Physical pragmatics: Inferring the social meaning of objects

Michael Lopez-Brau, Julian Jara-Ettinger

Department of Psychology, Yale University

Abstract

Beyond words and gestures, people have a remarkable capacity to also communicate through objects: A hat on a chair means it is occupied, rope hanging across an entrance means we should not cross, and objects placed in a closed box means they are not ours to take. How do people generate and interpret the social meaning of objects? We propose that this capacity emerges from our ability to reason about the motives behind other agents' actions. We show that a computational model that infers mental states based on how agents manipulate their environment spontaneously gives rise to the ability to communicate through objects. As predicted by our model, we found that people from both industrialized and non-industrialized societies can infer the communicative meaning of objects, even in the absence of a pre-existing convention. Moreover, we show that, after inferring the meaning of an unconventional object, people can store and retrieve its meaning in subsequent encounters, revealing a mechanism behind the widespread use of conventional communicative objects. Our model sheds light on how humans use their ability to interpret other people's behavior to embed social meaning into the physical world.

Key words: Computational modeling, Pragmatics, Social objects, Theory of Mind

1. Introduction

Humans have a remarkable capacity to communicate through objects, even ones we do not usually think of as conveying meaning. A hat on a chair reveals that a seat is taken; rope surrounding a patch of grass tells us not to walk through; and, during snowy Boston winters, a plastic chair on a shoveled parking spot signals that it is not up for grabs (Figure 1a-c). These objects do little to physically constrain our actions, yet they affect our behavior because

8 we recognize the meaning they convey. Despite their ubiquity, little is under-
9 stood about the representations that allow humans to invent and reason about
10 communicative objects.

11 Our goal in this paper is to provide a first step in addressing central questions
12 of this phenomenon: What cognitive capacities enable humans to use objects
13 to communicate? How do we recognize when an object is communicative and
14 infer its meaning? And how does this meaning become conventionalized? To
15 answer these questions, we focus on a family of communicative objects like those
16 shown in Figure 1. These objects are pervasive in society, but to our knowledge,
17 have received little attention in cognitive science. While these objects do not
18 capture the full scope of common communicative objects, we believe they reveal
19 fundamental principles of how people come to use objects to communicate. We
20 return to this point in the Discussion.

21 We propose that seeing objects as communicative emerges from our ability
22 to reason about the psychological causes behind other people’s behavior—our
23 *Theory of Mind* (ToM; Wellman 2014; Gopnik et al. 1997). The central idea
24 in our proposal is that, if we can infer other people’s mental states by reason-
25 ing about why they manipulated objects in a specific way, then agents can also
26 strategically manipulate objects with the purpose of eliciting mental-state infer-
27 ences in people who encounter them. Imagine, for instance, walking through an
28 unfamiliar building and finding a branching hallway with a broom positioned
29 diagonally across it. Intuitively, it is easy to recognize that (1) someone in-
30 tententially placed the broom here, (2) that doing so probably did not require
31 much effort, and (3) that the broom makes it slightly more inconvenient for us
32 to go into that hallway. When considering why someone would leave a broom
33 blocking a hallway entrance, one possibility that comes to mind is that they
34 wanted to prevent people from walking through. But if that were the case, why

35 did the person not put more effort into physically blocking the way? Intuitively,
36 this is because the agent’s goal was not to create an insurmountable constraint,
37 but rather to place an object that would prompt this reasoning in other agents,
38 allowing them infer that they ought to stay away.

39 Past research shows that the capacities necessary for this kind of inference
40 emerge early in development. From early childhood, people can infer the pres-
41 ence of a hidden agent based on the structure of the environment [Newman
42 et al., 2010; Keil & Newman, 2015; Saxe et al., 2005; Ma & Xu, 2013]; we can
43 estimate the difficulty associated with fulfilling different tasks [Gweon et al.,
44 2017; Bennett-Pierre et al., 2018; Yildirim et al., 2019]; and we can explain be-
45 havior in terms of unobservable mental states like beliefs, desires, and intentions
46 [Wellman, 2014; Gopnik et al., 1997]. Nonetheless, combining these capacities
47 to imbue objects with meaning is no simple feat. As our example above shows,
48 seeing an object as communicative requires not only representing the commu-
49 nicator’s beliefs and desires about the world, but also their expectation about
50 the beliefs that we would form upon encountering the communicative object.

51 In this paper we provide a computational model that uses Theory of Mind
52 to recognize the communicative meaning of objects. We show that a basic ca-
53 pacity to estimate the effort associated with different actions, combined with an
54 ability to infer mental states, spontaneously gives rise to the creation and use of
55 communicative objects. Our model makes a number of qualitative predictions
56 that are consistent with examples of everyday communicative objects (like those
57 shown in Figure 1), and we also confirm these predictions empirically, showing
58 that people can infer the meaning of objects that are not typically seen as com-
59 municative. Moreover, because our model is built upon universal components
60 of human cognition, it predicts that the way we recognize objects as commu-
61 nicative ought to be independent of culturally-variable experiences. To this end,

we replicate the key features of our predictions with the Tsimane’—a farming-foraging group living in the Bolivian Amazon. Finally, we use our computational model to generate quantitative predictions that enable us to rigorously test our account, and we show that people’s inferences about the meaning of objects can be quantitatively predicted by our model, but not by closely-related lesioned models.

Figure 1: Real-world examples of people communicating through objects. (a) A hat on a chair indicating that someone intends to return. (b) Rope a few inches above the grass so that people know not to walk through. (c) Chairs along the side of the street in South Boston to reveal someone shoveled and claimed this parking spot. (d) A traffic cone in front of some stairs signaling limited access. (e) A bucket along the side of the street in central Mexico indicating that the parking spot is reserved. (f) An easy-to-cross fence marking a property limit. (g) A stanchion across a stairwell revealing access may be restricted to certain individuals. (h) Belt barriers at the airport telling passengers that they should form a line (and where). (i) An unconventional ironing board along the side of the street indicating that the parking spot is taken. (j) A wooden pole and two small benches in a store in Bolivia indicating that the owner is not available. (k) A small rope along a sidewalk asking people not to walk near a construction site. (l) A pair of traffic posts preventing people from using this walkway.

2. Computational Framework

Our computational model is structured around an assumption that agents act to maximize their subjective utilities—the difference between the costs that they incur and rewards that they obtain. This assumption is at the heart of

human mental-state inferences in adults [Jern et al., 2017; Baker et al., 2017; Jara-Ettinger et al., 2019b], it emerges early in development [Liu et al., 2017; Jara-Ettinger et al., 2016; Gergely & Csibra, 2003; Lucas et al., 2014], and supports high-level cognition from moral reasoning [Ullman et al., 2009; Jara-Ettinger et al., 2015; Kleiman-Weiner et al., 2017], to language understanding [Grosse et al., 2010; Hanna & Tanenhaus, 2004; Jara-Ettinger et al., 2019a] and pedagogy [Bridgers et al., 2020].

Our framework instantiates Theory of Mind as a form of simple recursive social reasoning, similar to models developed to understand pedagogical demonstrations [Ho et al., 2016; Shafto et al., 2014], pragmatics [Frank & Goodman, 2012; Goodman & Frank, 2016], and mental-state inferences [Ullman et al., 2009]. For simplicity and clarity, we describe our model in the context of a simple event similar to the ones we use in Experiments 1-3. Here, agents encounter two doors and must decide which one to walk through. Before they do, an actor has the opportunity to position objects in front of either door to affect their choices. The actor can build physical constraints by stacking multiple objects. Our model thus consists of (i) an actor that modifies scenes—collections of objects, each associated with a position in space—with the goal of modifying other people’s behavior, and (ii) an observer that reacts to the scene by thinking about why the actor modified it. Critically, we assume that the actor and observer are never in the scene at the same time, and we give our model no concept of communication or notion that objects can be used to communicate.

