order. The variances were equal ($F(1, 59) = 1.1, p = .8$, ratio of variances 1.1), so we conducted an independent-samples t-test with homogeneous variances to compare the mean planning performance of the two samples. We found no significant improvement in planning performance of participants who did MST Search first and those who did MST Search second, $t(118) = 2, p = .05$, suggesting that performing MST Attribution does not train planning performance in MST Search (the marginally significant p-value suggests the opposite effect).

5.4.4 Discussion

In the third experiment we replicated the results of Experiment 1 and 2 with stimuli controlling for the length of path. This result established that our findings hold even when planning has no effect on outcome, given that the starting positions of the optimal agents are handicapped.

Figure 19: MELM planning and outcome weights of participants based on task order in Experiment 3. Participants who completed the MST Search task first, before the MST Attribution task, tended to have higher planning weights in the Attribution task, relative to participants who completed the Search task second, after the Attribution task.

Specifically we found that, (1) both outcome and planning contribute to attributions of intelligence, (2) individual planning weights in MST Attribution can be predicted from participants' CRT scores, and from participants planning performance in MST Search, and (3) planning weights in MST Attribution can be increased by a context manipulation of engaging in MST Search. The results of Experiment 3 also verify that our results generalize to different sets of stimuli, used for MST Attribution and MST Search.

We found that MST Search increases attributing intelligence to planning in MST Attribution, but has only a marginally significant effect on outcome-based attributions. We also found that MST Attribution does not improve planning performance in MST Search (and on the contrary, may have decreased it). These findings suggest that the outcome-bias was reduced by up-weighting planning. This reduction of outcome bias could happen by increasing salience of the task constraints, or by priming attention to planning, but is unlikely to have been caused by training planning ability. Taken together, these results suggest that attributions of intelligence are flexible – influenced by the extent to which the observer can evaluate planning, and can relate to the agent's task.

6 General Discussion

In this paper we examined how people attribute intelligence to plans and outcomes when observing an agent’s behavior. We used a new experimental method, the Maze Search Task (MST). We used the MST both as a Search task to measure how people plan, and as an Attribution task, to measure the attribution of intelligence. The MST is easy to learn, closely resembles natural planning behaviors, and allows us to dissociate between planning and outcome more cleanly than many real-world scenarios. While in natural contexts outcomes may be locally informative about planning – since the goal of planning is optimizing outcomes over time – the MST paradigm allows us to dissociate plans and outcomes by manipulating how close the agent starts to the goal. We modeled planning-based attributions of intelligence in MST by treating the agent as an approximately rational planner, for which attributed intelligence is inversely proportional to the agent’s degree of deviation from the optimal plan. Outcome-based attributions can be modeled as inversely proportional to the incurred cost.

Across three experiments, we found that intelligence attributions relied on both plans and outcomes, with significant individual differences in the extent to which people relied on these two factors, which correlated with an individual’s cognitive reflection. Individuals higher on cognitive reflection weighted plans more relative to outcomes when attributing intelligence. We also found that for some participants, attribution strategies were influenced by a contextual manipulation. Engaging in MST Search before MST Attribution made people slightly more likely to weight planning during attribution. This suggests that attribution strategies are flexible, and could be influenced by several factors. Attributing intelligence to planning could depend on planning abilities, as measured by MST Search, or on cognitive reflection abilities measured by CRT, such as a tendency to correct initial impressions (Reeder et al., 1992). Further, MST Search could increase the salience of planning, by priming the agent’s perspective, or task constraints.

Taken together, our results suggest an important implication. Outcome-centered theories suggest that outcome is the more attractive metric by which to attribute dispositions to an agent’s perfor-

mance. If true, then the ability to plan, understand the agent’s constraints, or engage in cognitive reflection should have no effect on attributions. Our findings suggest otherwise, demonstrating that people make attributions of intelligence based on planning, provided sufficient context and cognitive resources to make the attribution. In the remainder of this section, we consider alternative interpretations of our results, the contexts that may rationally favor reliance on outcome or planning, the generalizability of our approach, and the broader implications of our results.

