ERROR AND EXPECTATION IN LANGUAGE LEARNING: THE CURIOUS ABSENCE OF *MOUSES* IN ADULT SPEECH

Michael Ramscar — *University of Tübingen*
Melody Dye — *Indiana University*
Stewart M. McCauley — *Cornell University*

As children learn their mother tongues, they make systematic errors. For example, English-speaking children regularly say *mouses* rather than *mice*. Because children's errors are not explicitly corrected, it has been argued that children could never learn to make the transition to adult language based on the evidence available to them, and thus that learning even simple aspects of grammar is logically impossible without recourse to innate, language-specific constraints. Here, we examine the role children's expectations play in language learning and present a model of plural noun learning that generates a surprising prediction: at a given point in learning, exposure to regular plurals (e.g. *rats*) can decrease children's tendency to overregularize irregular plurals (e.g. *mouses*). Intriguingly, the model predicts that the same exposure should have the opposite effect earlier in learning. Consistent with this, we show that testing memory for items with regular plural labels contributes to a decrease in irregular plural overregularization in six-year-olds, but to an increase in four-year-olds. Our model and results suggest that children's overregularization errors both arise and resolve themselves as a consequence of the distribution of error in the linguistic environment, and that far from presenting a logical puzzle for learning, they are inevitable consequences of it.*

Keywords: learning, morphology, prediction, negative evidence, nativism, noun plurals, overregularization

> Gregory: 'Is there any other point to which you would wish to draw my attention?'
> Holmes: 'To the curious incident of the dog in the nighttime.'
> Gregory: 'The dog did nothing in the nighttime.'
> Holmes: 'That was the curious incident.'
>
> ('Silver Blaze', by Sir Arthur Conan Doyle)

1. Introduction. A racehorse vanishes, its trainer murdered. Sherlock Holmes lights upon a crucial piece of evidence: a dog has remained silent throughout (Gregory 2007). The fact that an expected event did not occur—the dog never barked—provides Holmes with a critical clue, enabling him to deduce that the culprit must be familiar with the dog. Holmes's deduction is a reminder that much can be learned from the discrepancy between what is expected and what actually occurs (Wasserman & Castro 2005). Here, we show how children use these discrepancies as an important source of evidence in learning, and that often, as in the curious incident of the dog in the nighttime, the nonoccurrence of expected events provides a rich and critical source of information.

The information offered by violations of expectation has often been marginalized or ignored in discussions of language learning (Brown & Hanlon 1970, Marcus 1993). It is claimed that this kind of 'indirect' negative evidence has little to offer a child engaged in a task as complex as language learning (Pinker 1984, 1989, 2004). There is, however, reason to believe that evidence acquired by expectation may be of more use to children than has often been supposed, because it is now commonly accepted that both positive evidence (the reinforcement of successful predictions) and negative evidence (unlearning as a result of prediction error) are necessary to account for even the most basic aspects of animal learning (Kamin 1969, McLaren & Mackintosh 2000, Pearce & Hall 1980, Rescorla 1968, Rescorla & Wagner 1972, Sutton & Barto 1998). As a result,

* This material is based upon work supported by the National Science Foundation under Grant Nos. 0547775 and 0624345 to Michael Ramscar, and Grant Nos. 2010083519 and DGE-0903495 to Melody Dye. We are grateful to Harald Baayen, Bradley Love, and Daniel Yarlett for many helpful discussions of these ideas, and to Rowan Goddard, Ian Goddard, and Johanna Moore, who inspired this project.

many researchers have wondered whether expectation might not also play a more substantial role in children's language learning (Bates & Carnevale 1993, Elman 1991, Hahn & Oaksford 2008, Johnson 2004, Lewis & Elman 2001, MacWhinney 2004, Prinz 2002, Pullum & Scholz 2002, Ramscar & Yarlett 2007, Ramscar, Yarlett, et al. 2010, Rohde & Plaut 1999, Seidenberg & MacDonald 1999).

In what follows, we show how a learning model that tunes its expectations according to the success or failure of its predictions exhibits the SAME trajectory of linguistic development in learning irregular plurals that children do, a pattern that has often been claimed to be incompatible with learning from the environment (Pinker 1989). Moreover, the model makes a novel empirical prediction: at early stages of learning, exposure to regular plurals can increase children's tendency to overregularize irregular plurals, while at a later stage, the exact same intervention will have precisely the opposite effect, such that learning about regulars will cause overregularization rates in older children to drop. Consistent with this, we find that memory testing for items that have regular plural labels increases the overregularization of irregular plurals in four-year-olds, but decreases it in six-year-olds. The model and results we present show how children's overregularization errors can arise as a natural consequence of the distribution of error in the linguistic environment, and subsequently are resolved as a natural consequence of the same learning mechanisms and the same distribution that give rise to them in the first place: rather than presenting a logical puzzle for learning, we show that overregularization errors are inevitable consequences of it.

2. THE LOGICAL PROBLEM OF LANGUAGE ACQUISITION. In the course of learning language, children often go through phases in which they make predictable errors. For example, English-speaking preschoolers often say *mouses* where their parents and older siblings would say *mice*. Because these errors are systematic, and because they are usually not explicitly corrected, it has been argued that children could never learn to make the transition to adult language based on experience alone. Accordingly, it is often claimed that learning even simple aspects of grammar is logically impossible in the absence of innate constraints on what is learned (this argument is often referred to as the 'logical problem of language acquisition', or LPLA; see Baker 1979).

A classical statement of the LPLA is given by Pinker (1984): in attempting to learn language, he argues, children must 'hypothesize the grammar of the adult language' (Figure 1). Strictly speaking, the child's task is to 'guess' the identity of the set of grammatical strings that makes up the language (Gold 1967).

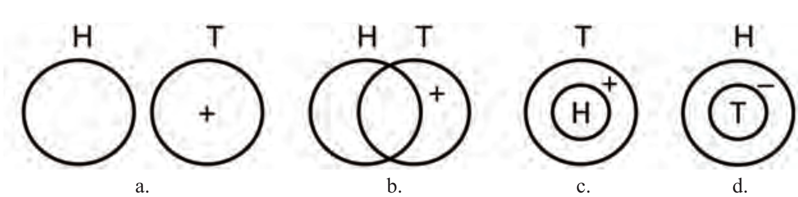

FIGURE 1. Four logical situations a child might arrive at while trying to 'learn' a language. Each circle represents the set of sentences in a language. H: child's hypothesized language; T: adult target language; +: grammatical sentence in the language the child is trying to learn; –: ungrammatical sentence (Pinker 1984).

Pinker depicts languages as circles that correspond to sets of word sequences and offers four logical possibilities for how a child's hypotheses might differ from adult language. In the first possibility (a), the child's hypothesis language, H, is disjoint from the language to be acquired (the target language, T). In terms of noun usage (our focus here), this corresponds to the state of a child learning English who cannot produce any well-formed irregular noun plurals (the child might say things like *mouses* but never *mice*). In (b), the sets H and T intersect, corresponding to a child who has correctly learned some irregular plurals, but not others (the child uses *mice* alongside incorrect forms like *gooses*). In (c), H is a subset of T, which means that the child has mastered usage of some but not all English noun plurals and never uses forms that are not part of English. Finally, in (d), H is a SUPERSET of T, meaning that the child uses English nouns correctly and also produces forms that are not part of the English language (i.e. the child uses both *mouses* and *mice* interchangeably).

A core assumption of this statement of the LPLA is that learners can only recover from erroneous superset inferences if they receive explicit corrective feedback from their parents or linguistic community (Pinker 1989). In the absence of such feedback, it is argued that all of the positive evidence children encounter will be consistent with the superset hypothesis they have made and will thus give them no reason to believe that this hypothesis is in error (Pinker 1984). Because children do not receive explicit corrective feedback about their mistakes (Brown & Hanlon 1970, but see also Bohannon & Stanowicz 1988, Schoneberger 2010), and because they do go through stage (d), it is claimed that children cannot learn the correct target language solely from experience—that is, on the basis of positive evidence alone.

It follows logically, then, that both the validity of the LPLA and the claim that the LPLA effectively disproves the idea that language can be learned without innate constraints (Baker 1979, Marcus et al. 1992, Pinker 1984, 1989, 2004) hinge on the idea that the kind of information that would allow children to correct their behavior is simply not present in the linguistic environment (Johnson 2004, Pinker 2004). Accordingly, if it can be shown that children CAN learn to correct themselves solely on the basis of evidence available in the environment, then clearly the argument does not hold: in that case, there would simply be no 'logical' problem of language learning (Johnson 2004, Pullum & Scholz 2002, Ramscar & Yarlett 2007).

3. MODELS OF LEARNING INFLUENCE CONCEPTIONS OF LEARNABILITY. In his 1989 book, *Learnability and cognition*, Steven Pinker raises—and dismisses—the possibility that 'indirect' negative evidence could provide a solution to the LPLA over the course of a single page. In a more recent article devoted to the LPLA (Pinker 2004), the matter is demoted to a footnote. This approach is not unusual; it reflects a set of beliefs that have come to dominate the study of children's language learning over the past half century (see Landauer & Dumais 1997 and Schoneberger 2010 for further discussion of this point).

To understand what is remarkable here, one has to step outside the realm of child language learning and venture into the humble world of the laboratory rat, because for the past forty years, psychologists studying animal behavior have been busy applying a fully fleshed-out theory of learning strategies to the study of RODENTS, and have shown that rats' expectations provide a critical source of evidence across a wide range of learning tasks. Strikingly, psychologists studying rats have found it impossible to explain the behavior of their subjects without acknowledging that rats are capable of learning in ways that are far more subtle and sophisticated than many researchers studying language tend to countenance in human children (Dayan & Daw 2008, Rescorla 1988).

Moreover, not only is it the case that animal learning models have been fleshed out in ways that embrace the idea that animals make extensive use of indirect evidence in learning, but the computational properties of these models have also been extensively explored (Dickinson 1980, Mackintosh 1975, Pearce & Hall 1980, Rescorla & Wagner

1972; see Danks 2003 for a review), and much progress has been made in understanding the biological underpinnings of these mechanisms (Daw & Shohamy 2008, Montague et al. 1996, Montague et al. 2004, Niv 2009, Schultz 1998, 2006, Schultz et al. 1997, Schultz & Dickinson 2000, Waelti et al. 2001).

Accordingly, while it is often claimed that animal models are insufficient to explain language learning and that some kind of domain-specific module or specialized set of learning principles is necessary to account for linguistic development, it is clear that many unlearnability arguments rely on inaccurate or outdated characterizations of learning (see also Pereira 2000). This is important, because questions about whether language is learnable from the environment (or whether animal models can offer insight into language learning) are best answered empirically (Gold 1967), by testing the predictions of well-specified learning models that have been trained on well-defined tasks and accurately characterized representations of the learning environment.

This is the approach to understanding the development of children's noun pluralization taken here: we show how a model of learning developed in the animal literature can be used to specify—and, critically, predict—the circumstances that can prompt a child to 'conclude that a nonwitnessed [form] is ungrammatical' (Pinker 1989:14). Explaining how children come to learn that some forms are more grammatical than others does not '[take] the burden of explaining learning out of the environmental input and [put] it back in the child', as has sometimes been claimed (Pinker 1989:14–15). Instead, we show how a proper understanding of both learning AND context—that is, the distribution of error in the child's environment—is critical to explaining how children learn language and understanding why they exhibit the characteristic patterns of linguistic development that they do.

We begin this explanation by briefly describing the picture of learning that has emerged from the study of animals.

4. THE ROLES OF EXPECTATION AND ERROR IN ANIMAL LEARNING. Although much of our contemporary understanding of animal learning has its origins in Ivan Pavlov's (1927) conditioning experiments, it is critical to note that the ideas about learning that people typically take from Pavlov's work are, in most ways, the OPPOSITE of the understanding of animal learning that has developed in the century since Pavlov's initial discoveries (Rescorla 1988). As is well known, Pavlov discovered that if he rang a bell as food was presented to a dog, the dog would later salivate upon hearing the bell, even if no food was on offer. This finding gave rise to a view of learning based on ASSOCIATION: animals were thought to learn to 'associate' previously unrelated things, such as bells and meals, by tracking the degree to which a STIMULUS (a bell) and a RESPONSE (salivation brought on by food) were paired.

Empirically, the naive view of Pavlovian conditioning, which sees learning as a simple process of recording cooccurrences that 'computes nothing more than correlations' (Santos et al. 2007:446), has been shown to be deeply mistaken (Rescorla 1988), as have two stubbornly popular—yet empirically false—beliefs pertaining to the necessary and sufficient conditions for learning: first, that explicit 'rewards' or 'punishments' are NECESSARY for learning; and second, that a simple cooccurrence between a 'stimulus' and a 'response' is SUFFICIENT for learning (i.e. if a bell is paired with food often enough, a dog will always learn the association). Although the results of animal experiments have long since shown these ideas to be wrong (Rescorla 1988), they still pervade the literatures in linguistics and cognitive science.

Rescorla (1968) provided one of the first clear demonstrations that these ideas are inadequate to explain the learning that occurs in animal conditioning: in a variant of the classic Pavlovian paradigm, a group of rats learned to associate a tone with a mild electric shock, according to the schedule of tones and shocks depicted in Figure 2.

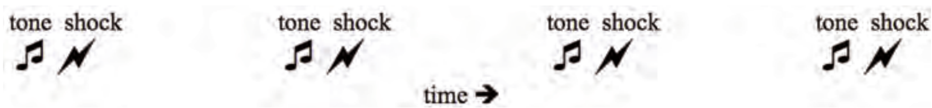

FIGURE 2. Schematic of a conditioning schedule used in Rescorla 1968. The rate of tones absent shocks here is zero.

Like Pavlov's dogs, these rats quickly learned to associate the tones with the shocks, freezing when a tone later sounded. However, a second group of rats that was exposed to an identical number of tone–shock pairings as the first group, but into which a number of tones that were not followed by shocks were interpolated (Figure 3), exhibited a very different pattern of learning.

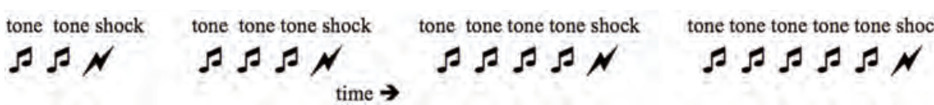

FIGURE 3. A training schedule with an increased background rate of tones without shocks: although the absolute number of tones leading to shocks is identical, approximately 70% of the tones are not followed by shocks, and the degree to which rats condition to the relationship between the tones and the shocks diminishes proportionally.

