Intentional Commitment as Goal Perseverance in Human Planning

Shaozhe Cheng[1]
shaozhecheng@zju.edu.cn

Minglu Zhao[3]
minglu.zhao@ucla.edu

Jingyin Zhu[1]
zhujy@zju.edu.cn

Jifan Zhou[1]
jifanzhou@zju.edu.cn

Mowei Shen[1]
mwshen@zju.edu.cn

Tao Gao[2,3]
tao.gao@stat.ucla.edu

[1] Department of Psychology and Behavioral Sciences, Zhejiang University
[2] Department of Communication, UCLA [3] Department of Statistics, UCLA

Abstract—**Human mind is a mosaic composed of multiple selves with conflicting desires. How can coherent actions emerge from such conflicts? The classical desire model of rational actions depends on maximizing the expected utilities evaluated by all desires. In contrast, the intention model suggests that humans regulate conflicting desires with an intentional commitment that constrains action planning towards a fixed goal. Here, through a 2D navigation game where humans were instructed to decide between two goals, we explored whether human actions will demonstrate a distinctive commitment to intention, as compared to a pure desire-driven agent who acts only to maximize the expected utility. Results indicated that humans spontaneously commit to an intention when facing conflicting desires. The pursuit is persistent even when unexpected disruptions make the prior intention sub-optimal, indicating the unique "goal perseverance" nature of intention. We further explored two functional hypotheses of intentional commitment: computational constraint hypothesis and social origin hypothesis. Followed experiments showed that, first, when humans were given enough time to plan, "goal perseverance" decreased but still existed, suggesting that intentional commitment may partially result from humans' limited resources, but it's not all. Second, when humans are involved in social contexts such as when being observed or when having another unrelated player in the game, "goal perseverance" was enhanced, supporting the social origin hypothesis of intentional commitment. These findings shed light on the uniqueness of human agency and its potential social origin.**

I. Introduction

Humans are purposive agents who act to fulfill their desires. While most Artificial Intelligence (AI) systems today are designed to solve one specific task, humans in real life typically need to decide between conflicting desires. As mentioned in the Odyssey, the ancient Greek hero Ulysses wanted to hear the Siren's song, yet he was also eager to get back to his homeland safely without being seduced by the Siren. In everyday life, we mundane people also constantly experience this contradiction within ourselves. People suffer from conflicting desires as if they have multiple selves: parts of you want longevity while another part is addicted to alcohol [26]. This multiple-selves dilemma has long been discussed in philosophy [27, 13] and psychology [16], yet, debates remain on how rational actions can be generated when people desire conflicting things.

A. Desire Model

The classical model of rationality defines rational actions as the ones that are expected to fulfill desires, as captured by Hume's famous claim that "Reason is and ought to be the slave of the passions". Following this tradition, the classical philosophical model of actions asserts that desires, despite their complexity and incompatibility, are sufficient for directly generating coherent actions when combined with beliefs [11, 2]. For example, my intention to turn on the light is adequately explained by my desire to have a clear view and my belief that I can see better with the light on. While this seems to be intuitive when the desire is simple, how could the desire model handle complicated situations with multiple desires in conflict with each other?

Modern decision theory based on the desire-belief model offers a solution using the probabilistic nature of expectation calculation. In the mathematical formulation, desires can be treated as a utility function that takes a state as input and outputs a scalar representing the desirability of the state. Multiple-desire problems can be considered as having multiple sources of utility for reaching different states. A rational agent does not need to make the hard choice among desires and can simply act to maximize expected utility (MEU), where the expectation is jointly evaluated by the probability of all future states and how well the states can satisfy all desires [39, 24]. In this tradition, there is no limitation on the complexity of desires, as long as the expected utility is fulfilled. Here, the critical part of generating rational actions is to fulfill a given desire, but not to choose one from multiple desires. This line of rationality echoes Aristotle's assertion that deliberation is always about means, never about ends.

