UC Merced
Proceedings of the Annual Meeting of the Cognitive Science Society

Title
Entropy, order and agency: The cognitive basis of the link between agents and order

Permalink
https://escholarship.org/uc/item/74c6105q

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 40(0)

Authors
Schachner, Adena
Kim, Min-Ju

Publication Date
2018

eScholarship.org Powered by the California Digital Library
University of California

Entropy, order and agency: The cognitive basis of the link between agents and order

Adena Schachner (schachner@ucsd.edu), Min-Ju Kim (minjukim@ucsd.edu)
University of California, San Diego, Department of Psychology
9500 Gilman Drive, La Jolla, CA, 92093-0109

Abstract

People often believe that orderly structures were created by agents. We examine the cognitive basis of this tendency, asking if learned associations or causal reasoning drives us to link order with agents. Causal reasoning predicts that knowledge of an alternative physical-mechanical cause should 'explain away' orderliness, weakening the link with agents. In a preregistered experiment, we manipulated the context to provide (or not provide) a physical-mechanical explanation for orderly outcomes, and participants judged if an object or agent had been present. We compared outcomes differing in (a)levels of orderliness and (b)whether context provided an alternative explanation. We found that environmental context 'explained away' orderliness, such that participants observing order inferred agency only when there was no alternative explanation. The link between order and agents is moderated by causal reasoning, and is malleable: It can be weakened by understanding alternative causal mechanisms by which order could arise.

Keywords: causal reasoning, inference, order, agency, animacy, music, event perception, social cognition, religion

Introduction

"The visible order of the universe proclaims a supreme intelligence." - Jean-Jacque Rousseau, *Emile,* 1762

Our human environments are full of orderly, complex structures: We hear orderly sounds, like music (Cross, 2003), create objects with orderly structure, like tools (Vaesen, 2012); and see order in nature (Paley, 1802; Dawkins, 1986). How do we perceive and reason about this 'orderliness', pattern or regularity?

People readily invoke intentional agents – both physical and supernatural – to explain the patterns and regularities they observe (Dawkins, 2006; Swinburne, 1968), and recent work has shown that this tendency arises early in life (Baer, Tam, Nancekivell & Friedman, 2013; Friedman, 2001; Ma, Berthiaume, Hoch, & Xu, 2013; Ma & Xu, 2013; Newman, Keil, Kuhlmeier, & Wynn, 2010). In effect, we appear to link orderly outcomes with intentional, animate agents – whether these agents be human, animal, or supernatural. This link has broad consequences: For example, it explains the intuitive nature of creationist beliefs, as these involve positing a supernatural agent to explain order in nature (Barrett, 2007; Bering, 2011; Kelemen, 2004).

What underlying cognitive processes cause us to link order with agents? This goes beyond the question of how we *detect* randomness versus regularity (Reimers, Donkin & LePelley, 2018; Griffiths, Daniels, Austerweil & Tenenbaum, 2018): Given that one *has* detected regularity, what cognitive processes lead us to believe it was produced by an agent, versus an inanimate force (i.e. a non-random but inanimate process)? Two broad types of cognitive accounts could underlie this tendency.

Associations with perceptual features account

Firstly, the tendency to invoke agents could be driven by a network of associations between agents and the perceptual features of orderly stimuli (Keil & Newman, 2015). Such associations could be learned from experience: We often observe other people organizing objects and producing orderly things like music and art. We do not often observe inanimate objects or forces producing orderly outcomes. With experience, then, we should come to associate people with the orderly outcomes we have observed. We could then generalize from 'people' to 'agents' as a class (Gelman & Opfer, 2007); and from the specific orderly outcomes we have seen to other perceptually similar outcomes.

This associative account has penetrated into a diverse set of literatures, and has been used to explain a wide range of phenomena. In music cognition, an associative account has been used to explain why musical patterns evoke agency in an automatic or obligatory way (Launay, 2015). In religious cognition, an associative account has been used to explain why order in nature evokes creationist beliefs, even after learning about evolution (Barrett, 2007; Bering, 2011; Cornelius & Lopez-Mobilia, 2013; Kelemen, 2004).