As the number of tones without shocks increased, rats came to associate the tones with the shocks less and less. Indeed, the degree to which the rats froze upon hearing the tone decreased in direct proportion to the BACKGROUND RATE of tones absent shocks. As the background rate increased, conditioning decreased, despite the fact that the rate at which the tones cooccurred with the shocks remained exactly the same.

This finding cannot be explained by the naive 'associative' conceptions of learning that we described above (Rescorla 1988). Given that there was no change in the tone–shock association rate between the groups of rats—only the background rate varied—it follows that the difference in what was learned must be due to the 'no shock' trials. The NONoccurrence of expected shocks after certain tones influenced the degree to which the rats conditioned to the tones that did precede shocks. It follows then that learning cannot simply be a process of tracking positive cooccurrences of cues and events.

Indeed, it has long been well established that there is more to learning than simply counting successful and unsuccessful predictions. The results of numerous experiments have revealed that animal learning is a process that can be seen, informally, to reduce UNCERTAINTY in an animal's developing understanding of the predictive structure of its environment (Rescorla 1988). Because uncertainty is reduced as cues are learned and reliable expectations are formed, learning is best understood as a competitive process: if an animal learns to predict an outcome from one cue, there will be less uncertainty to drive the learning of another. CUE COMPETITION is thus a simple statistical consequence of uncertainty reduction and can be illustrated by the results of BLOCKING experiments (Kamin 1969), in which learning about the predictive value of a novel cue is effectively 'blocked' by the presence of an already well-learned cue.

For example, if a rat has learned that it will be shocked when it hears a tone, and a light is subsequently paired with the tone in training, any learning of the light as an additional predictive cue will be inhibited. Because the tone is already fully informative about the upcoming shock, the information provided by the light is redundant and is therefore ignored. Prior learning about the tone blocks subsequent learning about the light. As numerous results like this demonstrate, rats do not learn simple 'associations' between stimuli and responses; rather, they learn the degree to which individual cues are systematically informative about the environment.

In cases where the informative cues to an event (or other aspect of the environment) have not yet been established, potentially predictive cues compete for relevance. As a result, cues that are more reliably informative are discriminated from cues that are less informative (Rescorla 1988). Cue competition uncovers positively informative relationships within an animal's environment by eliminating the influence of less informative relationships. Since there are invariably far more UNinformative coincidences in the environment than informative ones, it follows that expectations that are wrong have more influence on the shape of learning than expectations that are right (for discussion, see Ramscar et al. 2011).

Given the logic of the foregoing discussion of error and expectation, one might ask: WHAT expectations? WHICH errors? Since the rats in Rescorla's experiment had no a priori knowledge about the relationship between the tones and the shocks, it is natural to wonder why it was only the background rate of the tones that mattered in predicting the upcoming shock. The answer is that, in principle, everything in the rat's local environment mattered (Rescorla 1988). However, just as the rat will learn to discount tones as predictive cues the less they appear with shocks, so it will have learned to discount the myriad other aspects of its environment that have often been present in the absence of shocks. Prior learning thus influences—and, indeed, is integral to—subsequent learning. What the rat learns in a given context can only be understood against the backdrop of what it has learned already. For the sake of simplicity, models and explanations tend to focus on informative cues, while ignoring cues whose high background rates are likely to render them largely irrelevant in competitive terms.[1] It is important to understand, however, that the novelty of a given cue is entirely relative and can only be computed in relation to the other potential cues that are available to a learner (Ramscar, Yarlett, et al. 2010, Rescorla 1988). (This helps clarify why learning is often related to a 'stimulus complex', rather than to individual stimuli; Rescorla & Wagner 1972.)

Finally, it is worth noting that the logic of discrimination learning suggests that at the outset, what a young learner encounters is best conceptualized as a large, undifferentiated set of cues connected to little or no environmental knowledge,[2] and that the perceptible variances and invariances in the environment, along with the learner's developing expectations about them, drive discrimination of the combination of the predictors that best capture that environment (Rescorla 1988). Interestingly, this is conceptually very similar to William James's (1890:488) suggestion that an infant first experiences the world as a 'blooming, buzzing confusion', and that the perception of variance leads her to learn to DISCRIMINATE its contents:

> the undeniable fact being that *any number of impressions, from any number of sensory sources, falling simultaneously on a mind* WHICH HAS NOT YET EXPERIENCED THEM SEPARATELY, *will fuse into a single undivided object for that mind*. The law is that all things fuse that *can* fuse, and nothing separates except

[1] In Rescorla's (1968) experiments, rats exposed to a high, random base rate of tones did not condition to the tone, but did condition to the experimental chamber.

[2] For modeling purposes, one might initially idealize this as making up no more than 'the environment', that is, $n = 1$.

> what must. … Although they separate easier if they come in through distinct nerves, yet distinct nerves are not an unconditional ground of their discrimination … The baby, assailed by eyes, ears, nose, skin, and entrails at once, feels it all as one great blooming, buzzing confusion; and to the very end of life, our location of all things in one space is due to the fact that the original extents or bignesses of all the sensations which came to our notice at once, coalesced together into one and the same space. (emphases in original)

Although James's 'blooming, buzzing confusion' is frequently mischaracterized in the literature—perhaps because the specifically discriminative conception of learning in which James situated these remarks is often ignored—for animals, at least, learning from expectation and error offers a fleshed-out account of the process through which the perception of variance can lead to learning about the world.

5. Prediction and language learning. The discovery that animals are perfectly capable of learning about predictive relationships even when they have no explicit access to the locus of their predictions contrasts with a critical assumption in the LPLA—and much of the language learning literature—that learned inferences can only be unlearned when explicit correction is provided (Baker 1979, Brown & Hanlon 1970, Marcus 1993, Marcus et al. 1992, Pinker 1984, 1989, 2004). If the logic of the LPLA were applied to rat learning, it would predict that rats could only learn about the relationship between a tone and an absent shock if they were provided with additional, explicit information about this relationship. Rescorla's—and countless other—experiments make clear that, for many species of animals, at least, this prediction is simply false.

Learning from prediction error is, of course, not the sole preserve of rats, pigeons, and dogs. Outside the domain of language, models that make assumptions similar to those just described have been successfully applied to the study of decision making, executive function, habitual learning, and response selection in humans (McClure 2003, Montague et al. 2004, Montague et al. 1996, Niv 2009, Schultz 1998, 2006, Waelti et al. 2001). Numerous behavioral studies have shown that human learning is sensitive to background rates at a high level of abstraction (for reviews, see Miller et al. 1995, Siegel & Allan 1996). In addition, a growing body of evidence provides compelling reason to believe that human children are sensitive to background rates in language learning tasks (Ramscar, Dye, & Klein 2013, Ramscar et al. 2011, Ramscar, Yarlett, et al. 2010; see also Saffran 2001, Saffran et al. 1996, Saffran et al. 1999).

Perhaps just as compellingly, there is now a substantial body of research showing that prediction is ubiquitous in language processing. As people listen to or read language, they build up a wealth of linguistic expectations, anticipating upcoming linguistic material at numerous levels of abstraction based on the structure and semantics of prior discourse (Altmann & Mirković 2009, Altmann & Steedman 1988, Balling & Baayen 2012, Chang et al. 2006, Kutas & Federmeier 2007, Levy 2008, MacDonald et al. 1994, MacDonald & Seidenberg 2006, Otten & Van Berkum 2008, Ramscar, Matlock, & Dye 2010, Tanenhaus & Brown-Schmidt 2008, Tanenhaus et al. 1995, Wicha et al. 2003). These findings suggest that indirect negative evidence is available to children, and thus that it may well play the same kind of role in their learning as it does in that of animals. Importantly, these findings suggest that paying closer attention to the predictive nature of children's learning can help us gain insight into the way linguistic understanding develops in learners.

6. A model of plural learning.

6.1. Overview. Given that children make linguistic predictions, and given too that they learn in response to prediction errors, an obvious question arises: are the mechanisms we have described sufficient to provide an account of the patterns of overregularization that have been observed in plural noun learning? To formally address this question, we constructed a model of the way a child might learn to name singular and plural objects.

6.2. WHY PLURALS PRESENT A PROBLEM. The question of what governs the inflection of linguistic forms has been a topic of heated debate in relation to the question of language learnability (McClelland & Patterson 2002, Pinker & Ullman 2002; see also Albright & Hayes 2003, Baayen & Moscoso del Prado Martín 2005, Clahsen 1999, Ernestus & Baayen 2004, Harm & Seidenberg 1999, Haskell et al. 2003, Joanisse & Seidenberg 1999, Justus et al. 2008, MacWhinney & Leinbach 1991, Marslen-Wilson & Tyler 2007, Pinker 1991, 1999, Pinker & Prince 1988, Plaut & Booth 2000, Plunkett & Marchman 1991, 1993, Prasada & Pinker 1993, Ramscar & Dye 2011, Ramscar & Yarlett 2007, Rumelhart & McClelland 1986, Taatgen & Anderson 2002, Tabak et al. 2010, Woollams et al. 2009). In the case of plural nouns, English-speaking children tend to overregularize irregulars—saying, for example, *mouses* instead of *mice*—and this behavior is rarely explicitly corrected. As children grow older, however, they come to produce only the adult form: *mice*. Since there is no obvious reason for them to stop saying *mouses*, it has been argued that this presents a logical puzzle: how could they LEARN to do this without feedback (Baker 1979, Pinker 1984, 2004; see also Clahsen 1999, Huang & Pinker 2010, Marcus 1993, 1995, Marcus et al. 1995, Pinker 1991, 1999, Pinker & Prince 1988, Prasada & Pinker 1993)?

In English, correct irregular plural marking is particularly difficult to acquire (Ramscar & Dye 2011), even in comparison to past-tense marking, another source of youthful error and the object of much prior study. This likely reflects the nature of the input. While irregular verbs are rare as types, they tend to have high token frequencies, such that in the Corpus of Contemporary American English (Davies 2009), the forty most frequent verb forms are all irregular. Moreover, in the Reuters corpus (Rose et al. 2002), just three irregular verbs (*be*, *have*, and *do*) account for fully a quarter of the attested verbs forms, with past-tense verb forms outnumbering base or present-tense verb forms. In learning the past tense, then, children are likely to encounter more past-tense verbs forms than uninflected forms, and more irregular past-tense forms than regular past-tense forms. Plurals are different: children generally encounter singular noun forms, and when they do encounter plural forms, they are highly likely to be regular. In the Reuters corpus, only around 30% of nouns occur in their plural form, and of these, the overwhelming majority in terms of both types and tokens are regular. This makes the learning problem substantively more difficult. However, the two problems may not be different in kind: as with the past tense, children's irregular plural production follows a U-SHAPED developmental trajectory, such that children who have been observed to produce *mice* in one context may still frequently produce overregularized forms such as *mouses* in another (Arnon & Clark 2011). Given the nature of the learning problem, there is much scope for experimental interventions to be made, and their effects to be measured, as children engage in the lengthy process of mastering plural forms (Ramscar & Yarlett 2007).

6.3. THE RESCORLA-WAGNER LEARNING RULE. The model described here was intended to have sufficient detail to allow predictions to be derived from the error-driven learning mechanisms we have outlined above, while being simple enough for the relationship between the mechanisms and the predictions to remain transparent. Plural learning was simulated using the learning rule from Rescorla and Wagner (1972), which treats learning as a process that enables a learner to better predict events in the world and, in particular, to weigh and assess the INFORMATIVITY of various cues in predicting relevant outcomes.

While the Rescorla-Wagner model cannot account for all of the phenomena observed in 'associative' learning, the model provides an accessible formalization of the basic principles of error-driven learning, and is sufficiently detailed to allow a straightforward testing of the analysis we present here. It should be noted, however, that the analysis is consistent with similar principles embodied in a wide range of learning models, in which equivalent simulations could be implemented (see e.g. Barlow 2001, Courville et al. 2006, Danks 2003, Dayan & Daw 2008, Gallistel 2003, Kruschke 2008, McLaren & Mackintosh 2000, Pearce & Hall 1980, Sutton & Barto 1998). Furthermore, because the model is mathematically very similar to a perceptron (Rosenblatt 1959), our employment of it allows for ready comparison with a popular discriminative approach in machine learning (e.g. Brill 1995, Collins & Koo 2005, Roark et al. 2007).

The Rescorla-Wagner model simulates changes in the associative strengths between individual cues C and an outcome as the result of discrete learning trials. If the presence of a cue or outcome X at time t is defined as present(X, t) and its absence as absent(X, t), then the predictive value V of a cue C_i for an outcome O_j after a learning event at time $t + 1$ can be stated as in 1.

(1) $V_{ij}^{t+1} = V_{ij}^{t} + \Delta V_{ij}^{t}$

The change (Δ) in the predictive value of C_i after t can be defined as in 2.

(2) $$\Delta V_{ij}^{t} = \begin{cases} 0 \; if\, ABSENT\,(C_i,t) \\ \alpha_i\beta_1(\lambda - \sum_{PRESENT(C_j,t)} V_{ij}) if\, PRESENT\,(Cj,t)\;\&\; PRESENT(O,t) \\ \alpha_i\beta_2(0 - \sum_{PRESENT(C_j,t)} V_{ij}) if\, PRESENT\,(Cj,t)\;\&\; ABSENT(O,t) \end{cases}$$

Thus, learning is governed by a DISCREPANCY FUNCTION where λ is the total value of the predicted event (i.e. the maximum amount of associative strength that an outcome j can support; here it is simply set to 1, indicating that an event is fully anticipated), and V_j is the predictive value for outcome j given the set of cues present at time t.

In trials in which there is POSITIVE EVIDENCE—that is, in which expected outcomes DO occur—the Rescorla-Wagner learning rule produces a negatively accelerated learning curve (the result of events being better predicted, which reduces the discrepancy between what is expected and what is observed) and asymptotic learning over repeated trials (as events become fully predicted). Conceptually, this happens because the model embodies the idea that the function of learning is to align our expectations with reality, and the better that alignment becomes over time, the less we need to learn (Anderson & Schooler 1991, Ebbinghaus 1913).