The desire model, like models based on Markov Decision Process (MDP) [7], has a profound impact on modern AI. Models such as deep reinforcement learning can solve MDPs approximately in high dimensions and generate complex intelligent behaviors, reaching human-expert level performance in games like Atari [22] and Go [31, 32]. The desire model is also prominent in cognitive psychology, particularly in the studies of Theory of Mind (ToM). It has been assumed that humans spontaneously explain others' actions by attributing them to a combination of beliefs and desires [15, 41, 42]. Under this framework, researchers take the desire model as the default model of the mind that needs to be inferred. More recently, ToM has been modeled by Bayesian inverse planning [3] with two components: a forward planning process that assumes an agent acts rationally based on its mental states, formalized as $P(action|mind)$; an inverse planning process that uses the Bayes rule to infer the mind as being able to explain the planning process of observed actions: $P(mind|action) \propto P(action|mind)P(mind)$. The key of this model is the planning engine that computes $P(action|mind)$.

B. Intention beyond Desire

However, contemporary philosophers [29, 19, 8] argue that there is a gap between desires and the generation of rational actions: a deliberate process irreducible to a simple complex of desires and beliefs. Searle claims that in general, only a small portion of rational actions are sufficiently induced by beliefs and desires. Consider a drug addict with a predominant desire to take heroin and takes whatever he believes to be heroin. Here the drug addict's belief and desire are sufficient to determine the action, but that is hardly a model of rationality.

To fill the gap, generation of rational actions must presuppose the process of deliberation. In psychology and neuroscience, the studies of volition [18], self-regulation [5], self-control [1, 14], and goal pursuit [30, 23] suggest that there is a unique process for the human mind beyond desires and passions, which lies at the core of human agency. It has been suggested that intention is a mental state distinct from desires, which is defined as the deliberate choice among potential desires and the commitment to a course of action [8]. In this belief-desire-intention model, desires do not directly drive human actions but are instead mediated by intentions [29, 19]. Intention-based actions do not consider the expectation of all future states evaluated by all desires but are committed to bringing about one fixed future [6, 10]. Therefore, any conflicting nature of desires must be "filtered out" before forming an intention to act: an agent is allowed to desire conflicting things but not to intend conflicting things [8, 28]. In other words, intention serves as a resolution settling the debate between conflicting desires: the course of actions is determined once the intention is established, and the execution of those actions will no longer be influenced by unchosen desires.

C. Why Commit to an Intention?

1) Computational constraint hypothesis: Humans are resource limited agents. Commitment may serve as a bounded rationality [33], allowing agents with limited computational resources, like humans, to apply a prior deliberation for a future conduct without being time-sliced agents who always start from scratch [8]. Empirical studies on commitment first explored by economists demonstrated that humans are not fully rational as utility-maximizers due to the fact that their preferences may change over time, referred to as the "changing tastes" problem [34]. To forestall the changing tastes, commitment has been proposed as a regulation device to deal with the temporal fluctuations of preferences [35, 25, 26, 9]. Intuitively, considering the expectation of all desires demands massive computational resources, while committing to a fixed intention helps reduce the computational burden of online decision making for rapid changes both of tastes and of the environment.

2) Social origin hypothesis: Humans are social animals, the unique aspects of human mind may be shaped by social interactions [40, 28, 20, 37]. Being committed means being predictable to others. This resonates with the evolutionary perspective that humans evolved with a high color contrast between the white sclera and the darker colored iris to better convey attention with different displacements of the gaze [21, 38]. Similarly, commitment makes humans' intentions more readable, suggesting a social origin hypothesis that humans' commitment emerges from a communicative intention as people try to better demonstrate themselves to others. This also supports the hypothesis that while ToM was originally developed to understand others, it has been internalized to monitor one's own actions due to the evolutionary pressure of cooperation and communication [40, 20], where an agent makes its own mind more explainable and predictable from an intentional stance [12]. This social perspective has also been explored in developmental psychology, showing that children demand commitment from their partners, regulate partners when they are being difficult, and are voluntarily constrained by joint commitment [37].

D. Present Study

Here we examine the commitment nature of intention using an individual planning task. The task was deliberately designed to follow an MDP setting, where commitment is not necessary and may even compromise the optimality. We started by exploring whether humans demonstrate distinctive commitment to intention, as compared to a baseline desire-MDP model that acts to maximize the expected utilities over many desires. Then, we examined two functional hypotheses of intention, computational constraint hypothesis and social origin hypothesis. To test the computational constraint hypothesis, we studied whether an increase in computational resources would make people less biased to commitment. To test the social origin hypothesis, we examined whether people's commitment to intention will be enhanced when they are involved in certain social contexts.