This associative account is consistent with both adult and developmental findings, e.g. that year-old infants expect only agents to produce orderly outcomes like neatly stacked blocks (Newman et al., 2010; Ma & Xu, 2013). Infants routinely observe people creating order (organizing toys, laundry), and can generalize between perceptually similar stimuli and from people to other kinds of agents (Gelman & Opfer, 2007). In this way, infants could learn an association between order and agents within the first year. Associative learning thus provides a plausible, intuitive, and influential account of the order-agent link.

Causal reasoning (rational inference) account

In contrast to this associative account, the link between order and agents may instead be a product of causal reasoning, as an inference to the best explanation (Keil & Newman, 2015; Lipton, 2004; Ma & Xu, 2013; Tenenbaum, Griffiths & Kemp, 2006). Under this account, to explain observed data (i.e. an orderly outcome), people consider multiple hypotheses about its cause (perhaps it was generated by a person; by an animal; by gravity; by wind; etc.). They then choose the hypothesis that is both plausible a priori and provides a good explanation of the observed

data (e.g. 'a person is the most probable cause'). This type of inferential reasoning occurs in multiple related domains, including physical-mechanical reasoning, social reasoning, and mental state inference (Baker et al., 2009; Battaglia, Hamrick & Tenenbaum, 2013; Teglas et al., 2011).

One signature of this type of inferential reasoning is 'explaining away' – the idea that a plausible alternative explanation should weaken evidence for other explanations. This phenomenon serves as evidence that the underlying cognitive representation is rich and structured (Gopnik & Sobel, 2000; Pearl, 2000; Tenenbaum et al., 2006). If structured causal reasoning drives the link, then knowledge of a plausible physical-mechanical cause should 'explain away' the orderliness, reducing the likelihood of inferring an agent as cause – even for the very same perceptual stimulus.

The current study

The current study aims to tease apart these two accounts, asking whether the link between order and agents is driven by associations, causal reasoning, or both processes simultaneously. This bears on the malleability of the link between agency and order (as noted by Meng, Griffiths & Xu, 2017): Can we weaken or break the link between order and agents, by understanding alternative causal mechanisms by which order could arise? If causal reasoning accounts for the link, we would expect this to be the case; if a learned association accounts for the link, we would expect it to be robust even in the face of conflicting high-level causal knowledge.

To test these ideas, we used methods modeled on previous work on order and agency (Newman et al., 2010; Ma & Xu, 2013): We presented participants with a particular scene, occluded the scene from view, and then showed that either an orderly or disorderly outcome was generated. We then asked participants to judge whether an inanimate object or animate character had been behind the occluder (i.e., causing the outcome). To tease apart the two theories, we manipulated whether or not there was a plausible physical-mechanical causal explanation for each outcome, and asked if this alternative explanation could 'explain away' the link between order and agents.

In particular, the scene took the form of a staircase-like xylophone, on which striking each stair produced a different musical note (see Figure 1). For some participants, the xylophone's bars were positioned such that a ball simply rolling down the staircase would produce a descending scale – a highly ordered sequence of sounds (as verified in a pilot study; see below). Previous work has shown that orderly sounds cue agency, just as orderly visual stimuli do (Ma et al., 2013). Thus, if associations with perceptual features drive the link between order and agency, orderly sounds (but not disorderly sounds) should lead observers to invoke an agent, regardless of their context. In contrast, if causal reasoning accounts for the tendency to invoke agents, then a context that provides an alternative physical-mechanical causal explanation should 'explain away' the orderliness, weakening evidence of the agents' involvement. This result would provide evidence that causal reasoning can modulate the link between order and agency.

