Intention regulates conflicting desires in human decision making

Shaozhe Cheng¹, Ning Tang¹, Yang Zhao¹, Jifan Zhou^{1*}, Mowei Shen^{1*} & Tao Gao^{2,3*}

¹Department of Psychology and Behavioral Sciences, Zhejiang University

²Department of Communication, UCLA

³Department of Statistics, UCLA

*Corresponding author. Email: jifanzhou@zju.edu.cn, mwshen@zju.edu.cn, or tao.gao@stat.ucla.edu.

Abstract

It is an ancient insight that human actions are driven by desires. This insight inspired the formulation that a rational agent acts to maximize expected utility (MEU), which has been widely used in psychology for modeling theory of mind and in artificial intelligence (AI) for controlling machines' actions. Yet, it's rather unclear how humans act coherently when their desires are complex and often conflicting with each other. Here we show desires do not directly control human actions. Instead, actions are regulated by an intention — a deliberate mental state that commits to a fixed future rather than taking the expected utilities of many futures evaluated by many desires. Our study reveals four behavioral signatures of human intention by demonstrating how human sequential decision-making deviates from the optimal policy based on MEU in a navigation task: “Disruption resistance” as the persistent pursuit of an original intention despite an unexpected change has made that intention suboptimal; “Ulysses-constraint of freedom” as the proactive constraint of one's freedom by avoiding a path that could lead to many futures, similar to Ulysses's self-binding to resist the temptation of the Siren's song; “Enhanced legibility” as an active demonstration of intention by choosing a path whose destination can be promptly inferred by a third-party observer; “Temporal leap” as committing to a distant future even before reaching the proximal one. Our results showed how the philosophy of intention can lead to discoveries of human decision-making, which can also be empirically compared with AI algorithms. The findings showing that to define a theory of mind, intention should be highlighted as a distinctive mental state in between desires and actions, for quarantining conflicting desires from the execution of actions.

Keywords

action, conflicting desires, intention, commitment, theory of mind, sequential decision-making

Intention regulates conflicting desires in human decision making

Humans, like all animals, act to fulfill their desires (Aristotle, 330BC; Hume, 1978). Nevertheless, these desires are not coherent as they are constantly advocating for different interests (Elster, 1987; Schelling, 1984; Searle, 2003; Sigmund, 1920; Tversky & Shafir, 1992). An everlasting debate between desires could lead to indecisiveness with wavering actions that lack the resolution to achieve any desired outcome. How do coherent actions emerge from these complex, even conflicting desires?

In philosophy, the classic desire theory argues that despite their complexity, desires, in combinations with beliefs, are sufficient for directly generating coherent actions (Audi, 1974; Davidson, 1963). To act intentionally is to act rationally for fulfilling desires (Dennett, 1989). In decision theory and artificial intelligence (AI), this rationality principle has been formalized as the maximization of expected utility (MEU) (Morgenstern & von Neumann, 1953; Russell & Norvig, 2009). Here utility is defined as a scalar measuring the desirability of an action's future outcomes. The complexity of desires can be addressed by taking the expectation of different possible outcomes, resulting in a summed utility of different futures weighted by their probabilities. The power of MEU has been demonstrated by modern AI, such as deep reinforcement learning, which has achieved super-human level performance in many challenging games (Silver et al., 2016; Vinyals et al., 2019). MEU has also been incorporated into cognitive science, especially theory of mind (ToM) — humans' spontaneous attribution of beliefs and desires to others' actions (Wellman, 2014). A Bayesian ToM (BToM) further models the attribution of mental states by inverting the rationality principle to infer the most likely belief-desire combinations given actions (Baker et al., 2009; Jara-Ettinger et al., 2016).

However, the desire model may lack a critical mental representation — intention, which is the deliberate state of mind between desire and action (Bratman, 1987; Harman, 1986; Searle & Willis, 1983). Acting intentionally is not only to fulfill desires, but also to executively regulate desires, so that their persisting conflicts could be quarantined from the execution of actions (Bratman, 1987). To support this theory, philosophical analysis of language shows that intention and desire are semantically different. Unlike desires

whose strength can be quantified weaker or stronger as measured by a utility function, intention is a constraint that can only be dichotomously satisfied or not (Brand, 1984). Moreover, in contrast to desires that can be fulfilled in a number of ways, intention must be “satisfied in the right way” — accidentally killing one's archenemy in a car crash doesn't satisfy the intention of murdering that enemy (Searle & Willis, 1983). These philosophical analyses have been supported by empirical human studies using introspections and self-reports (Malle & Knobe, 2001; Perugini & Bagozzi, 2004; Schult, 2002).

