Beyond rationality: We infer other people's goals by learning agent-variable expectations of efficient action

Joan Danielle K. Ongchoco (joan.ongchoco@yale.edu)
Department of Psychology, Yale University

Julian Jara-Ettinger (julian.jara-ettinger@yale.edu)
Department of Psychology, Yale University

Abstract

Our ability to make sense of goal-directed behavior is central to social reasoning. From infancy, this capacity is structured around an assumption that agents act efficiently. But agents are often inefficient and how we move is affected by our emotional states and personal idiosyncrasies. How, then, does an assumption of efficiency allow us to accurately interpret people's actions? We hypothesized that people expect agents to move efficiently relative to an agent-specific baseline rather than to an objective notion of efficiency. Consistent with this, we found that people can quickly learn and subtract agent-idiosyncratic movements when interpreting goal-directed action (Experiment 1). Moreover, in a free-response task, people's propensity to explain superfluous movement in terms of goals depended on the agent's relative efficiency rather than on the path's objective efficiency (Experiment 2). Our results show that people flexibly adjust their expectations of efficiency by attending to how agents typically move.

Keywords: Movement reasoning; efficiency; action understanding; mental states; Theory of Mind

Introduction

One of the greatest complexities of human behavior is our capacity for movement. We can move in infinite combinations of facial, arm, and leg movements to pursue a wide range of physical, epistemic, and communicative goals. But beyond the complexity and functionality of physical movement, perhaps even more striking is how such movements can reveal more intangible states—what people are thinking and feeling. A fast and steady gait reveals confidence while an uneasy one betrays hesitation; an arm leaping forward suggests the person will grab an object, while one oscillating sideways indicates a greeting; and a person tapping their finger can reflect a mindless tic, or a song stuck in their head.

The assumption of efficiency

In simple contexts, inferences about other people's behavior are driven by an expectation that agents move efficiently in space (Gergely & Csibra, 2003; Jara-Ettinger, Gweon, Schulz, & Tenenbaum, 2016; Liu & Spelke, 2017; Scott & Baillargeon, 2013). This notion of efficiency tends to be agent-invariant—agents will move in the shortest possible path given environmental constraints. When agents deviate from the most efficient plan, we try to come up with explanations that justify the apparent inefficiency. If, for instance, we watch someone extend their hands upwards while in front of a bookcase, we can infer that their goal is to get a book from one of the top shelves. And if instead, we watched someone repeatedly move books around the bookcase until finally taking one from the top shelf, we can conclude that they did not know which book they wanted, or they were unsure about where to find it. In other words, when watching seemingly inefficient behavior (such as reshuffling books), we infer the mental states under which the actions can be conceived as efficient (e.g. this was an efficient way to search for the book).

Mental-state inferences from apparent inefficiencies explain how we infer other people's goals (Baker, Saxe, & Tenenbaum, 2009), preferences (Jara-Ettinger, Sun, Schulz, & Tenenbaum, 2018), beliefs (Baker, Jara-Ettinger, Saxe, & Tenenbaum, 2017), and even communicative intent (Royka, Aboody, & Jara-Ettinger, 2018). Moreover, even infants can perform efficiency-based inferences, suggesting that they are foundational to human action-understanding (Gergely, Nádasdy, Csibra, & Bíró, 1995; Liu & Spelke, 2017; Liu, Ullman, Tenenbaum, & Spelke, 2017).

But movements are inefficient

Consider, however, how such inferences would work in a more realistic situation. Imagine sitting across someone who is mindlessly bouncing their leg, or fiddling with a pen. Clearly, their behavior is meaningful, perhaps revealing restlessness or anxiety, but these movements are not directed towards any external goal. If the person suddenly began to act towards an external goal, such as reaching for an object, goal-recognition would require us to dismiss the leg movement as irrelevant and focus uniquely on the arms, allowing us to find a goal that renders the arm movements as efficient, without having to explain other superfluous movements.