We also aimed to test whether the link between order and agency is driven by causal reasoning alone, or whether perceptual features of orderly stimuli have an additional effect, above and beyond the effect of causal reasoning. If both of these factors contribute to the link between agents and order, we would expect that a stimulus that was both perceived as orderly and also causally implausible without an agent would be more strongly linked with agents than a stimulus that was causally implausible but not orderly. In contrast, if causal reasoning alone accounts for the link between order and agents, then the percept of orderliness should not result in increased agency attributions above and beyond the effect of causal reasoning.

We test this prediction by comparing three different outcomes, including one that is orderly (descending scale), one that is causally implausible without an agent (scrambled tones), and one that is *both* causally implausible and also highly orderly (ascending scale). If both causal reasoning and orderly perceptual features drive the link with agency, we should find that the ascending scale is more strongly linked with agents than are the scrambled tones. In contrast, if the ascending scale is linked with agency to an equal or lesser extent as the scrambled tones, this would suggest that only causal inference plays a role, and that the perceptual features of orderly stimuli do not play any additional role in connecting orderly outcomes with animate agents.

Method

Design

The current study design, sample size, and analyses were pre-registered before data collection (www.aspredicted.org). The experiment followed a 2 (environmental context) x 3 (sound outcome) within-subject design (see Figure 1). Each participant completed six trials, composed of two blocks (by environmental context: bars-descending xylophone, bars-scrambled xylophone) of three test trials each (sound outcome: descending scale, scrambled tones, ascending scale). The order of the blocks, and the order of test trials within a block, were counterbalanced between subjects.

On each trial, the scene was covered with an animated curtain; one of the sound outcomes occurred; and participants were asked to judge whether the inanimate ball or animate agent (named 'Fred') had been present behind the curtain, on a -2 to +2 scale (Definitely the ball, Probably the ball, Could be either one, Probably Fred, Definitely Fred).

We manipulated the environmental context such that the context either did or did not provide a physical-mechanical explanation for the orderly sounds. In the bars-descending environmental context, the xylophone's bars were positioned in a descending sequence such that a ball simply rolling downward would produce an orderly descending scale. In contrast, in the bars-scrambled environmental context, the bars were positioned such that a ball rolling

down could not produce an orderly sound sequence, but would instead play a scrambled sequence of tones.

We tested three sound outcomes: A descending scale, scrambled tones, and an ascending scale (see Stimuli). These stimuli were chosen to manipulate whether the sounds were orderly; causally implausible without an agent (given the environmental context); or both.

To verify our manipulation of the sounds' orderliness, we asked a separate sample of adult participants to judge which of each pair of sound sequences was more orderly, or if they were equally orderly (N=40, none of whom participated in the main experiment). As expected, participants judged the descending and ascending tones as more orderly than the scrambled tones (all p's<0.01, binomial tests), and rated the descending and ascending scales as equally orderly.

Participants

Sixty participants were recruited from Amazon Mechanical Turk website and tested online (Age: *M* = 35.07, Range = 19 to 61; 37 males). Participants were required to have had at least 93% of their previous work on the MTurk website judged as acceptable; this criterion was determined a-priori based on criteria used in the literature (Berinsky, Huber & Lenz, 2012; Goodman, Cryder & Cheema, 2012; Mason & Suri, 2011). All participants gave informed consent, and were paid $0.50 for approximately 5 minutes of their time. 34 additional participants were tested, but excluded following our preregistered exclusion criteria, for leaving questions blank (2 participants), failing to identify sounds in the initial sound check (1), having done a previous similar experiment (6), or failing the memory check questions (24).

Stimuli

Animated videos and accompanying sounds were constructed using Apple Keynote '09, Apple GarageBand '09, and Apple QuickTime Pro 7 software.

