

8-Month-Old Infants Spontaneously Learn and Generalize Hierarchical Rules

Denise M. Werchan, Anne G. E. Collins, Michael J. Frank, and Dima Amso

Brown University

Psychological Science

1–11

© The Author(s) 2015

Reprints and permissions:

sagepub.com/journalsPermissions.nav

DOI: 10.1177/0956797615571442

pss.sagepub.com



Abstract

The ability to extract hierarchically organized rule structures from noisy environments is critical to human cognitive, social, and emotional intelligence. Adults spontaneously create hierarchical rule structures of this sort. In the present research, we conducted two experiments to examine the previously unknown developmental origins of this hallmark skill. In Experiment 1, we exploited a visual paradigm previously shown to elicit incidental hierarchical rule learning in adults. In Experiment 2, we used the same learning structure to examine whether these hierarchical-rule-learning mechanisms are domain general and can help infants learn spoken object-label mappings across different speaker contexts. In both experiments, we found that 8-month-olds created and generalized hierarchical rules during learning. Eyeblink rate, an exploratory indicator of striatal dopamine activity, mirrored behavioral-learning patterns. Our results provide direct evidence that the human brain is predisposed to extract knowledge from noisy environments, and they add a fundamental learning mechanism to what is currently known about the neurocognitive toolbox available to infants.

Keywords

cognitive development, cognition, cognitive neuroscience

Received 10/7/14; Revision accepted 1/15/15

Seminal work has shown that infants exploit the statistical properties of the environment to learn visual and auditory information (Kirkham, Slemmer, & Johnson, 2002; Saffran, Aslin, & Newport, 1996). Infants can also learn simple algebraic rule structures in patterned sequences of speech sounds, for example, where the abstract rule maps onto an arbitrary rather than a concrete class of items (M. C. Frank, Slemmer, Marcus, & Johnson, 2009; Marcus, Fernandes, & Johnson, 2007; Marcus, Vijayan, Rao, & Vishton, 1999). However, especially in early postnatal life, infants are faced with complex changing and noisy environments that require learning and action. What mechanisms are available to help young infants transform these signals into organized behavior in the absence of repetitive patterns, cues, or incentives? Here, we provide evidence that younger infants exploit latent hierarchical-rule-learning mechanisms that to date have been considered characteristic of more mature learning.

This work is part of a long line of theoretical and experimental research examining hierarchical action and

reinforcement learning. In reinforcement learning, an agent selects among multiple actions in response to stimuli to learn stimulus-action-outcome contingencies. In a hierarchical framework, these contingencies depend on a higher-order rule set, which can be cued by multiple contexts. Thus, a hierarchical agent can select the valid stimulus-action-outcome contingencies in a context-appropriate fashion and can transfer those contingencies to novel contexts without having to learn them anew (Collins, Cavanagh, & Frank, 2014; Collins & Frank, 2013; Donoso, Collins, & Koechlin, 2014; M. J. Frank & Badre, 2012; Monsell, 2003).

This hierarchical framework is domain general and can apply to rule learning across multiple classes of stimulus inputs. For example, children growing up in a

Corresponding Author:

Denise M. Werchan, Department of Cognitive, Linguistic, and Psychological Sciences, Brown University, 190 Thayer St., Providence, RI 02912

E-mail: denise_werchan@brown.edu

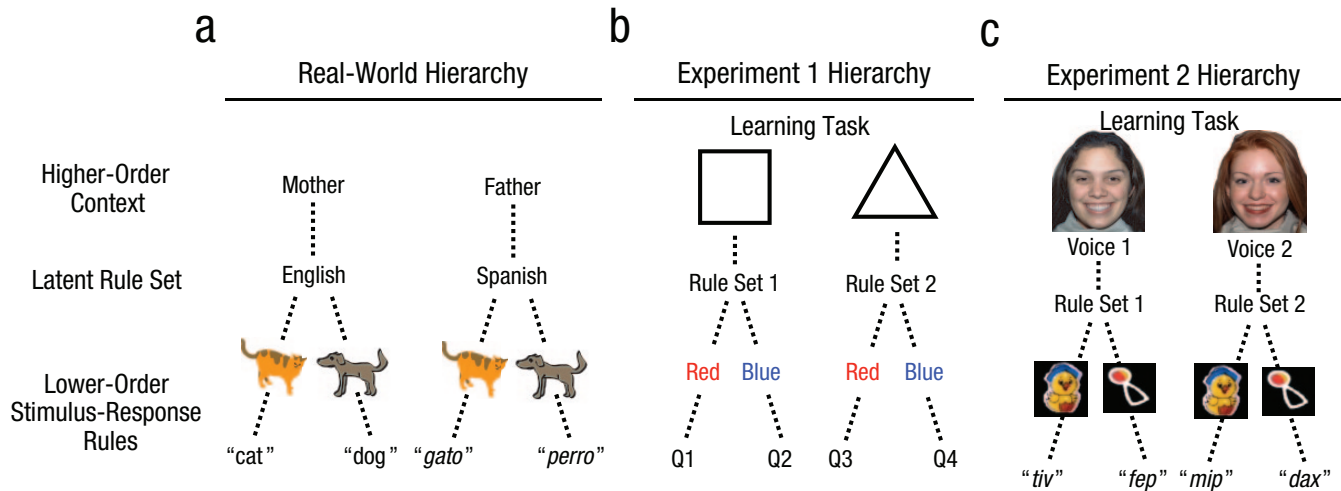


Fig. 1. Examples of hierarchical structures in (a) a real-world context and in the learning tasks from (b) Experiment 1 and (c) Experiment 2. During development, children may learn that specific higher-order contexts are associated with distinct rule sets that determine lower-order stimulus-response rules. For example, a child raised in a bilingual environment may come to expect that each parent will speak in a different language and, therefore, different words will be used to label the same objects. This mechanism was manipulated in two experiments. Experiment 1 used a visual hierarchical structure, in which two higher-order shapes each cued a separate rule set that dictated which quadrant (Q) of the screen the shape would appear in, given its color. Experiment 2 used a word-learning hierarchical structure, in which two higher-order face-voice combinations each cued a separate rule set that dictated which artificial words a pair of animated toys were associated with.

bilingual environment may learn that when they are with their mother (context), they should expect to hear English and respond in English (one rule set), but when they are with their father, they should expect to hear Spanish and respond in Spanish (another rule set). Thus, such children may use a higher-order context (mother or father) to determine the appropriate rule set to use (language that specifies object-label mappings; see Fig. 1a). This mechanism could then help infants learn and separate multiple languages without having to experience every word in each speaker context. In this framework, rule sets are distinct from the contexts that cue them. Therefore, infants may learn that the context “grandmother” is also associated with the “Spanish” rule set: When they hear their grandmother use the Spanish word “*botella*” when referring to a bottle, they can immediately infer that their father, but not their mother, will also respond to “*botella*,” even if they have never encountered a bottle with their father. Note that this example describes a different type of hierarchy than those found in linguistic structures, such as embedded clauses in syntax (Chomsky, 1988). Instead, it describes a domain-general rule-learning approach based on higher-order contexts governing lower-level rule structures.

