Précis:
Nature and origins of intuitive psychology in human infants

Shari Liu

Ph.D. 2020, Department of Psychology, Harvard University

1 Chapter One: Introduction

Every day, we look out into the world and see scenes of people in motion, and we understand such scenes by invoking knowledge of the social and physical world. We appreciate the mental world of people, including their desires, percepts, and beliefs, and we also see these actors as solid bodies acting in a physical world, who can exert forces and navigate themselves through their physical environment. How do our minds get so much meaning from this input? This question has motivated many research programs in cognitive science, including research investigating the computational basis for social and physical understanding (Baker et al., 2009; Piloto et al., 2018; Rabinowitz et al., 2018; Ullman et al., 2017), the neural origins of these computations (Adolphs, 2009; Fischer et al., 2016), their evolutionary origins (Martin & Santos, 2016; Povinelli et al., 2000; Premack & Woodruff, 1978), how this understanding develops in the minds and brains of infants and children (Baillargeon et al., 2016; Gopnik & Wellman, 1992; Wilcox & Biondi, 2015), and how this understanding varies across cultures and differences in experience (Barrett et al., 2016; Koster-Hale et al., 2014; Landau et al., 2009).

This dissertation contributes to this multidisciplinary enterprise of understanding social intelligence by revealing its origins in human minds. How do infants learn to understand what people want, or which actions are hard and easy? How does their limited motor experience constrain this understanding? The main thesis of this dissertation is that infants hold a coherent theory of other people's actions: a precursor to a representational theory of mind. In support of this argument, I present two main pieces of evidence. First, infants' knowledge includes quantitative abstract variables, which relate to each other causally, and allow infants to both infer hidden causes from limited observations and to make predictions in new situations. Second, infants can go beyond their observations and experiences to reason about novel agents engaging in actions that infants themselves are not capable of performing. This theory of action may be a key foundation for our prodigious ability to reason about and learn from other minds.

1.1 Previous work

As human adults, we understand other people's actions and minds by invoking a folk theory of psychology, an abstract and coherent system of knowledge that causally relates the actions of other people to their beliefs, knowledge, percepts and goals, and specifies how the physical world constrains what actions people can take (Dennett, 1987; Gopnik & Meltzoff, 1998; Perner, 1991; Wellman & Woolley, 1990). While concepts like proportional belief, on some views, take 4 or so years to develop (Wellman et al., 2001), a narrower understanding of other minds, that represents people as intentional agents who pursue their goals efficiently, is available even to infants. In the first year of life, infants represent agents' goals (Woodward, 1998), physical constraints (Saxe et al., 2006), perceptual states (Luo & Johnson, 2009), the cost of their actions (Gergely et al., 1995), and their causal powers (Muentener & Carey, 2010). However, a key open question is what form this understanding takes: Do infants hold a coherent theory of other people that integrates across what they want, see, and make happen? And how does this kind of intuitive theory work, in computational terms, in people of any age?

1.1.1 Understanding other minds by inverting their plans

Research in adults and older children shows that we understand other people's actions and minds by appealing to the internal mental states that best explain people's overt behaviors (Baker et al., 2009; Jara-Ettinger et al., 2016). We assume that other agents' actions result from a balance between cost and reward–how much actions cost, and how much reward agents place over states of the world. Rather than only acting to carry out their desires, or only acting to minimize effort, we expect agents to carry out actions that maximize reward while minimizing cost. Cost and reward, which enter into a utility function, are thus separable causes of action: When an agent chooses an apple, easily in reach, from a banana, sitting all the way across the room, it could be because the agent likes bananas, but not enough to take a costly action to get them, or because the agent likes apples and bananas the same, and chooses the easier action. The proposal that we understand other people's actions by inverting their plans, which are composed of the agent's goal states and the cost of the actions that will bring about those goal states, is termed the Naive Utility Calculus (NUC) (Baker et al., 2017; Jara-Ettinger et al., 2016), and has gained empirical support in studies of young children (Bridgers et al., 2019; Jara-Ettinger et al., 2017; Jara-Ettinger et al., 2015a; Jara-Ettinger et al., 2015b) and adults (Baker et al., 2017).

What is the content and form of action understanding in the first year of life, before infants accumulate the kinds of experiences older children and adults possess? Because infants are built to learn, and their theories of other people change and develop in childhood (Gopnik & Wellman, 1992; Perner, 1991; Wellman, 2002), it is possible either that they share the computational capacities to represent other minds that older children have, or that this is built over development. I investigate this question in Chapter 2 (Do infants hold continuous representations of action cost?) and Chapter 3 (Do infants expect other people's plans to be composed of costs and rewards?).

1.1.2 Learning about other minds via first-person experience

Infants are agents. They pursue goals by moving their bodies, and learn to explore the world in increasingly sophisticated ways over development. A prevailing theory in developmental psychology is that this wealth of first-person motor experience teaches infants to understand other people's actions (Gallese et al., 2004; Piaget, 1954). For example, this could explain why it takes infants 5 months to understand other people's reaching actions as intentional and goal-directed: it takes infants, on average, 4 to months of figure out how to reach for and grasp objects (von Hofsten, 1989). Furthermore, lab interventions on infants' motor abilities help their understanding. Giving 3-month-old, so-called 'pre-reaching', infants training acting on objects with Velcro-covered ('sticky') mittens causes them to see other people's actions as goal-directed (Sommerville et al., 2005) and physically constrained (Skerry et al., 2013). This work shows that infants learn about action through motor experience, but leaves open *what* they learn.

In Chapter 4, I test the hypothesis that 3-month-old infants already see other people as causal agents who pursue goals through costly actions. However, without experience reaching for objects (e.g. through sticky mittens training, or experience reaching built through motor development), 3-month-olds are unsure of whether reaching counts as a causal intentional action. I test this hypothesis by asking whether we can give 3-month-olds this insight–that reaching causes rewarding events to happen–without motor training.