We created two videos to demonstrate how the xylophone produced sound, one for each environmental context (the bars-scrambled xylophone and the bars-descending xylophone). Both videos showed a similar scene: A staircase with colored bars, a pipe at the top of the hill, and small 'digital display' on the top left corner of the video, used to scaffold pitch memory (this displayed a visual record of the notes that had been played, as numbers; see Figure 1). In both videos, a mallet entered the scene and played a sequence of sounds before exiting the scene. The mallet moved in the same path and struck the same positions in both videos; but this resulted in a different sequence of pitches across the two videos. For the bars-descending xylophone, it played an ascending scale (1-2-3-4-5), two disorderly five-note sequences (5-2-4-3-1; then 1-3-4-2-5), and a descending scale (5-4-3-2-1). For the bars-scrambled xylophone, it played a disorderly five-note sequence (5-2-4-3-1), an ascending scale (1-2-3-4-5), a descending scale (5-4-3-2-1), and another disorderly five-note sequence (1-3-4-2-5). The sequences were designed so that the sounds played at test would each be equally novel relative to the sounds participants had previously heard in the experiment. These videos thus included each of the test sequences one time, along with an additional five-note sequence.

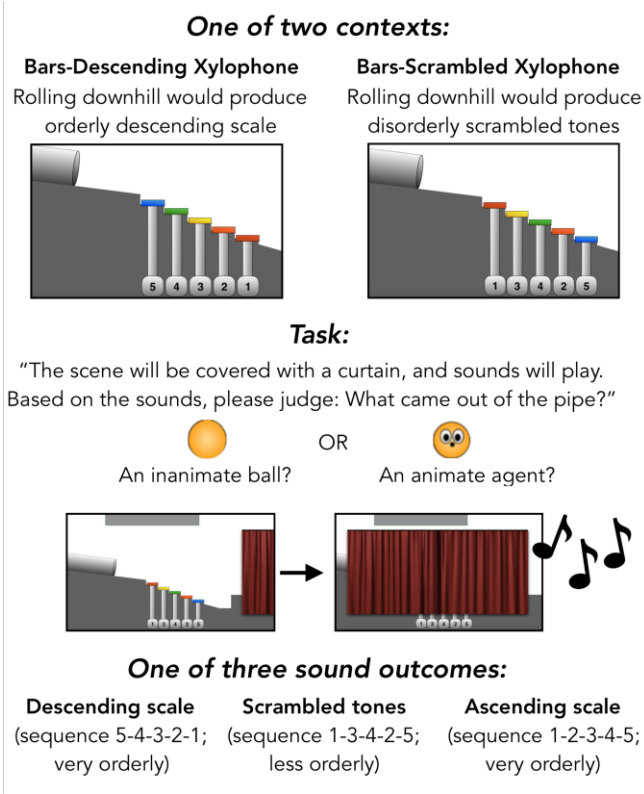

Figure 1: Method. We manipulated the environmental context, such that it either did or did not provide a physical-mechanical explanation for the orderly sounds (bars-descending, bars-scrambled xylophones). On each trial, the scene was covered with a curtain; one of three sound outcomes occurred; and participants judged whether the inanimate ball or animate agent had been present behind the curtain. The causal reasoning and perceptual features accounts make different predictions about when each sound outcome will be linked with the agent.

The sounds indirectly showed that for the bars-descending xylophone, the stairs were arranged to play the following pitches from left to right: G, F, E, D, C (or 5, 4, 3, 2, 1, with fundamental frequencies of 778, 702, 647, 569, and 511 Hz), For the bars-scrambled xylophone, the video indirectly showed that the bars were arranged to play G, D, F, E, C (5, 2, 4, 3, 1).

During each trial, an additional video was shown in which the xylophone and pipe was covered by an animated curtain, and one of three sound outcomes played: a descending scale (pitch sequence 5-4-3-2-1), an ascending scale (1-2-3-4-5), and a scrambled sequence of the same tones (1-3-4-2-5). All tone sequences were of equal duration.

Procedure

Participants were first asked to report age and gender, before clicking a link to submit these answers and move to a second page. Throughout the experiment, participants could not return to previous pages after moving forward to a new page, preventing revision of answers based on subsequent information. On a second page, participants were asked to identify two simple sounds, to ensure that all participants' computers could play sounds.