The necessity of intention has also been analyzed for the purpose of pragmatic decision-making. In analogy to a parliament, intention functions to resolve the debate between opposing desires. Unlike desire, intention doesn't consider the expectation of many outcomes evaluated by many desires, but a proactive commitment to a plan to achieve one fixed future (Bandura, 2001). Thus, intention must be “admissible” for executing a sequence of actions — an agent may have conflicting desires but cannot have conflicting intentions (Bratman, 1987). The feasibility of this philosophical theory of intention has been shown in a few modeling studies. Early work of logical AI formalizes intention as the selection of a goal for persistent pursuit (Cohen & Levesque, 1990). More recently, intention has also been modeled as the ordering of reaching multiple goals that can be inferred by BToM (Jara-Ettinger et al., 2020).

Here we explore whether humans indeed regulate conflicting desires with intention in their decision-making. Our behavioral experiments focused on how the prospect of future events influences humans' current action: taking the expected desirability of many futures or intentionally committing to one of them. Our paradigm was inspired by the paradox of Buridan's ass, an ancient thought experiment speculating that an ass, placed in the middle of two equally desirable piles of hay, may end up starving to death due to indecisiveness. In our experiments, the task of an agent, either controlled by a human or a MEU model (see Supplement, MEU model), was to navigate to one of two equally desirable restaurants (destinations) which were located apart from each other in a 2D map. It is noteworthy that in contrast to the initial philosophical analysis, a purely desire-driven MEU agent will not “starve to death”. It can break the tie of conflicting desires by sampling an action from a stochastic policy, formulated as

Markov Decision Process (MDP) that can be solved by classic control algorithms such as value iteration (Bellman, 1966; Sutton & Barto, 1998). Hence, instead of focusing on the overall task performance, we engineered critical moments in the navigation, at which the agent would act differently depending on whether it followed a plan to a fixed future or a MEU policy that maximizes the expected utility of many futures. Following this strategy, we explored four behavioral signatures of intentions in three experiments.

Experiment 1: Disruption resistance

Experiment 1 explored the “disruption resistance” nature of intention as the persistent pursuit of a future despite unexpected disruptions and setbacks that have made that future suboptimal. A disruption was introduced as a “drift” that nullified the agent’s action by placing the agent in one of the nearby cells except its intended position (see Fig. 1A).

Method

There were 10 trials in total. In the first nine trials, just as participants were instructed, the disruption occurred randomly with a 10% probability at every timestep, resulting in roughly 1 drift per trajectory. In the last trial, without participants' awareness, the disruption was not random but deliberately engineered — it was triggered when the agent first revealed its destination by executing an action towards one destination while away from the other. This deliberate disruption was against the agent’s action, by placing it in the cell in the opposite direction of its action (see Fig. 1B). As a result, the agent ended up closer to the destination not revealed by its action. An intention-driven agent should commit to the original destination by fighting against the drift and getting back to its planned course of action. In contrast, a desire-driven agent should just move towards the other destination as it brings higher expected utilities.

Participants

Across all experiments, sample sizes were determined before data collection, no additional data were collected after experiment began. As there is no similar prior study, Experiment 1 used a suggested sample size in social science (Simmons et al., 2013). We

recruited fifty undergraduate and graduate students (27 females, $M_{age} = 21.32$, $SD = 2.07$) from the Zhejiang University participants' pool, they participate experiments for credits or payments. Samples were not intended to be representative of any population, because we assume the intentional nature of actions applies to all populations. This study and the following studies were approved by the Zhejiang University. All participants were given informed consent.

Fig. 1. Humans commit to the original destination despite disruptions. (A) Design of the random disruptions. Both the time step and the direction of the disruptions were randomly sampled. (B) Design of the deliberate disruption. Both the time step and the direction of the disruptions were deliberately designed, to push the agent away from the destination the moment it was revealed. (C) Percentage of reaching the original destination with different types of disruptions. $*P < 0.05$; $***P < 0.001$. Error bars indicate 95% confidence intervals.

Procedure

Participants were instructed that they were players controlling a hungry agent to reach any of the two restaurants (destinations) as soon as possible. They could control the agent with the four arrow keys (up, down, left, and right) on a standard keyboard. They were clearly explained that the environment was not fully deterministic: at every step, there was a 10% probability that the agent's action could be disrupted by a random drift, which could randomly push the agent to a nearby cell. A trial ended once the agent reached a destination, immediately followed by a new trial with a new map (see Supplement for detailed design of maps).

Across all experiments, participants performed the task individually in a single room using laboratory computers. Researchers who collected the data were blind to study hypothesis during data collection.

Results

As predicted, with the deliberate disruption, humans' averaged percentage of reaching the original destination ($N = 50$; 70%; 95% CI [0.57,0.83]) was significantly higher than that from MEU (0%, averaged from simulations equal to the number of human trials) (see Fig. 1C, independent t -test, $t(98) = 10.69$, two-tailed $P < 0.001$, Cohen's $d = 2.14$). With random disruptions, both humans and MEU agents reached the original destination with a high percentage. Still, humans' percentage (96%; 95% CI [0.93,0.98]) was significantly higher than that of MEU (92%; 95% CI [0.89,0.94]) ($t(98) = 2.43$, $P = 0.017$, $d = 0.49$). These results collectively demonstrate the “disruption resistance” nature of human intention as committing to a destination in an uncertain environment. MEU was far less likely to commit to a destination and demonstrated less 'destination perseverance' than shown by human participants.