2 Chapter Two: Costly, more or less[1]

Many previous studies show that infants expect agents to be efficient. For instance, infants understand that, when faced with an obstacle, agents tend to move around it, but without the obstacle, agents tend to move directly to their goals (Gergely et al., 1995; Phillips &

[1]The research from this chapter appears in the literature as: Liu, S., & Spelke, E. S. (2017). Six-month-old infants expect agents to minimize the cost of their actions. *Cognition*, *160*, 35–42.

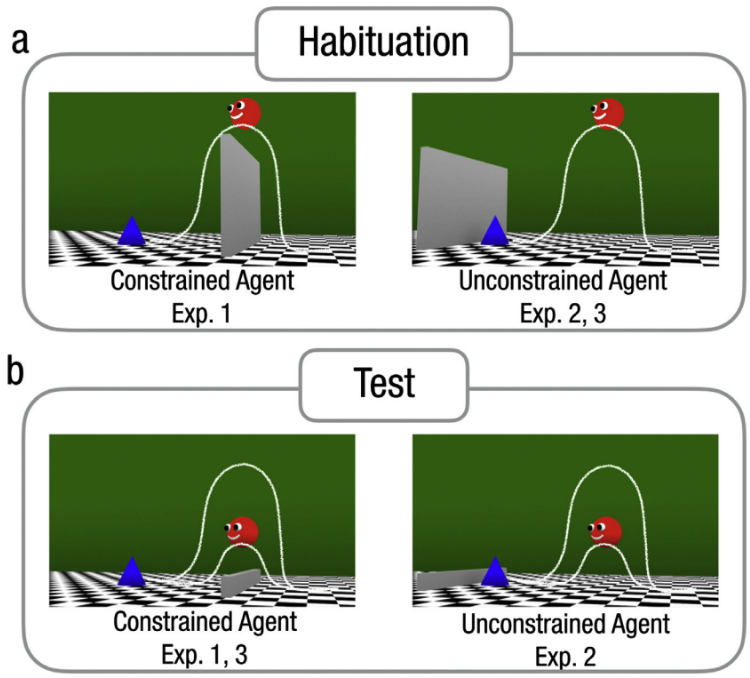

Figure 1: Trial structure for Experiments 1–3, including (a) habituation to an agent leaping over tall barriers efficiently (left, Experiment 1) or performing identical motions without a physical constraint (right, Experiments 2 and 3) and (b) test, with the agent performing low and high jumps over a novel barrier (left, Experiments 1 and 3) or no barrier (right, Experiment 2). White lines indicate trajectories of motion.

Wellman, 2005; Skerry et al., 2013). Nevertheless, this research leaves open exactly how the infant mind solves this problem: Do infants have (1) a simple expectation that agents move around obstacles, but not open space, towards their goals, without representing the relative cost of different actions, or do they (2) have a richer expectation that agents maximize efficiency (i.e. minimize a continuous variable, action cost) with respect to these obstacles and goals?

In Liu and Spelke (2017), we show that 6-month-old infants expect other agents to minimize the cost of their actions. After being habituated to an animated agent leaping over a barrier towards a goal object, infants saw a low barrier come between the agent and its goal. Infants looked longer when this character takes an unnecessarily high leap over this barrier, relative to when it jumps just high enough to clear the barrier, suggesting that they were more surprised by the former action. In Experiment 2, we showed that this effect is not because infants are more interested in longer or faster jumps. If the barrier was moved out of the agent's way throughout the experiment, infants no longer had systematic expectations about the height of the agent's jump. In Experiment 3, we showed that this effect is not the

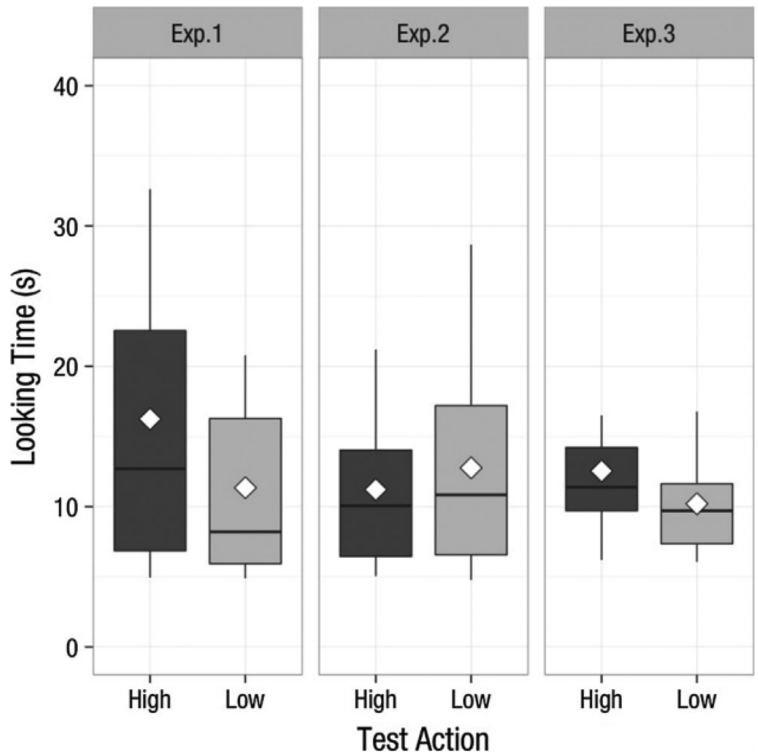

Figure 2: Boxplots for raw looking times in seconds to test events in Experiments 1–3 (N = 20 per experiment). Boxes indicate interquartile ranges, bold horizontal lines indicate medians, and white diamonds indicate means.

result of infants' learning that the agent jumps just high enough to clear the barrier. When the barrier was moved out of the agent's way initially, and then suddenly blocks her path, infants still expected the agent to take a minimally costly jump, looking longer at the high inefficient jump, even though they never observed this agent circumventing a barrier until that point in the experiment. In sum, the work in Chapter 2 shows that infants represent a continuous notion of action costs, and expect agents to minimize the cost of its actions the very first time they see this specific agent navigate around an obstacle. However, this work leaves open what action cost is to an infant (path length, time expended, or force applied over a path, among other possibilities), and whether this variable participates in more complex reasoning about other people's action plans.