On the third page, participants began the main task. They were instructed to watch a video of a "musical staircase", which "demonstrates how it works". Participants then viewed one xylophone demonstration video, and were asked to describe what happened. To check that participants remembered how the notes mapped to the bars of the xylophone, participants were asked two multiple-choice 'memory check' questions ("Which part of the staircase plays the highest note?"; "Which part of the staircase plays the lowest note?"). Participants then moved to a fourth page.

On the fourth page, participants were asked to notice the pipe on the left-hand side of the scene, and were told that one of two things could come out of the pipe: An inanimate ball, or a cartoon character. To provide evidence of the ball's inanimacy and the character's animacy, brief videos were shown (of the character moving in a self-propelled way, or the ball rolling with gravity), and they were identified in ways consistent with inanimacy and animacy: As a billiard ball, or as cartoon character with a name (Fred).

Participants then completed three test trials, each presented on a separate page. On each test trial, participants watched a video where the xylophone was occluded, and a brief sound sequence played (either the descending scale, scrambled tones, or ascending scale). Participants were then asked to judge: "What came out of the pipe? Was it the ball? Or Was it Fred?", on a -2 to +2 scale (Definitely the ball, Probably the ball, Could be either one, Probably Fred, Definitely Fred), and asked to explain their answer ("What made you think that?").

After completing three test trials, participants read similar instructions for the second block (the other environmental context condition), with a new but similar scene introduced as "a different kind of musical staircase". They then viewed the other of the two xylophone demonstration videos, and completed the same questions and parallel three test trials as in the first block, but with the videos showing the new environmental context. Participants lastly were asked if the ideas in the questions were clear, and to guess what the experiment was about, before submitting their answers.

Results

We examined participants' judgments of whether an agent or object had been present, asking whether environmental context affected participants' judgments (i.e. whether context 'explained away' the link between orderly stimuli and the agent). We thus conducted a 2 (environmental

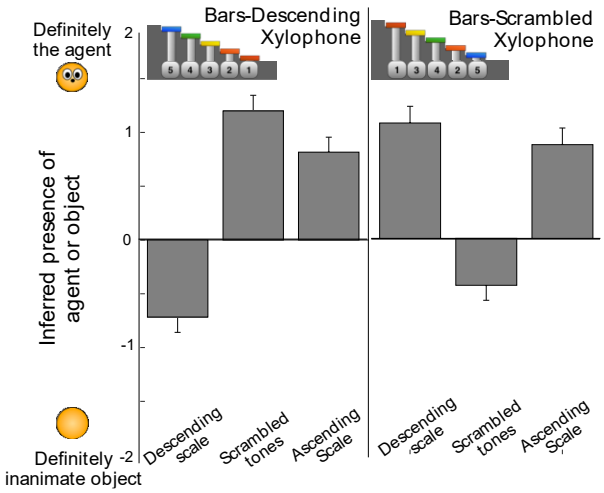

Figure 2: Results. Participant's judgments of whether an inanimate object (a ball) or an animate agent had been behind the curtain when the sounds occurred, for two environmental contexts (bars-descending, bars-scrambled xylophones), and three sound outcomes (orderly descending scale; disorderly scrambled tones; orderly ascending scale). Causal reasoning (not perceptual features) predicts and explains when each stimulus was linked with the agent.

context: bars-descending vs. bars-scrambled xylophone) by 3 (sound outcome: descending scale vs. scrambled tones vs. ascending scale) repeated measures ANOVA with the rating of likelihood that the agent vs. object was present as the dependent measure (-2 to +2 probability scale). As predicted, we found a significant interaction of sound outcome with environmental context, $F(2,118)=46.85$, $p<0.001$, $\eta_p^2= 0.44$ (also a main effect of sound outcome, $F(2, 118) = 8.16$, $p < 0.001$, $\eta_p^2= 0.12$; and no main effect of the environmental context, $F(1, 59)=0.72$, $p=0.40$, $\eta_p^2=0.01$). Thus, participants' judgment of whether the agent or inanimate object had been present differed not only based on the orderliness of the sound itself; but also depending on the context in which each sound outcome occurred.