2.1. Actor model

Suppose that an actor wants to ensure that observers take action a in a situation that is initially in configuration s_0 (e.g., ensuring that agents take the second door upon encountering two initially unobstructed doors with some objects between them). The actor’s utility for moving objects to transform scene

99 s_0 into scene s is given by

$$U_A(s; a, s_0) = R_A(a)p_O(a|s) - C_A(s, s_0). \quad (1)$$

100 Here, $R_A(a)$ is the actor’s reward when observers take action a (i.e., capturing
 101 the actor’s motivation to ensure that observers take action a), and $p_O(a|s)$
 102 is the probability that observers will indeed take desired action a when they
 103 encounter scene s .¹ The probability that the observer takes the desired action
 104 a is computed by integrating the space of unknown reward functions $R_O \in R$
 105 that the observer may have, such that

$$p_O(a|s) = \sum_{R_O \in R} p_O(a|s, R_O)p_O(R_O) \quad (2)$$

106 where $p_O(a|s, R_O)$ is estimated through a mental model of observers (described
 107 in the sub-section below).

108 The expected reward in Eq. 1 ($R_A(a)p_O(a|s)$) is then balanced against the
 109 cost $C_A(s, s_0)$ of transforming scene s_0 into scene s . Combined, the first term
 110 leads the actor to prefer scenes where they expect observers to take the desired
 111 action a (regardless of their unknown rewards R_O), and the second term leads
 112 the actor to favor minimal, but effective, scene changes over costly ones.

113 Given the actor’s utility function, we define the probability of transforming
 114 scene s_0 into scene s as the softmax of Eq. 1, given by

$$p_A(s; a, s_0) \propto \exp(U_A(s; a, s_0)/\tau) \quad (3)$$

115 where $\tau \in (0, \infty)$ is the model’s temperature parameter. As $\tau \rightarrow 0$, the agent

¹For clarity, here we present a formulation where there is only one action that the actor wants observers to take. The framework can be trivially extended to cases where there are multiple actions that satisfy the actor’s desires.

116 converges towards optimal actions, always selecting those with the highest util-
 117 ity. As $\tau \rightarrow \infty$, the softmax function transforms the optimal policy into a
 118 uniform distribution over actions, regardless of the utilities. In our model, we
 119 use softmaxing to provide a small relaxation to the assumption that agents are
 120 strictly optimal [as is typical in Theory of Mind models; Baker et al., 2009, 2017;
 121 Jara-Ettinger et al., 2019b]. Throughout our analyses and experiments, we set
 122 $\tau = 0.1$ prior to all data collection and did not fit the parameter to participant
 123 judgments.

124 2.2. Observer model

125 Computing the utility associated with each scene requires the actor to predict
 126 how the observer will react to different situations ($p_O(a|s, R_O)$ in Eq. 2). We
 127 define this probability as the softmax of the observer’s utility function, given by

$$U_O(a; s, \phi) = R_O(a, s) - C_O(a, s) + \phi \left[\sum_{R_A \in R} R_A(a) p_A(R_A|s, s_0) \right]. \quad (4)$$

128 $R_O(a, s)$ and $C_O(a, s)$ capture the observer’s egocentric rewards and costs, re-
 129 spectively, for taking action a in scene s . The last term then represents the
 130 actor’s expected reward, weighted by a cooperation parameter ϕ that captures
 131 the observer’s motivation to pursue, act against, or ignore the actor’s prefer-
 132 ences. When $\phi = 0$ the observer decides what to do based only on how much
 133 they value different outcomes, and how difficult they are to pursue. When ϕ
 134 is positive or negative, the observer infers the absent actor’s desires based on
 135 scene s , and integrates them into their choice.

136 Because the observer does not know the actor’s reward function, this term
 137 is inferred given the knowledge that the actor chose to modify scene s_0 into
 138 scene s . In line with computational models of Theory of Mind [Baker et al.,
 139 2017, 2009; Jara-Ettinger et al., 2019b], we assume that the observer infers the

140 actor’s reward function R_A via Bayesian inference:

$$p_A(R_A|s, s_0) \propto \ell_A(s|R_A, s_0)p_A(R_A), \quad (5)$$

141 with the likelihood given by the softmax of Eq. 1 (also using a $\tau = 0.1$ through-
142 out).

143 Note that our model assumes that, when an observer encounters scene s ,
144 they also know the scene’s original state s_0 . This allows observers to infer
145 which objects an agent moved (by comparing the difference between scene s and
146 scene s_0). In reality, it is more likely that observers have a prior distribution
147 over scenes $p(s_0)$ rather than perfect knowledge about the initial scene s_0 . For
148 instance, when encountering a broom placed across an entrance, an observer
149 may not know where the broom was situated before an agent moved it, but
150 they may believe it was more likely that it was positioned elsewhere. Modeling
151 prior expectations about scene distributions is beyond the scope of our model,
152 but we return to the implications of this assumption in the Discussion.

153 The actor and observer model specify a recursive relation where the actor’s
154 utility depends on the observer’s utility (Eq. 1) and the observer’s utility de-
155 pends on the actor’s utility (Eq. 4). After three levels of depth [which is well
156 within human capacity; O’Grady et al., 2015; Goodman & Frank, 2016], we
157 ground the recursive structure in a model of an observer with no Theory of
158 Mind (using $U_O(a; s) = R_O(a, s) - C_O(a, s)$ in place of Eq. 4).

159 2.3. Model Analysis

160 We begin by analyzing our model’s performance in a hypothetical situation
161 that shows how mental-state reasoning gives rise to the creation and under-
162 standing of communicative objects. We then present a set of simple qualitative
163 studies designed to validate core predictions and model assumptions. Finally,

we present a quantitative experiment where we generate and test our model’s graded predictions against participant judgments.

We illustrate how our model spontaneously uses objects in a communicative manner in a simple domain with two doors—door *A* and door *B*—and an actor that wants to ensure that observers always select door *A*. We assume that the initial scene s_0 has a set of objects between the two doors, such that the objects do not initially block either door and placing an object in front of either door is equally costly. For simplicity, we also assume that the cost the actor incurs in placing an object in front of a door equals the cost that the observer pays to move that object out of the way (see Experiment 4 for cases where we relax this assumption and have asymmetrical costs for the actor and the observer).

Figure 2: (a) Actor model spontaneously using objects to communicate. Here, an actor wants to ensure that observers avoid a target door. The x-axis shows the observer’s preference (negative values indicating that observers already disprefer walking through the target door) and the y-axis shows the number of objects the actor chooses to stack in front of the target door (negative values indicating stacking objects on the opposite door). When the observer lacks a Theory of Mind (blue line), the actor builds a physical constraint as a function of the observer’s preference, effectively building a barrier. When the observer has a Theory of Mind (yellow line), the actor places a single object in front of the target door, even when the observer has a strong preference for going through it. (b) Observers reacting to objects in front of their preferred door. Observers lacking a Theory of Mind (blue line) become more likely to avoid the door as a function of how many objects are blocking it. Observers with a Theory of Mind (yellow line) recognize the meaning of a single object and show a sharp discontinuity in how they adjust their behavior. The temperature parameter was set to $\tau = 0.1$, the actor’s reward was set to $R_A = 9$, and the observer’s cooperation parameter was set to $\phi = 10.0$.

Figure 2 shows our model’s performance in a situation where the observer’s rewards are known. When the actor expects the observer to prefer door *A*

177 (negative observer preferences in Figure 2a), the rational choice is to do nothing,
 178 as the observer will naturally act the way the actor hopes that they will. As
 179 the observer shifts their preference to door B , the actor begins to modify the
 180 scene, producing two different types of behavior as a function of the observer's
 181 Theory of Mind. When the observer lacks a Theory of Mind (removing the last
 182 term in Eq. 4), the observer is expected to respond only to their egocentric costs
 183 and rewards. Consequently, the actor stacks objects in front of door B , placing
 184 the minimum number of objects necessary to dissuade the observer (as moving
 185 objects is also costly for the actor). This is captured in Figure 2a, where the
 186 blue line shows how the actor stacks more objects as the observer's preference
 187 becomes stronger.