3 Chapter Three: What's worth the effort[2]

Are the powerful computational resources that guide action understanding in adults and older children (Baker et al., 2009; Jara-Ettinger et al., 2017) present in the infant mind? The work in Chapter 3 tests whether infants represent cost and reward in an integrated utility function. As our case study, we focus on the problem of learning what people want from how hard they are willing to try. We ask whether infants appreciate that agents try harder for goals they value more, and use this information predict what agents will choose in a new situation.

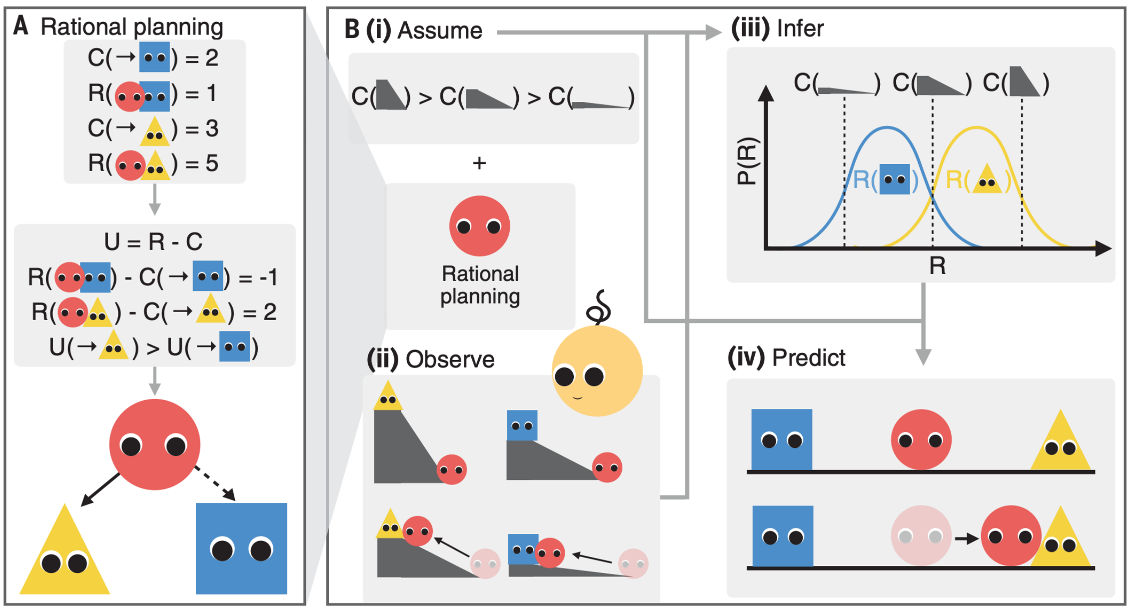

Figure 3: A schematic of our computational model. (A) The forward direction defines the agent as a rational planner that calculates the utilities of different actions from their respective costs and rewards and then selects an action stochastically in proportion to its utility. In this case, the overall utility for approaching the triangle is higher than for approaching the square, so the central agent (circle) will likely choose triangle over square. (B) An observer (i) assuming this model and priors over the costs of different actions can (ii) observe a series of actions and then (iii) infer a posterior distribution over the hidden values of an agent's costs and rewards given its actions. (iv) These posteriors can then be used to predict the actions of the agent in a new situation, in which the same goal states can be reached with different actions.

Across 3 experiments, Liu et al. (2017) show that infants use continuous representations of action cost in order to learn about the preferences of other people. In Experiment 1,

[2]The research from this chapter appears in the literature as: Liu, S., Ullman, T. D., Tenenbaum, J. B., & Spelke, E. S. (2017). Ten-month-old infants infer the value of goals from the costs of actions. *Science*, *358*(6366), 1038–1041.

after seeing a central protagonist character jump a low (but not a medium) barrier for a yellow agent and jump a medium (but not a high) barrier for a blue agent, 10-month-old infants expected the character to choose the blue over the yellow agent in a new situation, when both target agents appeared in the scene with no barriers. Specifically, they looked longer when the central protagonist chose the yellow agent (i.e. the one it was only willing to jump a low barrier for) than when the protagonist chose the blue agent (i.e. the one it was willing to jump a medium barrier for).