Alternative interpretations. We defined a measure of outcome as the number of steps relative to the shortest possible number of steps in a given environment, to account for the fact that different environments have different scale, and different possible shortest paths. However, our results also can be replicated with outcome measured as an unnormalized number of steps. Thus, people could interpret outcome by either mistaking lucky trajectories for a hidden skill, or by mistaking a lucky environment for a hidden skill, in addition to lucky trajectories. Notably, environmental luck alone would not explain outcome bias in MST – for example – two environments with an identical map, and differing only in scale, would likely elicit similar attributions of intelligence, even though the path lengths would differ between these environments. However, our results allow for a possibility that people’s perception of outcome combines environmental luck and trajectory luck.

In a world where the reward placement is biased, it may be rational to look for spatial biases, and expect that agents who correctly identify such biases would consistently achieve good performance. Under this assumption people could reason that the agent is able to figure out the pattern by which the exit is placed, and rate the agents highly if they appear to have figured out the pattern. To avoid this interpretation, we instructed participants that the reward is *equally likely to be anywhere in the maze*, and asked them to report any strategies they used to perform the MST tasks. Only four out of 240 participants reported noticing biases in the distribution of exit locations. A small number of biased runs are expected in randomly generated sequences, and we consider this to be a small number of participants. Further, participants were instructed that each trial shows *a different agent*. Agent appearance varied between trials to reinforce this point. Thus, while we cannot definitely rule out that some participants may unconsciously assume a spatial bias, there is no current support

for that notion.

The assumption that agents rationally maximize rewards and minimize costs is consistent with interpreting mental effort as a cost. Different costs, such as steps and mental effort, can be traded-off against each other (Kool et al., 2018). Even though participants were instructed that agents want to get to the exit in the fewest moves possible, some people could implicitly assume that thinking is costly. We would expect such people to score low on cognitive reflection due to avoiding mental effort, as well as make less outcome-based attributions, since trading off mental cost against steps assumes that step cost is negligible compared to mental effort. This interpretation is unlikely, given that we observe more outcome-based attributions in people with low CRT scores.

Lastly, the Naïve Utility Calculus interprets optimal action as *the best available action given an agent's knowledge and goals*. It is possible that some people interpret optimal action as an action that would have been optimal *in retrospect* (e.g. De Freitas & Johnson, 2018), after the location of the exit has been observed. Attributing intelligence based on this assumption would be equivalent to attributing intelligence based on outcome. Previous work suggests that people are unlikely to make this interpretation in partially observable grid-worlds (Baker et al., 2017). However, it is possible that children who fail to pass the false belief test think this way. To investigate this possibility further, future studies could use MST Attribution in conjunction with a False Belief tests to study whether a simpler implementation of the rationality principle may develop before full theory of mind.

Outcome may sometimes be a rational cue to intelligence. MST Attribution creates a favorable context for evaluating planning. However, the rationality principle is difficult to apply when critical information is missing (Jara-Ettinger et al., 2016), or when mental resources are limited, e.g. under cognitive load. Thus, outcome-based attributions of intelligence may sometimes be resource-rational, especially if the observer has a prior belief that the agent is intelligent. For example, a person playing chess against an AI agent would likely see an unexplained move as smart, if they expect that AI agent is good at such games. Under such circumstances unexplained success

could reinforce the belief that the agent will do well in the future, and cause attributions of high intelligence. It is possible that children observing adults think this way. For example, children are known to imitate causally irrelevant adult actions as much as relevant actions (Keupp et al., 2018), which can facilitate learning by modeling social behavior, and secondarily learning instrumental skills (McGeer, 2007; Zawidzki, 2013). Boyd (2017) further suggests that human cultural evolution is driven by learning through imitating successful models, more than by figuring out how things work. As long as outcomes are highly correlated with competence, it may be sufficient to care about outcomes without explicitly considering the process that drives them.