These intuitions suggest that action-understanding involves a signal decomposition problem where we do not expect all movement to be efficiently directed towards a goal. Instead, we readily identify which movements reflect 'undirected' mental states such as anxiety and agent-specific idiosyncratic movements. Having identified these basic movements, we then infer agents' goals by expecting them to be efficient relative to their baseline movement, rather than relative to an objective and agent-independent standard of efficient movement.

The current study

Here we present two experiments that provide initial evidence for our proposal. In the first experiment, we took one of the most classical demonstrations of this expectation of efficiency—people tend to infer unobservable goals or constraints to explain inefficient movement (Baker et al., 2009; Gergely et al., 1995; Jara-Ettinger, Schulz, & Tenenbaum, 2019)—but added a critical prior step. Participants first watched agents of varying baseline efficiency. We varied this by adding more or less noise to their paths (see Figure 1a). Participants then watched a test action (see Figure 1b), and they were simply asked whether the agent pursued an unobservable goal (see Figure 1c). If people infer goals relative to an objective level of efficiency, they should make identical inferences for identical movements, independent of the agent's baseline motion. By contrast, if people infer goals relative to an agent-specific level of efficiency, then their inferences should depend on the agent's inefficiency at baseline, so they should feel less inclined to attribute inefficiencies to unobservable constraints or goals.

In a second experiment, instead of asking participants to infer whether the agent pursued an unobservable goal, we now asked them to simply explain the agent's movements in their own words. This way of probing people's reasoning can reveal whether people are indeed less likely to explain away observed inefficiencies by appealing to unobservable goals or constraints when they see an inefficient agent at baseline; but it can also reveal how people do reason about such inefficiencies, if not by defaulting to these goals or constraints (e.g., explaining the inefficiency by appealing to emotional or dispositional states).

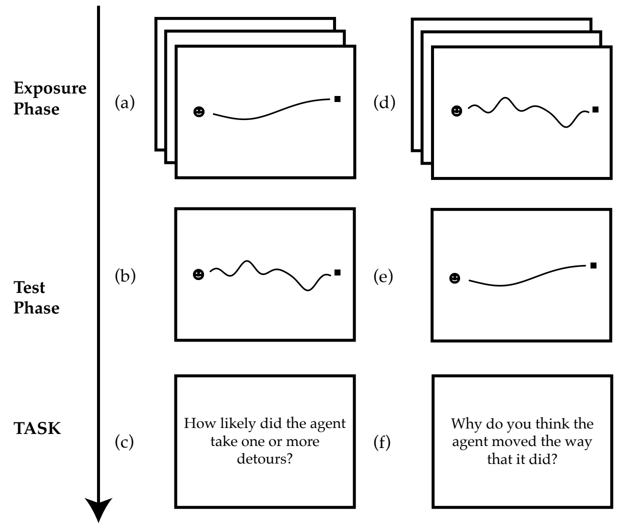

Figure 1: A caricatured depiction of the sample trial in Experiments 1 and 2. The two experiments were identical in the first two phases and only differed in the test question. In the first two phases we varied whether the agent was efficient or inefficient. (a) The agent moves to a goal in an efficient path (i.e. minimally inefficient). (b) The agent moves towards a goal in an inefficient path. (c) Experiment 1 task. (d) The agent moves in an inefficient path. (e) The agent moves in an efficient path. (f) Experiment 2 task.

Experiment 1a

Do we interpret inefficient motion relative to an agent-specific level of efficiency? Participants were first exposed to varying agent baseline efficiencies, and were asked to infer whether the agent completed unobservable goals at test.

Method

Data and code for all experiments reported here are available on `https://osf.io/h42kc`.

Participants. 120 participants were recruited through Amazon Mechanical Turk (Crump, McDonnell, & Gureckis, 2013). This sample size was split across four conditions (30 participants per condition), and was chosen before data collection began.