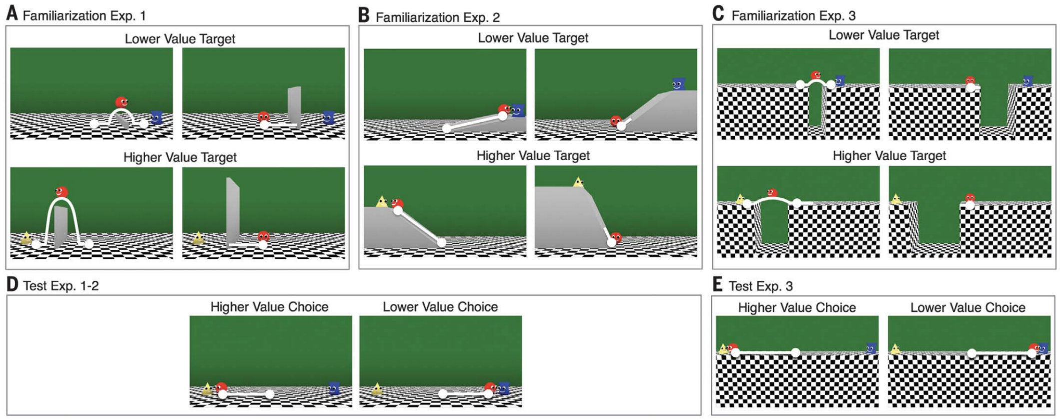

Figure 4: Structure of experiments 1 through 3. (A to C) During familiarization, the central agent (circle) accepted a low cost and refused a medium cost for the lower-value target (square) and accepted a medium cost and refused a high cost for the higher-value target (triangle). Other than the sizes of the barriers, ramps, and trenches, and the consequent trajectories of motion, the pairs of events displaying approach or refusal of approach to the two targets were identical. (D and E) At test, the agent looked at each of the two targets and chose either the lower- or higher-value target. White circles indicate start- and end points of action, and white lines indicate trajectories.

We then replicated this finding in 2 other experiments. Regardless of whether infants saw an agent jump higher barriers (Experiment 1), climb steeper inclines (Experiment 2), or jump wider gaps (Experiment 3) for one goal than another, we found that infants expected the agent to choose the goal for which performed higher-cost actions, looking longer when the agent chose the lower-cost goal instead (Figure 5). Across all of these studies, we show that infants' conceptions of physical effort are not reducible to any single path feature like speed (equated for in Experiment 2), time (equated for across Experiments 1-3), path length (equated for in Experiment 2), vertical motion (equated for in Experiment 3), or horizontal motion (equated for in Experiment 1) (Figure 4).

Finally, we tested an array of computational models on these and past findings (Gergely et al., 1995; Liu & Spelke, 2017; Woodward, 1998), and found that only a generative model that maximizes utility over cost (conceived as physical effort), and reward (conceived as

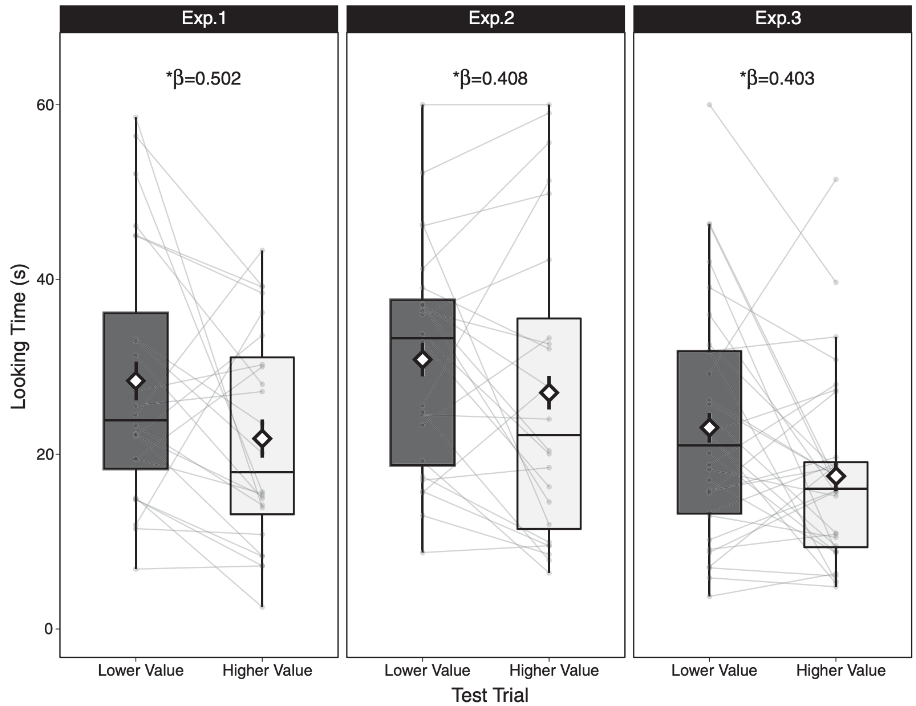

Figure 5: Boxplots of average looking time toward the higher- and lower-value choice during test in experiments 1 through 3. White diamonds indicate means, with error bars indicating within-subjects standard errors. Horizontal lines indicate medians, boxes indicate middle quartiles, and whiskers indicate points within 1.5 times the interquartile range from the upper and lower edges of the middle quartiles. Light gray points connected across boxes indicate looking times from individual participants. Beta coefficients indicate effect sizes in standard deviations, and asterisks indicate significance relative to prespecified (experiments 1 and 2) and preregistered (experiment 3) alphas ($*P < 0.05$). Statistical analyses are provided in the text and supplementary materials.

value) fully accounted for all of them. When we lesioned the model to represent cost as path length (e.g. moving five meters costs the same regardless of whether the agent moves horizontally or vertically), reward as frequency of selection (e.g. how often an agent chooses a goal indicates its value to the agent), or limited the model to represent cost and reward separately rather than integrated together in a utility function, the model failed to account for a subset of the current and past empirical findings. Overall, this research show that infants not only expect others to be efficient (Gergely et al., 1995; Liu et al., 2017) and to have preferences (Woodward, 1998), but they put these two concepts (cost and reward) together to explain, learn about, and predict other people's actions (Baker et al., 2009; Jara-Ettinger et al., 2016). Thus, by ten months of age, infants appear to see other people as utility-maximizing intentional agents whose plans include abstract variables like cost and reward.