In trials in which there is NEGATIVE EVIDENCE—that is, in which an expected outcome fails to occur—λ_j (the expected outcome) takes a value of zero because it did not occur. In such cases, the discrepancy function ($\lambda_j - V_j$) produces a negative value, resulting in a reduction in the associative strength between the cues present on that trial and the absent outcome j. Conceptually, these prediction errors can be thought of as violations of expectation that allow the model to learn from negative evidence.

The total amount of predictive (cue) value any given outcome can support in learning is finite. (Informally, we can think of this as capturing the idea that if predictive confidence keeps rising, it must eventually reach a point of relative certainty.) As a result, cues compete with one another for relevance, and this produces learning patterns that often differ greatly from those that would arise by simply recording the correlations between cues and outcomes (i.e. simply tracking base rates—a common misconstrual of learning); see Figure 4.

The rate of change (Δ) at t is determined by two factors: the overall learning rate β (where $0 \leq \beta \leq 1$), and the individual saliency of cues α_i (where $0 \leq \alpha \leq 1$). Because we were interested in how learning affects the relative value of cues, α_i was set to 1, eliminating its influence on our simulations. Lambda was set at $\lambda = 100\%$ for each word, and the beta β_j learning rate took the default value in the Rescorla-Wagner implementation contained in the ndl package (a library of the R statistical programming language).

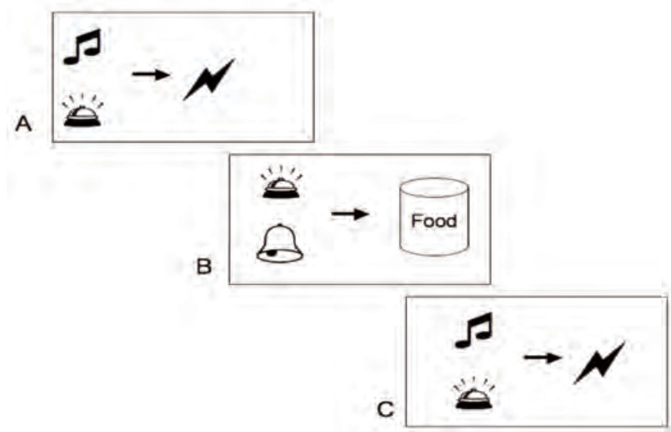

FIGURE 4 (TOP). Consider a rat being conditioned to expect either shocks or food. A light shines just before both food and shocks (A, B, C), while an accompanying bell only ever sounds before food (B), and an accompanying tone only ever sounds before shocks (A, C). In order to best anticipate when shocks and food will be forthcoming, the rat must learn to attend to the cues that are most informative about each outcome. In trial (A), it learns that both the tone and the light predict shocks. Because the light indiscriminately predicts both shocks and food, the rat incorrectly predicts a shock in trial (B). As a result, the strength of the association between light and shock decreases, even though no shock is present on this trial. The converse occurs in trial (C), when light incorrectly predicts food. In this trial, the strength of the light–food association decreases.

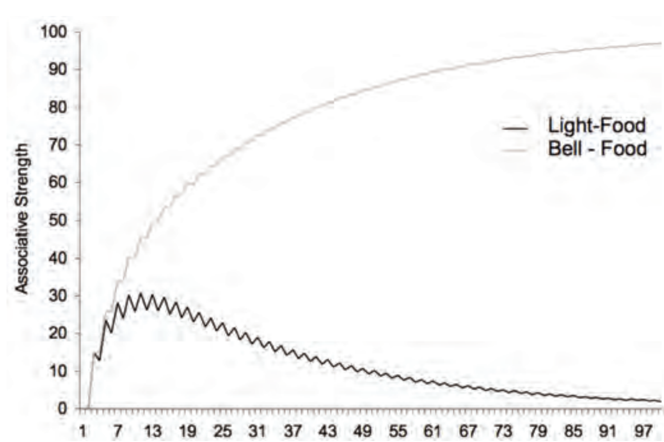

FIGURE 4 (BOTTOM). A simulation of error-driven learning of the relationship between bell and food and light and food in this scenario. The graph shows the cue values developing in the Rescorla-Wagner (1972) model. The errors produced by light cause it to lose out in cue competition with bell so that the association between bell and food is emphasized, while the association between light and food is devalued. Though bell and food cooccur with exactly the same frequency as light and food in this scenario, learning effectively dissociates light as an uninformative cue.

6.4. Implementation of the model. Our simulations make three key assumptions about the learning environment.

- Children do not learn their native languages in formal teacher-pupil settings (Chomsky 1959, Pinker 1984).
- Children learn words, at least initially, by hearing them used in context (Smith & Yu 2012, Tomasello 2003).
- The distribution of error in the early linguistic environment—that is, the combined value of both positive and negative evidence—favors the appropriate mappings. For example, a child learning the word *mice* will hear the word used in a way that makes it most informative about mice, or depictions of them, and must learn to associate the appropriate cues in the environment—mouse-things—with the word (Quine 1960, Wittgenstein 1953). Conceptually, this assumption reflects the idea that adult speakers use language in informative ways, and hence, that a mouse ought to be more informative about the English word *mouse*, and mice more informative about the word *mice*, than they are about other words such as *rat*, *chair*, *moon*, or *allele*.

Notably, the way our model learns from this environment differs markedly from many previous models, which envisage a child as learning to transform a 'word stem'—*cat*—into an inflected form—*cat***s** (cf. MacWhinney & Leinbach 1991, Plunkett & Marchman 1991, 1993, Rumelhart & McClelland 1986). By contrast, our model learns to predict word forms from semantics (i.e. the environment), a process that much more closely approximates the situation of the child learner (see also Ambridge et al. 2009, Andrews et al. 2009, Cottrell & Plunkett 1994, Durda et al. 2009, Goldberg 2011, Moscoso del Prado Martín et al. 2004, Ramscar 2002, Ramscar & Yarlett 2007).

In addition, our simulations were shaped by a number of working assumptions about the nature of the learning task:

(i) The model assumes that when a child is asked to name a picture of mice, the child has some prior experience of mice, and this results in activation of the word *mice*, because this is the phonological form the child has learned to associate with the semantic representation of mice (Meyer & Schvaneveldt 1971, Ramscar & Yarlett 2007). What the child actually says, however, is contingent on both the strength of the representation of *mice*, and the degree to which other forms interfere with *mice* production.[3]

(ii) The model assumes that a child must learn to discriminate between single and multiple items in naming, and that set size serves as a cue to whether forms are singular or plural (Ramscar et al. 2011).

(iii) The model assumes that the phonological forms of regular singular and plural (+*S*) nouns are distinguished temporally, by the occurrence (in plurals) or nonoccurrence (in singulars) of a sibilant after a common form (see Ramscar & Dye 2011 for converging evidence). While this ignores the many differences between the single and plural forms of regular nouns—such as different sibilant allomorphs, coarticulation effects, and so forth—it captures the idea that regular plurals resemble one another with respect to their key pho-

[3] Recent discussions of reinforcement learning distinguish between model-based learning, in which a model—or map—of the states that best predict relevant environmental information is acquired, based on an intermediate representation of candidate actions, and model-free learning, in which learning simply reflects the difference between actual and expected events (see e.g. Gläscher et al. 2010). We assume that language learning is a model-based process.

netic indicator of plurality (the sibilant), whereas irregular plurals resemble neither regular plurals nor one another (Ramscar & Dye 2011). For a child who has heard a large number of regular plurals and relatively few irregulars, and who is still learning to discriminate many of these items, this knowledge will support the expectation of a sibilant after a common form, leading to correct regular plural production (*rats*), but interfering with irregular plural production (*mouses*).

(iv) The model assumes that the strength of this expectation arises out of two factors:
- the degree to which the other word forms the child knows about are activated by the cues present on mice trials (as Fig. 3 shows, learners will come to ignore these cues over time as they better discriminate specific items).
- the overall learned values of those other forms.

These assumptions reflect the idea that children will be learning to categorize objects at the same time that they are learning to name them (Swingley & Aslin 2007), and that early exposure to mice in the context of hearing *mice* will not only support mice as an informative cue to *mice*, but will also support less well-discriminated cues (such as stuff, or multiple-items, or mousiness). Until these alternatives are discriminated, they will interfere with the production of *mice*, as they will serve to cue other, competing forms the child has learned (such as other plurals, or the singular form *mouse*).

6.5. SIMULATING PLURAL LEARNING. The model simulates how cues to the irregular plural *mice*, its singular form *mouse*, and a set of twenty-eight other nouns that have regular plural forms are learned and discriminated. These forms were represented in proportion to the frequency distribution of singular and plural noun forms in English, such that *mouse* was twice as frequent as *mice*, and the proportion of singular to plural forms of the regular nouns was 10 : 6 (see Ramscar & Dye 2011 for detailed analyses).

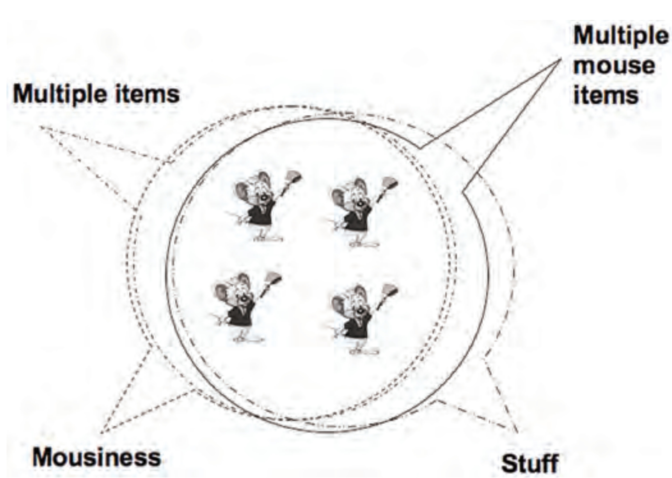

FIGURE 5. Four cues that will all be supported by a child's exposure to the word *mice* in the context of mice. Although these cues always cooccur with the word *mice*, their covariance with other singular and plural nouns—and thus the distribution of error associated with them—differ such that the balance of evidence favors the multiple-mouse-items → *mice* mapping. (Note that while the cues are separated out for explanatory convenience here, they could be ranges of values on continuous perceptual dimensions as far as the model is concerned.)

Figure 5 illustrates the four environmental cues that consistently covary with *mice*, and that are most relevant to (and informative about) plural mouse naming. These cues represent the idea that over the course of learning, information about the world—initially a mass of undifferentiated stuff—is gradually discriminated, as learning uncovers the relevant cues to objects, events, affordances, and so forth. At the outset of learning, all and any kind of 'stuff' in the world is potentially informative about concrete nouns like *mouse* and *mice*, such that learning to discriminate the correct cues to *mouse* and *mice* involves discriminating the 'mousey stuff' associated with *mouse* and *mice* from the other kinds of stuff associated with nouns. At the same time, learning to discriminate *mice* from *mouse* requires discriminating the specific mousey stuff that best predicts *mice* as opposed to *mouse* (i.e. the presence of multiple mouse objects as opposed to a single mouse object). Finally, learning to use *mice* correctly simultaneously also involves learning to discriminate the appropriate KIND of multiple items associated with *mice* (mouse-items) from other sets of items in the world.[4]

Crucially, because all four of these cues—stuff, multiple-items, multiple-mouse-items, and mousiness—are present whenever mice are seen and *mice* is heard, all of these cues will receive identical support, meaning that a child could never hope to discriminate the cue(s) appropriate to naming mice on the basis of positive evidence alone. Because the DISTRIBUTION OF ERROR associated with each cue differs, however, children should still be able to learn the correct association between *mice* and multiple mouses. This becomes clear when we consider the background rates of each cue. Since *mice* is frequently heard when mouse-items are present (e.g. 'look at those mice!') and infrequently when they are not, there will be little error in the relationship between mouse-items and *mice*. Conversely, since there will be many occasions when stuff and other items are present in the child's environment and *mice* is NOT heard (e.g. *cups* or *daddy* might be heard instead), these cues will generate a great deal of error as cues to *mice*. Similarly, whenever a single mouse is present, and *mouse* is heard, the presence of mousiness in the absence of multiple-mouse-items will generate erroneous expectations of *mice*, which will allow the meaning of *mice* to be discriminated from the meaning of *mouse*. Thus from a discriminative learning perspective, the fact that stuff, multiple-items, mousiness, and mouse-items provide identical positive evidence for *mice* is not an impediment to learning because their background rates—and thus, the negative evidence each provides—differ dramatically (Figure 6).

In the model, overregularization occurs on mice trials because the cues to stuff and multiple-items, which gain support when *mice* is heard in the presence of mice, also gain support whenever the (usually regular) labels for other plural items are learned. Because of this, further encounters with mice will lead not only to the expectation of the label *mice*, but also to the expectation of other noun forms (Figure 7; see also Ramscar & Yarlett 2007), leading to competition between the responses. This competition yields an initial bias toward overregularization errors, a product of the distribution of regular and irregular plural forms in English and the cues to them in the environment.

To simulate how response competition will affect the production of correct irregular forms over learning, we examined the likelihood that the model would produce the label *mice* when presented with mice at each point in learning (thereby allowing for a

[4] It is worth noting that while for the purposes of exposition, we describe these different dimensions in discrete terms, we assume that these dimensions will be largely undifferentiated prior to learning. The degree to which they are actually experienced as discrete (i.e. the degree to which they are actually discriminated from one another) will depend on what has actually been learned up to that point. The current learned status of any 'discrete' response can only be evaluated in relation to an overall system of responses.

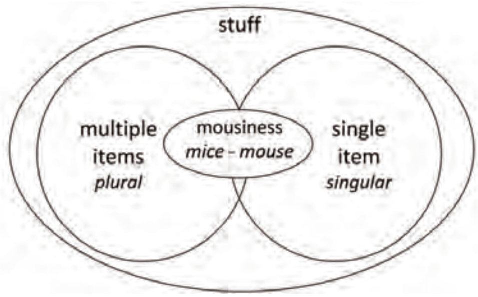

FIGURE 6. The relative specificity of the four cues: while the generality of the less specific cues (stuff and mousiness) will support their positive reinforcement early on in learning, that generality will also generate a high degree of error relative to the more uniquely informative cues. As a result, the influence of less specific cues on more specific responses will wane over time.