II. Experiment

The experimental task of an agent was to navigate to one of two equally desirable destinations located apart from each other on a 2D map. We introduce unexpected disruptions during the agent's navigation in order to explore whether people would show commitment to one fixed destination even when disruptions made it sub-optimal. A disruption was introduced as a "drift" that nullified an agent's action by placing the agent in one of the nearby cells except for its intended position (see Fig. 1). In cases when a carefully engineered drift places the agent closer to the destination that the agent was not pursuing, we were particularly interested in whether the agent will continue with its previous goal or change its goal. A pure desire-driven agent would always move towards the other destination as it brings higher expected utilities. However, if human planning involves a deliberate process beyond desires, they would show commitment to the original destination by fighting against the drift to return to the planned course of action. Moreover, we are interested in to what extent will the commitment, if it exists, be affected by its potential functional roles. Here we explored two functional hypotheses of intentional commitment: computational and social functions. If commitment to intention serves to reduce computational burden, an increase in the time resource for deliberation would make people less committed. If commitment to intention serves to support potential social interactions, people in social contexts would show more commitment than when they act alone.

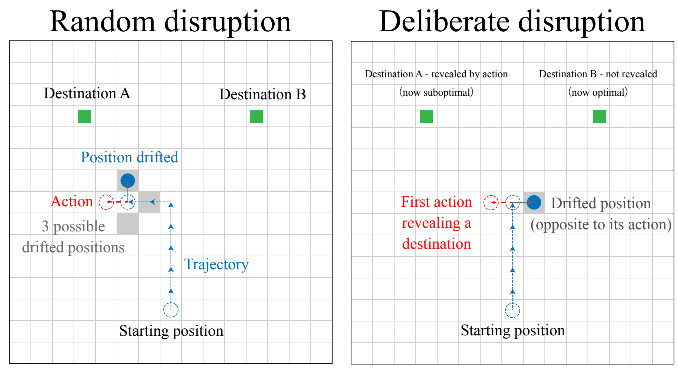

Fig. 1: **Design of disruption.** For random disruption, both time step and direction of the disruptions were randomly sampled. For deliberate disruption, both were designed to push the agent away from the destination the moment it was revealed.

A. Design and Procedure

There were 10 trials in total. Each trial consisted of a 15×15 grid map with one agent and two destinations. The two destinations were placed so that Manhattan distances between the agent's starting position and two destinations were equal, ranging from 7 to 13 with 10 as the mean. In the first 9 trials, as participants were instructed, the disruption occurred with a 10% probability at every time step, resulting in roughly 1 drift per trajectory. Across trials, these "random" 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 disruptions were distributed within the 3 trials. 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 fights against the agent's action by placing it in a cell to the opposite side of its action (Figure 1). As a result, the agent would end up being closer to the destination not revealed by its action. Participants were explicitly informed that the environment was not deterministic: at every step, there was a 10% probability that the agent's action could be disrupted by a random drift that pushes it to a nearby cell. A trial ended once the agent reached a destination, immediately followed by a new trial with a new map.

We designed 5 between-subject conditions for participants.

- Individual Real-time Condition. Participants were instructed to control an agent to reach any of the two destinations as soon as possible with the least number of steps using the four arrow keys (up, down, left, and right) on a standard keyboard. In this condition, participants performed the task alone in a single room with laboratory computers. There was no time delay after disruptions.
- Time Delay Conditions. The design and procedure were the same as the Individual real-time condition except there was a pause after each disruption (both random and deliberate) occurred, allowing for more time to re-plan. During the pause, participants would see a blank phase with the instruction "please wait" presented in the center of the screen. Participants were given clear explanations on this time duration before the experiment. Duration of the pause was manipulated in two conditions:
  - Individual-2s-delay: Duration of the pause was 2s.
  - Individual-5s-delay: Duration of the pause was 5s.
- Social Context Conditions. The design and procedure were the same as the Individual real-time condition except participants were grounded in two kinds of social context:
  - Observed real-time condition. Another person sat behind the participant by a distance. Participants were told that he/she was an experiment assistant to help with any potential computer problems.
  - Dual real-time condition. A second independent player was present in the game and performed the same task as in Individual real-time condition. In this task, one agent was controlled by the participant and the other controlled by a desire-driven agent. There was no interference between two agents: they could enter the same cell or go to the same destination without conflicts. Pairs of participants unfamiliar with each other were recruited to attend the experiment together. The two people were introduced to each other and read the instructions together. They were then assigned to two separate rooms with computers visibly connected via a network cable, making them believe they were doing the task with the other person. An experimenter unfamiliar with the participants impersonated the absent one if anyone in the pair did not attend the experiment as arranged.