188 When the observer is cooperative and has a Theory of Mind (using Eq. 4
 189 with $\phi > 0$), the actor no longer attempts to build purely physical constraints.
 190 Instead, the actor now places a single object in front of door B , even when the
 191 observer has a strong preference for that door, as the actor expects the observer
 192 to infer the object's meaning and react accordingly. This is captured in Figure
 193 2a, where the yellow line shows the actor using at most one object, independent
 194 of the observer's preference.

195 Figure 2b shows how our observer model recognizes the meaning of commu-
 196 nicative objects. Observers with no Theory of Mind (blue line) respond to the
 197 physical cost alone, becoming more likely to abandon their preferred door as
 198 a function of the number of objects blocking it. By contrast, observers with a
 199 Theory of Mind (yellow line) show a sharp discontinuity, recognizing a single ob-
 200 ject's meaning, and adjusting their behavior accordingly. Together, these results
 201 show how our model gives rise to actors who use objects in a communicative
 202 manner and observers who recognize the meaning of these objects.

203 Our model analyses show proof-of-concept that Theory of Mind, combined

204 with an ability to estimate the cost of different actions, can give rise to the
205 creation and recognition of communicative objects. Yet, creating and recog-
206 nizing objects are asymmetrical activities in two ways. First, according to our
207 model, the capacity to recognize the meaning of objects is cognitively simpler
208 than creating them, requiring one fewer level of recursion (i.e., the actor model
209 that invents communicative objects is at the top level, while the observer model
210 that recognizes them is one level below; see SI for full details). Second, for com-
211 municative objects to become ubiquitous, the ability to recognize them must
212 be widespread, while the ability to invent them can be restricted to a few in-
213 dividuals. Thus, our experiments focus on testing people’s ability to recognize
214 communicative objects, infer their meaning, and invoke this meaning in subse-
215 quent encounters.

216 Experiments 1-3 are qualitative. Experiments 1 and 2 first test how people
217 react to objects that impose a low cost to observers. In Experiment 1, partici-
218 pants are asked to infer which of two objects has a communicative purpose, and
219 in Experiment 2, participants are asked to infer an object’s meaning. Experi-
220 ment 3 then presents evidence that people store the inferred meaning of objects
221 and retrieve it in subsequent encounters. Finally, Experiment 4 is quantita-
222 tive, where we test communicative inferences in a graded inference task. We
223 compare participant judgments to our model as well as to lesioned models that
224 highlight the contribution of different components of our account. All studies
225 were approved by Yale’s IRB (protocols “Culture and Cognition” #2000022403
226 and “Online reasoning” #2000020357).

227 **3. Experiment 1**

228 In Experiment 1, we began by testing whether people believe that ob-
229 jects that impose a low cost to observers are more likely to be communicative.

Through this, we also test an alternative possibility to our account: Commu-
 nicative objects may gain and sustain their meaning through convention alone,
 with no Theory of Mind at play. As children, for instance, most of us likely
 ignored strap barriers at banks, movie theatres, and DMVs, and we only began
 to see their social meaning after watching how other people reacted to these
 objects. If the meaning of objects were purely conventional, then the mapping
 from objects to meanings should be arbitrary. By contrast, our model pre-
 dicts that only objects that impose a cost to observers should be perceived as
 communicative.

Figure 3: (a-b) Example stimuli from Experiment 1. In both panels, the left side shows the low-cost door and right side shows the no-cost door. (c) Experiment 1 results. The blue bar represents the percentage of participants that inferred that the low-cost door had the communicative object. Error bars are bootstrapped 95% CIs.

3.1. Participants

80 U.S. participants (as determined by their IP address) were recruited using Amazon Mechanical Turk ($M = 37.84$ years, $SD = 12.22$ years). 15 additional participants were recruited and replaced for failing our inclusion criteria (see Results).

3.2. Stimuli

Our stimuli consisted of 16 images of doors, each with a single object in front of it, like in Figure 3a-b (see SI for full stimuli set). Each object was

247 associated with two different doors: a “low-cost” door, where the object was
248 placed directly in the middle of the doorway (e.g., Figure 3a, left), or a “no-
249 cost” door, where the object was placed next to the door, not blocking the way
250 (e.g., Figure 3a, right). We used eight different objects across participants that
251 were unconventional in their usage: a plant, a chair, a pile of books, a pile of
252 cinderblocks, some tape, some meter sticks, a hat, and a fishbowl tied to a tack
253 on a door frame. Half of these doors were open and the other half were closed.

254 3.3. Procedure

255 Participants read a brief cover story that asked them to imagine leaving an
256 office and encountering a pair of doors with objects in front of both of them. Af-
257 ter describing the objects to participants, they had to answer a single multiple-
258 choice attention check question (“What objects are in front of the doors?”).
259 Participants that failed to answer the attention check question correctly were
260 sent to the beginning of the cover story and not permitted to access the experi-
261 ment until they answered the question correctly. In the experiment, participants
262 completed a single trial where they saw a low-cost door and a no-cost door, both
263 with the same object in front of them (see Figure 3a-b for two examples; door or-
264 der randomized). Participants were then asked three test questions (sequentially
265 and only one at a time so that they could not revise their previous answers):
266 “Which door do you think someone was trying to tell you something?” (pos-
267 sible responses: “The door on the left” or “The door on the right”), “Which
268 door requires more work to walk through?” (possible responses: “The door on
269 the left,” “The door on the right,” or “Equally easy”), and “Do you think you
270 would be able to walk through this door if you wanted to?” (possible responses:
271 “Yes” or “No”).

272 3.4. Results

273 Participants who said either door was impossible to walk through were
274 excluded from the study and replaced ($n = 15$). Figure 3c shows the re-
275 sults from Experiment 1. When asked which of the two objects was meant
276 to be communicative, 72.50% of participants selected the low-cost door ($CI_{95\%}$:
277 62.50% – 82.50%; $p < 0.001$ from a two-tailed binomial test), suggesting that
278 participants indeed saw the cost as having a communicative intent, rather than
279 performing at chance because the object was not a conventional communicative
280 object.

281 Because our interpretation of door difficulty might be different than those
282 of participants (i.e., they might find some of our no-cost doors as more costly
283 than their low-cost counterpart), we ran a follow-up analysis that instead used
284 participants’ judgments of door difficulty based on their responses to our second
285 test question. Using these labels, we found that 76.20% of participants selected
286 the door they identified as having a lower cost ($CI_{95\%}$: 66.25% – 85.00%; $p <$
287 0.001 from a two-tailed binomial test), reinforcing our previous analysis.

288 4. Experiment 2a

289 The results from Experiment 1 show that people were more likely to judge
290 an object as communicative if it imposed a low cost, relative to an object that
291 did not impose a low cost. However, participants in this experiment were not
292 asked to judge what the object meant. In Experiment 2a, we tested a further
293 qualitative prediction our model makes that a convention-based account does
294 not: people should infer that an object that imposes a low cost is signaling
295 observers to avoid that path, even if the object is unconventional.

296 4.1. *Participants*

297 160 U.S. participants (as determined by their IP address) were recruited via
298 Amazon Mechanical Turk ($n = 80$ per condition; $M = 39.50$ years, $SD = 17.68$
299 years). 17 additional participants were recruited and replaced for failing our
300 inclusion criteria (see Results).

301 4.2. *Stimuli*

302 Our stimuli consisted of 16 images of pairs of doors, each with a single ob-
303 ject in front of one of the doors, like in Figure 4a-b. Each object was associated
304 with two different pairs of doors: a “low-cost” pair, where the object was placed
305 directly in the middle of one of the doorways (e.g., Figure 4a; door order ran-
306 domized), or a “no-cost” pair, where the object was placed next to one of the
307 doors, not blocking the way (e.g., Figure 4b; door order randomized). We used
308 eight objects identical to those used in Experiment 1. Half of these pairs of
309 doors were open and the other half were closed.