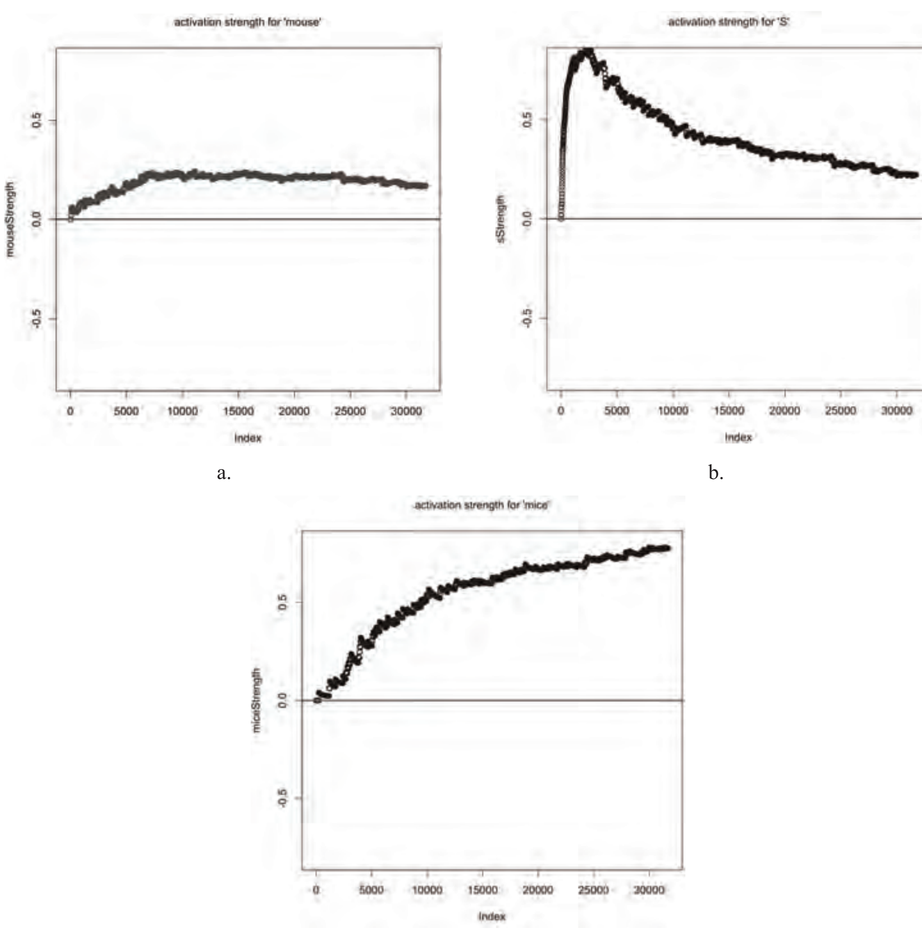

FIGURE 7. The relative strength of each response across learning (learned strengths are represented by the height of each line): (a) *mouse*, (b) +*S*, and (c) *mice*. Early in learning, less specific cues that are shared across the responses generate interference that then diminishes as these uninformative cues are unlearned over cue competition.

fully incremental evaluation of the model's predictions to be made; cf. McCauley & Christiansen 2011). To estimate these response propensities, we calculated the activation each response received from the cues to *mice* and then calculated an INTERFERENCE VALUE—the activation of *mouse* plus the activation of *+S* at the end of a common form—which was subtracted from the activation of the appropriate response, *mice*. If the interference value is greater than the activation of *mice*, this subtraction yields a negative value, indicating a bias toward overregularization. Conversely, when the activation of *mice* is greater than the summed activations of the competing responses, the bias is to produce the correct form (Figure 8).

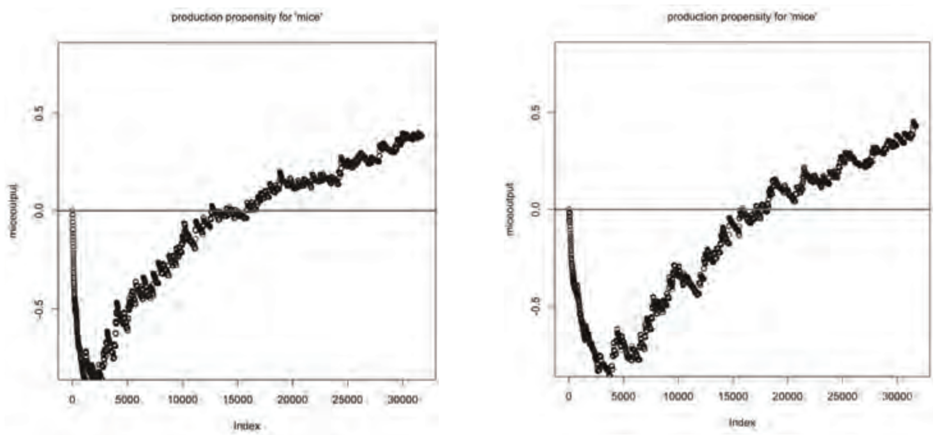

a. b.

FIGURE 8. Panel (a) plots development of irregular plural production in the model, showing its response propensity at each point in time when the cues to *mice* are present. Negative values favor overregularized responses; positive values favor correct irregular plural responses. To illustrate the relative robustness of this result, panel (b) plots the same pattern of development in a second implementation of the model in which the ratio of regular singular forms to plurals was 70 : 30, as observed in the Reuters corpus. Consistent with U-shaped learning, both models produced initial periods in which correct forms precede overregularizations.

Although this simple model ignores a range of factors that will influence specific instances of overregularization (e.g. linguistic context also influences the predictability—and overregularization—of irregular forms; Arnon & Clark 2011), it successfully captures how the tendency toward overregularization first arises as a result of the frequency of different word forms and the frequency and distribution of the cues to them, and then later diminishes as a function of the distribution of error among those same cues. (The R code required to implement this version of the model is included in the appendix; exploration will reveal that so long as a representation of the learning problem respects the distribution of cues and lexical outcomes, this pattern of performance is robust.) This developmental trajectory exhibits the classic U-shaped learning pattern—where production mixes correct and incorrect forms prior to settling on the correct form—previously noted in the development and resolution of children's overregularization (Brown 1973, Marcus et al. 1992).

6.6. SIMULATING PLURAL LEARNING WITH NATURALISTIC INPUT. In order to test the scalability of the model as well as its performance when exposed to naturalistic input, we extracted nouns from a corpus of child-directed speech taken from the CHILDES database (MacWhinney 2000). In order to compensate for data sparsity resulting from the low frequency of irregular nouns in individual corpora, the entire American English portion of CHILDES was aggregated after being reordered chronologically by the age of the target child in each recording session.[5] To maintain a naturalistic developmental trajectory, files that included speech directed to multiple target children of different ages were excluded. Each noun token was extracted from the resulting aggregated corpus and lemmatized, using the CELEX database (Baayen et al. 1995), and then attached to a corresponding cue bundle. For example, the singular noun *cat* was attached to the cue bundle of stuff, single-item, cattiness, and cat-item, while the plural noun *mice* was attached to stuff, multiple-items, mousiness, and mouse-items.

With the order of the aggregated corpus preserved, each utterance was treated as a separate learning trial, with the cue bundles corresponding to each noun in the utterance treated as a single compound conditioned stimulus, and each noun's word form treated as a separate unconditioned stimulus. As an example, the utterance 'the cat chases the mice' would result in the compound stimulus of stuff, single-item, multiple-items, cattiness, cat-item, mousiness, mouse-items, which the word forms *cat* and *mice* would be conditioned to. The alpha, beta, and lambda parameters of the model were identical to those used in the initial simulations.

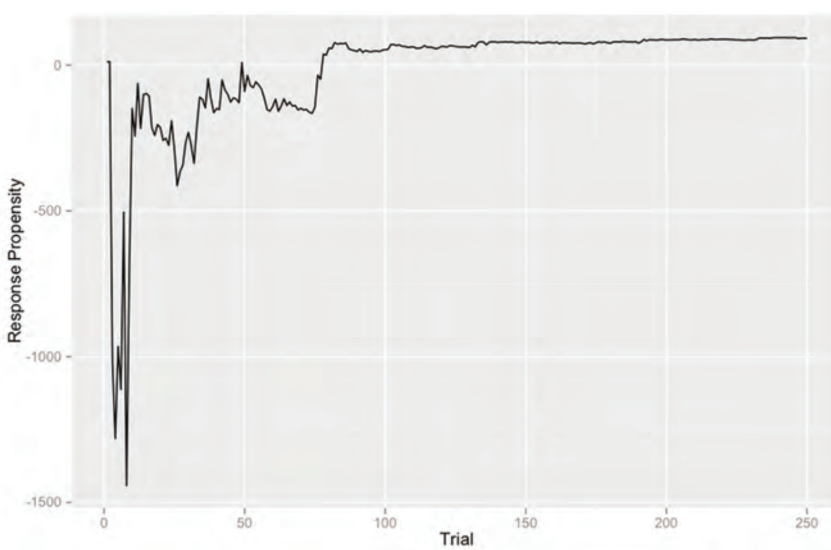

FIGURE 9. Response propensity of the model during a single pass through the entire American English portion of the CHILDES database. Negative values favor overregularized responses; positive values favor correct irregular plural responses. The first 250 production attempts are shown (one trial every 1000th utterance).

This version of the model allowed for fully incremental predictions to be made. At each point in learning, attempts to produce the plural form *mice* were simulated by calculating the difference between the activation of *mice* (given the cues stuff, multiple-items, mousiness, and multiple-mouse-items) and the activation of *mouse* and +*S* (given the same cues), based on the learned values of the cues and responses at any given point

[5] The idea of an aggregated CHILDES corpus, ordered by the target child age in each recording file, was originally proposed by Morten Christiansen in the context of a different modeling project.

in time. A negative value on this difference measure represents a higher association for *mouse* and +*S* than for *mice*, indicating a propensity to overregularize (i.e. produce the singular form + sibilant combination *mouses*).

When trained on a naturalistic data set, the model again produces the U-shaped pattern of learning observed in the idealized simulation (Figure 9). Here again, the initial tendency to overregularize arises out of the frequency of different word forms and the frequency and distribution of the cues to them, before resolving itself as a result of the distribution of error among these same cues.

6.7. GENERATING NOVEL PREDICTIONS FROM THE MODEL. The formal properties of the model allow for detailed predictions to be made about the circumstances that might lead to an increase or decrease in the rate of overregularization in young children, depending on their prior learning. Figure 10 illustrates the effect of exposure to the same mixture of regular and irregular plurals at different junctures in the model's training: early in learning and then later on in learning.

Conceptually, these interventions might be expected to have a broadly similar effect: given that children are initially learning to discriminate between the semantic cues to regulars and irregulars, they should have some expectation of irregulars on regular trials. Thus whenever children incorrectly expect an irregular form, this will result in prediction error (negative evidence), which will raise the error rate of unreliable cues (such as stuff and multiple-items). Over the long run, this will help young speakers discriminate the appropriate semantic cues to irregulars. This is the big picture. Importantly, however, because discrimination learning is always systematic—that is, the overall effects of learning and unlearning can only be established in relation to whatever else a learner knows—the LOCAL effect of such interventions can differ dramatically depending on how they interact with the learner's prior knowledge. This idea is easily captured by looking at how exposure to regular plurals can have different effects on overregularization at different stages in learning.

In the model, production of a given form is the result of a competitive process based on the degree of support for each possible response given the evidence available, and the overall degree to which a given response has already been learned. Because of the different frequencies of regular plural forms, and irregular singular and plural forms, irregular plurals are learned and discriminated more slowly than the forms they compete with. Early in plural learning, the rate at which support for the +*S* regular response is growing far outstrips that at which the (erroneous) cues supporting that response are weakening, resulting in an increase in the likelihood that an overregularized form will be produced (Fig. 10a). As learning about these other responses begins to asymptote, however, and as the cues to *mice* become better discriminated, the exact same sequence of training trials will yield the opposite result, and exposure to regulars will actually increase the likelihood of a correct irregular response (Fig. 10b). Finally, at the point that cue competition has effectively eliminated the influence of the erroneous cues, the trial-to-trial effects of learning will have little impact on the likely response, as support for the +*S* response is now so weak that local fluctuations will not affect production (Fig. 10c).

It is worth noting that this pattern of learning can potentially arise in any situation where the items that need to be discriminated from one another differ greatly in their frequency. It also further underlines the point that learning is SYSTEMATIC, and depends not only on the information currently available to the learner but also on the information the learner has accrued through previous experience. Talking about the 'information' available to a learner makes sense only in relation to what the learner already knows, because it is that prior knowledge that determines both HOW informative any new 'information' is and IN WHAT WAY it is informative.

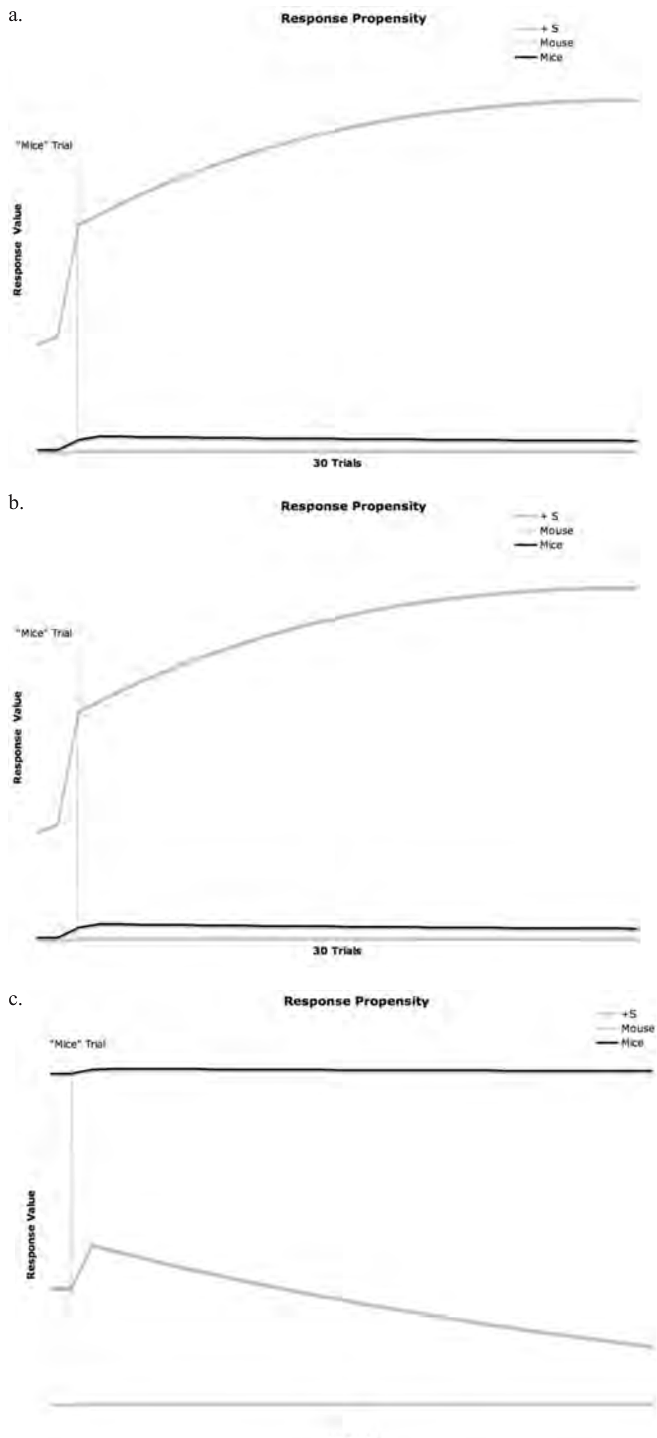

FIGURE 10. The effects of learning about mice (i.e. the effect of positive evidence about the cues to *mice*) at different stages of the simulation depicted in Fig. 6. The sequence of training trials is identical across all three plots and comprises a regular plural trial, followed by a mice trial, followed by an additional twenty-eight regular plural trials. Each plot line represents the strength of cues on each trial (the summed value of the cues normalized by the learned strength of each response) and thus represents the relative likelihood of a particular response at each point in learning.