B. Participants

Human participants were adults recruited from the Zhejiang University participants' pool. The sample size was set to $n = 50$ for each condition. Data is still being collected for Time Delay Conditions, with 20 participants for the two Time Delay Conditions as of now. A total of 190 participants joined the experiments for course credits or monetary payments (10 RMB). All participants were given informed consent before the experiments. All studies were pre-reviewed and approved by the institutional review board at the Department of Psychology and Behavioral Sciences, Zhejiang University. Researchers who collected the data were blind to the hypotheses of the study during data collection.

C. Desire-MDP model

We adopted the pure desire-driven agents as a baseline. We employed MDP as an implementation of the desire-model following the MEU principle. Desires are defined as sources of rewards, and an agent acts to maximize its expected long-term future reward. The definition of an MDP includes a state space S, an action space A, a transition function $T(s,a)$, and a utility function $R(s,a)$. Solution to an MDP is an optimal policy π, which takes state s as input, and outputs a probabilistic distribution of actions, $P(a|s)$. An agent acts by sampling an action from this distribution. The above definition and the solution to an MDP do not involve a formulation of intention.

1) State space: The agent's state was its location, defined as a tuple with 2D coordinates $(X_{coordinate}, Y_{coordinate})$. The size of the map is 15x15, composing a state space of size 225.

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

3) Transition function: We used a stochastic transition function that takes state s and action a as input and outputs $P(s')$, a probability distribution over the next state s'. In the experiment, the two destinations were set as the terminal states. During navigation, the agent can only reach 4 nearby states. The agent moved to the cell in the direction of its action with probability 9/10, with the other nearby cells evenly splitting the rest of the 1/10 probability.

4) Reward function: The reward function takes the state s and action a as input and outputs a scalar as the short-term reward. Here the reward has two components: 30 for reaching any destination, and -1/30 for every movement on the map.

5) Solving MDP: The optimal policy of MDP was solved by value iteration using the Bellman optimality equation [7] of the value function V, which is 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 achieved when this iterative process converges.

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

Here γ is the discount factor that is fixed to be 0.9 in all MDP simulations, which is a commonly used value.

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

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

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

Then, we derive a Boltzmann policy for the probability of taking an action a given state s as proportional to $Q^*(s,a)$:

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

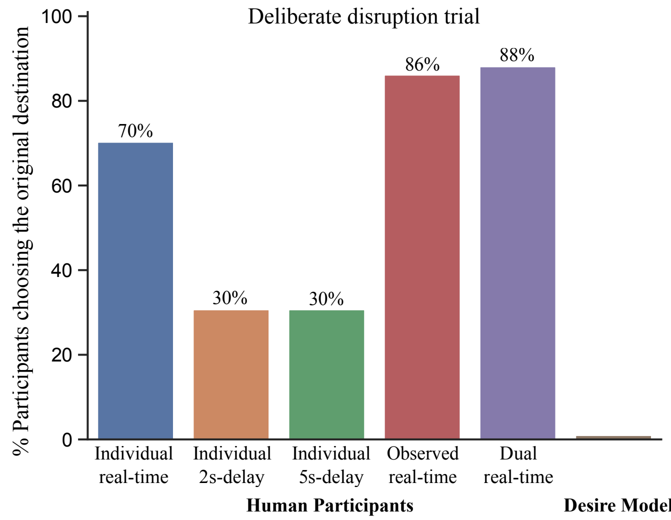

Fig. 2: **Experimental results.** Percentage of participants who reached the original destination in the last trial of deliberate disruption, with different conditions and the desire-MDP Model baseline.