Previous work with adults shows that hierarchical organization has a dual learning benefit (Collins et al., 2014; Collins & Frank, 2013). First, using higher-order contextual information to specify lower-order rule sets helps adults structure learning and behavior in such a way that learning new information does not interfere with behaviors learned in other contexts. For example, in

the case of a child in a bilingual home, receiving the label “cat” and the label “*gato*” should not create interference as long as the labels are governed by unique higher-order contexts (as in Fig. 1a).

Simultaneously, the rule sets are latent: They are not tied to a specific higher-order context and can thus be transferred to novel contexts when useful. Further, novel stimulus-action-outcome contingencies can be appended to latent rule sets (e.g., one can always learn a new object-word label in an existing language). Hierarchical organization of this sort (Collins et al., 2014; Collins & Frank, 2013) is incidental and automatic during learning, which raises the possibility that it may be functional early in life. However, computational models and electroencephalographic data suggest that this type of learning depends on hierarchical nesting of dopamine-innervated frontostriatal loops (Frank & Badre, 2012; Collins et al., 2014; Collins & Frank, 2013), the anterior components of which are involved in motor-action selection and thought to be underdeveloped in infancy (von Hofsten, 2004). We thus assessed whether infants exhibit a predisposition for organizing behavior into latent rule sets using an oculomotor task that does not require motor-action selection. Because infants are capable of attention-guided oculomotor control by roughly 6 to 8 months of age (Amso & Johnson, 2006, 2008), we predicted that in environments that involve oculomotor responses, infants might also automatically apply hierarchical structure to facilitate learning and generalization across contexts.

We adapted a canonical adult paradigm for assessing incidental hierarchical rule learning (Collins et al., 2014;

Collins & Frank, 2013). In Experiment 1, we used this task to investigate whether 8-month-olds spontaneously apply hierarchical structure to organize visual information (Fig. 1b). In Experiment 2, we used an identical hierarchical learning structure (Fig. 1c) to test whether this mechanism is useful for word learning, a relevant domain for young infants (Xu, Cote, & Baker, 2005). We tested the idea that if hierarchical rule sets are latent, then one should be able to append a novel object label to an existing rule set (in this case, a language) and then transfer it back to other speakers associated with that language. The juxtaposition of the two experiments allowed us to investigate whether these hierarchical-rule-learning mechanisms operate across inputs from multiple domains including visual, auditory, and multisensory information. As noted, hierarchical-structure learning is thought to depend on dopamine-innervated frontostriatal loops (Collins et al., 2014; Collins & Frank, 2013). Therefore, we also measured infants' eyeblink rate as an exploratory measure of dopamine activity. Eyeblink rate is thought to be an indirect marker of striatal dopamine activity in infants (Bacher & Smotherman, 2004) and adults (Karson, 1983) and has been implicated in similar cognitive-control rule-learning tasks in adults (Dreisbach et al., 2005; Müller et al., 2007).

Experiment 1

Method

Participants. Twenty healthy 8-month-old infants (8 females, 12 males; mean age = 8.5 months, $SD = 1.00$) were recruited via advertisements and by identifying potential candidates using birth records from the state department of health. Sample size was determined on the basis of similar studies in our lab that used the same age group. We continued collecting data until we reached this target sample size. An additional 9 infants were tested, but data were discarded because of fussiness or crying ($n = 5$), technical or experimenter error ($n = 3$), or parental interference ($n = 1$). All families were compensated for time and travel to our lab.

Materials. We used eye tracking apparatus to streamline calculation of speed of eye movements, or reaction times (RTs), to target locations. Infants' eye movements were recorded using remote eye tracking software (RED system; SensoMotoric Instruments, or SMI; Teltow, Germany), and the task was presented using E-Prime software (Version 2.0; Schneider, Eschman, & Zuccolotto, 2001).

Procedure

Task overview. Eight-month-old infants participated in a learning task and a generalization task, during which

they saw cue/target-location pairings (Fig. 2a). The cues were presented in the center of the screen and varied by both shape (e.g., square or triangle) and color (red or blue); the target was an animated toy presented in one of four quadrants on the screen (Fig. 2b). These pairings could be learned simply as individual associations between the central cues and the target locations. Alternatively, infants could apply a hierarchical structure to learn the pairings (as depicted in Fig. 2c) as adults have been previously observed to do (Collins et al., 2014; Collins & Frank, 2013). In this case, one dimension (shape) would be used as a higher-order context that cues a latent rule set, which then groups together simpler rules between the lower-order feature (color) and the target location.

After the learning task, infants saw two novel shape contexts during the subsequent generalization task. The task was designed such that if infants learned latent rule sets, they could subsequently transfer these rule sets to novel contexts (e.g., novel shapes) during the generalization task—such transfer would be evidenced by faster learning of a rule set that analogously groups together the same color-location associations in an existing set, compared with a control rule set that also involves previously experienced color-location pairings but not in a coherent set.

The mappings between rule sets, shapes, and target locations were counterbalanced. The dependent variables were (a) mean RT to the location of the target (animated toy) and (b) mean eyeblink rate per trial. RT was defined as the time between trial onset (presentation of the center cue) and the time the point of gaze arrived at the target location. Eyeblink rate was defined as the average number of eye blinks per trial. We predicted that if infants' RTs decreased with trial exposure, this would indicate that they were learning to correctly predict or anticipate the target location after the onset of the cue (Canfield & Haith, 1991).

Before the study began, infants' point of gaze was calibrated by presenting two target stimuli, one in the upper left and one in the lower right corner of the monitor. The point of gaze was validated by presenting one stimulus in each of the four corners of the monitor. Target locations were defined in the native SMI software-analysis package BeGaze and encompassed the target location stimulus (Fig. 2a). Average eye blink rates per trial were computed using SMI RED and software native to the RED system. Infants sat on their parents' laps approximately 75 cm from a 22-in. monitor in a dark room.