7. Training experiment.

7.1. Overview. Our learning model predicts that as a result of the distribution in English, learning about regular plurals will have different behavioral consequences for children's irregular plural production, depending on each child's prior experience. Training on regular plurals will increase overregularization rates for irregular plurals early in learning, but DECREASE rates of overregularization later on. To test this counterintuitive prediction, we recruited four- and six-year-old children to take part in a simple training experiment. We employed a semantic old-new task to expose children to plural forms, and a test-train-test paradigm to compare baseline rates of overregularization with posttraining rates (Ramscar & Yarlett 2007).

7.2. Participants. Thirty-eight four-year-old and forty six-year-old children were recruited from a database of volunteers living in the vicinity of Palo Alto, California. The average ages were four years and six months for the four-year-olds, and six years and seven months for the six-year-olds. Children of these ages have fully mastered regular plural inflection (Brown 1973, de Villiers & de Villiers 1973), but often overregularize irregular plural nouns (Graves & Koziol 1971, Ramscar & Yarlett 2007). The children were randomly assigned to two groups: an experimental condition and a control condition.

7.3. Methods and materials.

Pretest. Both groups of children were pretested on plural production that exposed them to correct singular forms and established a baseline rate of overregularization for each child. In the pretest, the children were asked to help a cookie monster puppet learn to name a series of plural nouns. The children were shown pictures of six regular and six irregular nouns, first singular and then plural depictions that were presented on a laptop computer. As each picture was shown, the children were asked to tell the monster the names of these items (i.e. they were made to retrieve the phonological response to the semantic cue). Regardless of the plural form the children produced, they were provided with encouraging feedback from the puppet. The six irregular items in the test were MOUSE-MICE, CHILD-CHILDREN, SNOWMAN-SNOWMEN, GOOSE-GEESE, TOOTH-TEETH, and FOOT-FEET; the six regular semantic matches were RAT, DOLL, COW, DUCK, EAR, and HAND. These items were chosen from each of the families of irregular plurals that young children reliably learn to master. Although children in this age range tend to overregularize these irregular plurals, they have reliable knowledge of their correct forms (Ramscar & Yarlett 2007).

Experimental condition. In the experimental condition, children were required to exercise their knowledge of plural nouns by telling a cookie monster whether depictions of regular plural noun-objects had the same name as items they had previously named in the pretest. The children were asked to tell the cookie monster 'yes' or 'no' to indicate that they had or had not, respectively, already seen these depictions. If the child saw something that had the same name as an item in the pretest, the child was asked to say 'yes', and if it did not have the same name as an item in the pretest, the child was asked to say 'no'. When a set of objects appeared, the experimenter asked the child to 'Look at those—did cookie monster see any of those before?'. Children who did not spontaneously respond were prompted 'Did cookie see any of these? Yes? No?'. If no response was forthcoming, the experimenter proceeded to the next item. Half of the presented items were new depictions of the regular items in the pretest, and half were foils. The children were thus tested on twelve new and twelve old items per block.

Notably, the absence of overt naming responses by children was intended both to reduce the effect of perseverative biases on posttest performance, and to subject our hypotheses about the effect of implicit expectation on children's discrimination learning to a particularly stringent test (see also Ramscar & Yarlett 2007). By simply having children provide 'yes' or 'no' answers in the training phase, we could increase our confidence that any changes to children's underlying REPRESENTATIONS of the plural forms of the objects they encountered in training were brought about by the implicit expectations that those objects evoked (i.e. since we were interested in the development of children's knowledge, we wished to limit the influence of factors that did not relate to that knowledge as best we could). All depictions of the 'old' items in training were novel, which required children to make categorization judgments to generate the correct answers, and children were told to base their category judgments on whether the items would be 'called by the same name' as previously presented items. Because words' phonological representations are cued by their semantics, these measures could be expected to result in reinforcement of the regular plural forms, as well as prediction errors and latent learning (Meyer & Schvaneveldt 1971). As Fig. 7 indicates, the behavioral consequences of this latent learning should vary depending on the prior experiences of learners.

Control condition. In the control condition, children were shown six color slides after the pretest, and then asked to tell the cookie monster whether they had seen that particular color, in an old-new task with an equal number of foils. To avoid cuing any notion of plurality, the colors were presented as solid blocks filling the screen. The total time to complete this condition was equated to that of the experimental training condition.

Posttest. Both sets of children then completed a posttest identical to the pretest.

7.4. Results and discussion. Children's performance in these tests supported the model's predictions. A 2 (pre- to posttest) × 2 (age) × 2 (condition) repeated-measures ANOVA analysis of the overregularized forms produced by each child in the pre- and posttests revealed a significant interaction between age, training type, and pre- to posttest performance ($F(1,58) = 4.701$, $p < 0.05$), and a significant interaction between age and pre- to posttest performance ($F(1,58) = 6.329$, $p < 0.001$). The older children in the experimental condition improved their irregular production, overregularizing less in the posttest ($M = 1.5$ overregularizations out of six) than the pretest ($M = 2.25$; $t(14) = 2.665$, $p < 0.01$), whereas rates of overregularization increased in the younger children (pretest $M = 2.54$; posttest $M = 3.27$; $t(14) = 1.761$, $p < 0.02$).There was little change in the performance of either age group in the control condition (see Figure 11).

The same results were obtained when the data were coded as per Ramscar & Yarlett 2007: 0 = failure to respond, 1 = overregularization, 2 = uninflected form, 3 = correct irregular. The same repeated-measures ANOVA revealed significant interactions between age, training type, and pre- to posttest performance ($F(1,58) = 4.996$, $p < 0.05$), and age and pre- to posttest performance ($F(1,58) = 11.559$, $p < 0.001$). In the experimental condition, older children's improvement ($t(15) = 2.992$, $p < 0.01$) and younger children's decline were both significant ($t(15) = 2.374$, $p < 0.05$).

Thus testing memory for regular plural nouns led to six-year-olds overregularizing plurals significantly less in the posttest, whereas the same training had the opposite effect on younger children. Testing memory for color words had no effect on either group. In line with the counterintuitive predictions of the model, then, the ability of the older children to produce plurals like *mice* and *feet* improved with training, even though none of these labels were actually present in the training trials.

8. General discussion. To the extent that the results we present here are surprising, it may be due to common misunderstandings of the way learning works (Rescorla 1988) and particularly to how prediction error provides a rich source of negative evi-

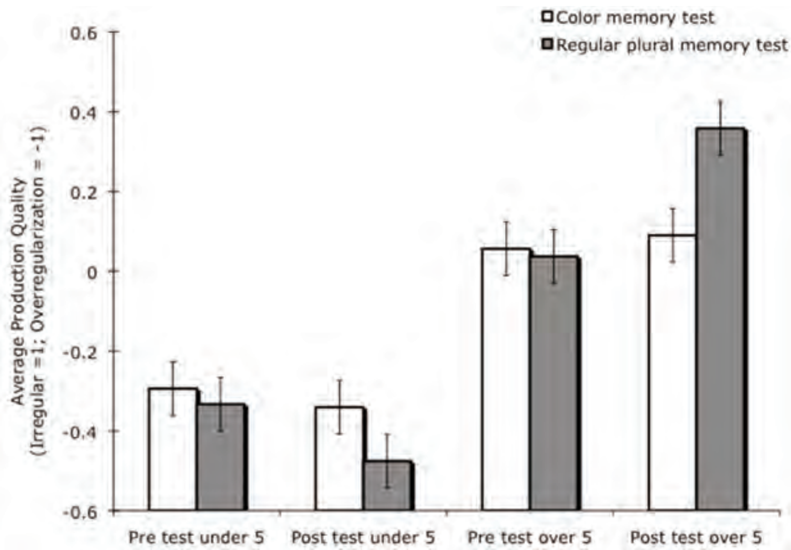

FIGURE 11. Pre- and posttest performance by age and condition. The data are plotted as the number of correct forms minus overregularized forms averaged across each pair of trials. Error bars denote standard error of the mean.

dence to learners. Overwhelmingly, research into language learning has preoccupied itself with the observable: with what a child hears or sees.[6] The underlying assumption has been—and largely remains—that a child can only learn about what is directly in front of her. This assumption is inconsistent with much of what we understand about animal (and human) learning.

While the idea that learning about a word can be thought about in terms of a 'single exposure' is common in the language learning literature, in formal theories of learning THERE IS NO SUCH THING AS LEARNING IN ISOLATION. Discrimination learning is SYSTEMATIC: it is a property of systems (see also Ramscar, Dye, & Klein 2013). What this means is that the learning that occurs at any given instant (on a trial in a learning experiment, or from 'a single exposure' to a word) is wholly contingent on what has already been learned in a given system—that is, everything the learner has already been exposed to—and can be influenced by anything else that a learning system might subsequently be exposed to (Rescorla 1988).

Because many researchers have assumed that children learn from 'positive evidence' alone (e.g. Brown & Hanlon 1970, Pinker 1984, 2004), linguistic theory has been guided by constraints imposed by the logical problem of language acquisition (Johnson 2004) and Gold's demonstration of the limitations of learning without negative evidence (Gold 1967). As Gold himself noted, however, his proof applied to an unrealistic formal model of language (Johnson 2004), which suggested EITHER that only the most

[6] This preoccupation is not the preserve of language researchers, but rather it is widespread in cognitive psychology. For example, the finding that testing for knowledge robustly improves the accuracy of its encoding in students has a clear parallel with the findings we report in children here (Roediger & Karpicke 2006, Karpicke & Roediger 2008, Karpicke & Blunt 2011). However, the mechanisms that give rise to 'testing effects' are poorly understood (see Roediger & Butler 2010 for a review). We suggest that attempts to explain testing effects could be much improved by conceiving of the memories under test as related—and even competing—components within larger systems of learned knowledge (i.e. in the same way as children appear to treat noun plurals).

trivial class of languages is learnable OR ELSE that children have access to negative evidence 'in a way we do not recognize' (Gold 1967:453). Since Gold's time, it has become clear that language processing involves prediction at every conceivable level (see Ramscar, Yarlett, et al. 2010 for a review) and that processes responsive to prediction error are ubiquitous in learning. It is also clear that the information available to children in the structure of linguistic distributions is evidently far richer than has traditionally been supposed (see Baayen et al. 2011, Landauer & Dumais 1997, McCauley & Christiansen 2011, Ramscar & Dye 2011, Reali & Christiansen 2005). Our results suggest that these predictive processes, in conjunction with the learning mechanisms that they drive, enable children to correct their own mistakes in learning language. It would seem that there simply is no logical problem in the way that children who say *mouses* manage, without explicit correction, to grow into adults who say *mice*.

In light of this, it is worth clarifying several points about the work described here, and in particular, the learning model used in these simulations. As we noted at the outset, the model we employed is not a new one. And, while it has limitations, these limitations are well known and did not prevent the model from serving the purpose of simulating and successfully predicting behavior in our task, suggesting that even stubbornly puzzling aspects of language learning may still be consistent with well-understood learning processes. This last point is important. When it comes to fitting behavioral data, the Rescorla-Wagner model is arguably more successful than any other learning formalism in the history of psychology (Miller et al. 1995, Siegel & Allan 1996). Further, as we noted earlier, there is much evidence that the mechanisms proposed by the model are neurally plausible (for a review, see Schultz 2010).

Moreover, the model is not confined to mere data fitting: Roberts and Pashler (2000) have argued, convincingly, that models need to be evaluated against data that they CANNOT be simply fit to, and that the clearest test case is to have the model make falsifiable predictions that can be evaluated empirically:

> Quantitative theories with free parameters often gain credence when they closely fit data. This is a mistake. A good fit reveals nothing about the flexibility of the theory (how much it cannot fit), the variability of the data (how firmly the data rule out what the theory cannot fit), or the likelihood of outcomes (perhaps the theory could have fit any plausible result), and a reader needs all three pieces of information to decide how much the fit should increase belief in the theory. The use of good fits as evidence is not supported by philosophers of science nor by the history of psychology; there seem to be no examples of a theory supported mainly by good fits that has led to demonstrable progress. (Roberts & Pashler 2000:358)

The Rescorla-Wagner model has generated a number of successful predictions in regard to animal learning (Kamin & Gaioni 1974, Kremer 1978), and the simulation and experiments reported not only show that the model (and our theory) can generate and gain support from this kind of 'strong testing' in the domain of human learning, but also that it can do so in the domain of language learning.