Experiment 2: Ulysses-constraint of freedom

Experiment 2 explored the “Ulysses-constraint of freedom” as the proactive constraint of one's freedom by avoiding a path that could lead to many futures, named after Ulysses, who bound himself to a mast to resist the temptation of the Siren's song (Elster, 2000). The opportunity of “self-binding” was presented at a crossroad with two paths: an open-ended path that could lead to two destinations, or a fixed-future path that leads to only one destination (see Fig. 2A for an exemplary map). A desire-driven agent should show no preference, as the expected utilities of taking these two paths were identical. An intention-driven agent may prefer the fixed future path by which the agent can demonstrate their commitment to an intention. This self-demonstration of intention can enhance the legibility of human trajectories. That is, from a third-party perspective, inferring the destination of a human trajectory should be easier than that of a machine trajectory. We tested this hypothesis by using a BToM model (see Supplement: Bayesian theory of mind) that can infer the destinations of trajectories in real time.

Fig. 2. Humans constrained their freedom with a bias towards the fixed-future path.

(A) An agent at a crossroad. Agents can either choose a fixed-future path that leads to one destination (yellow arrow) or choose an open-ended path that could lead to both destinations (green arrow). This is a 4-steps condition defined as the length of the shortest path (blue arrow) between the starting position and the crossroad. (B) Design of the deliberate disruption. Once the agent reveals its destination, it will be immediately pushed to a position equally distanced from the two destinations. (C) Averaged of the percentage of choosing the fixed-future paths across steps-to-crossroad conditions. (D) Percentage of choosing the fixed-future paths as a function of steps-to-crossroad. (E) The posterior of the BToM inference of the actual destination by an agent as a function of steps. The error bars/shading reflect 95% confidence intervals. $**P < 0.01$, $***P < 0.001$.

Method

As the intentional commitment may take time and effort (Bratman, 1987; Tversky & Shafir, 1992), we also manipulated when the agent would face the crossroad. From the starting positions, the steps of action required to reach the crossroad varied from [0, 1, 2,

4, 6] steps, with 48 trials for each step (see Supplement, Experiment 2; see Supplement Fig. S1 for sample maps). Similar to Experiment 1, a random drift could appear with 6.7% probability at each step (roughly 1 drift per trajectory), to push the agent to one of 7 nearby cells.

The task and environment of Experiment 2 were identical to Experiment 1, except the variety of maps was greatly increased by both introducing barriers and manipulating the agent's distances to the two destinations (see Supplement: Experiment 2 for detailed design of maps). We further replicated the “disruption resistance” with these richer set of maps. A special trial with a deliberate disruption was appended to the end of Experiment 2. When the agent arrived at a position where its destination was first revealed, a drift would drag the agent back to a position equally distanced from the two destinations (Fig. 2B).

Participants

Despite a large effect size (Cohen's $d = 2.14$) observed in Experiment 1, we chose a conservative effect size (Cohen's $d = 0.8$) to determine the sample size in Experiment 2. A power analysis (power = 0.8, alpha = 0.05) indicated that 20 human participants were sufficiently needed in Experiment 2. Thus, we recruited 20 participants (15 females, $M_{age} = 20.85$, $SD = 2.03$) in the same fashion as in Experiment 1. All participants were given informed consent.

Results

Ulysses-constraint of freedom

Overall, humans ($N = 20$) preferred the fixed-future path (63%; 95% CI [0.56, 0.69]) over the open-ended path (37%) (one sample t -test with a 50% baseline, $t(19) = 4.17$, two-tailed $P < 0.001$, $d = 0.93$; Fig. 2C). By contrast, as expected, MEU showed no preference ($t(19) = -0.51$, $P = 0.62$). Furthermore, humans' preferences developed over time, plateauing at 4-steps (see Fig. 2D). For humans, the main effect of steps-to-crossroad was significant ($F(4, 195) = 25.43$, $P < 0.001$, $\eta^2 = 0.34$). They chose the fixed-future path more at steps 2, 4, 6, as revealed by one sample t -test with a 50% baseline. Post-hoc analyses further indicated this preference was much stronger at steps 4

and 6, compared with step 2 (both $P < 0.001$, see Supplement, Table S1). In contrast, MEU showed no bias to the fixed future path at all steps. We also replicated this effect when the distances to the two destinations varied slightly (see Supplement, Fig. S2). These results showed that humans prefer a path locked to a fixed destination. Moreover, this preference was not instantly revealed but developed gradually. This is consistent with the theory that unlike desires often arise effortlessly, intention is a deliberate process that requires time and effort (Harman, 1986).

Enhanced legibility

The posterior of BToM inference of the actual destination (finally reached) was plotted in Fig. 2D. We focus on the 6-steps condition, as it had sufficient time steps for humans to establish a commitment. Humans revealed their actual destination much faster than MEU (cluster-based permutation tests identified significant gaps from steps 6 to 11, all $P < 0.05$). As this is an individual navigation task without an observer at all, the enhanced legibility of human trajectory can be viewed as a self-demonstration of intention.