Apparatus. The experiment was conducted using the Qualtrics online survey platform (`http://www.qualtrics.com`), and custom software written using CSS and HTML. Individuals could not participate more than once or in any of the experiments reported here. Because our experiment required the viewing of embedded videos, each participant also completed a simple browser compatibility check to ensure that they could do this (this also served as our attention check). Only participants who passed this check were allowed to continue.

Stimuli. Animations were created via custom software written in Python with the PsychoPy libraries (Peirce et al., 2019) and were presented as embedded videos in .mp4 format (800 by 600 pixels). Stimuli consisted of simple animations of an agent (grey disc, 0.7° in diameter, with a happy face) moving in a two-dimensional space (see Fig. 1). In each animation, the agent navigated towards a goal (a 0.9° image of a red jewel) in an efficient or inefficient way. These paths were generated by sampling two functions from the range $[0,1]$ from a Gaussian process with a squared exponential kernel (Seeger, 2004) with variance $\sigma^2 = 0.1$, and lengthscale $\ell = 0.1$ for inefficient paths, and $\ell = 0.3$ for efficient paths ($n = 4$ paths generated for each set). The benefit of generating paths this way is that we can now test for efficiency beyond the dichotomy of 'straight' vs. 'curved' paths, and so none of our 'efficient' paths were ever straight.

Design. A participant could be assigned to one of four possible conditions in a 2x2 design. They could either be exposed to efficient or inefficient agents, and then tested with either efficient or inefficient paths. Thus, in some conditions, the efficiency of the agent in the Exposure and Test phases matched (e.g. Expose Inefficient - Test Inefficient), and in other conditions, they did not (e.g. Expose Inefficient - Test Efficient).

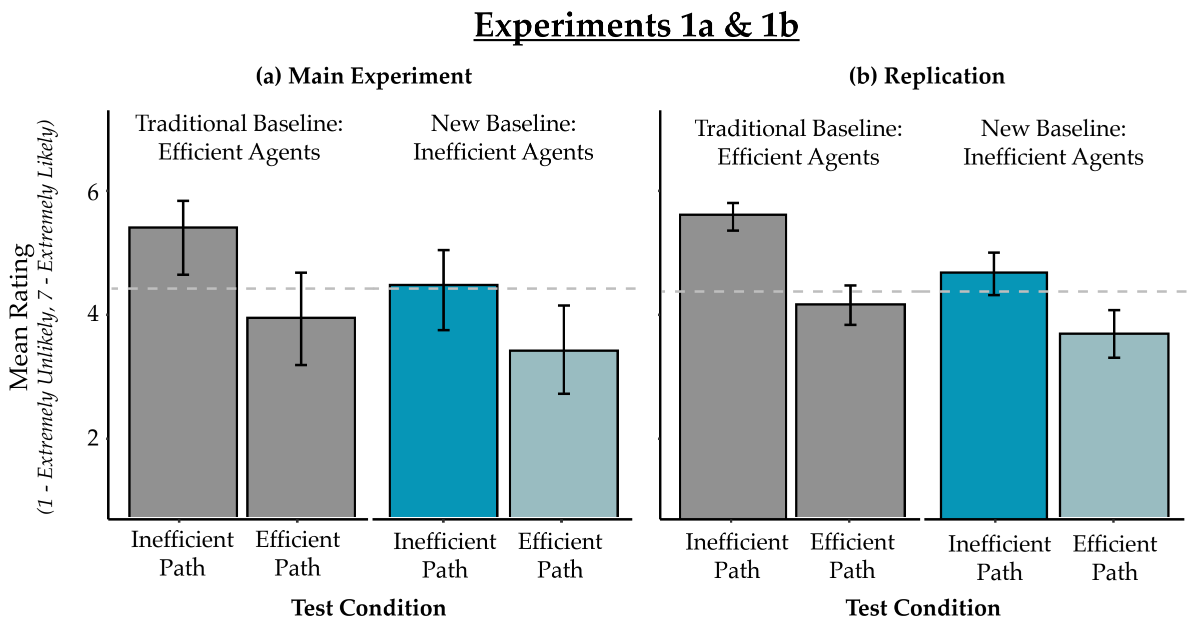

Figure 2: Results from Experiments 1a and 1b. The bars depict mean ratings. Error bars reflect 95% corrected and accelerated bootstrapped confidence intervals. The dashed line reflects the baseline rating for efficient paths with no exposure phase.