Participants' agency judgments differed by context for the same sound outcome, and this only occurred when these sounds differed in causal plausibility. In the case of the descending scale, they judged it more likely that the animate agent had been present in the context of the bars-scrambled xylophone (M=1.07, SEM =0.17), and more likely that the inanimate ball had been present in the context of the bars-descending xylophone (M=-0.72, SEM=0.14), Mean difference=-1.78, d=-1.03; $t(59)$ = -7.99, $p<0.001$). In the case of the scrambled tones, the participants judged it more likely that the animate agent had been present in the context of the bars-descending xylophone (M=1.2, SEM=0.14), and more likely that the inanimate ball had been present in the context of the bars-scrambled xylophone, where a ball rolling with gravity would result in these tones, M=-0.43, SEM=0.15, Mean diff.=1.63, d=0.91; $t(59)$=-7.05, $p< 0.001$.

In contrast, animacy ratings on the ascending tones did not differ across the two environmental contexts, as predicted by the causal reasoning account, since they were causally implausible in both cases (bars-descending xylophone, ascending scale: M=0.82, SEM=0.17; bars-scrambled xylophone, ascending scale: M=0.87, SEM =0.18; *Mean difference*=-0.05, d =-0.03; $t(59)$=-0.25, p=0.80).

Overall, in the context of the bars-descending xylophone, participants judged it more likely that the agent (not the ball) had been present after hearing the scrambled tones (M=1.2, $t(59)$=8.44, p<0.001) or ascending scale (M=0.82, $t(59)$=4.84, p<0.001); and judged it more likely that the ball (not the agent) had been present after hearing the orderly descending scale (M=-0.72, $t(59)$=-5.02, p< 0.001).

In contrast, in the context of the bars-scrambled xylophone, the pattern was reversed: Participants judged it more likely that the agent had been present after hearing the ascending scale (M=0.87, $t(59)$=4.77, p<0.001) or descending scale (M=1.07, $t(59)$=6.23, p<0.001); and judged it more likely that the ball had been present after hearing the disorderly scrambled tones (M=-0.43, $t(59)$=-2.91, p=0.005). This reversal shows that causal reasoning can modulate the link between order and agents.

We next examined if whether the outcome was perceptually orderly or disorderly had an additional effect on the agency judgments (in addition to the effect of causal plausibility). The context of the bars-descending xylophone provided a crucial test case: If orderly perceptual features have an additional effect, then the ascending scale should be more strongly linked with agents than the scrambled tones. In contrast, we found that the two conditions were not significantly different, and that judgments in the ascending condition were numerically lower (i.e., less linked with the agent) than the scrambled tones condition (Ascending: M=0.82, SEM=0.17; Scrambled tones: M=1.2, SEM=0.14, $t(59)$=-1.69, p=0.10).

Discussion

Humans tend to link orderliness with animate agents, and this idea holds broad implications for our representations of orderly stimuli in music, art, and nature. We asked whether this link was malleable, and tested two accounts of the link's cognitive basis, asking whether causal reasoning or learned associations with perceptual features leads people to invoke agents when observing order. We found that an alternative, physical-mechanical explanation 'explained away' orderliness, such that it no longer led people to infer the presence of an agent. When the very same orderly outcome occurred in a context that was similar, but did not provide an alternative explanation, participants readily inferred the presence of an agent from the same orderly stimulus.

The percept of orderliness did not lead to increased agency attributions beyond the effect of causal reasoning: We compared outcomes that were orderly (descending scale), causally implausible without an agent (scrambled tones), and *both* causally implausible and also highly orderly (ascending scale). If both causal reasoning and orderly perceptual features drive the link with agency, we should find that the ascending scale is more strongly linked with agents than are the scrambled tones. In contrast, we found that this outcome was judged numerically *less* linked to the agent than an outcome that was less orderly, but similarly causally implausible (scrambled tones).