Learning task. During the learning task, we presented infants with four cue/target-location pairings, in which the centrally presented cues varied in shape (e.g., square or triangle) and color (red or blue). An animated toy

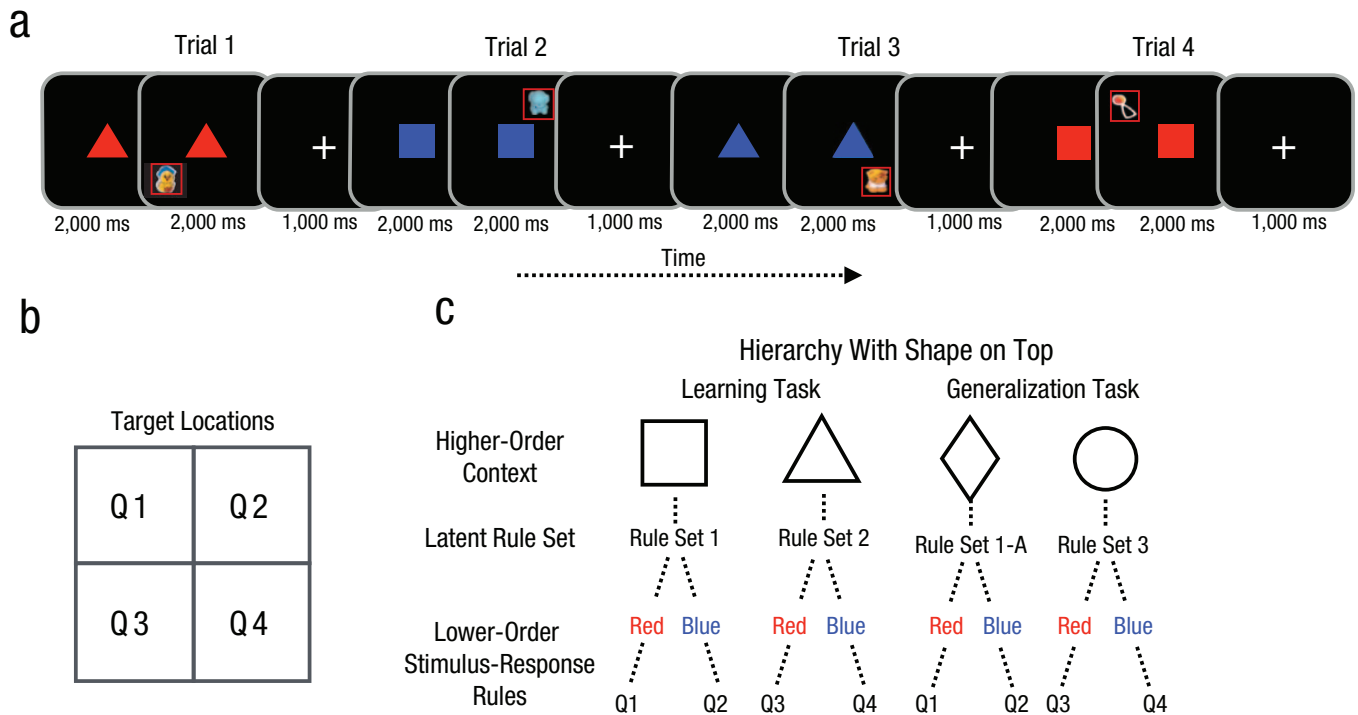


Fig. 2. Sample trial sequence and paradigm from Experiment 1. Each trial in the learning task (a) began with a centrally presented cue that varied in color (red or blue) and shape (square or triangle). Then an animated toy (the target) appeared in one of four quadrants of the computer screen (b). Eye movements were measured to determine how quickly infants looked toward the quadrant containing the target stimulus (highlighted here by the red box). Infants could use shape as a higher-order context to cluster the pairings into latent rule sets specifying lower-order color/target-location rules (c). The generalization task was similar to the learning task, except that the shapes were a diamond and a circle. The color pairings for one shape were the same as in the learning task, but the color pairings for the other shape required a new rule set.

appeared in the target location, which could be in any one of four screen quadrants. In principle, the cue/target-location pairings could be learned efficiently as four separate shape/color/target-location rules with no latent hierarchical structure. In such a case, the two dimensions (color and shape) of each cue would be used in conjunction as a single state, with no privilege given to either shape or color. Thus, infants might learn the following rules: A red triangle means the toy will appear in Quadrant 1, a blue triangle means Quadrant 2, a red square means Quadrant 3, and a blue square means Quadrant 4.

Alternatively, infants could apply a hierarchical structure to learn the pairings, as adults have been previously observed to do (Collins et al., 2014; Collins & Frank, 2013). In this case, one dimension (e.g., shape) is used as a higher-order context that cues a latent rule set, which then groups together simpler rules between the lower-order feature (e.g., color) and the target location. Thus, infants might learn sets of rules during the learning task as follows: If the higher-order context “shape” is a square, then the color red means the toy will appear in Quadrant 1 and blue means the toy will appear in Quadrant 2 (Rule Set 1, or RS1). If the higher-order context “shape” is a triangle, then different rules apply: Red predicts the toy

will appear in Quadrant 3 and blue predicts Quadrant 4 (Rule Set 2, or RS2). If infants learned in this hierarchical format, we predicted that they would more likely use shape than color as a higher-order context; we made this prediction on the basis of pilot data and the known shape bias in infants and children (e.g., Graham & Diesendruck, 2010; Landau, Smith, & Jones, 1988). The learning task was designed such that there were no clues or incentives offered to structure the input in a hierarchical format. While there is no immediate benefit to creating this sort of hierarchical structure, computational models and work with adult subjects have shown that it affords future generalization opportunities (Collins et al., 2014; Collins & Frank, 2013).

Infants received 10 trials per rule set during learning (each trial consisted of presentation of the two stimuli from the rule set). The presentation order of the stimuli was intermixed and pseudorandomized, with the constraint that the randomization resulted in an equal number of trials in which the shape changed from one trial to the next and trials in which the color changed from one trial to the next. During each stimulus presentation, the central cue was shown for 2,000 ms, after which an animated toy associated with the central cue appeared for

2,000 ms (Fig. 2a). The central cue remained on screen while the animated toy was presented. There was a 1,000-ms intertrial interval. We binned every two consecutive trials per rule set to create five learning blocks for each rule set. We defined learning as an increase in RT with trial exposure.

Generalization task. Immediately after the learning task, infants saw four new cue/target-location pairings. These new pairings were associated with the same colors (red or blue), but they were composed of new shapes (e.g., diamond or circle). These novel pairings could again be grouped by shape to form rule sets (Fig. 2c). One such rule set (RS1-A) had the same set of rules governing color/target-location pairings as a rule set from the first task (RS1, in which red was associated with Quadrant 1 and blue with Quadrant 2). The other rule set (Rule Set 3, or RS3) consisted of two color/target-location rules that had both been experienced individually before, but across different rule sets (RS1 and RS2), which allowed us to control for simple low-level stimulus-response learning (Collins & Frank, 2013). Infants again received 10 pseudorandomized trials per rule set, as in the learning task, and RTs to the target locations from cue onset were measured. We again binned every two consecutive trials to create five learning blocks per rule set. We defined learning as the change in infants' RTs with trial exposure.