4 Chapter Four: Knowing before doing[3]

How do infants' experiences acting on the world influence their action understanding? On one view of development (Piaget, 1952), all of our complex knowledge about the world comes from sensorimotor experience: Acting on the world provides new sensorimotor information and prompts infants to posit new models of the world that they have not considered before to explain it. On another view (Carey, 2009; Spelke & Kinzler, 2007), infants start out with some understanding of agents and actions, which could support their subsequent motor and social development. In support of Piaget's view, babies come to understand simple actions like reaching and grasping, and more complex, multi-step actions, like pulling a cloth in order to reach a toy sitting on top of the cloth, at around the same time that infants learn to successfully perform these actions (Loucks & Sommerville, 2012; Sommerville & Woodward, 2005; Woodward, 1998). Furthermore, lab interventions show that infants indeed learn by doing. The most striking evidence for this comes from studies of 3-month-old, pre-reaching infants, who do not yet reach for, grasp, or act intentionally on objects. When 3-month-olds are trained to pick up objects using Velcro-covered mittens ('sticky mittens' training), this experience produces subsequent changes in infants' action understanding. Just two or three minutes of experience of batting at and picking up objects helps infants appreciate that when mittened hands reach towards objects, they're reaching towards those objects, rather than whatever object is at that location (Gerson & Woodward, 2014b; Sommerville et al., 2005), and that reaching is costly and constrained by barriers (Skerry et al., 2013). This work shows that first-person experience acting on the world indeed causes changes in infants' interpretation of other people's actions.

Infants learn something from their motor experience, but what? In the work presented in Chapter 4, we test the hypothesis that infants at 3 months already see other people as goal-directed agents who act intentionally to carry out their goals, and learn from sticky mittens that reaching is a causal intentional action. Three-month-olds, without ever having reached for or grasped an object, may be puzzled by how reaching an object could possibly cause it to be entrained by the hand grasping it. They could then learn during sticky mittens training that simply contacting a specific set of objects with the mittened hand can result in this effect. This view is consistent with findings that infants who are trained on sticky mittens succeed only at understanding actions performed with a similar mitten, and not with a bare hand (Woodward, 2008), and that training infants on one set of objects, and testing them on a different set, takes away the benefit of training (Gerson & Woodward,

[3] The research from this chapter appears in the literature as: Liu, S., Brooks, N. B., & Spelke, E. S. (2019). Origins of the concepts cause, cost, and goal in prereaching infants. *Proc. Natl. Acad. Sci. U. S. A.*, *116*(36), 17747–17752.

2014a).

This interpretation predicts that showing infants *simpler actions with clear contact relations* should boost their sensitivity to the cost of these actions, even in the absence of motor experience with mittens. In five experiments, our findings were consistent with this prediction: even without motor training, 3-month-old infants looked longer at inefficient than efficient reaches when these actions caused an object to light up. In Experiment 3, after being habituated to person reaching over a barrier, and the barrier was removed, infants looked longer when the person reached in the same curved path they did before (an inefficient action) than when the person reached in a new direct path (an efficient action). We found that this effect was stronger when infants saw the hand touch the object and cause it to illuminate, than when infants saw the hand grasp and pick up the object (Experiments 1-2, as in Skerry et al., 2013).

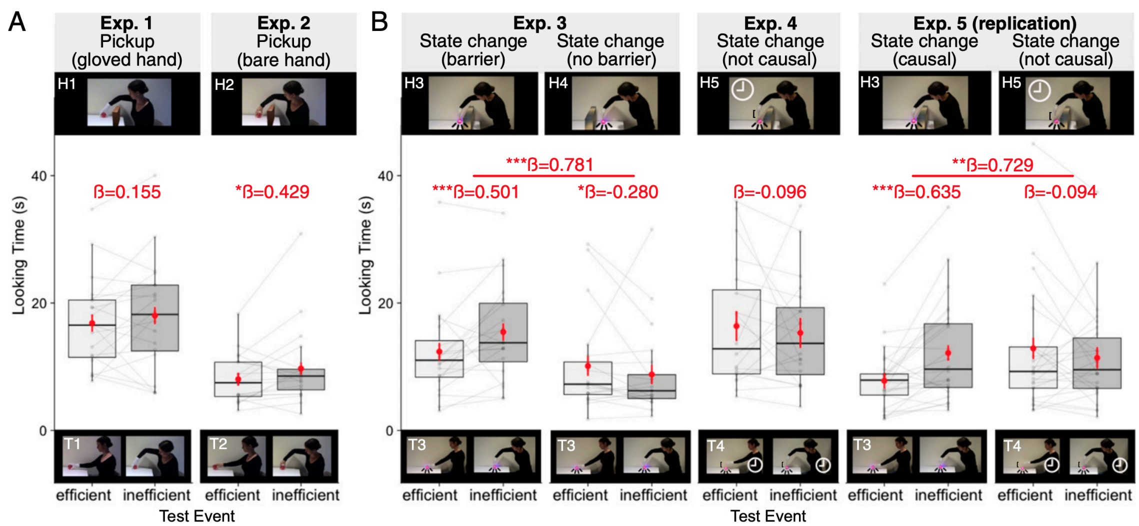

Figure 6: Looking time in seconds toward the efficient versus inefficient reach at test across Exps. 1 to 5 (n = 152), for both (A) pick-up events (Exps. 1 and 2) and (B) state-change events (Exps. 3 to 5). Images indicate video displays used during the habituation phase (above each graph) and test phase (below each graph) for each experiment. Red dots and error bars indicate means and within-subjects 95% CIs. Pairs of connected points indicate data from a single participant. Horizontal bars within boxes indicate medians, and boxes indicate the middle 2 quartiles of data. Upper whiskers indicate data up to 1.5 times the interquartile range above the third quartile, and lower whiskers indicate data up to 1.5 times the interquartile range below the first quartile. Beta coefficients (β) list effect sizes in SD units for each condition. $*P < 0.05$, $**P < 0.01$, $***P < 0.001$, two-tailed, except for the causal condition in Exp. 5, which was preregistered as a one-tailed test.