We note that successful outcomes can also act as a cue attracting the observer’s attention to the planning strategy used by the agent. One participant’s response to debriefing questions reflects learning a search strategy by observation: *‘I was systematically visiting rooms left to right, but I saw someone going into bigger rooms before going into the smaller rooms. That was smart!’*. More experiments are needed to investigate under which conditions people could see unexplained behavior with good outcomes as a cue that initiates an evaluation of planning strategy in this way.

Generalizability to other domains. Our framework for intelligence attribution is based on the proximity of an agent’s plan to the optimal plan. This way of attributing intelligence can, in theory, generalize to other domains in which the observer has a rational model of the agent, and sufficient observations to make an evaluation. The strategy of evaluating behavior by its optimality would be conceptually similar in all these domains, but differ in implementation details.

For example, to evaluate intelligence in verbal communication, one can assume a pragmatic model of communication. In such models rationality is expressed by choosing utterances that minimize the cost of communication, while maximizing the speaker’s informative intent (e.g. Pu et al., 2020). Such a model could be inverted to obtain an evaluation of the speaker’s intelligence with respect to optimal pragmatic communication.

To evaluate intelligence in physical domains, one can assume and invert a hierarchical symbolic and physical planner (Shu et al., 2020). For example, when we see a child re-purpose a dictionary

to reach a shelf, we could evaluate the quality of the planning solution. This evaluation may include the physical stability of the structure, and its counterfactual effectiveness in achieving a given goal.

Further, one can attribute intelligence to agents engaging in active learning by observing whether the agent asks good questions. Models of information gain and active learning can provide an optimal benchmark specifying which questions are best to ask, and can be inverted to measure how closely is the agent tracks the optimal question. In active learning scenarios the outcome (whether the student has gained the right knowledge) can also be dissociated from planning, since the observer can evaluate the efficiency of the student's queries before the outcome is achieved.

Our methodology can also be extended to include prior beliefs into Bayesian updating. To the extent that prior beliefs can be quantitatively measured, they can be used to initialize Bayesian inference. One way to measure such priors would be to withhold information about outcomes, plans, or other cues to agent rationality, and record the attributions of intelligence in this context. For example, initial beliefs could be informed by the agent's appearance (Todorov et al., 2005), prior experience with a similar type of agent, and anthropomorphism (Epley et al., 2007).

Limitations of our approach. The optimality-based approach has important limitations, since it requires sufficient observations of behavior, and sufficient mental computational resources to evaluate them. Everyday attributions of intelligence likely combine rational inference, outcomes, prior beliefs and perceptual cues, and the latter can be weighted more strongly whenever observations and computation are limited. For example, to attribute intelligence by talking to a person, someone could evaluate the rationality of arguments, quality of vocabulary, or appearance of the speaker. Some cues – such as "large vocabulary" – are hard to separate from an evaluation of rationality, since an observer could reason that an intelligent learning strategy was required to learn all those words. Others – such as "wearing a suit" – are likely explained by statistical generalization such as, "experts wear suits". Different weighting of each type of cue may be reasonable, depending on the quality of evidence and computational resources available to the observer.

Information about outcomes may be limited as well. For example, it may be hard to evaluate the outcome without knowing an agent's utilities, which may have a circular dependence on the agent's constraints (Anderson, 2001). It may also be hard to evaluate outcome when the objective utility is undefined, such as in creative domains. For example, evaluating the outcome of a drawing is relatively easy when the artist's objective is 'copy a photograph', compared to when the objective is 'to make unique art'. So, the observer's ability to evaluate outcome depends on the reliability and availability of information about an agent's objective, and evidence of whether it was achieved.