If infants learned latent rule sets that were not tied to the particular shape contexts that they were learned in, then we predicted that RTs would be faster (positive transfer) for the analogous rule set (RS1-A) and slower (negative transfer) for the novel rule set (RS3). If infants did not learn a hierarchical structure, we expected to find no differences between learning of these rule sets during the generalization task. That RS3 preserved the same color/target-location rules from the learning task was a control that lower-level stimulus-response learning did not drive generalization performance. That is, it ensured that any difference in learning RS1-A and RS3 could only be due to transfer of the set of color/target-location rules rather than individual color/target-location rules. Therefore, any benefit to learning RS1-A over RS3 can be attributed only to participants having created latent rule sets during learning that could then be generalized across shape contexts.

It is also critical to note that generalization could occur only if infants used one dimension (shape) as a higher-order context that cues a latent rule set, which then groups together a set of lower-order color/target-location pairings. If infants used only shape, then there would be no generalization at test because both shapes were entirely novel. If infants used only color, then generalization would occur in both conditions at test, because both conditions preserved the color/target-location associations from the learning task. Thus, if we observed only

generalization in an analogous rule set, we could confidently adopt a model in which infants created a latent hierarchical structure during learning and then reused this structure to support learning in a novel context.

Results

Behavioral performance in the learning task.

Across subsequent trials, RTs significantly decreased for both ostensibly formed RS1 and RS2, which indicates that infants were anticipating or predicting the correct quadrant after cue presentation more quickly with exposure to both rule sets, $F(4, 76) = 6.221, p < .001, \eta_p^2 = .247$ (Fig. 3a). Notably, the majority of infants showed evidence of learning a hierarchical rule structure as in Figure 2c, rather than individual rules involving shape, color, and target location. In analogous tasks in adults, RT costs are commonly observed when the higher-order rule switches on a trial-by-trial basis and thus has to be updated into working memory (Collins et al., 2014; Collins & Frank, 2013; Monsell, 2003). We reasoned that if infants learned rule sets based on shape, then RTs should be slower (i.e., a switch cost should be evident) when the shape rule switched on consecutive trials (indicating a switch to a different rule set or group of color/target-location pairings) relative to when the shape rule repeated (indicating the same rule set or lower-level color/target-location pairings as the previous trial), regardless of color.

We calculated two switch-cost values, one assuming a higher-order shape structure (RT shape switch – RT shape repeat) and, as a fidelity check, a second assuming a higher-order color structure (RT color switch – RT color repeat). Fifteen (of 20) infants had a greater (more positive) cost to shape-rule switches than to color-rule switches (sign test, $p = .041$). Additionally, these shape-rule switch costs were significantly greater than zero, $t(14) = 2.657, p = .019$; mean switch cost = 27.04 ms. These data provide the first clue that infants may be establishing a hierarchical rule structure from ambiguous input, as indicated by a selective RT cost related to updating of higher-order rules into working memory.

Eyeblink rate in the learning task.

Neuroscience and computational-modeling research provide mechanistic evidence that the formation of hierarchical rule structures is supported by interactions between the prefrontal cortex (PFC), striatum, and their dopaminergic innervation in adults (Collins et al., 2014; Collins & Frank, 2013; M. J. Frank & Badre, 2012). These models posit that frontostriatal loops are hierarchically nested, such that a context cues a higher-order level that selects the appropriate rule structure, which in turn constrains a lower stimulus-response selection level (Collins & Frank, 2013; M. J.

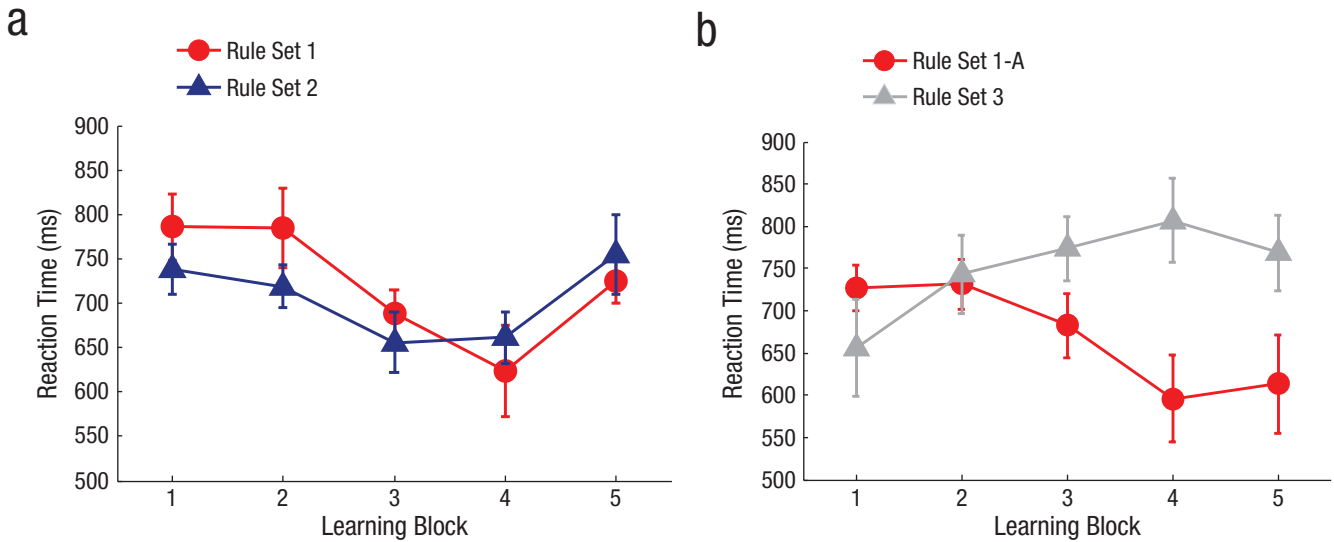


Fig. 3. Results from Experiment 1: reaction time as a function of block and rule set during (a) the learning task and (b) the generalization task. Error bars indicate ± 1 SEM.

Frank & Badre, 2012). Learning which rule structures apply is thought to rely on dopaminergic signals in frontostriatal pathways. We used this information to generate a prediction about a physiological indicator of striatal dopamine function, namely eyeblink rate (Blin, Masson, Azulay, Fondarai, & Serratrice, 1990; Karson, 1983; Kleven & Koek, 1996; Taylor et al., 1999).

In adults, higher eyeblink rate is correlated with better performance in cognitive control tasks that require updating rule representations into working memory (Dreisbach et al., 2005; Müller et al., 2007), where this same updating function is related to striatal activity (Collins & Frank, 2013; M. J. Frank & Badre, 2012; M. J. Frank, Loughry, & O'Reilly, 2001; McNab & Klingberg, 2007). Infants also show increased eyeblink rate during feeding and presentation of novel stimuli (Bacher & Smotherman, 2004), both of which are modulated by dopamine agonists (Pitts & Horvitz, 2000). These data hint at a link between eyeblink rate and dopamine activity as early as infancy. Therefore, we used this information to predict that this eyeblink-rate measure would be engaged only on precise trial types relevant to switching the higher-order dimension. We compared infants' eyeblink rate on trials in which the shape switched (and the color stayed the same) with infants' eyeblink rate on trials in which the shape repeated (and the color again stayed the same). We controlled for color switches in this way to ensure that any difference in eyeblink rate was only due to changes in the higher-order shape rule and not to factors related to color switches, such as a change in luminance.