But did infants succeed in this task because (1) our stimuli conveyed a clear and simple causal relation between the reaching action and its effect, or because (2) they featured an exciting toy that lights up? To address this question, we tested 3-month-old infants in a control condition (Experiment 4) with identical stimuli as Experiment 3, except that the person's hand stopped short of the object, which then illuminated after a 0.5s delay. If infants see other people as causal intentional agents at 3 months, then this manipulation should take away the impression that the person caused the object to change (a signature of causal perception in older infants and adults; Michotte, 1963; Muentener & Carey, 2010), and also the impression that the person acted with the goal of causing this change. Thus, if infants' expectations about other people's action are conditioned on them understanding other people as causal intentional actors, then they should look equally at the inefficient and efficient versions of this action at test. That is exactly what we observed in Experiment 4: given this subtle change in the stimuli that removed impressions of causal agency, infants looked equally at the test events. We then replicated this finding (longer looking at the inefficient reach, conditioned on the action being causal) in a pre-registered direct replication (Experiment 5).

In summary, the work in Chapter 4 provides an existence proof that infants do not require first-person action experience with specific actions, like object-directed reaching, in order to see these actions, when performed by other people, as physically constrained. Three-month-old infants understand that obstacles impose constraints on object-directed reaching, and display this understanding most clearly when they see simple events where object-directed reaches result in action on contact, followed immediately by the object changing state. These findings challenge theories that root our intuitive psychology in our own sensorimotor experiences (Piaget, 1952), in favor of theories that grant infants a few early-emerging abstract concepts like cost, cause, and goal (Carey, 2009; Spelke & Kinzler, 2007). Given these minimal, early-emerging concepts, infants still have to learn what people want, which actions are hard versus easy, and how acting causes people to achieve their goals. Abstract concepts at the core of action understanding may precede and enable this learning.

5 Chapter Five: Conclusion and open questions

Human infants appreciate other agents' goals (Woodward, 1998), physical constraints (Saxe et al., 2006), perceptual states (Luo & Johnson, 2009), the cost of their actions (Gergely et al., 1995), and their causal powers (Muentener & Carey, 2010). In this dissertation, I investigate the content and form of this knowledge. I show that in the first year of life, infants analyze other people's actions in terms of abstract concepts like cost, cause and reward, organized into an intuitive theory of psychology. Infants hold expectations about key elements of this system even before mastering simple goal-directed actions, like object-directed reaching. These findings suggest that infants hold systematic theories of other people's actions. That is, infants interpret other people's actions in terms of continuous abstract variables (Chapter 2) and represent coherent plans over those variables (Chapter 3), including plans that go beyond infants' own specific experiences (Chapters 2-4). Like adults and older children, infants are able to stretch beyond the actions they have observed and experienced and reason about unfamiliar agents and novel actions. They can see someone act and infer why, and make predictions about what they will do in a novel situation. In this concluding chapter, I summarize the key findings from this thesis, lay out the broader theoretical questions that this work raises, and suggest ways to test them in future research.

5.1 How are intuitive psychology and intuitive physics organized in the infant mind?

The work in this thesis, as well as the work preceding it, suggests that infants see agents as acting in accord with their mental states, and also as physical bodies acting in a physical world. For example, infants appreciate that agents direct their actions towards physical objects, and are constrained by physical obstacles like barriers (Csibra, 2008; Csibra et al., 2003; Gergely et al., 1995; Liu et al., 2019; Liu & Spelke, 2017; Skerry et al., 2013; Woodward, 1998). Infants also connect the physical actions and physical constraints of agents to psychological variables, like what agents can see (Csibra & Volein, 2008; Luo & Johnson, 2009), what emotional states they hold (Skerry & Spelke, 2014), and what they prefer (Liu et al., 2017).

While it is clear that infants are able to reason about agents' actions towards and over physical objects, and to connect agents' physical actions to mental states, these findings raise questions about how the core domains of physics and psychology relate to each other in the first year of life. I believe that future studies focusing on representations of action cost can shed light on this question. One possibility is that intuitive physics and intuitive psychology are distinct domains in infancy, with computations proprietary to each domain. Under this possibility, infants make use of distinctively psychological concepts, like exhaustion and bodily exertion, to define action cost. Another possibility is that these two domains of knowledge are more deeply integrated: Infants could represent agents as a special kind of physical object, capable of causing its own motions[4]. Under this possibility, infants could define the cost of agents' actions in terms of the underlying mechanics or dynamics that apply to its physical body, and then understand psychological dimensions of cost, like bodily exhaustion, in terms of these physical representations.

[4]Another version of this hypothesis is that objects are inert agents, whose behavior is only defined in terms of physical dynamics or mechanics.

These two alternatives make distinct predictions about whether and how infants' understanding of physics relates to their understanding of physical effort. Whereas first proposal makes no prediction that infants' understanding of object mechanics or dynamics relates to their representation of exhaustion or exertion, the second, integrated proposal strongly predicts that infants' knowledge in the physical domain will predict and impact their intuitions about the cost of agent's actions, because representations of cost are defined in terms of these physical variables. Future work testing these predictions can reveal the nature and relation between these two fundamental domains of knowledge, and their origins in infant minds.

5.2 What is the nature of infants' causal representations?

Michotte (1963) proposed that all causal representations are rooted in our understanding of physical events like launching and entertainment. One interpretation of the work from Chapter 3 is that its findings refute Michotte, because I reported that 3-month-old infants are sensitive to causal relations in the context of agents' actions, 3 months earlier than they demonstrate a causal understanding for interactions between inanimate objects. (Previous findings suggest that prior to 6 months of age, infants do not interpret events like launching and entrainment as causal; Bélanger & Desrochers, 2001; Cohen & Amsel, 1998; Desrochers, 1999). Because the experiments in Chapter 4 differ in many ways from the experiments that test for infants' understanding of physical causation, I think this interpretation is premature, but testable with future work. For example, we can borrow methods from past work on causal perception in infants (Muentener & Carey, 2010), and repeat these experiments in 3-month-olds to ask whether infants at this age are indeed sensitive to contact relations in causal events featuring intentional agents, prior to being sensitive to contact relations in causal events involving physical objects.