Moreover, this is not the only strong test of the model in this domain: Ramscar, Yarlett, et al. 2010 shows that the model correctly predicts very different patterns of performance in category learning, depending on the temporal sequence of category labels and exemplars (see also Ashby et al. 2002). The model has also lent insight into how to optimally sequence information to facilitate color and number learning in two- and three-year-olds, and verbal rule learning in a card-sorting task with the same age group (Ramscar, Dye, et al. 2013, Ramscar et al. 2011, Ramscar, Yarlett, et al. 2010). In a particularly provocative set of results, Ramscar, Dye, and Klein (2013) show that while the model successfully predicts toddlers' behavior in a cross-situational word learning task, when a sample of developmental psychologists specializing in language learning were asked to predict the children's behavior, their intuitive predictions were consistently wrong. While the psychologists correctly predicted undergraduate performance on the task, this varied systematically from that of the two-and-a-half-year-olds. These successful tests of the model's surprising—and falsifiable—predictions on different aspects of language learning are worth noting both because they serve as an important check on our intuitions, and because, as Roberts and Pashler (2000) point out, it is comparatively rare to find instances of models being used to generate novel empirical predictions in psychology and linguistics.

This leads us to another advantage Rescorla-Wagner offers: simplicity. As Roberts and Pashler (2000) note, the more free parameters a model employs, the less clear its predictions are: the more the danger of overfitting grows, and the less falsifiable the model becomes. The implementation of Rescorla-Wagner we used here has one free parameter: a learning rate, which we held constant throughout the simulations. A number of recent studies have shown how simple models based on the Rescorla-Wagner learning rule often outperform more complicated (and more recent) models when it comes to fitting and predicting human data. For example, Gureckis and Love (2010) present evidence that a simple Rescorla-Wagner implementation produces better fits to human sequential learning data than a more complex model designed specifically to simulate this task (the simple recurrent network; Elman 1990).

Perhaps even more surprisingly, Baayen and colleagues (2011) show that when the task of learning is analyzed in terms of discrimination, a version of the Rescorla-Wagner model that allows for a great number of learned weights to be estimated efficiently (Danks 2003), trained on a relatively 'small' linguistic corpus (~11 million two- and three-word phrases), provides good fits to human data on a wide range of effects documented for lexical processing, including frequency effects, morphological family size effects, and relative entropy effects. For monomorphemic words, the model provided excellent fits with NO free parameters, and for morphologically complex words, Baayen and colleagues had only to add a few free parameters to enable the model to fit a broader range of data more closely and parsimoniously than other models in the literature that were designed specifically for the task (e.g. Norris 2006). The model also captures frequency effects for complex words and provides good fits to data revealing phrase frequency effects, despite it not having explicit representations of either complex words or phrases (Baayen et al. 2013).

We suggest that these findings are representative of a more general—and very successful—trend emerging in computational approaches to learning: that of focusing on understanding the structure of the learning task, and then using relatively simple but effective learning algorithms to DISCOVER structure in data (Halevy et al. 2009, Recchia & Jones 2009), rather than seeking to second-guess the structure of those data in advance. Our own findings can be seen as illustrating the merits of looking at human development from the same perspective: seeing a child as equipped by nature to discover the structure of the world by discrimination learning (for further discussion, see Ramscar 2010).

For this approach to work, it is essential that the relationship between the learner and the world be properly understood in terms of the way that the information available to learners is structured. Children acquire language in context, usually without any explicit instruction, and as such, they may never encounter situations in which forms like *walked* or *mice* are explicitly derived in real time from *walk* and *mouse* (Tomasello 2003). Yet the assumption that this rote conjugation process will somehow be the outcome of learning is common to models of all persuasions (e.g. MacWhinney & Leinbach 1991, McClelland

& Patterson 2002, Pinker 1984, 1989, 2004, Pinker & Ullman 2002, Plunkett & Marchman 1991, Rumelhart & McClelland 1986). Traditional 'connectionist' models of inflection (e.g. Rumelhart & McClelland 1986) have simply sought to account for how a particular conception of language might be learned, rather than using the logic of discrimination learning to reconceptualize the task that actually faces the learner, as in the approach taken here. It seems likely that it is just this kind of unexamined theoretical presupposition that makes the task of explaining how language is learned appear far harder than it actually is.

This point also applies to the widespread acceptance by linguists, philosophers, and psychologists that negative 'learnability' arguments warrant the conclusion that various aspects of our linguistic knowledge are innately specified. Negative learnability arguments do not and cannot warrant this conclusion. All one can conclude from a negative learnability argument is that its author is unable to conceive of how it is that something is learned given a particular conception of learning (as, indeed, Gold (1967) explicitly notes). It is always possible that either the characterization of learning or its outcome—the knowledge or cognitive ability that it is claimed cannot be learned—is simply wrong (Johnson 2004).

In the end, questions about whether cognitive capacities are learnable will not be decided by proclamation; they will be resolved by the formulation of convincing scientific accounts of how those capacities develop, which either explain how the information-processing architectures that underlie those capacities arise through learning, or explain how they develop otherwise. The simplicity of the Rescorla-Wagner model is helpful in this regard precisely because the modeler is forced to attend to the actual predictive and discriminative relationships that children learn from: the relationships between linguistic gestures and the objects they abstractly come to represent, as well as the systems of relationships present in linguistic systems themselves. The model predicts that representations of linguistic forms will become increasingly discriminated over the course of experience, making it highly unlikely that children process language in the same way as adults (Stemberger 2004, Stemberger & Middleton 2003, Tabak et al. 2010; see also Baayen et al. 2011, Bannard et al. 2009, Ramscar, Dye, & Klein 2013), or that younger adults process language in the same way as older adults (Ramscar, Hendrix, et al. 2013).

In the same vein, it is important to note that while Rescorla-Wagner is a useful model of a specific implicit-learning process, there is far more to learning than error monitoring. Humans, especially adult humans, are not the passive observers of the environment that Rescorla-Wagner idealizes them to be. They are agents with goals and desires who can direct their attention, and rerepresent their views of their worlds, and all of these cognitive behaviors will in turn have an effect on what they learn and how they learn it. At the same time, these agentive aspects of cognition appear to develop slowly in humans (see Ramscar & Gitcho 2007, Thompson-Schill et al. 2009 for reviews). There is reason to believe that this makes young children and infants more likely to sample the error in their environments in very similar ways (Ramscar, Dye, & Klein 2013), and that this makes it more likely that children who are exposed to cultural systems that embody probabilistic conventions will come to learn and represent the patterns of information in them in appropriate, conventionalized ways (Hudson Kam & Newport 2009, Newport 1990, Singleton & Newport 2004, Thompson-Schill et al. 2009).

Of course, human infants do not just differ from other animals in the way that they sample the environment: the SOCIAL environments that they sample and learn from are markedly different as well (Akhtar & Tomasello 2000, Tomasello 2003, 2008). A human infant is not just a qualitatively different learner from an infant rat; she is born into a qualitatively different environment as well. While social learning and 'associative learning' have often been painted as being in opposition to one another (e.g. Akhtar & Tomasello 2000), it is likely that—as with the LPLA—this opposition hinges on a flawed view of what learning is. As Quine (1960) noted, learning language does not merely require that a child master the relationships in a conventionalized system of sound tokens (Gold 1967), but that the child also learn what the tokens and their relationships mean.

Being able to do so appears to hinge on learning to share subjectivity; the child must somehow master the shared point of view of her community (see also Akhtar & Tomasello 2000, Tomasello 2003, Wittgenstein 1953). How human infants come to discriminate the 'intersubjectively available cues as to what to say and when' (Quine 1960:ix) is an incredibly complex task, yet it is clear that children manage to do this, and that they do so by learning. As a result, it may be that learning models that sample the environment in relatively simple ways are particularly well suited to capture the content and quality of children's social learning. Triesch and colleagues (2006) demonstrate, for example, that gaze following can emerge naturally from domain-general learning mechanisms, provided that a child has access to a caregiver that tends to look at things in ways that the infant finds informative. A concrete, mechanistic account such as this is scientifically preferable to competing explanations of human social development that assume that gaze following is determined by an unspecified innate mechanism that exists solely in order to glean information from a caregiver's gaze (see e.g. Spelke & Kinzler 2007), both because the accuracy of the former is easier to establish, and because this means in turn that even the discovery that it is inaccurate will advance our scientific understanding of development.

In this work, we have shown how a 'simple learning model' can provide a principled account of the specific pattern of data associated with children's learning of a much debated linguistic convention: plural inflection (for compatible approaches relating to verb argument structure, see Ambridge 2012, Ambridge et al. 2009, Boyd & Goldberg 2011, Goldberg 2011). These results establish that noun pluralization conventions are not in principle unlearnable, and that accounting for children's patterns of acquisition of them does not mandate the positing of innate computational mechanisms. We have also shown how a formal learning model—the mechanisms of which are well understood computationally and well supported by other empirical evidence—can make surprising predictions about children's overregularization errors and their eventual recovery from them. We take the success of these predictions as evidence that children do in fact LEARN the conventions of plural marking from the language that they encounter, in much the same way that they learn about many other aspects of the rich cultural environments into which they are born.

Appendix: R code and training set

A. Code required to implement the basic model using the ndl package in the R statistical programming language.

```
# load the naive discrimination learning package into R
library(ndl)

# load the file describing the training set (cues, outcomes, and their frequencies) into R
cuesOutcomes<-read.table("singplur.txt",T,stringsAsFactors=FALSE)

# set sampling of the training set to random
randomOrder = sample(1:sum(cuesOutcomes$Frequency))

## Estimate learning in Rescorla-Wagner
```

```
#mice outcome
mouseitems2mice = RescorlaWagner(cuesOutcomes,
traceCue="mouseitems", traceOutcome="mice",  randomOrder=randomOrder)
mousiness2mice = RescorlaWagner(cuesOutcomes,
traceCue="mousiness", traceOutcome="mice",  randomOrder=randomOrder)
items2mice = RescorlaWagner(cuesOutcomes,
traceCue="items", traceOutcome="mice",  randomOrder=randomOrder)
stuff2mice = RescorlaWagner(cuesOutcomes,
traceCue="stuff", traceOutcome="mice",  randomOrder=randomOrder)

#s outcome
mouseitems2s = RescorlaWagner(cuesOutcomes,
traceCue="mouseitems", traceOutcome="s", randomOrder=randomOrder)
mousiness2s = RescorlaWagner(cuesOutcomes,
traceCue="mousiness", traceOutcome="s",  randomOrder=randomOrder)
items2s = RescorlaWagner(cuesOutcomes,
traceCue="items", traceOutcome="s",  randomOrder=randomOrder)
stuff2s = RescorlaWagner(cuesOutcomes,
traceCue="stuff", traceOutcome="s",  randomOrder=randomOrder)

#mouse outcome
mouseitems2mouse = RescorlaWagner(cuesOutcomes,
traceCue="mouseitems", traceOutcome="mouse",  randomOrder=randomOrder)
mousiness2mouse = RescorlaWagner(cuesOutcomes,
traceCue="mousiness", traceOutcome="mouse",  randomOrder=randomOrder)
items2mouse = RescorlaWagner(cuesOutcomes,
traceCue="items", traceOutcome="mouse",  randomOrder=randomOrder)
stuff2mouse = RescorlaWagner(cuesOutcomes,
traceCue="stuff", traceOutcome="mouse",  randomOrder=randomOrder)

# Calculate the response propensities
sStrength <- (mouseitems2s$weightvector + mousiness2s$weightvector + items2s$weightvector +
stuff2s$weightvector)
miceStrength <- (mouseitems2mice$weightvector + mousiness2mice$weightvector +
items2mice$weightvector + stuff2mice$weightvector)
mouseStrength <- (mouseitems2mouse$weightvector + mousiness2mouse$weightvector +
items2mouse$weightvector + stuff2mouse$weightvector)
interference <- mouseStrength + sStrength
miceoutput <- miceStrength - interference

# Plot the strength of "mouse" when the cues to "mice" are present across training
plot(mouseStrength, ylim=c(-0.8, 0.8), col="blue")
mtext("activation strength for 'mouse'", 3, 1.5)
abline(h=0, col="red")

# Plot the strength of "S" when the cues to "mice" are present across training
plot(sStrength, ylim=c(-0.8, 0.8), col="blue")
mtext("activation strength for '+S'", 3, 1.5)
abline(h=0, col="red")

# Plot the strength of "mice" when the cues to "mice" are present across training
plot(miceStrength, ylim=c(-0.8, 0.8), col="blue")
mtext("activation strength for 'mice'", 3, 1.5)
abline(h=0, col="red")

# Plot the response propensities for "mice" when the cues to "mice" are present across training
plot(miceoutput, ylim=c(-0.8, 0.8), col="blue")
mtext("production propensity for 'mice'", 3, 1.5)
abline(h=0, col="red")
```

B. Training set for the basic model in R ('singplur.txt').

CUES	OUTCOMES	FREQUENCY
items_stuff_arminess_armitems	arm_s	120
items_stuff_beakeriness_beakeritems	beaker_s	120
items_stuff_beariness_bearitems	bear_s	120
items_stuff_bookiness_bookitems	book_s	120
items_stuff_bottleiness_bottleitems	bottle_s	120
items_stuff_bowliness_bowlitems	bowl_s	120
items_stuff_boyiness_boyitems	boy_s	120
items_stuff_cariness_caritems	car_s	120
items_stuff_catiness_catitems_catitem	cat_s	120
items_stuff_chairiness_chairitems	chair_s	120
items_stuff_cupiness_cupitems_cupitem	cup_s	120
items_stuff_doginess_dogitems	dog_s	120
items_stuff_duckiness_duckitems	duck_s	120
items_stuff_faceiness_faceitems	face_s	120
items_stuff_forkiness_forkitems	fork_s	120
items_stuff_froginess_frogitems	frog_s	120
items_stuff_girliness_girlitems	girl_s	120
items_stuff_handiness_handitems	hand_s	120
items_stuff_houseiness_houseitems	house_s	120
items_stuff_leginess_legitems	leg_s	120
items_stuff_plateiness_plateitems	plate_s	120
items_stuff_ratiness_ratitems	rat_s	120
items_stuff_spooniness_spoonitems	spoon_s	120
items_stuff_stooliness_stoolitems	stool_s	120
items_stuff_toyiness_toyitems	toy_s	120
items_stuff_traininess_trainitems	train_s	120
items_stuff_truckiness_truckitems	truck_s	120
items_stuff_tviness_tvitems	television_s	120
item_stuff_arminess_armitem	arm	120
item_stuff_beakeriness_beakeritem	beaker	200
item_stuff_beariness_bearitem	bear	200
item_stuff_bookiness_bookitem	book	200
item_stuff_bottleiness_bottleitem	bottle	200
item_stuff_bowliness_bowlitem	bowl	200
item_stuff_boyiness_boyitem	boy	200
item_stuff_cariness_caritem	car	200
item_stuff_catiness_catitem	cat	200
item_stuff_chairiness_chairitem	chair	200
item_stuff_cupiness_cupitem	cup	200
item_stuff_doginess_dogitem	dog	200
item_stuff_duckiness_duckitem	duck	200
item_stuff_faceiness_faceitem	face	200
item_stuff_forkiness_forkitem	fork	200
item_stuff_froginess_frogitem	frog	200
item_stuff_girliness_girlitem	girl	200
item_stuff_handiness_handitem	hand	200
item_stuff_houseiness_houseitem	house	200
item_stuff_leginess_legitem	leg	200
item_stuff_plateiness_plateitem	plate	200
item_stuff_ratiness_ratitem	rat	200
item_stuff_spooniness_spoonitem	spoon	200
item_stuff_stooliness_stoolitem	stool	200
item_stuff_toyiness_toyitem	toy	200
item_stuff_traininess_trainitem	train	200

item_stuff_truckiness_truckitem	truck	200
item_stuff_tviness_tvitem	television	200
items_stuff_mousiness_mouseitems	mice	100
item_stuff_mousiness_mouseitem	mouse	200

REFERENCES

Akhtar, Nameera, and Michael Tomasello. 2000. The social nature of words and word learning. *Becoming a word learner: A debate on lexical acquisition*, ed. by Roberta Michnick Golinkoff and Kathryn Hirsh-Pasek, 115–35. Oxford: Oxford University Press.