This shows for the first time that the link between order and agents is malleable, and can be eliminated by fully understanding an alternative causal mechanism by which order could arise. The link between order and agency is not obligatory, but instead appears to be a rational inference, the product of rich mental theories of how the physical and social world works (Battaglia et al., 2013; Gelman & Opfer, 2007; Gopnik, 2012; Tenenbaum et al., 2006).

This has broad implications for how we reason about the orderly, complex structures in our environment: When we observe orderly sounds (like music), objects with orderly structure (like tools), or order in nature, this may not 'cue' agent representations in a direct, associative way. Instead, we may make reasoned inferences about what we believe to be the most likely cause – leading us to infer the involvement of an agent when we have no better explanation.

The causal reasoning account thus predicts and explains over-attribution of agency in some contexts: Agent attribution should occur when the observer wrongly *believes* there is no alternative explanation, e.g. due to not knowing or fully understanding the underlying physical-mechanical mechanism. Causal reasoning is thus consistent with the idea that people over-attribute of agency to explain order in nature (i.e., as a result of a failure to fully understand natural selection as an alternative explanation).

In the situation tested here, there was little ambiguity about the cause of the outcomes – the context was known, and two potential causes were specified as hypotheses. In more naturalistic situations, there may be more ambiguity about the cause of orderly or disorderly stimuli. Do our findings generalize to these cases? This is an open empirical question; however, the framework we apply here – that of rational inference and causal reasoning – applies equally well to these more uncertain cases. When we have more uncertainty about the cause, rather than weighing the likelihood of each outcome under the single cause, people could instead weight the likelihood of each outcome under a variety of causes, and then average over those (weighted by each cause's prior probability; Griffiths & Tenenbaum, 2009). Using such a strategy of marginalizing over possible causes would allow for the same kind of reasoning process as we demonstrate here in the case of more uncertain causes.

Overall, we find that causal reasoning drives the link between order and agency, such that the link is malleable and orderliness can be 'explained away' by an alternative physical-mechanical explanation. This finding provides a new framework for predicting when and whether orderly stimuli will be linked with agency, and lays a foundation for better understanding of high-level representations of a wide

range of orderly stimuli, including music, artifacts like tools and visual art, and order in nature.

References

Baer, C., Tam, C., Nancekivell, S., & Friedman, O. (2013). Order and disorder in children's ownership judgments. Poster presented at the meeting of the *Society for Research in Child Development*, Seattle, WA.

Baker, C.L., Saxe, R., & Tenenbaum, J.B. (2009). Action understanding as inverse planning. *Cognition*, 113, 329-349.

Barrett, J. L. (2007). Cognitive science of religion: What is it and why is it? *Religion Compass, 1,* 768–786. doi:10.1111/j.1749-8171.2007.00042.x

Battaglia, P. W., Hamrick, J. B., & Tenenbaum, J. B. (2013). Simulation as an engine of physical scene understanding. *Proceedings of the National Academy of Sciences*, *110*(45), 18327–18332. doi:10.1073/pnas.1306572110

Bering, J. (2011). *The Belief Instinct: The Psychology of Souls, Destiny, and the Meaning of Life.* New York: W.W. Norton & Co.

Berinsky, A. J., Huber, G. A., & Lenz, G. S. (2012). Evaluating Online Labor Markets for Experimental Research: Amazon.com's Mechanical Turk. *Political Analysis*, *20*(3), 351–368. doi:10.1093/polana/mpr057/-/DC1

Cornelius, C., & Lopez-Mobilia, G. (2013). Relations between anthropomorphism, predictability, and patterned outcomes. Poster presented at the Eighth Biennial meeting of the *Cognitive Development Society*, Memphis, TN.

Cross, I. (2003). Music as a biocultural phenomenon. *Annals of the New York Academy of Sciences*, *999*, 106–111.

Dawkins, R. (1986). *The Blind Watchmaker*. New York: W.W. Norton & Co.

Dawkins, R. (2006). *The God Delusion*. Random House.

Friedman, W. J. (2001). The development of an intuitive understanding of entropy. *Child Development*, *72*, 460-473.