Disruption resistance

The result of “disruption resistance” was consistent with Experiment 1, humans demonstrated stronger destination perseverance (95%; 95% CI [0.85,1] by reaching the original destination much more often than MEU (55%; 95% CI [0.31,0.79]) (independent t -test, $t(19) = 3.21$, $P = 0.003$, $d = 1.02$).

Experiment 3: Temporal leap

Experiment 3 explored the “temporal leap” nature of intention as committing to a distant future even before reaching the proximal one. An intention promises to bring one future about. This stability of the future enables an agent to concatenate multiple intentions, with one starting from the fixed future promised by the previous one. Therefore, the agent can form a partial plan with a long horizon by leaping forward from one promised future to the next one without concerning the gaps between them. It implies

that the agent will be biased towards the very next promised future in the chain, even when new emerging opportunities at that time have made that future suboptimal.

Fig. 3. Navigating to a chain of two destinations in experiment 3. (A) At any moment there were always two destinations (green squares) for an agent (blue circles) to choose. (B) Once the agent reached a destination, the destination (A) disappeared, the other destination (B) stayed now as the old destination, and a new destination (C) appeared at a new location. The distance difference in this map is -5 (5-10). (C) The agent chose to pursue the old destination B, (D) alternatively, the agent chose to pursue the new destination C.

Method

We tested the temporal leap hypothesis with a Pac-Man like task, requiring participants to pursue a stream of destinations, with the constraint that at each moment there were always exactly two destinations to choose from. A trial ended when the agent reached one destination, immediately followed by another trial, with the old destination leftover from the previous trial and the presentation of a new destination (see Fig. 3).

A temporal leap implies a bias towards reaching the old destination as it has already been planned into the chain. We explored this bias by manipulating the old and the new destination's relative distances to the agent. The difference in their distances was evenly chosen from [-5, -3, -1, 0, 1, 3, 5], with positive values indicating the old was closer (see Supplement, Experiment 3). There were 45 trials for each condition.

Participants

The sample size was determined in the same way as in experiment 2. A total of 20 participants (12 females, $M_{age} = 21.3$, $SD = 2.13$) joined this experiment. Participants were recruited in the same way as in Experiments 1 & 2. All participants were given informed consent.

Results

Temporal leap

For the equal-distances (0 difference) condition, humans ($N = 20$; 61.3%; 95% CI [0.54,0.68]) reached the old destination significantly more often than MEU (48.3%; 95% CI [0.44,0.52]) (independent t -test, $t(38) = 3.31$, $P = 0.0023$, $d = 1.05$). Across all conditions, the averaged distances difference was 0, yet humans reached the old destination (56.9%; 95% CI [0.52,0.62]) significantly more often than MEU (49.7%; 95% CI [0.49,0.50]) (Fig. 4a, $t(38) = 2.79$, $P = 0.0083$, $d = 0.88$). Furthermore, we modeled the percentage of reaching the old destination as a function of distance differences with logistic regression (Fig. 4b). For humans, the intercept (0.42; 95% CI [0.36,0.49]) was significantly higher than 0 ($P < 0.001$), demonstrating a bias toward the old destination. For MEU, the intercept was not significantly different from 0 (-0.08; 95% CI [-0.21,0.06], $P = 0.26$). Similar bias was also observed in reaction time, across all conditions, humans took less time to initiate the first action towards the old destinations (176ms; 95% CI [113,239]) than the new destinations (206ms; 95% CI [134,279]) (paired t -test, $t(19) = 4.27$, two-side $P < 0.001$, $d = 0.14$).

Fig. 4. Humans commit to a distant future with a bias towards reaching old destinations. (A) Percentage of trials in which agents reached the old destination, from the equal-distances condition and averaged across all distance conditions. (B) Result of all distance conditions. Percentage of choosing the old destination and the curves fitted by logistic regressions. (C) The posterior of BToM inference of the destination finally reached by an agent over the temporal course of trajectories from the equal-distances condition. The error bars/shading reflect 95% confidence intervals.

Enhanced legibility

The human bias to the old destination was also revealed by the legibility of trajectories. By applying BToM to the equal-distances condition, we found that humans revealed their destinations more quickly when their destinations were the old ones compared with the new ones (Fig. 4C, cluster-based permutation tests identified significant

gaps from 6.7% to 46.7% of humans' trajectory, all $P < 0.05$). MEU showed no such bias. BToM of all distances-difference conditions showed similar patterns, albeit with a larger variance (see Supplement, Fig. S3). These results collectively demonstrate the temporal leap nature of human intention as establishing a bias to a distant future, even before reaching a proximal one.