Future research. In using the ideal planner to quantitatively measure the difference between 'better' and 'worse' plans we do not claim to accurately describe all manifestations of human intelligence. Many complex real-world situations are currently out of scope of our work, and may require a different methodological approach. Notably, important aspects of intelligence that are currently outside the scope of our methodology is learning, adaptability (Legg & Hutter, 2007), and ability to learn good concepts (Chollet, 2020). We may think that agents who learn more and faster, without over-fitting to local properties of the environment are more intelligent. Thus, one direction of future work is investigating whether human attributions of intelligence track the outcomes of learning – such as the rate of improvement at certain tasks – as well as structure of the learned knowledge – such as learning compositional and generalizable concepts.

Another direction for future research is how intelligence attribution relates to other situations where people consider outcome and intent, such as the moral domain. For example, multiple studies have reported that intuitions about punishment are determined by the outcome of the agent's action more than by the intent, which is known as moral luck (Martin & Cushman, 2016; Nagel, 2013; Williams, 1981). If people are consistent in how they treat outcome, then individual differences in processing moral luck may be associated with individual differences in attributing intelligence. Notably, people who score highly on CRT also place less weight on moral luck when judging moral outcomes (Cushman & Greene, 2012; Schwitzgebel & Cushman, 2012), which suggests that reasoning about outcomes in moral luck and in MST both rely on cognitive reflection.

Previous work has shown that in multi-stage bandit tasks mental effort can be traded-off against reward, so that people invest less mental effort when the effect of extra effort on reward is marginal (Kool et al., 2018). One direction of future work is to investigate people’s assumptions about trade-offs between mental and physical effort as they attribute intelligence in MST by using stimuli with identical maps at different scales. If observers assume that it is intelligent to trade off physical and mental costs, they should expect more minimization of mental effort when navigating smaller maps, and more minimization of physical effort in bigger mazes, as the physical cost of deviating from the shortest path increases. Alternatively, if the observers expect that for intelligent agents the costs of mental effort are negligible, then they should expect agents to minimize steps equally well, regardless of the difference in reward.

Lastly, the assumption of the Naïve Utility Calculus that humans approximate optimal plans leaves open the possibility that people use a heuristic planner different from the ideal planner that we describe (Jain et al., 2019; Zhi-Xuan et al., 2020). If so, participants could engage in more planning-based attributions of intelligence than we can currently detect. Notably, our current experiments use relatively simple mazes, on which many planning models could approximate the optimal planner. One possible direction of future work would be to design mazes that would bring out differences between the possible approaches to solving MST, for example, more complex mazes. These stimuli could then be used to test whether the models that best predict individual planning are also best at predicting how people evaluate the intelligence of others. In ongoing work, we are studying the empirical planning models that people may use to navigate in MST Search to explore this possibility further.

7 Conclusion

Our work takes a step toward understanding how people attribute intelligence to others, using a rigorous, quantitative framework. We have proposed that humans may intuitively reason about intelligence by evaluating the optimality of an underlying planning process that guides observed behavior. We evaluated this hypothesis in the domain of spatial search using a novel Maze Search

Task, and a Bayesian inverse planning computational model. Our results demonstrate that it is possible to model human intuitions about intelligence as a combination of inferred plans and observed outcomes, and explained disagreements with respect to attributing intelligence expressed in previous work.

Our approach explains not only lay people’s intuitions about whether someone is intelligent, but also individual differences in such intuitions, which may drive fundamentally different theories of social perception based on the rationality principle or based on luck. Looking beyond the simple setting of maze worlds, our approach provides a method for understanding planning and plan evaluation in environments and contexts similar to those encountered in daily life. The capacity to engage with and model an internal decision-making mechanism is one of the strengths of our framework, and we are currently exploring how it might explain a broader range of behaviors and tasks.

8 Acknowledgements

We thank the reviewers for their helpful comments and suggestions, and Stefano Anzelotti, Lucas Tian, Suhyoun Yu, and Jacob Browning for valuable discussions. This material is based upon work supported by the Center for Brains, Minds and Machines (CBMM), funded by NSF STC award CCF-1231216.