We found that trials in which the shape switched—indicating a switch to a new rule set—elicited more

eyeblinks than trials in which the shape rule repeated, specifically during the second half of the learning task, $F(1, 19) = 11.262, p = .003$. Eyeblink rate for shape-switch ($M = 0.59$ blinks per trial, $SD = 0.62$) versus shape-repeat ($M = 0.71$ blinks per trial, $SD = 0.68$) trials was not significantly different during the first half of learning, $t(19) = 1.259, p = .223$. However, by the second half of learning, when the rule sets were learned, the eyeblink rate for shape-switch trials ($M = 1.47, SD = 1.11$) was significantly higher than the eyeblink rate for shape-repeat trials ($M = 0.75, SD = 0.57$), $t(19) = 3.951, p = .001$. As a control, we ran the same analysis assuming a higher-order context of color and controlling for changes in shape, and we found no differences in eyeblink rate between color-switch and repeat trials, $F(1, 19) = 0.531, p = .475$. This exploratory measure suggests that the neural system supporting this learning in infants may engage dopamine-dependent mechanisms.

Behavioral performance on the generalization task. Finally, and most important, we found that infants treated these hierarchical rule sets as latent, meaning that the rule sets were not tied to the particular shape contexts and could be generalized to novel contexts. Infants indeed reliably transferred the rule structure from RS1, as indicated by faster learning of the analogous rule set (RS1-A) relative to the novel rule set (RS3), $F(4, 76) = 4.102, p = .005, \eta_p^2 = .178$ (Fig. 3b). This positive transfer is consistent with the prediction that infants built rule sets during learning and reused one of these rule sets to support learning in a novel context. In contrast, the relative slowing of RTs for RS3 may be indicative of negative

transfer (Collins & Frank, 2013): RS3 pairings involved individual rules that reminded them of either RS1 or RS2; hence, an incidental tendency to apply hierarchical structure would lead to incorrect predictions and slower RTs.

These results provide the first evidence that infants create hierarchical rule structures during incidental learning. Although it is possible that infants could have learned the pairings using alternate mechanisms, such as statistical learning, this is an unlikely explanation, as our input did not contain a statistical or patterned structure that infants could exploit to facilitate generalization in novel contexts. Infants also could not have used simple associative mechanisms to facilitate learning in a novel context, because learning of the analogous and novel rule sets would have been equivalent during the generalization task if this were the case. That infants learned the analogous rule set faster than the novel rule set, along with the fact that there was a significant RT cost to higher-order rule switches, is strong evidence that infants created and reused a hierarchical rule structure.

Experiment 2

Experiment 2 was designed to replicate the results of Experiment 1 using the same hierarchical structure but with different learning and response requirements. In this experiment, we examined whether this mechanism is useful for word learning, a domain relevant to 8-month-olds, and whether such a mechanism would support the ability to append novel lower-order contingencies (object-label pairings) to existing latent rule sets (languages).

Infants use several mechanisms to facilitate language acquisition and word learning, including statistical learning to segment words from strings of syllables (Kirkham et al., 2002; Saffran et al., 1996) and abstract rule-based mechanisms to form simple rules from syllable sequences (M. C. Frank et al., 2009; Marcus et al., 2007; Marcus et al., 1999). However, infants have difficulty extracting statistical regularities when more than one artificial grammar is presented in the same sequence without explicit cues to indicate a change to a novel statistical structure (Gebhart, Newport, & Aslin, 2009). Monolingual 12-month-old infants are also unable to simultaneously learn two separate abstract rule structures (e.g., AAB and ABA patterns) from syllable sequences using simple first-order rule-learning mechanisms (Kovács & Mehler, 2009). Yet bilingual infants are capable of learning multiple languages (e.g., Genesee & Nicoladis, 2007; Pearson, Fernandez, & Oller, 2003) and appear to reach language-acquisition milestones at similar ages as their monolingual counterparts (e.g., Petitto et al., 2001). This suggests that infants might have access to additional learning mechanisms that help them learn and separate multiple languages across contexts. In Experiment 2, we examined whether

hierarchical-rule-learning mechanisms serve this goal. We tested 8-month-olds using an identical hierarchical structure as in Experiment 1. We designed Experiment 2 to be similar with respect to the hierarchical-learning structure that could be formed, but unique with respect to the response requirements as well as the information to be learned. We sought to verify the domain generality of this mechanism and especially that it was not constrained by the visuospatial dimensions and oculomotor-response requirements of Experiment 1.

Method

Participants. The final sample consisted of 22 healthy 8-month-old infants (9 females, 13 males; mean age = 8.5 months, $SD = 1.03$). An additional 5 infants were tested, but their data were discarded because of fussiness or crying ($n = 4$) or parental interference ($n = 1$). All families were compensated for their time and travel to our lab.

Procedure

Task overview. Infants were familiarized with several trials that consisted of a face followed by audiovisual toy-word pairings during a learning task and a generalization task (Fig. 4). Infants saw a face on the left half of the screen, followed rapidly by a toy on the right of the screen being labeled by a female voice. The faces were two discriminable female faces (taken from the NimStim Face Stimulus Set; Tottenham et al., 2009), and the toys were two different animated toys. Four monosyllabic pseudowords (“*jic*,” “*mip*,” “*dax*,” and “*tiv*”) were used (e.g., Xu et al., 2005), with a separate word assigned to the same object in each of the two rule sets. In addition, each unique word was spoken by one of two female speakers.

The learning task was constructed such that infants could use the face-voice mappings as higher-order contexts to create two rule sets. Critically, the same two toys were used (e.g., cartoon duck and rattle) in both rule sets. However, each face-voice higher-order context labeled the toys using different words, thereby creating RS1 and RS2, akin to learning in a bilingual environment. As in Experiment 1, infants could simply learn four associations, but this would not allow them to transfer rule sets or pass the upcoming inference test.