Discussion

Our study shows that the human mind regulates conflicting desires through intention with a commitment to a fixed future, making human actions more predictable and explainable, at the cost of deviating from an optimal policy for maximizing the expected utility. Our empirical results showed that humans are surprisingly inflexible — they cling to prior inertia and resist re-planning, even when the environmental changes have made their intention suboptimal. This may in part be because humans are bounded by a limited computational resource. Sticking to a plan can certainly reduce the budget of online decision-making. Still, the extent of the inflexibility is striking considering the simplicity of the navigation tasks we used here — Experiments 1 and 3 used maps without barriers, which can hardly be further simplified. This suggests that the inflexibility may also serve certain purposes other than saving computational cost. Being inflexible means being predictable. This was supported by our results showing that BToM could more effectively predict the destination of human navigation, indicating the legibility of human actions was better than those of MEU. Such predictability can be a great advantage in real life. Unlike AI often designed to solve a specific task, humans often face tasks embedded in a complex network. Being predictable can greatly facilitate the coordinating of these intertwined tasks, both intrapersonal and interpersonal.

For intrapersonal coordination, humans always need to coordinate multiple temporal events with their future selves. Being predictable enables a partial plan with temporal gaps between sub-plans (Bratman, 1987), supported by the temporal leap results in Experiment 3. One seemingly paradoxical feature of partial planning is that a distant future could be more predictable than the proximal future. For instance, while I don't have a plan for the next week, I have already decided to fly to Paris for the next Olympics followed by a visit to the Musée du Louvre. This temporal leap nature is in sharp contrast

with algorithms based on MDP, which demand that the planned trajectories must be continuous. The contrast between humans' partial plans and MDP's continuous plans highlights the importance of intention for intrapersonal coordination of tasks with temporal gaps.

For interpersonal coordination, the predictability enabled by intentions allows one to coordinate with others both synchronically and diachronically. Intention demands commitment, which is essential for teamwork — in hunting a lion together, we won't stand a chance unless we both commit to it simultaneously and persistently. Any of my partners' flexibility of that commitment will place me in peril. For these reasons, commitment has been mostly studied in the context of collaboration (Gilbert, 2013; Tomasello et al., 2005). Under this context, one not only needs to form an intention with commitment, but also to demonstrate that intention to others. This demonstration of intention has been explored in models of multi-agent interaction (Dragan et al., 2013; Ho et al., 2016; Shafto et al., 2014), in which a demonstrator picks an action to facilitate the observers' Bayesian inference of the demonstrator's mind. Interestingly, we found a similar demonstration of intention in Experiment 2. Despite the fact it was a purely individual task, the experiment showed that the destinations of human trajectories were much easier to predict than those of MEU agent with the only concern of reaching a destination as soon as possible. In the effect of the Ulysses-constraint of freedom, humans prefer an inflexible path that forces the agent to pursue a fixed destination — as a demonstrator would do to convince an observer of their commitment to an intention. Indeed, with this constraint of freedom, BToM is able to predict the humans' destination more effectively. This self-demonstration of intention supports the hypothesis that while ToM is originally developed to understand others, due to evolutionary pressures of cooperation and communication, it has been internalized to monitor one's own actions (Vygotsky, 1980), so that the agent makes its own mind more explainable and predictable to a third-party observer from an intentional stance (Dennett, 1996).

Our results showed how the philosophy of intention can lead to discoveries of human decision making, which can also be empirically compared to an AI algorithm. These results collectively demonstrated that in a theory of mind, intention should be clearly defined as a distinctive mental state for quarantining conflicting desires from the

execution of actions. Such a theory of mind can also support a more human compatible AI in the future, which acknowledge the complexity of desires (Russell, 2019) and regulate them with intentions for more predictable and explainable interactions with humans.

References

- Aristotle. (2000). *The Nicomachean Ethics* (K. Ameriks & D. M. Clarke, Eds.). Cambridge University Press (Originally published around 330 B.C.).
- Audi, R. (1974). Intending. *The Journal of Philosophy*, 70(13), 387–403.
- 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(4), 1–10.
- Baker, C. L., Saxe, R., & Tenenbaum, J. B. (2009). Action understanding as inverse planning. *Cognition*, 113, 329–349.
- Bandura, A. (2001). Social cognitive theory: An agentic perspective. *Annual Review of Psychology*, 52(1), 1–26.
- Bellman, R. (1966). Dynamic programming. *Science*, 153(3731), 34–37.
- Brand, M. (1984). *Intending and acting: Toward a naturalized action theory*. MIT press.
- Bratman, M. (1987). *Intention, plans, and practical reason*. Harvard University Press Cambridge, MA.
- Cohen, P. R., & Levesque, H. J. (1990). Intention is choice with commitment. *Artificial Intelligence*, 42(2–3), 213–261.
- Davidson, D. (1963). Actions, reasons, and causes. *The Journal of Philosophy*, 60(23), 685–700.
- Dennett, D. C. (1989). *The intentional stance*. MIT press.
- Dennett, D. C. (1996). *Kinds of minds: Toward an understanding of consciousness*. Basic Books.
- Dragan, A. D., Lee, K. C. T., & Srinivasa, S. S. (2013). Legibility and predictability of robot motion. *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 301–308.
- Elster, J. (1987). *The multiple self*. Cambridge University Press.
- Elster, J. (2000). *Ulysses unbound: Studies in rationality, precommitment, and constraints*. Cambridge University Press.
- Gilbert, M. (2013). *Joint commitment: How we make the social world*. Oxford University Press.
- Harman, G. (1986). *Change in view: Principles of reasoning*. MIT Press.