Procedure. In the 'Exposure' phase, each participant watched three animations one at a time (randomly drawn from either the Inefficient or Efficient set of paths). Participants were told that they were watching three animations of the same agent move towards a red jewel. In the 'Test' phase, participants watched one final animation. Participants never saw the same path twice. When the Exposure and Test phase matched (e.g., Expose Efficient - Test Efficient) the set of paths was split into two, where three of the four paths were used in the Exposure phase and the remaining other path was then used for the Test phase. When they did not match, paths for the Exposure phase were drawn from one set, and the Test path was simply randomly drawn from the other set. In the Test phase, participants were told that they would watch the same agent move towards one more jewel, but this time, there may or may not be other jewels hidden in the display. Participants were told that, although they could not see the extra jewels, the agent could always see and collect all jewels. Their task was to rate how likely there was to have been other jewels with the prompt, 'How likely did the agent take one or more detours?' (with 1 being Extremely Unlikely, and 7 being Extremely Likely).

To interpret our results, we also obtained a baseline measure of people's propensity to infer unobservable goals from a single efficient path (without an exposure phase). 120 unique participants were recruited were simply asked to rate how likely an agent might have collected a hidden jewel, given a path drawn from the Efficient set. Thus, in our main task, values above this baseline would show that people are more likely to infer an unobservable goal relative to when they see a single agent move efficiently in space. Conversely, values above this baseline would show that people are less likely to infer an unobservable goal relative to when they see a single agent move efficiently in space. Finally, values that are not significantly different from this baseline would suggest that people's propensity to infer an unobservable goal is similar to the default expectation that they have when seeing efficient action.

Results and Discussion

The average baseline rating for generally efficient paths was 4.41 (depicted by the grey line in Figure 2). Average ratings across the four relevant conditions were compared against this baseline (see Figure 2a). When participants saw efficient agents (as revealed in the Exposure phase) and efficient paths (shown in the Test phase), their ratings were at 'floor', and there was indeed no difference in their ratings from the baseline (3.93 vs. 4.41, $t(29)=1.23$, $p=.229$, $d=.22$). Conversely, when participants saw efficient agents take inefficient paths, they were more likely to explain the inefficiency with the detours the agent might have taken (5.40 vs. 4.41, $t(29)=3.37$, $p=.002$, $d=0.62$). These results suggest that when participants expect the agent to be efficient, they reason about the efficiency of paths the way we might expect from previous work—deviations from the path towards a goal are costly, unless we can explain this by the presence of other goals or constraints (Baker et al., 2009; Jara-Ettinger et al., 2019; Gergely et al., 1995).

The critical conditions, however, were when we challenged this rationality assumption, and participants should have no longer expected agents to be efficient. When participants saw inefficient agents take inefficient paths, they were less likely to attribute this inefficiency to detours, and more likely to say the paths were generally 'efficient'—returning to baseline (4.47 vs. 4.41, $t(29)=0.17$, $p=.861$, $d=0.03$). And when participants saw inefficient agents take efficient paths, they were even less likely to attribute inefficiency to detours (3.40 vs. 4.41, $t(29)=2.70$, $p=.011$, $d=0.49$).

To analyze the data as a whole, but still with a focus on our primary question, we ran a linear regression predicting participant judgments as a function of change in efficiency, coded as three levels: -1 for a decrease in inefficiency (going from inefficient to efficient), 0 for no change in inefficiency, and 1 for an increase in inefficiency (going from efficient to inefficient). Change in efficiency significantly predicted responses ($\beta = 1.00$, $p < 0.001$).