References

- Aczel, B., Palfi, B., & Kekecs, Z. (2015). What is stupid? people’s conception of unintelligent behavior. *Intelligence*, 53, 51–58.
- Ambady, N., & Rosenthal, R. (1993). Half a minute: Predicting teacher evaluations from thin slices of nonverbal behavior and physical attractiveness. *Journal of personality and social psychology*, 64(3), 431.

- Amodio, P., Boeckle, M., Schnell, A. K., Ostojić, L., Fiorito, G., & Clayton, N. S. (2018). Grow smart and die young: Why did cephalopods evolve intelligence? *Trends in ecology & evolution*.
- Anderson, E. (2001). Symposium on amartya sen's philosophy: 2 unstrapping the straitjacket of 'preference': a comment on amartya sen's contributions to philosophy and economics. *Economics and Philosophy*, 17(1), 21.
- Asch, S. E. (1946). Forming impressions of personality. *The Journal of Abnormal and Social Psychology*, 41(3), 258.
- Baker, C. L., Jara-Ettinger, J., Saxe, R., & Tenenbaum, J. B. (2017). Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nature Human Behaviour*, 1, 0064.
- Baker, C. L., Saxe, R. R., & Tenenbaum, J. B. (2011). Bayesian theory of mind: Modeling joint belief-desire attribution. In *Proceedings of the thirty-second annual conference of the cognitive science society* (pp. 2469–2474).
- Baker, C. L., & Tenenbaum, J. B. (2014). Modeling human plan recognition using bayesian theory of mind. *Plan, activity, and intent recognition: Theory and practice*, 177–204.
- Baron, J., & Hershey, J. C. (1988). Outcome bias in decision evaluation. *Journal of personality and social psychology*, 54(4), 569.
- Boyd, R. (2017). *A different kind of animal: how culture transformed our species*. Princeton University Press.
- Caplan, R. A., Posner, K. L., & Cheney, F. W. (1991). Effect of outcome on physician judgments of appropriateness of care. *Jama*, 265(15), 1957–1960.
- Chandler, J., Mueller, P., & Paolacci, G. (2014). Nonnaïveté among amazon mechanical turk workers: Consequences and solutions for behavioral researchers. *Behavior research methods*, 46(1), 112–130.

- Chollet, F. (2020). A definition of intelligence for the real world. *Journal of Artificial General Intelligence*, 11(2), 27–30.
- Cushman, F., & Greene, J. D. (2012). Finding faults: How moral dilemmas illuminate cognitive structure. *Social neuroscience*, 7(3), 269–279.
- Daw, N. D., O’Doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, 441(7095), 876–879.
- De Freitas, J., & Johnson, S. G. (2018). Optimality bias in moral judgment. *Journal of Experimental Social Psychology*, 79, 149–163.
- Dennett, D. C. (1989). *The intentional stance*. MIT press.
- Duff, R. (2017). Thank god i failed. In *Moral puzzles and legal perplexities: Essays on the influence of larry alexander*. Cambridge University Press.
- Epley, N., Waytz, A., & Cacioppo, J. T. (2007). On seeing human: a three-factor theory of anthropomorphism. *Psychological review*, 114(4), 864.
- Fiske, S. T., Cuddy, A. J., & Glick, P. (2007). Universal dimensions of social cognition: Warmth and competence. *Trends in cognitive sciences*, 11(2), 77–83.
- Frank, R. H. (2016). *Success and luck: Good fortune and the myth of meritocracy*. Princeton University Press.
- Frederick, S. (2005). Cognitive reflection and decision making. *The Journal of Economic Perspectives*, 19(4), 25–42.
- Gao, T., & Scholl, B. J. (2011). Chasing vs. stalking: interrupting the perception of animacy. *Journal of experimental psychology: Human perception and performance*, 37(3), 669.
- Gergely, G., & Csibra, G. (2003). Teleological reasoning in infancy: The naive theory of rational action. *Trends in cognitive sciences*, 7(7), 287–292.