The Boltzmann policy takes β as a rationality parameter. When $\beta \to 0$, the agent acts almost randomly; when $\beta \to \infty$, the agent chooses the action greedily based on the optimal Q-value. Here we chose β = 2.5 following previous studies modeling human action with MDP [3, 4]. With this value, the action selection will be dominated by the maximum $Q(s,a)$, but still deviates from it with a small probability, to capture the fact that human decision-making is not entirely optimal.

D. Results

In the last trial with a deliberate disruption, human participants across all five conditions tended to choose the original destination (see Fig. 2 Individual real-time condition: 70%; Individual 2s-delay condition: 30%; Individual 5s-delay condition: 30%; Observed real-time condition: 86%,; Dual real-time condition: 88%) compared to the desire-MDP model that always chose the closer destination (0%) (all $ps < .001$ by Fisher's exact test).

In the Individual Time Delay Conditions, participants in both 2s-delay and 5s-delay conditions were less likely to choose the original destination than participants in real-time condition with no time delay (30% vs 70%, χ^2 = 9.42, p = .002, Cramer's ϕ = 0.37).

In the Social Context Conditions, participants in the Dual condition were more likely to choose the original destination than participants in the Individual condition (88% vs 70%, χ^2 = 4.88, p = .027, Cramer's ϕ = 0.22). Participants in Observed condition were also more inclined to choose the original destination as compared to the Individual condition, with a borderline significance level (86% vs 70%, χ^2 =3.74, p = .054, Cramer's ϕ = 0.19). There was no significant difference in commitment between the Observed and the Dual condition.

These results demonstrate that, compared to a desire-driven agent who acts only to maximize expected utilities, humans act while spontaneously committing to their intention. This self-commitment was influenced in two ways: people showed less commitment when they had more time and resources to make a decision, and more commitment when they were under the social context of being observed.

III. Discussion

Our empirical results showed that humans spontaneously commit to their intentions. They cling to prior inertia and resist re-planning even when environmental changes have made their intention suboptimal. This self-commitment may result from humans' limited computational resources since sticking to one plan reduces the budget for online decision-making. Thus, people are more committed to their intentions when having fewer computational resources.

Still, the extent of the commitment is striking considering the simplicity of the navigation tasks we used here—even with enough time to re-plan for more optimal decisions as the environment changes, people still stick to their prior intention. This suggests that commitment may also serve certain purpose other than saving computational costs.

Our results that people are more committed in social context provide hints on the origin of human intentional commitment. Humans are born in a social world. Being a committed agent can greatly facilitate interpersonal coordination during social interactions. Imagine our ancestors hunting a lion together, they won't stand a chance unless they both commit to it simultaneously and persistently. Any of their flexibility of that commitment will place each other in peril. For these reasons, commitment has been mostly studied in the context of collaboration as an obligation that binds each other [17, 36]. This social constraint of joint commitment may be eventually internalized as a self-constraint that one uses to regulate one's own mind. Even when people act alone, they still spontaneously demonstrate their commitment as if they are making their minds predictable from a third-party perspective.

References

[1] George Ainslie. Specious reward: a behavioral theory of impulsiveness and impulse control. *Psychological bulletin*, 82(4):463, 1975.

[2] Robert Audi. Intending. *The Journal of Philosophy*, 70 (13):387–403, 1974.

[3] Chris L Baker, Rebecca Saxe, and Joshua B Tenenbaum. Action understanding as inverse planning. *Cognition*, 113 (3):329–349, 2009.

[4] Chris L Baker, Julian Jara-Ettinger, Rebecca Saxe, and Joshua B Tenenbaum. Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nature Human Behaviour*, 1(4):1–10, 2017.

[5] Albert Bandura. Social cognitive theory of self-regulation. *Organizational behavior and human decision processes*, 50(2):248–287, 1991.

[6] Albert Bandura. Social cognitive theory: An agentic perspective. *Annual review of psychology*, 52(1):1–26, 2001.

[7] Richard Bellman. A markovian decision process. *Journal of mathematics and mechanics*, 6(5):679–684, 1957.