These results suggest that when reasoning about the efficiency of actions, the mind does not only take into account agent-independent factors like the cost or degree of deviation from a straight path towards a goal. Moreover, perhaps most surprising is that our expectations of an agent's rationality, whereas in previous work seen as a 'default' assumption, may actually be more flexible than we think. And our expectations for the baseline efficiency of an agent influences how we reason about its goal-directed actions.

Experiment 1b

Method

This experiment was a pre-registered replication of Experiment 1a `https://aspredicted.org/blind.php?x=nh3w3d`. We ran a power analysis on the results from Experiment 1a and determined a sample size of 480 participants (120 participants per condition).

Results

The pattern of results in this replication exactly mirrored that of Experiment 1a (see Figure 2b). When participants saw efficient agents (as revealed in the Exposure phase) and efficient paths (shown in the Test phase), there was indeed no difference in their ratings from the baseline (4.08 vs. 4.41, $t(119)=1.98$, $p=.050$, $d=.18$). Conversely, when participants saw efficient agents take inefficient paths, they were more likely to explain the inefficiency away with the 'detours' the agent might have taken (5.65 vs. 4.41, $t(119)=10.75$, $p<.001$, $d=0.98$). When participants saw inefficient agents take inefficient paths, they were again less likely to attribute this inefficiency to detours, and more likely to say the paths were generally 'efficient'—returning to baseline (4.71 vs. 4.41, $t(119)=1.73$, $p=.086$, $d=0.16$). And when participants saw inefficient agents take efficient paths, they were even less likely to attribute inefficiency to detours (3.68 vs. 4.41, $t(119)=3.68$, $p<.001$, $d=0.34$). Change in efficiency again significantly predicted participant responses in a linear regression ($\beta = 0.99$, $p < 0.001$).

Experiment 2

Results from previous experiments were surprising in that the movements that people have deemed inefficient across so much of the work on action reasoning (Baker et al., 2009; Gergely et al., 1995; Jara-Ettinger et al., 2019; Liu & Spelke, 2017) were significantly *less* inefficient by mere (and incredibly sparse!) exposure to the statistics of how an agent typically moves. This is the first demonstration, to our knowledge, of a more sophisticated signal decomposition mechanism that separates what we know about the agent from what we are watching the agent performing *now*. If people are less likely to judge movements as inefficient when the baseline is inefficient, how are they then explaining such inefficiencies at test? In Experiment 2, we asked people to explain an agent's movements in a free-response version of the task. We also included a "violation-of-expectation" prompt, where we asked people how surprised they were by the movement at test.

Method

The experiment was identical to Experiments 1a and 1b, except where noted. 120 participants were recruited through Amazon Mechanical Turk. After watching the animations of the agent's baseline movements, participants now watched the test animation, after which they were asked to explain 'why the agent moved the way it did' (see Figure 2). Participants were also asked how surprised they were by the agent's movements in a 7-point Likert scale (where 1 was for 'Not at all surprised' and 7 for 'Extremely surprised').

Results and Discussion

Two coders blind to the experimental conditions coded all free responses into three possible categories: (1) responses where people inferred hidden goals or obstacles; (2) responses where people did not infer hidden goals or obstacles, and just described the movement; (3) or responses that could not be categorized in either of the two. A third coder served as a tie breaker in cases where the two coders disagreed (n=10 trials). Figure 3a shows the proportion of participants who spontaneously inferred an unobservable goal or obstacle, revealing a similar pattern of results with those of Experiment 1. People were most likely to mention unobservable goals or constraints when the agent was efficient at baseline, but inefficient at test (such that the leftmost bar in the plot was the highest). Conversely, people were least likely to mention unobservable goals or constraints when the agent was inefficient at baseline, but efficient at test (such that the rightmost bar in the plot was the lowest). At the critical conditions, when the movement of the agent at baseline matched their movements at test, the likelihoods of mentioning unobservable goals were not different (such that the middle bars in the plot were roughly equal). When they didn't infer unobservable goals, participants either simply inferred