310 4.3. *Procedure*

311 Participants read a brief cover story that asked them to imagine leaving
312 an office and encountering a pair of doors with an object in front of one of
313 them. After describing the object to participants, they had to answer a single
314 multiple-choice attention check question (“What is the only difference between
315 the two exits?”). Participants that failed to answer the attention check question
316 correctly were sent to the beginning of the cover story and not permitted to
317 access the experiment until they answered the question correctly. Participants
318 then saw a single trial containing a pair of doors: one door with an object in front
319 of it and the other door with nothing near it. The object was either directly in
320 front of the door, blocking the way (the low-cost condition; e.g., Figure 4a), or
321 next to the door, not blocking the way (the no-cost condition; e.g., Figure 4b).

322 Participants in both conditions were told that the object was communicative
 323 and were asked three test questions (sequentially and only one at a time so that
 324 they could not revise their previous answers): “What do you think someone was
 325 trying to tell you about the door with the OBJECT?” (possible responses: “You
 326 **should** walk through the door with the OBJECT” or “You **should not** walk
 327 through the door with the OBJECT”), “Which door requires more work to walk
 328 through?” (possible responses: “The door on the left”, “The door on the right”,
 329 or “Equally easy”), and “Do you think you would be able to walk through this
 330 door if you wanted to?” (possible responses: “Yes” or “No”).

331 *4.4. Results*

332 Participants who said the unobstructed door was harder to walk through
 333 were excluded from the study and replaced ($n = 17$). Figure 4c shows the
 334 results from Experiment 2a. Consistent with our model, participants in the low-
 335 cost condition inferred that the object meant that the door should be avoided
 336 (88.75%; $CI_{95\%}$: 81.25% – 95.00%; $p < 0.001$ from a two-tailed binomial test).
 337 Participants also confirmed that the object did not physically prevent them
 338 from walking through the door (87.50%; $CI_{95\%}$: 80.00% – 93.75%), showing
 339 that their response was indeed driven by a communicative inference and not by
 340 the perception that the object was an insurmountable physical constraint. By
 341 contrast, only 60.00% of participants in the no-cost condition inferred that the
 342 door should be avoided ($CI_{95\%}$: 48.75% – 71.25%; $p = 0.093$ from a two-tailed
 343 binomial test; 83.75% of participants confirmed that object was not blocking
 344 the door; $CI_{95\%}$: 75.00% – 91.25%).

345 As in Experiment 1, we ran a follow-up analysis using participant judgments
 346 of door difficulty instead of our own. Using these labels, we found that 83.65%
 347 of participants in the low-cost condition inferred that the object meant that the
 348 door should be avoided ($CI_{95\%}$: 75.96% – 90.38%; $p < 0.001$ from a two-tailed

binomial test). By contrast, only 57.14% of participants in the no-cost condition inferred that the door should be avoided ($CI_{95\%}$: 44.64% – 69.64%; $p = 0.350$ from a two-tailed binomial test), again reinforcing our previous analysis.

5. Experiment 2b

Although our experiments used objects that are not typically used to communicate, it is possible that participants found the objects reminiscent of communicative objects that are common in the U.S., like construction tape or traffic cones, and responded accordingly. To test this possibility, we replicated a variation of Experiment 2a with the Tsimane’—a farming-foraging group native to the Bolivian Amazon. The Tsimane’ live in non-industrialized communities along the Maniqui river and have little exposure to market-integrated communities. Recent research has found that the Tsimane’ differ from WEIRD populations (Western, educated, and from industrialized, rich, and democratic countries; Henrich et al. 2010) in several cognitive domains including number-word learning [Piantadosi et al., 2014; Jara-Ettinger et al., 2017], color-word vocabulary [Conway et al., 2020; Gibson et al., 2017], and music perception [McDermott et al., 2016].

5.1. Participants

150 Tsimane’ adults were recruited in their local communities in the Bolivian Amazon ($M = 33.13$ years, $SD = 15.31$ years). 7 additional participants were recruited and replaced for failing our inclusion criteria (see Results).

5.2. Stimuli

Our stimuli consisted of six images of doors, each with a single object in front of it, like in Figure 3a-b. Each object was associated with two different doors: a “low-cost” door, where the object was placed directly in the middle

of the doorway (e.g., Figure 3a, left), or a “no-cost” door, where the object was placed next to the door, not blocking the way (e.g., Figure 3a, right). Here we used a subset of the objects used in Experiments 1 and 2a that our Tsimane’ participants were still familiar with while remaining unconventional as communicative deterrents: a plant, a chair, and a pile of cinderblocks. Half of these doors were open and the other half were closed.

5.3. Procedure

Participants were read a brief cover story that asked them to imagine trying to enter a friend’s house through one of two possible doors with objects in front of them. Participants were then shown a low-cost door and a no-cost door sequentially (order counterbalanced, with different objects used in each door). Participants were asked to infer whether the object meant they should enter through or avoid the door in each trial. Specifically, we asked: “Do you think the owner wants you to enter or stay away?”

5.4. Results

Participants that did not complete both trials were excluded from the study ($n = 7$). Like U.S. participants, Tsimane’ participants inferred that they should avoid a door when the object was minimally blocking the door (80.42%; $CI_{95\%}$: 73.43% – 86.71%; $p < 0.001$ from a two-tailed binomial test), but not when the object was on the side of the door (25.87%; $CI_{95\%}$: 18.88% – 32.87%; $p < 0.001$ from a two-tailed binomial test). These results build upon our previous findings, suggesting that people can not only detect which objects are more likely to be communicative, but can also spontaneously infer their meaning.

6. Experiment 3

Experiments 1 and 2 provide qualitative evidence for our account, and they suggest that the detection of communicative objects is not driven by convention

Figure 4: (a) Example stimuli from the low-cost condition in Experiment 2a. (b) Example stimuli from the no-cost condition in Experiment 2a. (c) Experiment 2a results. The blue bars represent the percentage of U.S. participants that selected the unobstructed door as a function of condition. (d) Experiment 2b results. The blue bars represent the percentage of Tsimane' participants that inferred that they should not go through the door as a function of door type. Error bars are bootstrapped 95% CIs.

alone. At the same time, these results do not imply that conventions are never at play. Because using Theory of Mind is arguably cognitively demanding [Birch & Bloom, 2007; Back & Apperly, 2010; Keysar et al., 2000] and people minimize cognitive effort [Shenhav et al., 2017; Kool & Botvinick, 2018], it is likely that Theory of Mind enables people to invent new communicative objects and infer their meaning, but that pedagogy and convention are the engines supporting the widespread use of these objects. This hypothesis would help explain why some communicative objects, like chairs to indicate shoveled parking spots, are used in a consistent manner, and is parsimonious with a resource-rational view of the mind [Lieder & Griffiths, 2019; Griffiths et al., 2015]. In Experiment 3 we sought initial evidence for this possibility.

6.1. Participants

160 U.S. participants (as determined by their IP address) were recruited using Amazon Mechanical Turk ($n = 80$ per condition; $M = 36.13$ years, $SD = 10.90$ years). 42 participants were recruited and replaced for failing our inclusion criteria (see Results).

Figure 5: (a) Example stimuli from the first trial of Experiment 3. (b) Example of a door from the congruent condition relative to (a). (c) Example of a door from the incongruent condition relative to (a). (d) Experiment 3 results. The blue bars represent the percentage of participants that inferred that they should avoid the door with the picture as a function of condition. (e) The orange bars represent the average confidence rating in their inference in the second trial of the experiment. Error bars are 95% bootstrapped CIs.