Albright, Adam, and Bruce Hayes. 2003. Rules vs. analogy in English past tenses: A computational/experimental study. *Cognition* 90.119–61.

Altmann, Gerry, and Jelena Mirković. 2009. Incrementality and prediction in human sentence processing. *Cognitive Science* 33.1–27.

Altmann, Gerry, and Mark Steedman. 1988. Interaction with context during human sentence processing. *Cognition* 30.191–238.

Ambridge, Ben. 2012. How do children restrict their linguistic generalizations? An (un)-grammaticality judgment study. *Cognitive Science* 3.508–43.

Ambridge, Ben; Julian Pine; Caroline Rowland; Rebecca Jones; and Victoria Clark. 2009. A semantics-based approach to the 'no negative evidence' problem. *Cognitive Science* 33.1301–16.

Anderson, John, and Lael Schooler. 1991. Reflections of the environment in memory. *Psychological Science* 2.396–408.

Andrews, Mark; Gabriella Vigliocco; and David Vinson. 2009. Integrating experiential and distributional data to learn semantic representations. *Psychological Review* 116.463–98.

Arnon, Inbal, and Eve V. Clark. 2011. When 'on your feet' is better than 'feet': Children's word production is facilitated in familiar sentence-frames. *Language Learning and Development* 7.107–29.

Ashby, Gregory; W. Todd Maddox; and Corey Bohil. 2002. Observational versus feedback training in rule-based and information-integration category learning. *Memory and Cognition* 30.666–77.

Baayen, R. Harald; Peter Hendrix; and Michael Ramscar. 2013. Sidestepping the combinatorial explosion: An explanation of *n*-gram frequency effects based on naive discriminative learning. *Language and Speech* 56.3.329–47.

Baayen, R. Harald; Petar Milin; Dušica Filipović Đurđević; Peter Hendrix; and Marco Marelli. 2011. An amorphous model for morphological processing in visual comprehension based on naive discriminative learning. *Psychological Review* 118.438–82.

Baayen, R. Harald, and Fermín Moscoso del Prado Martín. 2005. Semantic density and past-tense formation in three Germanic languages. *Language* 81.666–98.

Baayen, R. Harald; Richard Piepenbrock; and Leon Gulikers. 1995. The CELEX lexical database. Philadelphia: Linguistic Data Consortium, University of Pennsylvania.

Baker, Chris. 1979. Syntactic theory and the projection problem. *Linguistic Inquiry* 10.533–81.

Balling, Laura, and R. Harald Baayen. 2012. Probability and surprisal in auditory comprehension of morphologically complex words. *Cognition* 125.80–106.

Bannard, Colin; Elena Lieven; and Michael Tomasello. 2009. Modeling children's early grammatical knowledge. *Proceedings of the National Academy of Sciences* 106.17284–89.

Barlow, Horace. 2001. Redundancy reduction revisited. *Network: Computation in Neural Systems* 12.241–53.

Bates, Elizabeth, and George Carnevale. 1993. New directions in research on language development. *Developmental Review* 13.436–70.

Bohannon, John Neil, and Laura Stanowicz. 1988. The issue of negative evidence: Adult responses to children's language errors. *Developmental Psychology* 24.684–89.

BOYD, JEREMY K., and ADELE E. GOLDBERG. 2011. Learning what NOT to say: The role of statistical preemption and categorization in *a*-adjective production. *Language* 87.55–83.

BRILL, ERIC. 1995. Transformation-based error-driven learning and natural language processing: A case study in part-of-speech tagging. *Computational Linguistics* 21.543–65.

BROWN, ROGER. 1973. *A first language: The early stages*. Cambridge, MA: Harvard University Press.

BROWN, ROGER, and CAMILLE HANLON. 1970. Derivational complexity and order of acquisition in child speech. *Cognition and the development of language*, ed. by John Hayes, 11–54. New York: Wiley.

CHANG, FRANKLIN; GARY DELL; and KATHRYN BOCK. 2006. Becoming syntactic. *Psychological Review* 113.234–72.

CLAHSEN, HARALD. 1999. Lexical entries and the rules of language: A multidisciplinary study of German inflection. *Behavioral and Brain Sciences* 22.991–1060.

CHOMSKY, NOAM. 1959. Review of *Verbal behavior*, by B. F. Skinner. *Language* 35.26–57.

COLLINS, MICHAEL, and TERRY KOO. 2005. Discriminative reranking for natural language parsing. *Computational Linguistics* 31.25–69.

CONAN DOYLE, ARTHUR. 1894. *The memoirs of Sherlock Holmes: The adventure of Silver Blaze*. London: G. Newnes.

COTTRELL, GARRISON, and KIM PLUNKETT. 1994. Acquiring the mapping from meanings to sounds. *Connection Science* 6.379–412.

COURVILLE, AARON; NATHANIEL DAW; and DAVE TOURETZKY. 2006. Bayesian theories of conditioning in a changing world. *Trends in Cognitive Sciences* 10.294–300.

DANKS, DAVID. 2003. Equilibria of the Rescorla-Wagner model. *Journal of Mathematical Psychology* 47.109–21.

DAVIES, MARK. 2009. The 385+ million word corpus of contemporary American English (1990–present). *International Journal of Corpus Linguistics* 14.159–90.

DAW, NATHANIEL, and DAPHNA SHOHAMY. 2008. The cognitive neuroscience of motivation and learning. *Social Cognition* 26.593–620.

DAYAN, PETER, and NATHANIEL DAW. 2008. Decision theory, reinforcement learning, and the brain. *Cognitive Affective and Behavioral Neuroscience* 8.429–53.

DE VILLIERS, PETER, and JILL DE VILLIERS. 1973. A cross-sectional study of the acquisition of grammatical morphemes in child speech. *Journal of Psycholinguistic Research* 2.267–78.

DICKINSON, ANTHONY. 1980. *Contemporary animal learning theory*. Cambridge: Cambridge University Press.

DURDA, KEVIN; LORI BUCHANAN; and RICHARD CARON. 2009. Grounding cooccurrence: Identifying features in a lexical cooccurrence model of semantic memory. *Behavior Research Methods* 41.1210–23.

EBBINGHAUS, HERMANN. 1913. *Memory: A contribution to experimental psychology*. New York: Teachers College Press.

ELMAN, JEFF. 1990. Finding structure in time. *Cognitive Science* 14.179–211.

ELMAN, JEFF. 1991. Distributed representations, simple recurrent networks, and grammatical structure. *Machine Learning* 7.195–225.

ERNESTUS, MIRJAM, and R. HARALD BAAYEN. 2004. Analogical effects in regular past tense production in Dutch. *Linguistics* 42.873–903.

GALLISTEL, CHARLES. 2003. Conditioning from an information processing perspective. *Behavioural Processes* 61.1–13.

GLÄSCHER, JAN; NATHANIEL DAW; PETER DAYAN; and JOHN O'DOHERTY. 2010. States versus rewards: Dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron* 66.585–95.

GOLD, E. MARK. 1967. Language identification in the limit. *Information and Control* 10.447–74.

GOLDBERG, ADELE. 2011. Corpus evidence of the viability of statistical preemption. *Cognitive Linguistics* 22.1–28.

GRAVES, MICHAEL, and STEPHEN KOZIOL. 1971. Noun plural development in primary grade children. *Child Development* 42.1165–73.

GREGORY, RICHARD. 2007. In retrospect: Review of Conan Doyle 1894. *Nature* 445.152.

GURECKIS, TODD, and BRADLEY LOVE. 2010. Direct associations or transformations of an internal state? Exploring the mechanism underlying sequential learning behavior. *Cognitive Science* 34.10–50.

HAHN, ULRIKE, and MIKE OAKSFORD. 2008. Inference from absence in language and thought. *The probabilistic mind: Prospects for Bayesian cognitive science*, ed. by Nick Chater and Mike Oaksford, 121–42. Oxford: Oxford University Press.

HALEVY, ALON; PETER NORVIG; and FERNANDO PEREIRA. 2009. The unreasonable effectiveness of data. *IEEE Intelligent Systems* 24.8–12.

HARM, MICHAEL, and MARK SEIDENBERG. 1999. Phonology, reading acquisition, and dyslexia: Insights from connectionist models. *Psychological Review* 106.491–528.

HASKELL, TODD; MARYELLEN MACDONALD; and MARK SEIDENBERG. 2003. Language learning and innateness: Some implications of compounds research. *Cognitive Psychology* 47.119–63.

HUANG, YI TANG, and STEVEN PINKER. 2010. Lexical semantics and irregular inflection. *Language and Cognitive Processes* 25.1–51.

HUDSON KAM, CARLA, and ELISSA NEWPORT. 2009. Getting it right by getting it wrong: When learners change languages. *Cognitive Psychology* 59.30–66.

JAMES, WILLIAM. 1890. *The principles of psychology*, vol. 1. New York: Henry Holt.

JOANISSE, MARC, and MARK SEIDENBERG. 1999. Impairments in verb morphology following brain injury: A connectionist model. *Proceedings of the National Academy of Sciences* 96.7592–97.

JOHNSON, KENT. 2004. Gold's theorem and cognitive science. *Philosophy of Science* 71. 571–92.

JUSTUS, TIMOTHY; JARY LARSEN; PAUL DE MORNAY DAVIES; and DIANE SWICK. 2008. Interpreting dissociation between regular and irregular past-tense morphology: Evidence from event-related potentials. *Cognitive, Affective, and Behavioral Neuroscience* 8. 178–94.

KAMIN, LEON. 1969. Predictability, surprise, attention, and conditioning. *Punishment and aversive behaviour*, ed. by Byron Campbell and Russell Church, 279–96. New York: Appleton-Century-Crofts.

KAMIN, LEON, and STEPHEN GAIONI. 1974. Compound conditioned emotional response conditioning with differentially salient elements in rats. *Journal of Comparative and Physiological Psychology* 87.591–97.

KARPICKE, JEFFREY, and JANELL BLUNT. 2011. Retrieval practice produces more learning than elaborative studying with concept mapping. *Science* 331.772–75.

KARPICKE, JEFFREY, and HENRY ROEDIGER. 2008. The critical importance of retrieval for learning. *Science* 319.966–68.

KREMER, EDWIN. 1978. The Rescorla-Wagner model: Losses in associative strength in compound conditioned stimuli. *Journal of Experimental Psychology: Animal Behavior Processes* 4.22–36.

KUTAS, MARTA, and KARA FEDERMEIER. 2007. Event-related brain potential (ERP) studies of sentence processing. *The Oxford handbook of psycholinguistics*, ed. by Gareth Gaskell, 385–406. Oxford: Oxford University Press.

KRUSCHKE, JOHN. 2008. Bayesian approaches to associative learning: From passive to active learning. *Learning & Behavior* 36.210–26.

LANDAUER, THOMAS, and SUSAN DUMAIS. 1997. A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review* 104.211–40.

LEVY, ROGER. 2008. Expectation-based syntactic comprehension. *Cognition* 106.1126–77.

LEWIS, JOHN, and JEFFREY ELMAN. 2001. A connectionist investigation of linguistic arguments from poverty of the stimulus: Learning the unlearnable. *Proceedings of the 23rd annual conference of the Cognitive Science Society*, 552–57.

MACDONALD, MARYELLEN; NEAL J. PEARLMUTTER; and MARK SEIDENBERG. 1994. The lexical nature of syntactic ambiguity resolution. *Psychological Review* 101.676–703.

MACDONALD, MARYELLEN, and MARK SEIDENBERG. 2006. Constraint satisfaction accounts of lexical and sentence comprehension. *Handbook of psycholinguistics*, ed. by Matthew J. Traxler and Morton Ann Gernsbacher, 581–611. London: Elsevier.

Mackintosh, Nicholas. 1975. A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review* 82.276–98.

MacWhinney, Brian. 2000. *The CHILDES project: Tools for analyzing talk, vol. 2: The database*. Mahwah, NJ: Lawrence Erlbaum.

MacWhinney, Brian. 2004. A multiple process solution to the logical problem of language acquisition. *Journal of Child Language* 31.883–914.

MacWhinney, Brian, and Jared Leinbach. 1991. Implementations are not conceptualizations: Revising the verb learning model. *Cognition* 40.121–57.

Marcus, Gary. 1993. Negative evidence in language acquisition. *Cognition* 46.53–85.

Marcus, Gary. 1995. Children's overregularization of English plurals: A quantitative analysis. *Journal of Child Language* 22.447–59.

Marcus, Gary; Ursula Brinkmann; Harald Clahsen; Richard Wiese; and Steven Pinker. 1995. German inflection: The exception that proves the rule. *Cognitive Psychology* 29.189–256.