- Gerstenberg, T., Ullman, T. D., Nagel, J., Kleiman-Weiner, M., Lagnado, D. A., & Tenenbaum, J. B. (2018). Lucky or clever? from expectations to responsibility judgments. *Cognition*, 177, 122–141.
- Gilbert, D. T., & Malone, P. S. (1995). The correspondence bias. *Psychological bulletin*, 117(1), 21.
- Gino, F., Moore, D. A., & Bazerman, M. H. (2009). No harm, no foul: The outcome bias in ethical judgments. *Harvard Business School NOM Working Paper*(08-080).
- Godfrey-Smith, P. (2017). *Other minds: The octopus and the evolution of intelligent life*. Harper-Collins UK.
- Grice, H. P., Cole, P., & Morgan, J. (1975). Logic and conversation. 1975, 41–58.
- Heider, F. (1958). *The psychology of interpersonal relations*. Psychology Press.
- Heider, F., & Simmel, M. (1944). An experimental study of apparent behavior. *The American Journal of Psychology*, 243–259.
- Jain, Y. R., Callaway, F., & Lieder, F. (2019). Measuring how people learn how to plan. In *41st annual meeting of the cognitive science society (cogsci 2019)* (pp. 1956–1962).
- Jara-Ettinger, J., Gweon, H., Schulz, L. E., & Tenenbaum, J. B. (2016). The naïve utility calculus: computational principles underlying commonsense psychology. *Trends in cognitive sciences*, 20(8), 589–604.
- 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.
- Kepka, E. J., & Brickman, P. (1971). Consistency versus discrepancy as clues in the attribution of intelligence and motivation. *Journal of Personality and Social Psychology*, 20(2), 223—229.
- Keupp, S., Behne, T., & Rakoczy, H. (2018). The rationality of (over)imitation. *Perspectives on Psychological Science*, 13(6), 678-687. (PMID: 30301424) doi: 10.1177/1745691618794921

- Kool, W., Gershman, S. J., & Cushman, F. A. (2018). Planning complexity registers as a cost in metacontrol. *Journal of cognitive neuroscience*, 30(10), 1391–1404.
- Kun, A., & Weiner, B. (1973). Necessary versus sufficient causal schemata for success and failure. *Journal of Research in Personality*, 7(3), 197–207.
- Lee, D. (2020). *Birth of intelligence: From rna to artificial intelligence*. Oxford University Press.
- Legg, S., & Hutter, M. (2007). Universal intelligence: A definition of machine intelligence. *Minds and machines*, 17(4), 391–444.
- Lerner, M. (1980). *The belief in a just world; a fundamental delusion*. New York: Plenum Press.
- Liu, S., Ullman, T. D., Tenenbaum, J. B., & Spelke, E. S. (2017). Ten-month-old infants infer the value of goals from the costs of actions. *Science*, 358(6366), 1038–1041.
- Louie, K., Grattan, L. E., & Glimcher, P. W. (2011). Reward value-based gain control: divisive normalization in parietal cortex. *Journal of Neuroscience*, 31(29), 10627–10639.
- Luce, R. D. (1959). *Individual choice behavior: A theoretical analysis*. New York: Wiley.
- Mackie, D. M., Worth, L. T., & Allison, S. T. (1990). Outcome-biased inferences and the perception of change in groups. *Social Cognition*, 8(4), 325–342.
- Martin, J. W., & Cushman, F. (2016). The adaptive logic of moral luck. *The Blackwell companion to experimental philosophy*, 190–202.
- McGeer, V. (2007). The regulative dimension of folk psychology. In *Folk psychology re-assessed* (pp. 137–156). Springer.
- Nagel, T. (2013). Moral luck. In *The philosophy of free will: Essential readings from the contemporary debates* (pp. 24–38). Oxford University Press.