[8] Michael Bratman. Intention, plans, and practical reason. 1987.

[9] Gharad Bryan, Dean Karlan, and Scott Nelson. Commitment devices. *Annu. Rev. Econ.*, 2(1):671–698, 2010.

[10] Philip R Cohen and Hector J Levesque. Intention is choice with commitment. *Artificial intelligence*, 42(2-3): 213–261, 1990.

[11] Donald Davidson. Actions, reasons, and causes. *The journal of philosophy*, 60(23):685–700, 1963.

[12] Daniel C Dennett. *Kinds of minds: Toward an understanding of consciousness*. Basic Books, 1996.

[13] Jon Elster. *The multiple self*. Cambridge University Press, 1987.

[14] Ayelet Fishbach and James Y Shah. Self-control in action: implicit dispositions toward goals and away from temptations. *Journal of personality and social psychology*, 90(5):820, 2006.

[15] Jerry A Fodor. A theory of the child's theory of mind. *Cognition*, 1992.

[16] S Freud. The ego and the id. *Standard Edition*, 19:3–59, 1923.

[17] Margaret Gilbert. *Joint commitment: How we make the social world*. Oxford University Press, 2013.

[18] Patrick Haggard. Human volition: towards a neuroscience of will. *Nature Reviews Neuroscience*, 9(12): 934–946, 2008.

[19] Gilbert Harman. *Change in view: Principles of reasoning*. The MIT Press, 1986.

[20] Cecilia Heyes, Dan Bang, Nicholas Shea, Christopher D Frith, and Stephen M Fleming. Knowing ourselves together: The cultural origins of metacognition. *Trends in Cognitive Sciences*, 24(5):349–362, 2020.

[21] Hiromi Kobayashi and Shiro Kohshima. Unique morphology of the human eye. *Nature*, 387(6635):767–768, 1997.

[22] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.

[23] Gordon B Moskowitz and Heidi Grant. *The psychology of goals*. Guilford press, 2009.

[24] Stuart Russell and Peter Norvig. Artificial intelligence: a modern approach. 2002.

[25] Thomas C Schelling. *The Strategy of Conflict: with a new Preface by the Author*. Harvard university press, 1980.

[26] Thomas C Schelling. *Choice and consequence*. Harvard University Press, 1984.

[27] Arthur Schopenhauer. *The World as Will and Idea-Vol. I.*, volume 1. Read Books Ltd, 2015.

[28] John R Searle. *Rationality in action*. MIT press, 2003.

[29] John R Searle, S Willis, et al. *Intentionality: An essay in the philosophy of mind*. Cambridge university press,

1983.
[30] James Y Shah. The automatic pursuit and management of goals. *Current Directions in Psychological Science*, 14(1):10–13, 2005.
[31] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529 (7587):484–489, 2016.
[32] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.
[33] Herbert A Simon. The sciences of the artificial, 1969. *Massachusetts Institute of Technology*, 1981.
[34] Robert Henry Strotz. Myopia and inconsistency in dynamic utility maximization. *The review of economic studies*, 23(3):165–180, 1955.
[35] Richard Thaler. Toward a positive theory of consumer choice. *Journal of economic behavior & organization*, 1 (1):39–60, 1980.
[36] Michael Tomasello. The moral psychology of obligation. *Behavioral and Brain Sciences*, 43, 2020.
[37] Michael Tomasello, Malinda Carpenter, Josep Call, Tanya Behne, and Henrike Moll. Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and brain sciences*, 28(5):675–691, 2005.
[38] Michael Tomasello, Brian Hare, Hagen Lehmann, and Josep Call. Reliance on head versus eyes in the gaze following of great apes and human infants: the cooperative eye hypothesis. *Journal of human evolution*, 52(3): 314–320, 2007.
[39] J Von Neumann and O Morgenstern. Theory of games and economic behavior (3rd editon), 1953.
[40] Lev Semenovich Vygotsky and Michael Cole. *Mind in society: Development of higher psychological processes*. Harvard university press, 1978.
[41] Henry M Wellman. *The child's theory of mind.* The MIT Press, 1992.
[42] Henry M Wellman. *Making minds: How theory of mind develops*. Oxford University Press, 2014.