Marcus, Gary; Steven Pinker; Michael Ullman; Michelle Hollander; T. John Rosen; and Fei Xu. 1992. Overregularization in language acquisition. *Monographs of the Society for Research in Child Development* 57.1–165.

Marslen-Wilson, William, and Lorraine Tyler. 2007. Morphology, language and the brain: The decompositional substrate for language comprehension. *Philosophical Transactions of the Royal Society B: Biological Sciences* 362.823–36.

McCauley, Stewart, and Morten Christiansen. 2011. Learning simple statistics for language comprehension and production: The CAPPUCCINO model. *Proceedings of the 33rd annual conference of the Cognitive Science Society*, 1619–24.

McClelland, James, and Karalyn Patterson. 2002. Rules or connections in past-tense inflections: What does the evidence rule out? *Trends in Cognitive Sciences* 6.465–72.

McClure, Samuel. 2003. *Reward prediction errors in human brain*. Houston, TX: Baylor College of Medicine dissertation.

McLaren, Ian, and Nicholas Mackintosh. 2000. An elemental model of associative learning: I. Latent inhibition and perceptual learning. *Animal Learning and Behavior* 28.211–46.

Meyer, David, and Roger Schvaneveldt. 1971. Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology* 90.227–34.

Miller, Ralph; Robert Barnet; and Nicholas Grahame. 1995. Assessment of the Rescorla-Wagner model. *Psychological Bulletin* 117.363–86.

Montague, P. Read; Peter Dayan; and Terrence Sejnowski. 1996. A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *Journal of Neuroscience* 16.1936–47.

Montague, P. Read; Steven Hyman; and Jonathon Cohen. 2004. Computational roles for dopamine in behavioural control. *Nature* 431.760–67.

Moscoso del Prado Martín, Fermín; Aleksandar Kostić; and R. Harald Baayen. 2004. Putting the bits together: An information theoretical perspective on morphological processing. *Cognition* 94.1–18.

Newport, Elissa. 1990. Maturational constraints on language learning. *Cognitive Science* 14.11–28.

Niv, Yael. 2009. Reinforcement learning in the brain. *Journal of Mathematical Psychology* 53.139–54.

Norris, Dennis. 2006. The Bayesian reader: Explaining word recognition as an optimal Bayesian decision process. *Psychological Review* 113.327–57.

Otten, Marte, and Jos J. A. Van Berkum. 2008. Discourse-based word anticipation during language processing: Prediction of priming? *Discourse Processes* 45.464–96.

Pavlov, Ivan. 1927. *Conditioned reflexes: An investigation of the physiological activity of the cerebral cortex*, trans. and ed. by Gelb Anrep. London: Oxford University Press.

Pearce, John, and Geoffrey Hall. 1980. A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review* 87.532–52.

Pereira, Fernando. 2000. Formal grammar and information theory: Together again? *Philosophical Transactions of the Royal Society A* 358.1239–53.

Pinker, Steven. 1984. *Language learnability and language development*. Cambridge, MA: Harvard University Press.
Pinker, Steven. 1989. *Learnability and cognition: The acquisition of argument structure*. Cambridge, MA: MIT Press.
Pinker, Steven. 1991. Rules of language. *Science* 253.530–35.
Pinker, Steven. 1999. *Words and rules*. New York: Basic Books.
Pinker, Steven. 2004. Clarifying the logical problem of language acquisition. *Journal of Child Language* 31.949–53.
Pinker, Steven, and Alan Prince. 1988. On language and connectionism: Analysis of a parallel distributed processing model of language acquisition. *Cognition* 28.73–193.
Pinker, Steven, and Michael Ullman. 2002. The past and future of the past tense. *Trends in Cognitive Sciences* 6.456–63.
Plaut, David, and James Booth. 2000. Individual and developmental differences in semantic priming: Empirical and computational support for a single-mechanism account of lexical processing. *Psychological Review* 107.786–823.
Plunkett, Kim, and Virginia Marchman. 1991. U-shaped learning and frequency effects in a multi-layered perceptron: Implications for child language acquisition. *Cognition* 38.43–102.
Plunkett, Kim, and Virginia Marchman. 1993. From rote learning to system building: Acquiring verb morphology in children and connectionist nets. *Cognition* 48.1–49.
Prasada, Sandeep, and Steven Pinker. 1993. Generalization of regular and irregular morphological patterns. *Language and Cognitive Processes* 8.1–56.
Prinz, Jesse. 2002. *Furnishing the mind: Concepts and their perceptual basis*. Cambridge, MA: Bradford Books/MIT Press.
Pullum, Geoffrey, and Barbara C. Scholz. 2002. Empirical assessment of stimulus poverty arguments. *The Linguistic Review* 19.9–50.
Quine, Willard. 1960. *Word and object*. Cambridge, MA: MIT Press.
Ramscar, Michael. 2002. The role of meaning in inflection: Why the past tense does not require a rule. *Cognitive Psychology* 45.45–94.
Ramscar, Michael. 2010. Computing machinery and understanding. *Cognitive Science* 34.966–71.
Ramscar, Michael, and Melody Dye. 2011. Learning language from the input: Why innate constraints can't explain noun compounding. *Cognitive Psychology* 62.1–40.
Ramscar, Michael; Melody Dye; Jessica Gustafson; and Joseph Klein. 2013. Dual routes to cognitive flexibility: Learning and response conflict resolution in the dimensional change card sort task. *Child Development* 84.4.1308–23.
Ramscar, Michael; Melody Dye; and Joseph Klein. 2013. Children value informativity over logic in word learning. *Psychological Science* 24.1017–23.
Ramscar, Michael; Melody Dye; Hanna Muenke Popick; and Fiona O'Donnell-McCarthy. 2011. The enigma of number: Why children find the meanings of even small number words hard to learn and how we can help them do better. *PLoS ONE* 6.e22501.
Ramscar, Michael, and Nicole Gitcho. 2007. Developmental change and the nature of learning in childhood. *Trends in Cognitive Sciences* 11.274–79.
Ramscar, Michael; Peter Hendrix; Cyrus Shaoul; Petar Milin; and Harald Baayen. 2013. Nonlinear dynamics of lifelong learning: The myth of cognitive decline. *Topics in Cognitive Science*, to appear.
Ramscar, Michael; Teenie Matlock; and Melody Dye. 2010. Running down the clock: The role of expectation in our understanding of time and motion. *Language and Cognitive Processes* 25.589–615.
Ramscar, Michael, and Daniel Yarlett. 2007. Linguistic self-correction in the absence of feedback: A new approach to the logical problem of language acquisition. *Cognitive Science* 31.927–60.
Ramscar, Michael; Daniel Yarlett; Melody Dye; Katie Denny; and Kirsten Thorpe. 2010. The effects of feature-label-order and their implications for symbolic learning. *Cognitive Science* 34.909–57.
Reali, Florencia, and Morten Christiansen. 2005. Uncovering the richness of the stimulus: Structure dependence and indirect statistical evidence. *Cognitive Science* 29.1007–28.

Recchia, Gabe, and Michael N. Jones. 2009. More data trumps smarter algorithms: Comparing pointwise mutual information to latent semantic analysis. *Behavior Research Methods* 41.657–63.

Rescorla, Robert. 1968. Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative and Physiological Psychology* 66.1–5.

Rescorla, Robert. 1988. Pavlovian conditioning: It's not what you think it is. *American Psychologist* 43.151–60.

Rescorla, Robert, and Allan Wagner. 1972. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. *Classical conditioning II: Current research and theory*, ed. by Abraham H. Black and William Frederick Prokasy, 64–99. New York: Appleton-Century-Crofts.

Roark, Brian; Murat Saraclar; and Michael Collins. 2007. Discriminative n-gram language modeling. *Computer Speech and Language* 21.373–92.

Roberts, Seth, and Harold Pashler. 2000. How persuasive is a good fit? A comment on theory testing. *Psychological Review* 107.358–67.

Roediger, Henry, and Andrew Butler. 2010. The critical role of retrieval practice in long-term retention. *Trends in Cognitive Sciences* 15.20–27.

Roediger, Henry, and Jeffrey Karpicke. 2006. Test-enhanced learning: Taking memory tests improves long-term retention. *Psychological Science* 17.249–55.

Rohde, Douglas, and David Plaut. 1999. Language acquisition in the absence of explicit negative evidence: How important is starting small? *Cognition* 72.67–109.

Rose, Tony; Mark Stevenson; and Miles Whitehead. 2002. The Reuters corpus volume 1—From yesterday's news to tomorrow's language resources. *Proceedings of the third International Conference on Language Resources and Evaluation*, 29–31.

Rosenblatt, Frank. 1959. Two theorems of statistical separability in the perceptron. *Mechanisation of thought processes: Proceedings of a symposium held at the National Physical Laboratory*, ed. by D. V. Blake and Albert M. Uttley, 419–56. London: Her Majesty's Stationery Office.

Rumelhart, David, and James McClelland. 1986. On learning past tenses of English verbs. *Parallel distributed processing, vol. 2: Psychological and biological models*, ed. by David Rumelhart and James McClelland, 216–71. Cambridge, MA: MIT Press.

Saffran, Jenny. 2001. The use of predictive dependencies in language learning. *Journal of Memory and Language* 44.493–515.

Saffran, Jenny; Richard Aslin; and Elissa Newport. 1996. Statistical learning by 8-month-old infants. *Science* 274.1926–28.

Saffran, Jenny; Elizabeth Johnson; Richard Aslin; and Elissa Newport. 1999. Statistical learning of tone sequences by human infants and adults. *Cognition* 70.27–52.

Santos, Laurie; Jonathan Flombaum; and Webb Phillips. 2007. The evolution of human mind reading. *Evolutionary cognitive neuroscience*, ed. by Steven Platek, 433–56. Cambridge, MA: MIT Press.

Schoneberger, Ted. 2010. Three myths from the language acquisition literature. *Analysis of Verbal Behavior* 26.107–31.

Schultz, Wolfram. 1998. Predictive reward signal of dopamine neurons. *Journal of Neurophysiology* 80.1–27.

Schultz, Wolfram. 2006. Behavioral theories and the neurophysiology of reward. *Annual Review of Psychology* 57.87–115.

Schultz, Wolfram. 2010. Dopamine signals for reward value and risk: Basic and recent data. *Behavioral and Brain Functions* 6.1–9.

Schultz, Wolfram; Peter Dayan; and P. Read Montague. 1997. A neural substrate of prediction and reward. *Science* 275.1593–99.

Schultz, Wolfram, and Anthony Dickinson. 2000. Neural coding of prediction errors. *Annual Review of Neuroscience* 23.473–500.

Seidenberg, Mark, and Maryellen MacDonald. 1999. A probabilistic constraints approach to language acquisition and processing. *Cognitive Science* 23.569–88.

Siegel, Shepard, and Lorraine Allan. 1996. The widespread influence of the Rescorla-Wagner model. *Psychonomic Bulletin and Review* 3.314–21.

Singleton, Jenny, and Elissa Newport. 2004. When learners surpass their models: The acquisition of American Sign Language from inconsistent input. *Cognitive Psychology* 49.370–407.

Smith, Linda, and Chen Yu. 2012. Embodied attention and word learning by toddlers. *Cognition* 125.244–62.
Spelke, Elisabeth, and Katherine Kinzler. 2007. Core knowledge. *Developmental Science* 10.89–96.
Stemberger, Joseph. 2004. Phonological priming and irregular past. *Journal of Memory and Language* 50.82–95.
Stemberger, Joseph, and Christine Middleton. 2003. Vowel dominance and morphological processing. *Language and Cognitive Processes* 18.369–404.
Sutton, Richard, and Andrew Barto. 1998. *Reinforcement learning*. Cambridge, MA: MIT Press.
Swingley, Daniel, and Richard N. Aslin. 2007. Lexical competition in young children's word learning. *Cognitive Psychology* 54.99–132.
Taatgen, Niels, and John R. Anderson. 2002. Why do children learn to say 'broke'? A model of learning the past tense without feedback. *Cognition* 86.123–55.
Tabak, Wieke; Robert Schreuder; and R. Harald Baayen. 2010. Producing inflected verbs: A picture naming study. *The Mental Lexicon* 5.22–46.
Tanenhaus, Michael, and Sarah Brown-Schmidt. 2008. Language processing in the natural world. *Philosophical Transactions of the Royal Society B* 363.1105–22.
Tanenhaus, Michael; Michael Spivey-Knowlton; Kathleen Eberhard; and Julie Sedivy. 1995. Integration of visual and linguistic information in spoken language comprehension. *Science* 268.1632–34.
Thompson-Schill, Sharon; Michael Ramscar; and Evangelia Chrysikou. 2009. Cognition without control: When a little frontal lobe goes a long way. *Current Directions in Psychological Science* 8.259–63.
Tomasello, Michael. 2003. *Constructing a language*. Cambridge, MA: Harvard University Press.
Tomasello, Michael. 2008. *Origins of human communication*. Cambridge, MA: MIT Press.
Triesch, Jochen; Christof Teuscher; Gedeon Deak; and Eric Carlson. 2006. Gaze following: Why (not) learn it? *Developmental Science* 9.125–47.
Waelti, Pascale; Anthony Dickinson; and Wolfram Schultz. 2001. Dopamine responses comply with basic assumptions of formal learning theory. *Nature* 412.43–48.
Wasserman, Edward, and Leyre Castro. 2005. Surprise and change: Variations in the strength of present and absent cues in causal learning. *Learning and Behavior* 33.131–46.
Wicha, Nicole; Elisabeth Bates; Eva Moreno; and Marta Kutas. 2003. Potato not Pope: Human brain potentials to gender expectation and agreement in Spanish spoken sentences. *Neuroscience Letters* 346.165–68.
Wittgenstein, Ludwig. 1953. *Philosophical investigations*. Oxford: Blackwell.
Woollams, Anna; Marc Joanisse; and Karalyn Patterson. 2009. Past-tense generation from form versus meaning: Behavioural data and simulation evidence. *Journal of Memory and Language* 61.55–76.

Ramscar
Department of Linguistics
University of Tübingen
[michael.ramscar@uni-tuebingen.de]

[Received 29 January 2012;
revision invited 10 July 2012;
revision received 21 March 2013;
accepted 25 March 2013]

Dye
Program in Cognitive Science
Indiana University
[meldye@indiana.edu]

Stewart M. McCauley
Department of Psychology
Cornell University
[smm424@cornell.edu]