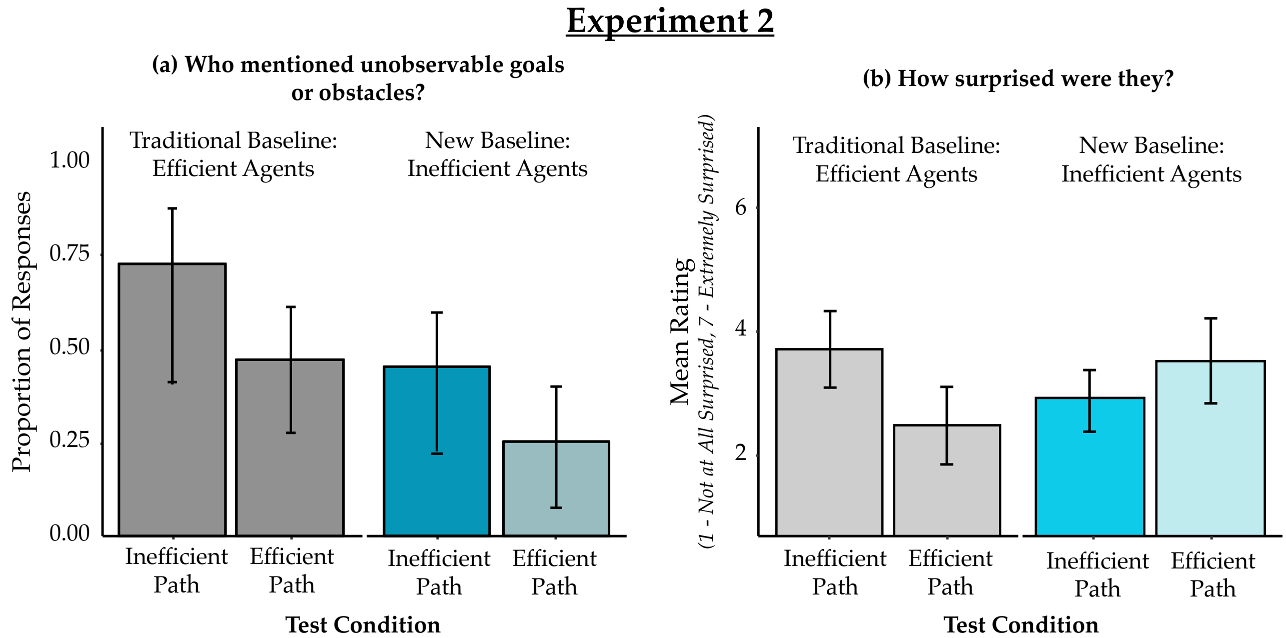

Figure 3: Results from Experiment 2. (a) The bars depict proportion of responses. (b) The bars depict the mean ratings. Error bars reflect 95% corrected and accelerated bootstrapped confidence intervals.

a single goal, or perhaps most intriguingly, noted instead the dispositions of the agent, describing the agent's movement for its 'energy' or speed (see Figure 4). These impressions were confirmed in a binomial logistic regression, where responses were coded as either 1 (for mentioning unobservable goals) or 0 (for not doing so), and predicted against change in efficiency, coded again as three levels: -1 for a decrease in inefficiency (going from inefficient to efficient), 0 for no change in inefficiency, and 1 for an increase in inefficiency (going from efficient to inefficient). Change in efficiency again significantly predicted responses ($\beta = 1.02$, $p = .002$).