- Ho, M. K., Littman, M. L., MacGlashan, J., Cushman, F., & Austerweil, J. L. (2016). Showing versus doing: Teaching by demonstration. *Advances in Neural Information Processing Systems* 29.
- Hume, D. (1978). *A treatise of human nature*. Oxford University Press.
- Jara-Ettinger, J., Gweon, H., Schulz, L. E., & Tenenbaum, J. B. (2016). The naive utility calculus: Computational principles underlying commonsense psychology. *Trends in Cognitive Sciences*, 20(8), 589–604.
- Jara-Ettinger, J., Schulz, L. E., & Tenenbaum, J. B. (2020). The naive utility calculus as a unified, quantitative framework for action understanding. *Cognitive Psychology*, 123, 101334.
- Malle, B. F., & Knobe, J. (2001). The distinction between desire and intention: A folk-conceptual analysis. In B. F. Malle, L. J. Moses, & D. A. Baldwin (Eds.), *Intentions and intentionality: Foundations of social cognition*. MIT Press.
- Morgenstern, O., & von Neumann, J. (1953). *Theory of games and economic behavior*. Princeton University Press.
- Perugini, M., & Bagozzi, R. P. (2004). The distinction between desires and intentions. *European Journal of Social Psychology*, 34(1), 69–84.
- Russell, S. (2019). *Human compatible: Artificial intelligence and the problem of control*. Penguin.
- Russell, S., & Norvig, P. (2009). *Artificial intelligence: a modern approach*. Prentice Hall.
- Schelling, T. C. (1984). *Choice and consequence*. Harvard University Press.
- Schult, C. A. (2002). Children's understanding of the distinction between intentions and desires. *Child Development*, 73(6), 1727–1747.
- Searle, J. (2003). *Rationality in action*. MIT press.
- Searle, J., & Willis, S. (1983). *Intentionality: An essay in the philosophy of mind*. Cambridge University Press.
- Shafto, P., Goodman, N. D., & Griffiths, T. L. (2014). A rational account of pedagogical reasoning: Teaching by, and learning from, examples. *Cognitive Psychology*, 71, 55–89.

- Sigmund, F. (1920). *A general introduction to psychoanalysis*. Createspace Independent Publishing Platform.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., & others. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587), 484–489.
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2013). Life after P-Hacking. *Meeting of the Society for Personality and Social Psychology*.
- Sutton, R. S., & Barto, A. G. (1998). *Introduction to reinforcement learning*. MIT Press.
- Tomasello, M., Carpenter, M., Call, J., Behne, T., & Moll, H. (2005). Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and Brain Sciences*, 28(5), 675–691.
- Tversky, A., & Shafir, E. (1992). Choice under conflict: The dynamics of deferred decision. *Psychological Science*, 3(6), 358–361.
- Vinyals, O., Babuschkin, I., Czarnecki, W. M., Mathieu, M., Dudzik, A., Chung, J., Choi, D. H., Powell, R., Ewalds, T., Georgiev, P., & others. (2019). Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575(7782), 350–354.
- Vygotsky, L. (1980). *Mind in society: The development of higher psychological processes* (M. Cole, V. John-Steiner, Scribner Sylvia, & Souberman Ellen, Eds.). Harvard University Press.
- Wellman, H. (2014). *Making minds: How theory of mind develops*. Oxford University Press.

Supplemental Materials

MEU model.

We employed the Markov decision process (MDP) as an implementation of the desire-model following the MEU principle, in which desires were defined as the reward function and the agent acted to maximize its expected long-term future rewards. The definition of an MDP includes a state space S , an action space A , a transition function, $T(s, a)$; a utility function, $R(s, a)$. The solution of MDP is an optimal policy π , which takes s as input, and outputs a probabilistic distribution of action given s , $P_\pi(a|s)$. The agent acted by sampling an action from this distribution. The above definition and the solution of MDP did not involve a formulation of intention.

State Space. The agent's state was its location, defined as a tuple with 2D coordinates $(X_coordinate, Y_coordinate)$. In Experiments 1 & 3, the size of the state space is 225, including every cell in a 15*15 map. In Experiment 2, cells occupied by a barrier were excluded from the state space.

Action Space. In all experiments, the agent can travel one cell in four directions: $a \in \{(0,1), (1,0), (0,-1), (-1,0)\}$.