If infants are tuned into the causal relations between agents and the physical world before they represent causal relations in purely physical contexts, these findings would suggest quite a different hypothesis about the origins of our causal concepts from Michotte (1963). While Michotte (1963) proposed that our sense of causation originates in perceptual input analyzers that form 'causal impressions' over the interactions of physical objects, this future work could reveal that these causal concepts are instead rooted in representations of intentional agents[5] that change the physical world through their actions. This understanding could stand separately from infants' later-developing causal understanding in physical situations, or even support the development of this understanding. Human adults clearly

[5]Paper 3 suggests the infant herself does not need to serve as this intentional agent: 3-month-old infants do not need to be capable of reaching to see these same actions, when performed by other people, as costly, causal, and intentional.

make use of this interventionist notion of causation in our scientific thinking (intervening on a variable, and measuring its effect), metaphor (the moon pulling the tides), and in our more mature understanding of desire and action (to want something is to hold in mind a state of the world, and to pursue that desire is to act so as to make that state real). Perhaps this interventionist notion of causation (Woodward, 2005) comes from our earliest understanding of people as causal agents.

5.3 In Summary

In this thesis, I argued that our social intelligence is rooted in an early-emerging, abstract, and coherent mentalistic theory of action. Even in infancy, we conceive other people as physical bodies with mental states, who circumvent obstacles, weigh the cost of acting against the rewards that these actions bring, and act in order to realize new states of the world. In the first year of life, infants make use of this knowledge to look beyond people's actions and make inferences about why they acted, and to form expectations about what they will do next. In this chapter, I raised two questions left open by this work, and proposed future work investigating infants' understanding of agents as physical bodies and catalysts of change. While we are far from a formal and explicit account of social intelligence at any age, this problem is tractable with the current tools and methods of cognitive science. The journey is long, but we are getting closer.

References

Adolphs, R. (2009). The social brain: Neural basis of social knowledge. *Annu. Rev. Psychol.*, *60*, 693–716.

Baillargeon, R., Scott, R. M., & Bian, L. (2016). Psychological reasoning in infancy. *Annu. Rev. Psychol.*, *67*(1), 159–186.

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*(3), 329–349.

Barrett, H. C., Bolyanatz, A., Crittenden, A. N., Fessler, D. M. T., Fitzpatrick, S., Gurven, M., Henrich, J., Kanovsky, M., Kushnick, G., Pisor, A., Scelza, B. A., Stich, S., von Rueden, C., Zhao, W., & Laurence, S. (2016). Small-scale societies exhibit fundamental variation in the role of intentions in moral judgment. *Proc. Natl. Acad. Sci. U. S. A.*, *113*(17), 4688–4693.

Bélanger, N. D., & Desrochers, S. (2001). Can 6-month-old infants process causality in different types of causal events? *Br. J. Dev. Psychol.*, *19*(1), 11–21.

Bridgers, S., Jara-Ettinger, J., & Gweon, H. (2019). Young children consider the expected utility of others' learning to decide what to teach. *Nature Human Behaviour*, *4*, 144–152.

Carey, S. (2009). *The origin of concepts*. Oxford University Press.

Cohen, L. B., & Amsel, G. (1998). Precursors to infants' perception of the causality of a simple event. *Infant Behav. Dev.*, *21*(4), 713–731.

Csibra, G. (2008). Goal attribution to inanimate agents by 6.5-month-old infants. *Cognition*, *107*(2), 705–717.

Csibra, G., Bıró, Z., Koós, O., & Gergely, G. (2003). One-year-old infants use teleological representations of actions productively. *Cogn. Sci.*, *27*(1), 111–133.

Csibra, G., & Volein, A. (2008). Infants can infer the presence of hidden objects from referential gaze information. *Br. J. Dev. Psychol.*, *26*(1), 1–11.

Dennett, D. C. (1987). *The intentional stance*. The MIT Press.

Desrochers, S. (1999). Infants' processing of causal and noncausal events at 3.5 months of age. *J. Genet. Psychol.*, *160*(3), 294–302.

Fischer, J., Mikhael, J. G., Tenenbaum, J. B., & Kanwisher, N. (2016). Functional neuroanatomy of intuitive physical inference. *Proc. Natl. Acad. Sci. U. S. A.*, *113*(34), E5072–81.

Gallese, V., Keysers, C., & Rizzolatti, G. (2004). A unifying view of the basis of social cognition. *Trends Cogn. Sci.*, *8*(9), 396–403.

Gergely, G., Nádasdy, Z., Csibra, G., & Bıró, S. (1995). Taking the intentional stance at 12 months of age. *Cognition*, *56*(2), 165–193.

Gerson, S. A., & Woodward, A. L. (2014a). The joint role of trained, untrained, and observed actions at the origins of goal recognition. *Infant Behav. Dev.*, *37*(1), 94–104.

Gerson, S. A., & Woodward, A. L. (2014b). Learning from their own actions: The unique effect of producing actions on infants' action understanding. *Child Dev.*, *85*(1), 264–277.

Gopnik, A., & Meltzoff, A. N. (1998). *Words, thoughts, and theories (learning, development, and conceptual change)*. MIT Press.

Gopnik, A., & Wellman, H. M. (1992). Why the child's theory of mind really is a theory. *Mind Lang.*, *7*(1-2), 145–171.

Jara-Ettinger, J., Floyd, S., Tenenbaum, J. B., & Schulz, L. E. (2017). Children understand that agents maximize expected utilities. *J. Exp. Psychol. Gen.*

Jara-Ettinger, J., Gweon, H., Schulz, L. E., & Tenenbaum, J. B. (2016). The naıve utility calculus: Computational principles underlying commonsense psychology. *Trends Cogn. Sci.*, *20*(8), 589–604.