Responses to the surprise prompt also suggest that participants were more surprised when the agent moved in a way that deviated from their baseline, i.e. more surprised when an inefficient agent moved in an efficient way, and an efficient agent moved in an inefficient way (see Figure 3b). Because we were interested in the interaction, we ran a linear regression predicting participant ratings as a function of the Exposure condition, Test condition, and the interaction between the two. The interaction significantly predicted surprise ratings ($\beta = -1.99$, $p = .003$), while beta-values for Exposure and Test conditions were not significant. This confirms that people were tracking the agent baseline, and were surprised when the agent deviated from this.

General Discussion

The work on movement perception and reasoning has so far been split into two dominant categories. There is the work on biological motion perception, and how we infer emotions and intentionality from motion patterns (Atkinson, Tunstall, & Dittrich, 2007; Dittrich & Lea, 1994; Troje, 2013), and there is the work on goal-directed action (Gergely & Csibra, 2003; Jara-Ettinger et al., 2016; Liu & Spelke, 2017; Scott & Baillargeon, 2013). People are certainly sensitive to how and why agents move in both senses, yet these bodies of work have mostly been explored independently of each other. How might they interact? Here we demonstrate that how any agent moves at baseline matters for how we reason about its goal-directed actions. Across two experiments, we found that people discounted inefficiencies of previously inefficient agents, and as a consequence, were less likely to resort to unobservable goals or constraints to explain away the inefficiency. (And in future work, we intend to run a pre-registered replication of the second experiment to further confirm these results).

Because we used unique stimuli involving unique movement paths and patterns, our results suggest that we track and pick-up agents' baseline movements quickly, even those we had never seen before. The tracking of this baseline efficiency relates to work in vision that explores how we recognize motion patterns despite each motion never being exactly quite the same, (e.g., a ball never really bounces the same way twice) (Cavanagh, Labianca, & Thornton, 2001). One open question is whether the filtering of 'noisy' inefficiencies (such as minor variability in one's movement) operates differently than the filtering of high-level structured inefficiencies (such as tapping one's foot).

The current study suggests that we extract and discount

	Why did the agent move the way it did? Representative answers
Decreased Efficiency from Baseline	I think it dipped down to collect an unseen jewel. Because there were no other jewels in his way I think it was the easiest way for him to catch the jewel. It was easier and a shorter route. Wanted to be slow and correct. Because the agent was not trying to look suspicious.
Stable Efficiency	It was energetic. To reach the single jewel. Maybe it needed to avoid something. Perhaps they thought it was the quickest way to get it. To gain momentum going foward.
Increased Efficiency from Baseline	There were invisible jewels that the agent was collecting. It wants more points. There are several hidden jewels in the path. I think he moved the way he did because he was getting passed the other jewels that were hidden. I think the agent stopped to pick up two unseen jewels before reaching the red one.

Figure 4: Participants answers to the question, "Why did the agent move the way it did?" in Experiment 2. Representative answers are categorized here according to the change in efficiency that resulted from the particular conditions they were assigned to.

an agent's baseline movements, allowing us to reason about the remaining goal-directed action under the assumption of efficiency. This does not mean, however, that the subtracted movements are then considered irrelevant or no longer carry information. This baseline can still reveal an agent's kinematics (e.g. inferring health from an inefficient limp) or their emotional state (e.g. inferring energy from excited ballistic movements). Previous work has found that even in the absence of a goal, people continue to reason about the intentions of the agent by resorting to 'movement-based goals', such that the agent may have intended to perform or carry out a sequence of actions (Schachner & Carey, 2013). If anything, our results similarly suggest that people reason in terms of these movement-based goals, but that they also go beyond this, referring in their responses to the properties of the movement, such as speed or energy, or even more subtle intentions (as in Figure 4, where one of our participants even answered that the agent was 'not trying to look suspicious'). Future work can explore the nature of these different kinds of inferences—of 'undirected' and goal-directed features of movements alike—and how they might relate to each other, or even be interpreted from within the same framework.