6.2. Stimuli

Our stimuli consisted of 16 images of pairs of doors, each with a single object in front of one of the doors, like in Figure 5a. Each object was associated with two different pairs of doors: a “low-cost” pair, where the object was placed directly in the middle of one of the doorways (e.g., Figure 5a), and a “symbol” pair, where a picture of the object was placed directly in the middle of one of the doorways (e.g., Figure 5b-c). We used eight objects identical to those used in Experiments 1 and 2a. Half of these pairs of doors were open and the other half were closed.

6.3. Procedure

Participants read a brief cover story that asked them to imagine leaving an office and encountering a pair of doors with an object in front of one of them. After describing the object to participants, they had to answer a single multiple-choice attention check question (“What is the only difference between the two exits?”). Participants that failed to answer the attention check question correctly were sent to the beginning of the cover story and not permitted to access the experiment until they answered the question correctly. In the first

433 trial, participants saw two doors: one with an object directly in front, blocking
 434 the way, and the other with nothing near it (e.g., Figure 5a). Participants were
 435 asked two test questions: What do you think NAME was trying to tell you about
 436 the door with the OBJECT? (possible responses: “You **should** walk through
 437 the door with the OBJECT” or “You **should not** walk through the door with
 438 the OBJECT”) and “How confident are you that that’s what NAME was trying
 439 to tell you?” (using a continuous slider ranging from “not confident at all” to
 440 “very confident”). Regardless of their first two answers, participants were told
 441 that their inference was correct. Then, in the second trial, participants saw a
 442 door with a picture in front of it instead of an object and were asked the same
 443 two test questions. Participants in the congruent condition saw a picture of
 444 the same object they saw in the first trial (e.g., Figure 5b if participants first
 445 saw Figure 5a) and participants in the incongruent condition saw a picture of a
 446 different object (e.g., Figure 5c if participants first saw Figure 5a). At the end
 447 of the experiment, participants were asked “Which door requires more work to
 448 walk through?” (“The door on the left”, “The door on the right”, or “Equally
 449 easy”) for both trials.

450 6.4. Results

451 Participants who did not respond that the object in the first trial imposed
 452 a cost were excluded from the study and replaced ($n = 42$), as our interest is in
 453 how participants generalized the inferred meaning of objects that they perceived
 454 as imposing a cost. We predicted that if participants treat the picture as a new
 455 signal, they should perform at chance in both conditions, as the picture does
 456 not impose a cost. By contrast, if participants assume that the signaler was
 457 treating the object as a newly-formed convention, they should infer that the door
 458 with the picture should be avoided in the congruent condition, but perform at
 459 chance in the incongruent condition. As predicted, participants in the congruent

condition judged that the door with the picture should be avoided, despite the picture not imposing a cost (67.50%; $CI_{95\%}$: 57.50% – 77.50%; $p < 0.01$ from a two-tailed binomial test). In the incongruent condition, only 43.75% inferred that the door should be avoided ($CI_{95\%}$: 32.50% – 55.0%; $p = 0.314$ from a two-tailed binomial test). Moreover, participants were significantly more confident in their interpretation of the picture in the congruent condition relative to the incongruent condition ($W = 3978$; $p < 0.01$ by a U -test; $\Delta = 11.63\%$; $CI_{95\%}$: 3.82% – 19.23%).

7. Experiment 4

Having found qualitative evidence for our proposal, we next used our computational model to generative quantitative predictions that we could compare to graded participant inferences.

7.1. Participants

80 U.S. participants (as determined by their IP address) were recruited using Amazon Mechanical Turk ($n = 40$ per condition; $M = 34.81$ years, $SD = 10.31$ years).

7.2. Stimuli

Our stimuli consisted of 27 gridworlds, like those in Figures 6a-c. Each gridworld was 10-by-10 squares in size and represented some farmland with two distinct fruit groves: pears and pomegranates. The stimuli was designed by parametrically varying two factors: the distance between the hiker (shown in blue) and the boulders (5, 7, and 9 squares away; i.e., the natural cost of the environment) and the number of boulders that the hiker must walk around (1, 2, or 3 boulders in front of a fruit grove; i.e., the artificial cost introduced by another agent). The hiker’s starting position was randomly selected to be at

one of the four corners, and the fruit groves were randomly placed on the two adjacent corners relative to the hiker.

7.3. Procedure

Participants read a cover story about hikers (corresponding to observers in our model) who sometimes walk through groves where farmers (corresponding to actors in our model) grow pears and pomegranates. Participants learned that, unbeknownst to hikers, farmers do not mind hikers taking pears, but they want to protect the pomegranates. To achieve this, farmers place boulders in front of their groves. Participants were then told they would see different farms with a hiker and a set of boulders positioned by a (now absent) farmer (see Figures 6a-c for examples). At the end of the cover story, participants were given a brief, five-question quiz (see SI for questions). Participants that gave incorrect answers to one or more questions were sent to the beginning of the cover story and not permitted to access the experiment until they answered all of the questions correctly.

Participants in the non-communicative condition were told that farmers did not expect hikers to realize that the boulders were placed intentionally, and would therefore treat them as natural constraints. In each trial, participants were asked to infer how much the farmer expected the hiker to like pomegranates. Participants in the communicative condition were told that farmers expected hikers to recognize that the boulders were placed intentionally. In each trial, participants answered two questions: “How much does this farmer think that hikers like pomegranates?” and “How cooperative does this farmer think hikers are?”, both using continuous sliders ranging from “not at all” to “very much.” Note that participants were not told that the objects might be communicative in either condition. All participants completed the same 27 trials where we varied both the initial cost of obtaining each type of fruit (by

Figure 6: (a-c) Example stimuli from Experiment 4. In these examples, both fruit groves were equally far and only varied on the number of boulders a farmer placed. (d-f) Model predictions and participant judgments from the non-communicative condition in purple and blue, respectively. (g-i) Model predictions and participant judgments from the communicative condition in purple and blue, respectively. Inference type is along the x-axis and the inferred value is along the y-axis. Error bars are bootstrapped 95% CIs.

manipulating the initial distance from the hiker), and the number of boulders that the farmer added (ranging from 1-3; see SI for details on cost estimations).

7.4. Results

Figure 7a shows the results of Experiment 4. Overall, our model and participant judgments had a correlation of $r = 0.97$ ($CI_{95\%}$: 0.95% – 0.98%). Figure 6a-c shows three example trials where the initial cost of getting pomegranates and pears was matched and the number of boulders varied from 1 to 3. Figure 6d-f shows model and participant reward inferences in the non-communicative condition (with each plot corresponding to the stimuli directly above it). Like

our model, participants assumed that, when objects lack communicative intent, the number of objects blocking the pomegranates directly reveals the hiker's expected reward. Thus, participants reported that farmers who placed more boulders also thought that hikers had a stronger preference.

Figure 6g-i shows how participants reacted to the same trials in the communicative condition. Overall, participants judged that the hiker was expected to be highly cooperative when a single boulder was blocking the pomegranates, and uncooperative when there were three boulders blocking the pomegranates. This pattern shows how participants recognized that the success of a single boulder depended on the hiker being cooperative and that, if the farmer put the effort of placing multiple boulders, it must have been because they expected the hiker to not care about the farmer's preferences.

Our model predicts that participants in the communicative condition should continue to infer stronger hiker rewards as a function of boulders in the scene, with a key difference: When a low-cost object is interpreted as non-communicative, its low cost reveals that farmers did not think hikers would want the pomegranates that much. By contrast, when an object is interpreted as communicative, its low cost no longer implies that the expected reward was lower (as it did in the non-communicative condition), because its purpose is to reveal the farmer's preferences rather than block the way. This predicts that the expected inferred reward in the cooperative condition should be higher than the inferred reward in the non-cooperative condition. However, as the number of boulders increases, participants should cease to infer that the object had a communicative purpose, and their reward inferences should converge to the ones from the non-communicative condition. This is indeed the pattern visible in Figure 6g-i (for the three example trials from Figure 6a-c).