Transition function. The transition function takes state s and action a as input, outputs $p(s')$, a probability distribution of the next states'. In all experiments, the two destinations were set as the termination state. In Experiment 1, the agent can only reach 4 nearby states. With the probability of 9/10, the agent moved to the cell in the direction of its action, with the other nearby cells evenly split the rest of the probability of 1/10 (transition noise). In Experiment 2, the noise can push the agent to 8 nearby cells. With the probability of 14/15, the agent moved to the cell in the direction of its action, with the other 7 nearby cells evenly split the rest of the probability of 1/15. Experiment 3 was identical to experiment 1, except the transition noise was 0.

Reward function. The reward function takes the state s and action a as input and outputs a scalar as the short-term reward. In all experiments, the reward had two components: 30 for reaching any of the destinations, -1/30 for every movement on the map.

Solving MDP. The optimal policy of MDP was solved by value iteration using the Bellman optimality equation (Bellman, 1966) of the value function V . It's an iterative bootstrapping process.

At time $t+1$, the new value function V_{t+1} is derived from the value function V_t . The optimal value function V^* can be found when this iterative process converges:

$$V^*(S) = \max_a \left[R(s, a) + \gamma \sum_{s' \in \mathbb{S}} p(s' | s, a) V^*(s') \right] \quad (1)$$

Where γ is the discount factor. In all MDP simulations, γ was fixed to 0.9, which is a value normally used.

Policy. The optimal policy can be derived from the optimal value function V^* in two steps:

First, we derived an optimal action-value Q^* from V^* :

$$Q^*(s, a) = R(s, a) + \gamma \sum_{s' \in \mathbb{S}} p(s' | s, a) V^*(s') \quad (2)$$

Then, a Boltzmann policy with the probability of taking an action a given its state s , was derived proportional to $Q^*(s, a)$.

$$P_{\pi}(a | s) \propto \exp(\beta Q^*(s, a)) \quad (3)$$

The Boltzmann policy takes β as a rationality parameter. When $\beta \rightarrow 0$, the agent will act in a more random way; when $\beta \rightarrow \infty$, the agent chooses the action greedy based on optimal Q-value. Here we chose $\beta = 2.5$ following previous studies modeling human action with MDP (Baker et al., 2009, 2017). With this value, the action will be dominated by the maximum $Q(s, a)$, but still deviate from it with a small probability, to capture the fact that human decision-making is not entirely rational.

Bayesian theory of mind (BToM). We used BToM (Baker et al., 2009) to infer the agent’s destination over time. As there were only two destinations, we only need to plot the posterior of the destination the agent actual reached $P_{reached}$. The posterior of the destination not reached was always $1 - P_{reached}$. Given an agent’s trajectory (the state-action pair up until $T \geq t$), the posterior of the agent’s destination was proportional to the product of the action likelihood and the prior probability of the destination:

$$P_{\beta}(\text{destination} \mid \text{action}_{1:T}, \text{state}_{1:T}) \propto \prod_{t=1}^T P_{\beta}(\text{action}_t \mid \text{destination}, \text{state}_t) * P(\text{destination}) \quad (4)$$

The action likelihood function $P_{\beta}(\text{action}_t \mid \text{destination}, \text{state}_t)$ was derived from an MDP policy, which was similar to the policy in Equation (3), except it only considered one destination as its goal. The initial $P_0(\text{destination})$ was set to 0.5 for both potential destinations.

Experiment 1.

Materials and Design.

Each trial consisted of a map with an agent and two destinations. Each map was randomly generated with two constraints. First, the Manhattan distances (the distance between two points measured along axes at right angles) between the agent’s starting position and two destinations were equal, ranging from 7 to 13, with 10 as the mean distance. Second, the agent’s starting position and the two destinations form an isosceles triangle, with the agent placed along the perpendicular bisector of the invisible line connecting the two destinations. This invisible triangle was randomly rotated by an angle randomly sampled from $[0, 90, 180, 270]$ degrees (see Fig. S4)

In the first nine trials, disruptions occurred randomly with a probability of 10%, resulting in roughly 1 drift per trajectory. In practice, the drifts were pseudo-randomly generated for each participant with one constraint: there were 3 disruptions in every 3 trials. There was no constraint on how these disruptions were distributed within these 3 trials.

Experiment 2.***Materials and Design.***

The maps in Experiment 2 were enriched in the following two ways:

First, depending on the design of the barriers, there were three types of maps: (a) Critical-crossroad maps, in which the barriers make certain cells a critical crossroad, from where an agent can choose between a fixed-future path or an open-ended path. Certain barriers were purposely placed so that (i) an agent cannot reach any destinations without encountering a critical crossroad; (ii) if there were more than one critical crossroad, the length of the shortest path to each critical crossroad was always identical; (iii) this length was manipulated according to the steps-to-crossroad condition introduced in the main text. To increase the variety of maps, the critical crossroads were created either by just one barrier or two barriers that aligned vertically or horizontally. For example, for the 6-steps-to-crossroad condition, it takes 7 barriers to create two one-barrier critical crossroads that participants could reach in 6 steps (see Fig. S5A). To cover the purpose of the experiment by making the critical crossroads less salient, we also added other barriers randomly scattered on the map with the constraints that they would not block the shortest path to each of the destinations (see Fig. S5B). There were 240 critical-crossroad maps, making up 5/6 of total trials. (b) Ambiguity-zone maps, in which the barriers were carefully placed to create an ambiguity-zone, within which the agent's action can not reveal its destination. It took at least 10 steps for a rational agent to move outside of this zone to reach its destination (see Fig. S5c). In the last trial, once the agent left this zone, a deliberate disruption would drag it back to the ambiguity-zone, therefore, the agent was once again placed in a position with equal distances to the two destinations. Before the last trial with the deliberate disruption, there were 24 ambiguity-zone maps but with random disruptions, intermingled with all other types of maps. The purpose of these maps was to cover the last trial of deliberate disruption with the same type of map and further increase the variety of the maps so that the critical crossroads were less salient. (c) Random-barriers maps, in which all barriers were placed randomly with only one constraint: the agent can reach at least one destination without being trapped by barriers. There were 24 random-barriers maps (1/12 of the total trials) to increase the variety of

maps so that both the critical crossroad and the ambiguity-zone became less salient. Across all types of maps, the total number of barriers in each map was fixed to 18.

Second, the positions of the agent and the two destinations were also systemically manipulated, so that the agent's distances to them vary from trial to trial, ranging from 10 to 20, with 15 as a mean distance. Unlike experiment 1, here the length of the shortest path for the agent to reach the two destinations was not always identical. The distance differences varied from $[0, 1, 2]$, to avoid participants using any heuristic to just pick one destination without even looking at the barriers of the map.

Experiment 2 consisted of 289 trials. The first 288 trials were designed to test the effect of Ulysses-constraint of freedom, with each of five steps-to-crossroad conditions consisting of 48 trials, plus another 24 trials of ambiguity-zone maps and 24 trials of random-barriers maps ($48 \times 5 + 24 + 24 = 288$). In all these trials, the random disruption occurred with a probability of $1/15$ (roughly 1 drift per trajectory). The last trial was designed to test the disruption resistance with a deliberate disruption, using an ambiguity-zone map.

Experiment 3.

Materials and Design.

Maps of Experiment 3 were similar to Experiment 1, with only one agent and two destinations on the map, without any barriers. The major difference was that this experiment required a continuous navigation task: once the agent reached one destination, the current map would not disappear, instead, it was updated with (a) the disappearance of the reached destination, and (b) the presence of a new destination at a new location, indicating the beginning of a new trial. Participants were told to reach as many destinations as possible during this continuous game.

At the beginning of the experiments, the two destinations were placed with equal distance to the agent. After this, the position of each newly presented destination was controlled so that its distances to the agent met the distances-difference condition. In addition, to maximize the divergence of the paths to each destination, so that the 'old destination - agent - new destination' angle was maximized, within the constraint of the

distances-difference. The transition function of the agent in this experiment was deterministic, without any random or deliberate disruptions.

There were 315 trials in total. Each trial was pseudo-randomly assigned to one of seven distances-difference conditions from $[-5, -3, -1, 0, 1, 3, 5]$, with each condition consisting of 45 trials.

Fig. S1. Sample maps with different steps-to-crossroad in Experiment 2.

Fig. S2. Humans' preference to a fixed future path, averaged across the distances-difference conditions in Experiment 2. Percentage of choosing the fixed-future paths as a function of steps-to-crossroad. The error bars indicate 95% confidence intervals. ** $P < 0.01$, *** $P < 0.001$.

Fig. S3. The human bias toward reaching old destinations averaged across all distances-difference conditions in Experiment 3. The posterior of BToM inference of the destination finally reached by an agent over the temporal course of trajectories. The error shading reflects 95% confidence intervals.

Fig. S4. Sampled maps used in Experiment 1. The agent (blue circle) and the two destinations (green squares) form an invisible isosceles triangle. Each map in Experiment 1 was randomly assigned one of the invisible isosceles triangles, whose orientation, width, and height were randomized across trials.

Fig. S5. Designs of barriers in Experiment 2. (A) The barriers form two critical crossroads. (B) Extra barriers scattered within the red region mark in the map. Barriers in this region will not block any shortest path to any destinations. (C) Barriers forming an ambiguity-zone which is marked in the yellow region. One sampled path within this zone is marked by the numbers. Within this zone, the path cannot reveal the agent's destination. Any deviation from this zone within 9 steps will lead to a suboptimal path. However, moving out of this zone at step 10 will clearly reveal the agent's destination.

Table S1. Post-hoc analysis showing humans' preference to the fixed-future path across steps-to-crossroad conditions in Experiment 2.

Steps-to-crossroad	1	2	4	6
0	0.01	0.21***	0.41***	0.35***
1	-	0.20***	0.40***	0.34***
2		-	0.20***	0.14**
4			-	-0.06

Post-hoc comparison using Tukey's HSD. Mean differences are shown. **indicates $P < 0.01$, ***indicates $P < 0.001$.