Perhaps the most significant contribution of this work though is its implications for our theory of 'rational action'. We note that the logic in our experiments has been shared in various versions of control experiments in past studies. At first pass, our results may seem intuitive—of course, people must constantly be tracking the statistics of another agent's movements. It was never clear, however, how prior patterns of movement would interact with people's reasoning, such that these would even change how people reason about efficiency. In fact, one particularly relevant finding from recent work is that prior movement patterns do *not* affect expectations of efficiency—infants continue to expect agents to be objectively efficient even when the agent is seen pursuing goals in an inefficient way (Liu & Spelke, 2017). While more research is needed, these results combined with our own, suggest that learning to adjust our inferences based on agent-specific inefficiencies may require experience watching how agents generally behave.

Our results ultimately demonstrate with adults that the prior efficiency of an agent matters. Previous work has suggested that exposure to 'non-rational' agents may lead us to abandon our 'intentional' stance, and thus not interpret the actions from an expectation of efficiency (Gergely et al., 1995). We suggest here that the mechanism may be more sophisticated than this, such that we don't just entirely abandon rationality when we don't find it. Rather, we simply estimate and use a different baseline level of efficiency, one that takes into account how people more naturally move, that is, in incredibly complex, meaningful, expressive, and crucially, *not always* goal-directed ways.

References

Atkinson, A. P., Tunstall, M. L., & Dittrich, W. H. (2007). Evidence for distinct contributions of form and motion information to the recognition of emotions from body gestures. *Cognition*, *104*, 59–72.

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., & Tenenbaum, J. B. (2009). Action understanding as inverse planning. *Cognition*, *113*, 329–349.

Cavanagh, P., Labianca, A. T., & Thornton, I. M. (2001). Attention-based visual routines: Sprites. *Cognition*, *80*, 47–60.

Crump, M. J., McDonnell, J. V., & Gureckis, T. M. (2013). Evaluating amazon's mechanical turk as a tool for experimental behavioral research. *PloS one*, *8*, e57410.

Dittrich, W. H., & Lea, S. E. (1994). Visual perception of intentional motion. *Perception*, *23*, 253–268.

Gergely, G., & Csibra, G. (2003). Teleological reasoning in infancy: The naıve theory of rational action. *Trends in Cognitive Sciences*, *7*, 287–292.

Gergely, G., Nádasdy, Z., Csibra, G., & Bíró, S. (1995). Taking the intentional stance at 12 months of age. *Cognition*, *56*, 165–193.

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*, 589–604.

Jara-Ettinger, J., Schulz, L., & Tenenbaum, J. (2019, Dec). *The naive utility calculus as a unified, quantitative framework for action understanding.* PsyArXiv. Retrieved from `psyarxiv.com/e8xsv` doi: 10.31234/osf.io/e8xsv

Jara-Ettinger, J., Sun, F., Schulz, L., & Tenenbaum, J. B. (2018). Sensitivity to the sampling process emerges from the principle of efficiency. *Cognitive Science*, *42*, 270–286.

Liu, S., & Spelke, E. S. (2017). Six-month-old infants expect agents to minimize the cost of their actions. *Cognition*, *160*, 35–42.

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*, 1038–1041.

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

Royka, A., Aboody, R., & Jara-Ettinger, J. (2018). Movement as a message: Inferring communicative intent from actions. In *Proceedings of the 40th annual conference of the cognitive science society.*

Schachner, A., & Carey, S. (2013). Reasoning about 'irrational'actions: When intentional movements cannot be explained, the movements themselves are seen as the goal. *Cognition*, *129*, 309–327.

Scott, R. M., & Baillargeon, R. (2013). Do infants really expect agents to act efficiently?: A critical test of the rationality principle. *Psychological Science*, *24*, 466–474.

Seeger, M. (2004). Gaussian processes for machine learning. *International journal of neural systems*, *14*, 69–106.

Troje, N. F. (2013). What is biological motion: Definition, stimuli and paradigms. *Social perception: Detection and interpretation of animacy, agency, and intention*, 13–36.