- Olson, K. R., Dunham, Y., Dweck, C. S., Spelke, E. S., & Banaji, M. R. (2008). Judgments of the lucky across development and culture. *Journal of Personality and Social Psychology*, 94(5), 757.
- Olson, K. R., Heberlein, A. S., Kensinger, E., Burrows, C., Dweck, C. S., Spelke, E. S., & Banaji, M. R. (2013). The role of forgetting in undermining good intentions. *PLoS ONE*.
- Pantelis, P. C., Gerstner, T., Sanik, K., Weinstein, A., Cholewiak, S. A., Kharkwal, G., . . . Feldman, J. (2016). Agency and rationality: Adopting the intentional stance toward evolved virtual agents. *Decision*, 3(1), 40.
- Pu, Y., Ellis, K., Kryven, M., Tenenbaum, J., & Solar-Lezama, A. (2020). Program synthesis with pragmatic communication. *Advances in Neural Information Processing Systems*, 33.
- Reeder, G. D., Pryor, J. B., & Wojciszke, B. (1992). Trait–behavior relations in social information processing. *Language, interaction and social cognition*, 37—57.
- Sacchi, S., & Cherubini, P. (2004). The effect of outcome information on doctors’ evaluations of their own diagnostic decisions. *Medical education*, 38(10), 1028–1034.
- Saxe, R., Carey, S., & Kanwisher, N. (2004). Understanding other minds: linking developmental psychology and functional neuroimaging. *Annu. Rev. Psychol.*, 55, 87–124.
- Schwitzgebel, E., & Cushman, F. (2012). Expertise in moral reasoning? order effects on moral judgment in professional philosophers and non-philosophers. *Mind & Language*, 27(2), 135–153.
- Shu, T., Kryven, M., Ullman, T. D., & Tenenbaum, J. B. (2020). Adventures in flatland: Perceiving social interactions under physical dynamics. In *42d proceedings of the annual meeting of the cognitive science society*.
- Surber, C. F. (1984). Inferences of ability and effort: Evidence for two different processes. *Journal of Personality and Social Psychology*, 46(2), 249.

- Todorov, A., Mandisodza, A. N., Goren, A., & Hall, C. C. (2005). Inferences of competence from faces predict election outcomes. *Science*, 308(5728), 1623–1626.
- Toplak, M. E., West, R. F., & Stanovich, K. E. (2011). The cognitive reflection test as a predictor of performance on heuristics-and-biases tasks. *Memory & cognition*, 39(7), 1275.
- Turing, A. M. (1950). Computing machinery and intelligence. In *Parsing the turing test* (pp. 23–65). Springer.
- Ullman, T., Baker, C., Macindoe, O., Evans, O., Goodman, N., & Tenenbaum, J. B. (2009). Help or hinder: Bayesian models of social goal inference. In *Advances in neural information processing systems* (pp. 1874–1882).
- Weisman, K., Dweck, C. S., & Markman, E. M. (2016). Varieties of experience: A new look at folk philosophy of mind. In *Proceedings of the thirty-eighths annual conference of the cognitive science society*.
- Welsh, M., Burns, N., & Delfabbro, P. (2013). The cognitive reflection test: How much more than numerical ability? In *Proceedings of the annual meeting of the cognitive science society* (Vol. 35).
- Williams, B. (1981). *Moral luck: philosophical papers 1973-1980*. Cambridge University Press.
- Wojciszke, B., Bazinska, R., & Jaworski, M. (1998). On the dominance of moral categories in impression formation. *Personality and Social Psychology Bulletin*, 24(12), 1251–1263.
- Zawidzki, T. W. (2013). Mindshaping. In *The oxford handbook of 4e cognition*.
- Zhi-Xuan, T., Mann, J., Silver, T., Tenenbaum, J., & Mansinghka, V. (2020). Online bayesian goal inference for boundedly rational planning agents. *Advances in Neural Information Processing Systems*, 33.