Jara-Ettinger, J., Gweon, H., Tenenbaum, J. B., & Schulz, L. E. (2015a). Children's understanding of the costs and rewards underlying rational action. *Cognition*, *140*, 14–23.

Jara-Ettinger, J., Tenenbaum, J. B., & Schulz, L. E. (2015b). Not so innocent: Toddlers' inferences about costs and culpability. *Psychol. Sci.*, *26*(5), 633–640.

Koster-Hale, J., Bedny, M., & Saxe, R. (2014). Thinking about seeing: Perceptual sources of knowledge are encoded in the theory of mind brain regions of sighted and blind adults. *Cognition*, *133*(1), 65–78.

Landau, B., Gleitman, L. R., & Landau, B. (2009). Language and experience: Evidence from the blind child.

Liu, S., Brooks, N. B., & Spelke, E. S. (2019). Origins of the concepts cause, cost, and goal in prereaching infants. *Proc. Natl. Acad. Sci. U. S. A.*, *116*(36), 17747–17752.

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*(6366), 1038–1041.

Loucks, J., & Sommerville, J. A. (2012). The role of motor experience in understanding action function: The case of the precision grasp. *Child Dev.*, *83*(3), 801–809.

Luo, Y., & Johnson, S. C. (2009). Recognizing the role of perception in action at 6 months. *Dev. Sci.*, *12*(1), 142–149.

Martin, A., & Santos, L. R. (2016). What cognitive representations support primate theory of mind? *Trends in cognitive sciences*, *20*(5), 375–382.

Michotte, A. (1963). *The perception of causality*. Routledge.

Muentener, P., & Carey, S. (2010). Infants' causal representations of state change events. *Cogn. Psychol.*, *61*(2), 63–86.

Perner, J. (1991). *Understanding the representational mind*. The MIT Press.

Phillips, A. T., & Wellman, H. M. (2005). Infants' understanding of object-directed action. *Cognition*, *98*(2), 137–155.

Piaget, J. (1952). *The origins of intelligence in children* (Vol. 8). International Universities Press New York.

Piaget, J. (1954). *The construction of reality in the child*. Routledge.

Piloto, L., Weinstein, A., TB, D., Ahuja, A., Mirza, M., Wayne, G., Amos, D., Hung, C.-c., & Botvinick, M. (2018). Probing physics knowledge using tools from developmental psychology. *arXiv preprint arXiv:1804.01128*.

Povinelli, D. J., Reaux, J. E., Theall, L. A., Giambrone, S., & Humphrey, N. (2000). *Folk physics for apes: The chimpanzee's theory of how the world works* (Vol. 7). Oxford University Press Oxford.

Premack, D., & Woodruff, G. (1978). Does the chimpanzee have a theory of mind? *Behav. Brain Sci.*, *1*(4), 515–526.

Rabinowitz, N. C., Perbet, F., Song, H. F., Zhang, C., Eslami, S., & Botvinick, M. (2018). Machine theory of mind. *arXiv preprint arXiv:1802.07740*.

Saxe, R., Tzelnic, T., & Carey, S. (2006). Five-month-old infants know humans are solid, like inanimate objects. *Cognition*, *101*(1), B1–B8.

Skerry, A. E., Carey, S. E., & Spelke, E. S. (2013). First-person action experience reveals sensitivity to action efficiency in prereaching infants. *Proc. Natl. Acad. Sci. U. S. A.*, *110*(46), 18728–18733.

Skerry, A. E., & Spelke, E. S. (2014). Preverbal infants identify emotional reactions that are incongruent with goal outcomes. *Cognition*, *130*(2), 204–216.

Sommerville, J. A., & Woodward, A. L. (2005). Pulling out the intentional structure of action: The relation between action processing and action production in infancy. *Cognition*, *95*(1), 1–30.

Sommerville, J. A., Woodward, A. L., & Needham, A. (2005). Action experience alters 3-month-old infants' perception of others' actions. *Cognition*, *96*(1), 1–11.

Spelke, E. S., & Kinzler, K. D. (2007). Core knowledge. *Dev. Sci.*, *10*(1), 89–96.

Ullman, T. D., Spelke, E. S., Battaglia, P., & Tenenbaum, J. B. (2017). Mind games: Game engines as an architecture for intuitive physics. *Trends Cogn. Sci.*

von Hofsten, C. (1989). Mastering reaching and grasping: The development of manual skills in infancy. In S. A. Wallace (Ed.), *Advances in psychology* (pp. 223–258). North-Holland.

Wellman, H. M. (2002). Understanding the psychological world: Developing a theory of mind. In U. Goswami (Ed.), *Blackwell handbook of childhood cognitive development* (pp. 167–187). Blackwell Publishers Ltd.

Wellman, H. M., Cross, D., & Watson, J. (2001). Meta-analysis of theory-of-mind development: The truth about false belief. *Child Dev.*, *72*(3), 655–684.

Wellman, H. M., & Woolley, J. D. (1990). From simple desires to ordinary beliefs: The early development of everyday psychology. *Cognition*, *35*(3), 245–275.

Wilcox, T., & Biondi, M. (2015). Object processing in the infant: Lessons from neuroscience. *Trends Cogn. Sci.*, *19*(7), 406–413.

Woodward, A. L. (1998). Infants selectively encode the goal object of an actor's reach. *Cognition*, *69*(1), 1–34.

Woodward, A. L. (2008). Infants' learning about intentional action. In A. L. Woodward & A. Needham (Eds.), *Learning and the infant mind*. Oxford University Press, USA.

Woodward, J. (2005). *Making things happen: A theory of causal explanation*. Oxford University Press, USA.