Specifically, as in Experiment 1, the generalization task was designed so that infants could reuse a rule-set structure that was identical to one shown during the learning task (e.g., RS1-A). This rule set could now either be transferred to a novel face-voice higher-order context or be relearned as a novel set of simple associations (Fig. 4). We also added a novel toy-word association to the rule set that was not previously experienced as part of the analogous RS1 during learning. The critical test in this experiment was whether infants would now transfer the

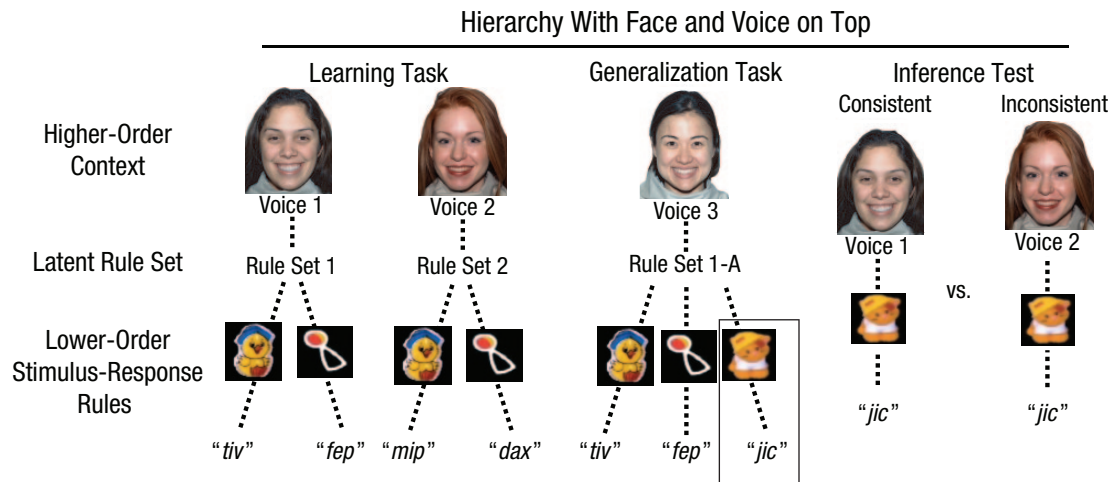


Fig. 4. Hierarchies in Experiment 2. In the learning task, infants could use a face-voice combination as a higher-order context to assign rule sets to pairings of toys with pseudowords. In the generalization task, infants were shown a learned rule set now associated with a novel face-voice context; an additional toy-word pairing was also added to the set. During the inference test, infants were shown pairings that were consistent and inconsistent with the rule-set structure. All mappings between faces, voices, toys, and words were counterbalanced.

novel word to the appropriate face-voice higher-order context originally experienced during the learning phase of the task—that is, whether they appended a novel association to an existing latent rule set. Thus, we included a final inference test trial in which we paired the higher-order face-voice contexts from learning with the novel toy-word association presented as part of RS1-A during generalization, and we examined looking time when the new toy-word pairing was paired with the consistent (RS1) versus the inconsistent (RS2) face-voice context from learning (Fig. 4). We predicted that if infants formed hierarchical rule sets using the face-voice as a higher-order context, then they would look longer at the inconsistent trials that violated the learned rule-set structure. If infants did not adopt a hierarchical rule-set structure, then we expected to find no difference in looking time between the consistent and inconsistent trials.

The mappings between faces, voices, toys, and words were counterbalanced. The dependent measures were the average looking time to the consistent relative to the inconsistent face-voice context during the inference-test trial and the average eyeblink rate during the learning task, as in Experiment 1. We used the same eye tracking software as in Experiment 1 to gather average eyeblink rate per trial.

Learning task. During the learning task, infants saw four different pairings of faces and voices with toys and words. Two female faces and voices, two toys, and four words were used in these pairings. All mappings between faces, voices, toys, and words were counterbalanced. The pairings were constructed such that each face was associated with a unique voice (e.g., Face 1 was always associated with words spoken by Voice 1, and

Face 2 was associated with words spoken by Voice 2). Both face-voice mappings were associated with the same two toys (e.g., both Face 1 and Face 2 were paired with a cartoon duck and a rattle); however, each face-voice mapping used different words to label the toys, as in a bilingual environment.

Infants received a total of 32 trials (8 trials per pairing). During each trial, infants would first see the face on the left side of the screen. After 500 ms, a toy appeared on the right side of the screen for an additional 1,500 ms, while a recorded female voice said the artificial word associated with the pairing. There was a 1,000-ms interval between trials. The pairings could be learned simply as individual associations between faces, voices, toys, and words, using simple associative-learning mechanisms. Alternatively, infants could use the face-voice mappings as higher-order contexts to learn the pairings as rule sets (RS1 and RS2) grouping together simpler toy-word rules or associations.

Generalization task. Immediately following the learning task, we presented infants with three new pairings of faces, voices, toys, and words. These pairings could again be grouped by face and voice to form a rule set identical to one experienced during the learning task (e.g., RS1-A); however, this rule set was now associated with a novel higher-order face-voice context. Additionally, one novel toy-word association was added to the rule set (Fig. 4). Infants again received 8 trials per pairing, as in the learning task, for a total of 24 trials.

Inference test. After the learning and generalization tasks, infants saw the faces and voices from the learning

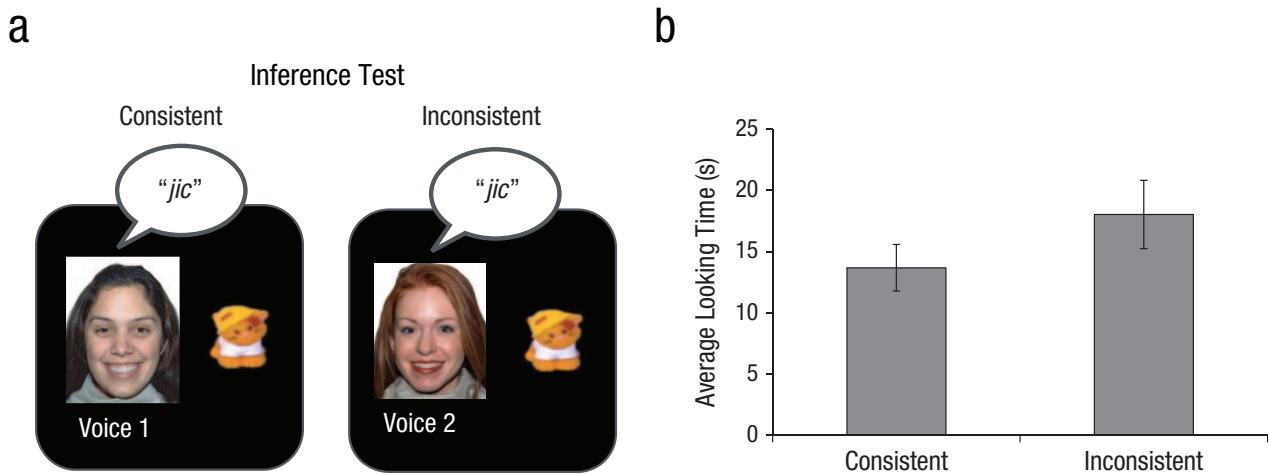


Fig. 5. Paradigm of (a) and results from (b) the inference test. During the inference test, infants saw pairings of faces and voices with toys; these pairings were either consistent or inconsistent with the hierarchical structure they had learned. The graph shows average looking time for consistent and inconsistent pairings. Error bars indicate ± 1 SEM.

task paired with the novel toy-word association from the generalization task. One of these pairings of faces, voices, toys, and words was consistent with the rule set structure formed during the tasks, whereas the other pairing was inconsistent with this rule-set structure (Fig. 5a). Infants received two consistent trials and two inconsistent trials during the inference test. The order of the consistent and inconsistent test trials was intermixed and counterbalanced across subjects. During each test trial, infants saw the face and toy while a recorded voice said the word associated with the toy once every 3 s. Each trial continued until the infant looked away for more than 2 s, for a maximum of 60 s. The dependent measure was the average looking time during the consistent trials compared with the average looking time during the inconsistent trials.