This predicted pattern also appeared more broadly when averaging reward

Figure 7: Experiment 4 results. Each point represents a judgment, with model predictions on the x-axis and participant judgments on the y-axis. Data includes the reward inferences made in the non-communicative (NC) condition and the reward and cooperation inferences made in the communicative (C) condition across all 27 trials. (a) Evaluation of our full model and participant judgments. (b) Evaluation of lesioned models on participant judgments. Model lesions include removing the influence of cost from the actor (left) and removing the influence of cost from the observer (right).

inferences across all trials as a function of boulders in the scene (Figure 8). In trials with a single boulder, participants inferred an average reward of 0.31 ($CI_{95\%}$: 0.20 – 0.43) in the communicative condition and 0.07 ($CI_{95\%}$: 0.01 – 0.18) in the non-communicative condition. This effect was weaker in the two-boulder trials, where participants inferred an average reward of 0.63 ($CI_{95\%}$: 0.54 – 0.71) in the communicative condition and 0.51 ($CI_{95\%}$: 0.40 – 0.60) in the non-communicative condition. Finally, this difference disappeared in the three boulder trials, with an inferred reward of 0.96 ($CI_{95\%}$: 0.85 – 1.00) in the communicative condition and 0.93 ($CI_{95\%}$: 0.78 – 0.99) in the non-communicative condition.

Beyond varying the number of boulders placed by farmers, our experiment also manipulated the initial distance between the hikers and each fruit grove. When boulders are not communicative, their purpose is to make the pomegranates difficult to obtain. Therefore, farmers should be more likely to block pomegranates when they are easy to reach, even if hikers don’t have a particular preference (and even more so when pears are in an inconvenient lo-

564 cation). However, farmers should block pomegranates that are hard to reach
 565 only when they expect hikers to really like them (but less so if the pears are in
 566 a convenient location). Consistent with this, a linear regression predicting in-
 567 ferred rewards in the non-communicative condition revealed a positive effect of
 568 boulders ($\hat{\beta} = 0.43, p < 0.001$), a positive effect of distance to the pomegranates
 569 ($\hat{\beta} = 0.05, p < 0.005$), and a negative effect of distance to the pears ($\hat{\beta} = -0.03,$
 570 $p < 0.05$).

571 When boulders can be used to communicate, their purpose may be to reveal
 572 the farmer’s desires, rather than to increase the cost of getting the fruit. As
 573 such, opening the possibility that objects are communicative should diminish the
 574 importance of the distances between the groves and the hiker. Consistent with
 575 this, a regression predicting inferred rewards in the communicative condition
 576 revealed a positive effect as a function of boulders ($\hat{\beta} = 0.32, p < 0.001$), but
 577 not as a function of the initial distances to the pears or pomegranates ($p > 0.16$
 578 in both cases; see SI for details).

579 Finally, we conducted a last analysis to test the critical components of our
 580 model. Our model’s inferences depend both on understanding that actors incur
 581 costs to position objects, and that these objects impose costs on observers. To
 582 evaluate the role that these assumptions played in our model, we implemented
 583 two model lesions. The first lesion removed the cost actors incur in positioning
 584 objects. In our experiment, this means that farmers do not incur costs when
 585 placing boulders. Figure 7b (left) shows how this lesioned model compares
 586 to participant judgments. This model was able to capture the broad pattern
 587 of inferences about observers ($r = 0.77$; $CI_{95\%}$: $0.66 - 0.85$), but performed
 588 worse than our main model ($\Delta r = 0.20$; $CI_{95\%}$: $0.12 - 0.30$). The second
 589 lesion removed the cost that observers incur in navigating around objects. In
 590 our experiment, this means that hikers do not consider the cost of navigating

591 around boulders. Figure 7b (right) shows how this model was no longer able to
 592 explain participant judgments ($r = -0.07$; $CI_{95\%}: (-0.27) - 0.12$), and was also
 593 reliably worse than our main model ($\Delta r = 1.04$; $CI_{95\%}: 0.84 - 1.23$). Together,
 594 these lesions show how considering the cost that actors incur in positioning
 595 objects, and how these objects also impose cost on observers, are critical for
 596 explaining how participants reasoned about objects in our experiment.

Figure 8: Comparison between participant judgments across the non-communicative (NC) and communicative (C) conditions in Experiment 4. On the x-axis is the number of boulders the farmer placed and on the y-axis is the reward participants inferred the hiker to have. Participants that saw one boulder inferred a significantly higher reward in the communicative condition than in the non-communicative condition. Error bars are bootstrapped 95% CIs.

597 8. Discussion

598 Human communication is remarkable in how we embed meaning into sounds,
 599 gestures, texts, and even objects. Our work shows how a capacity to represent
 600 and infer other people’s mental states spontaneously gives rise to our ability to
 601 use objects to communicate. In particular, our work shows how, if we can rea-
 602 son about the desires based on how an agent chose to arrange different objects,
 603 then we can also re-arrange objects with the purpose of revealing our desires
 604 to agents who encounter these objects. As predicted by our account, partic-
 605 ipants in Experiments 1 and 2 judged that an object had a communicative
 606 purpose when it imposed a low cost to observers, and they inferred its meaning
 607 accordingly. Experiment 3 further showed that people can store an object’s

608 meaning and retrieve it in future encounters, reducing the constant need for
609 any mental-state inference. Finally, Experiment 4 showed that people can make
610 quantitative mental-state inferences when encountering objects in novel situa-
611 tions. Combined, these results suggest that mental-state inferences support the
612 creation and recognition of novel communicative objects, while pedagogy and
613 convention drive their widespread use (see Figure 1 for everyday examples of
614 these objects).

615 *8.1. Model assumptions and study limitations*

616 Our model was built following several computational principles, each of
617 which enjoys strong empirical support: social interactions involve agents think-
618 ing about each other’s mental states [Goodman & Frank, 2016; Frank & Good-
619 man, 2012]; mental-state reasoning is structured around an assumption that
620 agents maximize utilities [Jara-Ettinger et al., 2016; Jern et al., 2017; Lucas
621 et al., 2014]; and inference about other people’s minds is performed via some
622 approximation of Bayesian inference [Baker et al., 2009, 2017; Jara-Ettinger
623 et al., 2019b]. While Experiment 4 provided additional support for each of
624 these principles, our experiments and model also made two further assumptions
625 based on related research.

626 Our first assumption was that, when encountering an object, agents can
627 identify whether it was intentionally placed by an agent or not. Formally, our
628 model achieved that by assuming that the observer knew both the initial scene
629 state s_0 and the final scene state s . While related research shows that some form
630 of this capacity emerges in infancy [Newman et al., 2010; Keil & Newman, 2015],
631 it is likely that in more realistic contexts people do not know the initial scene
632 state s_0 and instead have a distribution over potential initial scenes. Our model
633 can be easily extended to account for this possibility. However, determining
634 how people learn and represent this distribution across contexts is an open

question. Note that, in addition, detecting intentional action from objects can be substantially more complex than what we considered here. Imagine, for instance, finding a coffee cup on a bench. How do we determine if an agent forgot their coffee, chose to leave it behind, or left it with the purpose of signaling their return? In cases like these, the evidence left by an object can be ambiguous, and observers must rely on additional cues to reconstruct what happened (e.g., we may be more likely to believe the agent chose to leave the cup behind if it is empty and crumpled than if it is full of coffee and still warm). More research is needed to understand how people recognize intentional action in the absence of agents.

Our second assumption was that people can estimate the cost associated with placing and moving objects. In our quantitative experiments, costs could be visually determined by looking at the stimuli (Figure 6), and we therefore did not need to integrate a more complex cost estimator into our model. Recent work shows that people can estimate the difficulty of constructing physical displays through a motor simulation using intuitive physics [Yildirim et al., 2019; Battaglia et al., 2013], and integrating this work into our model may enable us to explain communicative objects in more complex situations. This is a direction we hope to pursue in future work.