Gelman, S. A., & Opfer, J. E. (2007). Development of the animate–inanimate distinction. *Blackwell Handbook of Childhood Cognitive Development*, 151–166.

Goodman, J. K., Cryder, C. E., & Cheema, A. (2012). Data Collection in a Flat World: The Strengths and Weaknesses of Mechanical Turk Samples. *Journal of Behavioral Decision Making*, *26*(3), 213-224. doi:10.1002/bdm.1753

Gopnik, A., & Sobel, D. M. (2000). Detecting blickets: How young children use information about novel causal powers in categorization and induction. *Child Development*, *71*(5), 1205-1222.

Gopnik, A. (2012). Scientific thinking in young children. *Science, 337,* 1623-1627.

Griffiths, T. L., Daniels, D., Austerweil, J. L., & Tenenbaum, J. B. (2018). Subjective randomness as statistical inference. *Cognitive Psychology*, *103*, 85–109.

Griffiths, T. L., & Tenenbaum, J. B. (2009). Theory-based causal induction. *Psychological Review*, *116*(4), 661–716. doi:10.1037/a0017201

Keil, F. C., & Newman, G. E. (2015). Order, order everywhere, and only an agent to think: The cognitive compulsion to infer intentional agents. *Mind and Language*, *30*(2), 117–139.

Kelemen, D. (2004). Are children "intuitive theists?" Reasoning about purpose and design in nature. *Psychological Science*, *15*(5), 295–301. doi:10.1111/j.0956-7976.2004.00672.x

Launay, J. (2015). Musical sounds, motor resonance, and detectable agency. *Empirical Musicology Review, 10* (1-2), 30-40.

Lipton, P. (2004). *Inference to the Best Explanation, 2nd edition.* Routledge Publishers.

Ma, L., & Xu, F. (2013). Preverbal infants infer intentional agents from the perception of regularity. *Developmental Psychology*, *49*(7), 1330–1337. doi:10.1037/a0029620

Ma, L., Berthiaume, V., Hoch, J., & Xu, F. (2013). Twelve-month-old infants infer intentional agents from the perception of auditory regularity. *Developmental Psychology*, *49*(7), 1330-1337.

Mason, W., & Suri, S. (2011). Conducting behavioral research on Amazon's Mechanical Turk. *Behavior Research Methods*, *44*(1), 1–23. doi:10.3758/s13428-011-0124-6

Newman, G. E., Keil, F. C., Kuhlmeier, V. A., & Wynn, K. (2010). Early understandings of the link between agents and order. *Proceedings of the National Academy of Sciences*, *107*(40), 17140–17145.

Pearl, J. (2000). *Causality: Models, reasoning, and inference*. New York: Cambridge University Press.

Reimers, S., Donkin, C., & Le Pelley, M. E. (2018). Perceptions of randomness in binary sequences: Normative, heuristic, or both? *Cognition*, *172*, 11–25.

Saxe, R. R., Tzelnic, T., & Carey, S. (2007). Knowing who dunnit: Infants identify the causal agent in an unseen causal interaction. *Developmental Psychology*, *43*(1), 149–158.

Saxe, R., Tenenbaum, J. B., & Carey, S. (2005). Secret agents: Inferences about hidden causes by 10-and 12-month-old infants. *Psychological Science*, *16*(12), 995.

Swinburne, R. G. (1968). The argument from design. *Philosophy*, *43*(165), 199-212.

Teglas, E., Vul, E., Girotto, V., Gonzalez, M., Tenenbaum, J. B., & Bonatti, L. L. (2011). Pure reasoning in 12-month-old infants as probabilistic inference. *Science, 332*(6033), 1054–1059.

Tenenbaum, J. B., Griffiths, T., & Kemp, C. (2006). Theory-based Bayesian models of inductive learning and reasoning. *Trends in Cognitive Sciences*, *10*(7), 309–318.

Vaesen, K. (2012). The cognitive bases of human tool use. *Behavioral and Brain Sciences*, *35*(4), 1–16. doi:10.1017/S0140525X11001452