Results

To determine whether there were any differences in looking time between the consistent and inconsistent trials, we conducted a two-tailed paired-samples *t* test, which indicated that infants looked significantly longer at the inconsistent pairing than at the consistent pairing, $t(21) = 2.461, p = .023$ (Fig. 5b).

We next examined differences in eyeblink rate for trials in which the higher-order face-voice rule switched from the previous trial—indicating the need to update the current rule set in working memory—compared with trials where the higher-order face-voice rule repeated during the learning task. We conducted a 2 (trial type: rule switch vs. rule repeat) \times 2 (time: first vs. second half of learning) repeated measures analysis of variance. Replicating the findings from Experiment 1, the results showed a

time-by-trial-type interaction, $F(1, 21) = 7.47, p = .013$. Eyeblink rate for face-switch ($M = 0.41, SD = 0.28$) versus face-repeat ($M = 0.34, SD = 0.30$) trials was not significantly different during the first half of the learning task, $t(21) = 1.16, p = .26$. However, by the second half of learning, eyeblink rate for face-switch trials ($M = 0.50, SD = 0.32$) was significantly higher than eyeblink rate for face-repeat trials ($M = 0.35, SD = 0.29$), $t(21) = 2.96, p = .008$.

General Discussion

Across two experiments, we showed that infants spontaneously apply hierarchical rule-set structures during incidental learning. Notably, our findings from both experiments suggest that these rule sets were not tied to a particular context but were instead latent, as evidenced by the fact that infants could generalize the sets to support learning in novel contexts. Eyeblink rate, an exploratory physiological indicator of dopamine activity in infants (Bacher & Smotherman, 2004) and adults (Karson, 1983), mirrored the behavioral findings.

Prior research shows that infants use statistical and algebraic rule-based mechanisms to drive learning in environments that have a statistical or patterned structure (M. C. Frank et al., 2009; Kirkham et al., 2002; Marcus et al., 2007; Marcus et al., 1999; Saffran et al., 1996). However, these mechanisms are unlikely to account for our findings, as neither of our experiments contained a statistical or patterned structure that infants could exploit to facilitate learning. Infants also could have learned the cue/target-location pairings in Experiment 1 and the pairings of faces and voices with toys and words in Experiment 2 using simple associative-learning mechanisms. However, if this were the case, then performance during the

generalization task should be equivalent in Experiment 1, and looking time during the inference test should be equivalent in Experiment 2. That infants exhibited faster learning of an analogous rule set in Experiment 1 is clear evidence that infants spontaneously constructed rule sets during initial learning and flexibly reused one of these sets to facilitate learning in a novel context. Experiment 2 replicated and extended these findings by showing that hierarchical rule sets are latent: 8-month-old infants were able to append a novel object label to an existing rule set during generalization and then reference the novel item back to a higher-order context associated with that rule set. This type of mechanism may thus help infants learn multiple languages without having to experience every word in each speaker context.

Evidence from computational modeling and neuroscience research suggests that hierarchical rule learning is supported by dopamine-innervated pathways between the PFC and striatum (Collins et al., 2014; Collins & Frank, 2013; Donoso et al., 2014; M. J. Frank & Badre, 2012). Clearly, it may well be that some other neural system supports the hierarchical rule learning observed in our infant sample. However, our data showing that infants have a greater switch cost for the higher-level than the lower-level dimension, paired with higher eyeblink rate specifically on rule-switch trials, is remarkably consistent with behavioral patterns traditionally associated with frontostriatal working memory updating mechanisms. While the PFC does not reach maturity until adolescence (e.g., Huttenlocher, 1979), the basal ganglia show relatively high functionality, as measured by glucose metabolism, compared with most of the cerebral cortex in the newborn period (Chugani, 1996). One hypothesis then is that frontostriatally mediated hierarchical rule learning may be dependent on frontostriatal loops in infancy in such a way that weights striatal involvement more heavily than prefrontal involvement. Thus, these frontostriatal loops perform similar computations across the life span, but on inputs that are appropriate to learning in infants' unique ecological niche (Rovee-Collier & Cuevas, 2009). As the individual's ecological niche changes and adapts across development, this mechanism may then be coopted to support increasingly complex tasks, such as cognitive control of complex thought and action. Another possibility is that this form of hierarchical learning may require less anterior frontal involvement than adult versions of the task. In adult work, participants must learn correct responses through reinforcement. In contrast, infants are simply shown the toys in the target locations, which directly indicate the correct actions (e.g., screen quadrants to direct gaze to). This form of hierarchical learning may require less anterior PFC processing than adult versions, because infants do not have to learn to select motor actions using reinforcement learning. This suggests that

the PFC may not need to be fully developed to support hierarchical rule learning as tested here. Future work using computational and neuroimaging tools appropriate to infants will bear directly on these questions.

Author Contributions

All authors helped conceive and design the experiments, which were based on models and tasks developed by A. G. E. Collins. D. M. Werchan collected and analyzed the data with input from A. G. E. Collins, M. J. Frank, and D. Amso. D. M. Werchan, M. J. Frank, and D. Amso wrote the manuscript.

Acknowledgments

We thank the infants and families who made this research possible.

Declaration of Conflicting Interests

The authors declared that they had no conflicts of interest with respect to their authorship or the publication of this article.

Funding

This work was supported by a National Science Foundation Graduate Research Fellowship under Grant No. DGE-1058262 to D. M. Werchan and by Grant No. R01 MH099078 from the National Institutes of Health to D. Amso and D. Badre.