Finally, our work has three main limitations. First, while our computational model can explain both the creation and recognition of communicative objects, our experimental work focused on the second component. We did so because, in our theory, all people ought to be able to detect and infer the meaning of communicative objects, while the ability to generate them can be limited only to a few individuals. We thus do not know the extent to which people can create ad-hoc communicative objects in new situations.

A second limitation of our work is that we did not consider additional forces

662 that shape communicative objects. When people create communicative objects,
663 the costs that these objects impose to the observer is not the only consideration
664 at play. An effective communicative object must also be noticeable to observers
665 and it must reveal that it was placed intentionally. Thus, a more complete
666 model would require modeling how actors decide to arrange objects so they are
667 noticeable and appear maximally intentional.

668 Our third limitation is that we focused our analysis on a specific class of
669 objects where agents reveal their intentions by imposing minimal costs on ob-
670 servers. These types of objects are widespread and easy to find in our everyday
671 lives (see Figure 1 for examples). Despite their ubiquity, to the best of our
672 knowledge, no work in cognitive science has attempted to formalize how peo-
673 ple come to use these objects in communicative manners. At the same time,
674 many objects contain meaning introduced by convention alone, such as how a
675 caduceus represents the medical field or how a light bulb can represent an idea.
676 The mechanism behind how these objects gain meaning may be simpler than
677 the one at work in ad-hoc communicative objects like the ones we explore here.
678 Nonetheless, observers must distinguish when an object is representative of a
679 message, and when it is nothing more than an object. In this case, reasoning
680 about the intention of the agent behind the objects may be critical. Future
681 work will investigate this question.

682 *8.2. Implications and open questions*

683 The process by which people use objects to communicate shares many sim-
684 ilarities with language and pragmatics [Grice, 1975] with the difference that
685 our medium is physical space instead of sound and our costs are physical effort
686 instead of utterance length. Indeed, formal analyses of meaning in language re-
687 semble how our model creates communicative objects. In particular, in Grice’s
688 classical analysis of meaning in human communication, he argued that an utter-

689 ance is meaningful if the speaker intends for it to induce a belief in the listener,
690 and also intends that the utterance is recognized as such [Grice, 1957]. The
691 similarities between the logic of communicative objects and accounts of mean-
692 ing in language suggest that the connection runs deep. Thus, one outstanding
693 question is whether the use of communicative objects emerges as an extension of
694 pragmatic reasoning in language, whether it constitutes a more primitive form
695 of proto-communication that can precede language, or whether the similarities
696 reflect only the shared mechanisms in reasoning about other agents goals and
697 costs.

698 More broadly, the interaction between physical reasoning and social cog-
699 nition extends beyond an ability to use objects to communicate. In particu-
700 lar, recent work has shown that people can also use Theory of Mind to infer
701 the transmission of ideas based on how different agents build similar artifacts
702 [Schachner et al., 2018; Hurwitz et al., 2019]. This “intuitive archeology” likely
703 shares a common basis with the inferences in our model, suggesting that hu-
704 man social reasoning provides rich inferences about the origins, structure, and
705 purpose of objects.

706 Our model implementation relied on an instantiation of Theory of Mind as
707 recursive Bayesian inference. While related work has shown that the amount of
708 recursion in our model is within human capacity [O’Grady et al., 2015; Good-
709 man & Frank, 2016], our results do not imply that this reasoning must be
710 implemented in a recursive manner. In particular, recent work has argued that,
711 in situations like the ones we explored, people can coordinate indirectly through
712 *virtual bargaining*—the idea that coordination without explicit communication
713 can be achieved by agents acting as if they had explicitly discussed and agreed
714 on how to coordinate, provided that such discussion would lead to a single clear
715 solution [Misyak et al., 2016; Misyak & Chater, 2014; Misyak et al., 2014]. Crit-

ically, the virtual bargaining account explains how an intuitive theory of how others think and act may support complex social interactions without multiple levels of recursion. It is possible that the same inferences in our model could be implemented via virtual bargaining, and our results are therefore only evidence that communicative objects emerge from Theory of Mind and intuitive physics at a computational level, without a strong commitment to our algorithmic implementation [Marr, 1982].

9. Concluding remarks

Research into human cognition has long sought to understand the power and limitations of intuitive psychology and intuitive physics: how these systems develop, support high-level cognition, and contribute to uniquely human behavior. Our work shows how people seamlessly infuse and extract social meaning from objects, integrating the two intuitive theories to go beyond separate understandings of the physical and social world. And yet, the types of meaning that we give objects is even broader than what we show here—a metal band can signify a lifelong vow, a chiselled stone can commemorate a lost loved one, and a menorah can reveal one’s metaphysical beliefs. Our work is a critical step towards the quantitative study of *social reality* [Searle, 1995]—the world we perceive as bearing social meaning that underlies our notions of abstract social constructs, like property, money, marriage, and government.

10. Acknowledgments

We thank Hyowon Gweon and Zenna Tavares for useful comments on a previous version of this manuscript. This work was supported by a Google Research Award to JJE.

740 11. Author Contribution

741 MLB and JJE conceptualized the model and designed the studies. MLB
742 implemented the models, collected the data and analyzed the results under the
743 supervision of JJE. MLB and JJE wrote the manuscript.

744 References

- 745 Back, E., & Apperly, I. A. (2010). Two sources of evidence on the non-
746 automaticity of true and false belief ascription. *Cognition*, 115, 54–70.
- 747 Baker, C. L., Jara-Ettinger, J., Saxe, R., & Tenenbaum, J. B. (2017). Rational
748 quantitative attribution of beliefs, desires and percepts in human mentalizing.
749 *Nature Human Behaviour*, 1, 1–10.
- 750 Baker, C. L., Saxe, R., & Tenenbaum, J. B. (2009). Action understanding as
751 inverse planning. *Cognition*, 113, 329–349.
- 752 Battaglia, P. W., Hamrick, J. B., & Tenenbaum, J. B. (2013). Simulation as an
753 engine of physical scene understanding. *Proceedings of the National Academy
754 of Sciences*, 110, 18327–18332.
- 755 Bennett-Pierre, G., Asaba, M., & Gweon, H. (2018). Preschoolers consider
756 expected task difficulty to decide what to do and whom to help. In *Cogsci*.
- 757 Birch, S. A., & Bloom, P. (2007). The curse of knowledge in reasoning about
758 false beliefs. *Psychological Science*, 18, 382–386.
- 759 Bridgers, S., Jara-Ettinger, J., & Gweon, H. (2020). Young children consider the
760 expected utility of others’ learning to decide what to teach. *Nature Human
761 Behaviour*, 4, 144–152.

- 762 Conway, B. R., Ratnasingam, S., Jara-Ettinger, J., Futrell, R., & Gibson, E.
763 (2020). Communication efficiency of color naming across languages provides
764 a new framework for the evolution of color terms. *Cognition*, *195*, 104086.
- 765 Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in
766 language games. *Science*, *336*, 998–998.
- 767 Gergely, G., & Csibra, G. (2003). Teleological reasoning in infancy: The naïve
768 theory of rational action. *Trends in cognitive sciences*, *7*, 287–292.
- 769 Gibson, E., Futrell, R., Jara-Ettinger, J., Mahowald, K., Bergen, L., Rat-
770 nasingam, S., Gibson, M., Piantadosi, S. T., & Conway, B. R. (2017).
771 Color naming across languages reflects color use. *Proceedings of the National*
772 *Academy of Sciences*, *114*, 10785–10790.
- 773 Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as
774 probabilistic inference. *Trends in Cognitive Sciences*, *20*, 818–829.
- 775 Gopnik, A., Meltzoff, A. N., & Bryant, P. (1997). Words, thoughts, and theories.
- 776 Grice, H. P. (1957). Meaning. *The Philosophical Review*, *66*, 377–388.
- 777 Grice, H. P. (1975). Logic and conversation. In *Speech acts* (pp. 41–58). Brill.
- 778 Griffiths, T. L., Lieder, F., & Goodman, N. D. (2015). Rational use of cognitive
779 resources: Levels of analysis between the computational and the algorithmic.
780 *Topics in Cognitive Science*, *7*, 217–229.
- 781 Grosse, G., Moll, H., & Tomasello, M. (2010). 21-month-olds understand the
782 cooperative logic of requests. *Journal of Pragmatics*, *42*, 3377–3383.
- 783 Gweon, H., Asaba, M., & Bennett-Pierre, G. (2017). Reverse-engineering the
784 process: Adults’ and preschoolers’ ability to infer the difficulty of novel tasks.
785 In *CogSci*.