References

- Amso, D., & Johnson, S. P. (2006). Learning by selection: Visual search and object perception in young infants. *Developmental Psychology, 42*, 1236–1245.
- Amso, D., & Johnson, S. P. (2008). Development of visual selection in 3- to 9-month-olds: Evidence from saccades to previously ignored locations. *Infancy, 13*, 675–686.
- Bacher, L. F., & Smotherman, W. P. (2004). Spontaneous eye blinking in human infants: A review. *Developmental Psychobiology, 44*, 95–102.
- Blin, O., Masson, G., Azulay, J. P., Fondarai, J., & Serratrice, G. (1990). Apomorphine-induced blinking and yawning in healthy volunteers. *British Journal of Clinical Pharmacology, 30*, 769–773.
- Canfield, R. L., & Haith, M. M. (1991). Active expectations in 2- and 3-month-old infants: Complex event sequences. *Developmental Psychology, 27*, 198–208.
- Chomsky, N. (1988). *Aspects of the theory of syntax* (Vol. 11). Cambridge, MA: MIT Press.
- Chugani, H. T. (1996). Neuroimaging of developmental nonlinearity and developmental pathologies. In R. W. Thatcher, G. R. Lyon, J. Rumsey, & N. Krasnegor (Eds.), *Developmental neuroimaging: Mapping the development of brain and behavior* (pp. 187–195). San Diego, CA: Academic Press.
- Collins, A. G., Cavanagh, J. F., & Frank, M. J. (2014). Human EEG uncovers latent generalizable rule structure during learning. *The Journal of Neuroscience, 34*, 4677–4685.
- Collins, A. G., & Frank, M. J. (2013). Cognitive control over learning: Creating, clustering, and generalizing task-set structure. *Psychological Review, 120*, 190–229.

- Donoso, M., Collins, A. G. E., & Koechlin, E. (2014). Foundations of human reasoning in the prefrontal cortex. *Science*, *344*, 1481–1486.
- Dreisbach, G., Müller, J., Goschke, T., Strobel, A., Schulze, K., Lesch, K. P., & Brocke, B. (2005). Dopamine and cognitive control: The influence of spontaneous eyeblink rate and dopamine gene polymorphisms on perseveration and distractibility. *Behavioral Neuroscience*, *119*, 483–490.
- Frank, M. C., Slemmer, J. A., Marcus, G. F., & Johnson, S. P. (2009). Information from multiple modalities helps 5-month-olds learn abstract rules. *Developmental Science*, *12*, 504–509.
- Frank, M. J., & Badre, D. (2012). Mechanisms of hierarchical reinforcement learning in corticostriatal circuits 1: Computational analysis. *Cerebral Cortex*, *22*, 509–526.
- Frank, M. J., Loughry, B., & O'Reilly, R. C. (2001). Interactions between frontal cortex and basal ganglia in working memory: A computational model. *Cognitive, Affective, & Behavioral Neuroscience*, *1*, 137–160.
- Gebhart, A. L., Newport, E. L., & Aslin, R. N. (2009). Statistical learning of adjacent and nonadjacent dependencies among nonlinguistic sounds. *Psychonomic Bulletin & Review*, *16*, 486–490.
- Genesee, F., & Nicoladis, E. (2007). Bilingual first language acquisition. In E. Hoff & M. Shatz (Eds.), *Handbook of language development* (pp. 324–342). Oxford, England: Blackwell.
- Graham, S. A., & Diesendruck, G. (2010). Fifteen-month-old infants attend to shape over other perceptual properties in an induction task. *Cognitive Development*, *25*, 111–123.
- Huttenlocher, P. R. (1979). Synaptic density in human frontal cortex—developmental changes and effects of aging. *Brain Research*, *163*, 195–205.
- Karson, C. N. (1983). Spontaneous eye-blink rates and dopaminergic systems. *Brain*, *106*, 643–653.
- Kirkham, N. Z., Slemmer, J. A., & Johnson, S. P. (2002). Visual statistical learning in infancy: Evidence for a domain general learning mechanism. *Cognition*, *83*, B35–B42.
- Kleven, M. S., & Koek, W. (1996). Differential effects of direct and indirect dopamine agonists on eye blink rate in cynomolgus monkeys. *Journal of Pharmacology and Experimental Therapeutics*, *279*, 1211–1219.
- Kovács, Á. M., & Mehler, J. (2009). Cognitive gains in 7-month-old bilingual infants. *Proceedings of the National Academy of Sciences, USA*, *106*, 6556–6560.
- Landau, B., Smith, L. B., & Jones, S. S. (1988). The importance of shape in early lexical learning. *Cognitive Development*, *3*, 299–321.
- Marcus, G. F., Fernandes, K. J., & Johnson, S. P. (2007). Infant rule learning facilitated by speech. *Psychological Science*, *18*, 387–391.
- Marcus, G. F., Vijayan, S., Rao, S. B., & Vishton, P. M. (1999). Rule learning by seven-month-old infants. *Science*, *283*, 77–80.
- McNab, F., & Klingberg, T. (2007). Prefrontal cortex and basal ganglia control access to working memory. *Nature Neuroscience*, *11*, 103–107.
- Monsell, S. (2003). Task switching. *Trends in Cognitive Sciences*, *7*, 134–140.
- Müller, J., Dreisbach, G., Goschke, T., Hensch, T., Lesch, K. P., & Brocke, B. (2007). Dopamine and cognitive control: The prospect of monetary gains influences the balance between flexibility and stability in a set-shifting paradigm. *European Journal of Neuroscience*, *26*, 3661–3668.
- Pearson, B. Z., Fernandez, S. C., & Oller, D. K. (1993). Lexical development in bilingual infants and toddlers: Comparison to monolingual norms. *Language Learning*, *43*, 93–120.
- Petitto, L. A., Katerelos, M., Levy, B. G., Gauna, K., Tétreault, K., & Ferraro, V. (2001). Bilingual signed and spoken language acquisition from birth: Implications for the mechanisms underlying early bilingual language acquisition. *Journal of Child Language*, *28*, 453–496.
- Pitts, S. M., & Horvitz, J. C. (2000). Similar effects of D1/D2 receptor blockade on feeding and locomotor behavior. *Pharmacology, Biochemistry & Behavior*, *65*, 433–438.
- Rovee-Collier, C., & Cuevas, K. (2009). Multiple memory systems are unnecessary to account for infant memory development: An ecological model. *Developmental Psychology*, *45*, 160–174.
- Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, *274*, 1926–1928.
- Schneider, W., Eschman, A., & Zuccolotto, A. (2012). *E-Prime 2.0 reference guide manual*. Pittsburgh, PA: Psychology Software Tools.
- Taylor, J. R., Elsworth, J. D., Lawrence, M. S., Sladek, J. R., Jr., Roth, R. H., & Redmond, D. E., Jr. (1999). Spontaneous blink rates correlate with dopamine levels in the caudate nucleus of MPTP-treated monkeys. *Experimental Neurology*, *158*, 214–220.
- Tottenham, N., Tanaka, J. W., Leon, A. C., McCarry, T., Nurse, M., Hare, T. A., . . . Nelson, C. (2009). The NimStim set of facial expressions: Judgments from untrained research participants. *Psychiatry Research*, *168*, 242–249.
- von Hofsten, C. (2004). An action perspective on motor development. *Trends in Cognitive Sciences*, *8*, 266–272.
- Xu, F., Cote, M., & Baker, A. (2005). Labeling guides object individuation in 12-month-old infants. *Psychological Science*, *16*, 372–377.