- 786 Hanna, J. E., & Tanenhaus, M. K. (2004). Pragmatic effects on reference res-
787 olution in a collaborative task: Evidence from eye movements. *Cognitive*
788 *Science*, *28*, 105–115.
- 789 Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the
790 world? *Behavioral and Brain Sciences*, *33*, 61–83.
- 791 Ho, M. K., Littman, M., MacGlashan, J., Cushman, F., & Austerweil, J. L.
792 (2016). Showing versus doing: Teaching by demonstration. *Advances in*
793 *Neural Information Processing Systems*, (pp. 3027–3035).
- 794 Hurwitz, E., Brady, T., & Schachner, A. (2019). Detecting social transmission
795 in the design of artifacts via inverse planning.
- 796 Jara-Ettinger, J., Floyd, S., Huey, H., Tenenbaum, J. B., & Schulz, L. E.
797 (2019a). Social pragmatics: Preschoolers rely on commonsense psychology
798 to resolve referential underspecification. *Child Development*, *91*, 1135–1149.
- 799 Jara-Ettinger, J., Gweon, H., Schulz, L. E., & Tenenbaum, J. B. (2016). The
800 naïve utility calculus: Computational principles underlying commonsense psy-
801 chology. *Trends in Cognitive Sciences*, *20*, 589–604.
- 802 Jara-Ettinger, J., Piantadosi, S., Spelke, E. S., Levy, R., & Gibson, E. (2017).
803 Mastery of the logic of natural numbers is not the result of mastery of count-
804 ing: Evidence from late counters. *Developmental Science*, *20*, e12459.
- 805 Jara-Ettinger, J., Schulz, L., & Tenenbaum, J. (2019b). The naïve utility cal-
806 culus as a unified, quantitative framework for action understanding. URL:
807 psyarxiv.com/e8xsv. doi:10.31234/osf.io/e8xsv.
- 808 Jara-Ettinger, J., Tenenbaum, J. B., & Schulz, L. E. (2015). Not so innocent:
809 Toddlers’ inferences about costs and culpability. *Psychological Science*, *26*,
810 633–640.

- 811 Jern, A., Lucas, C. G., & Kemp, C. (2017). People learn other people’s prefer-
812 ences through inverse decision-making. *Cognition*, 168, 46–64.
- 813 Keil, F. C., & Newman, G. E. (2015). Order, order everywhere, and only an
814 agent to think: The cognitive compulsion to infer intentional agents. *Mind &
815 Language*, 30, 117–139.
- 816 Keysar, B., Barr, D. J., Balin, J. A., & Brauner, J. S. (2000). Taking per-
817 spective in conversation: The role of mutual knowledge in comprehension.
818 *Psychological Science*, 11, 32–38.
- 819 Kleiman-Weiner, M., Saxe, R., & Tenenbaum, J. B. (2017). Learning a com-
820 mon-sense moral theory. *Cognition*, 167, 107–123.
- 821 Kool, W., & Botvinick, M. (2018). Mental labour. *Nature Human Behaviour*,
822 2, 899–908.
- 823 Lieder, F., & Griffiths, T. L. (2019). Resource-rational analysis: understand-
824 ing human cognition as the optimal use of limited computational resources.
825 *Behavioral and Brain Sciences*, 43, 1–85.
- 826 Liu, S., Ullman, T. D., Tenenbaum, J. B., & Spelke, E. S. (2017). Ten-month-
827 old infants infer the value of goals from the costs of actions. *Science*, 358,
828 1038–1041.
- 829 Lucas, C. G., Griffiths, T. L., Xu, F., Fawcett, C., Gopnik, A., Kushnir, T.,
830 Markson, L., & Hu, J. (2014). The child as econometrician: A rational model
831 of preference understanding in children. *PLOS ONE*, 9, e92160.
- 832 Ma, L., & Xu, F. (2013). Preverbal infants infer intentional agents from the
833 perception of regularity. *Developmental psychology*, 49, 1330–1337.
- 834 Marr, D. (1982). *Vision: A Computational Investigation into the Human Rep-*
835 *resentation and Processing of Visual Information*. MIT Press.

- 836 McDermott, J. H., Schultz, A. F., Undurraga, E. A., & Godoy, R. A. (2016).
837 Indifference to dissonance in native amazonians reveals cultural variation in
838 music perception. *Nature*, *535*, 547–550.
- 839 Misyak, J. B., & Chater, N. (2014). Virtual bargaining: A theory of social
840 decision-making. *Philosophical Transactions of the Royal Society B: Biological*
841 *Sciences*, *369*, 20130487. doi:10.1098/rstb.2013.0487.
- 842 Misyak, J. B., Melkonyan, T., Zeitoun, H., & Chater, N. (2014). Unwritten
843 rules: virtual bargaining underpins social interaction, culture, and society.
844 *Trends in Cognitive Sciences*, *18*, 512–519.
- 845 Misyak, J. B., Noguchi, T., & Chater, N. (2016). Instantaneous conventions:
846 The emergence of flexible communicative signals. *Psychological Science*, *27*,
847 1550–1561.
- 848 Newman, G. E., Keil, F. C., Kuhlmeier, V. A., & Wynn, K. (2010). Early under-
849 standings of the link between agents and order. *Proceedings of the National*
850 *Academy of Sciences*, *107*, 17140–17145.
- 851 O’Grady, C., Kliesch, C., Smith, K., & Scott-Phillips, T. C. (2015). The ease and
852 extent of recursive mindreading, across implicit and explicit tasks. *Evolution*
853 *and Human Behavior*, *36*, 313–322.
- 854 Piantadosi, S. T., Jara-Ettinger, J., & Gibson, E. (2014). Children’s learning
855 of number words in an indigenous farming-foraging group. *Developmental*
856 *Science*, *17*, 553–563.
- 857 Saxe, R., Tenenbaum, J., & Carey, S. (2005). Secret agents: Inferences about
858 hidden causes by 10- and 12-month-old infants. *Psychological Science*, *16*,
859 995–1001.

860 Schachner, A., Brady, T., Oro, K., & Lee, M. (2018). Intuitive archeology:
861 Detecting social transmission in the design of artifacts.

862 Searle, J. R. (1995). *The construction of social reality*. Simon and Schuster.

863 Shafto, P., Goodman, N. D., & Griffiths, T. L. (2014). A rational account of
864 pedagogical reasoning: Teaching by, and learning from, examples. *Cognitive*
865 *Psychology*, 71, 55–89.

866 Shenhav, A., Musslick, S., Lieder, F., Kool, W., Griffiths, T. L., Cohen, J. D.,
867 & Botvinick, M. M. (2017). Toward a rational and mechanistic account of
868 mental effort. *Annual Review of Neuroscience*, 40, 99–124.

869 Ullman, T., Baker, C., Macindoe, O., Evans, O., Goodman, N., & Tenenbaum,
870 J. B. (2009). Help or hinder: Bayesian models of social goal inference. In
871 *Advances in Neural Information Processing Systems* (pp. 1874–1882).

872 Wellman, H. M. (2014). *Making minds: How theory of mind develops*. Oxford
873 University Press.

874 Yildirim, I., Saeed, B., Bennett-Pierre, G., Gerstenberg, T., Tenenbaum, J.,
875 & Gweon, H. (2019). Explaining intuitive difficulty judgments by modeling
876 physical effort and risk.