

## Expectation and error distribution in language learning: The curious absence of “mouses” in adult speech


Michael Ramscar & Melody Dye  
*Department of Psychology, Stanford 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 aren’t 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.<sup>1 2 3</sup> Here, we examine the role children’s expectations play in language learning. Expectation and prediction-error are important and uncontroversial components of animal learning theory,<sup>4 5</sup> but their applicability to language learning is the subject of much debate. We present an error-driven model of the task facing children learning noun-plurals that generates a surprising prediction: that at an appropriate point in learning, exposure to regular plurals alone will cause children’s tendency to over-regularize irregular plurals to *decrease*. Interestingly, the model predicts that the same exposure should have an opposite effect earlier on in learning. Consistent with this, we found that memory testing for items with regular plural names decreased irregular plural over-regularization in six-year-olds, but increased it in four-year-olds. Our model and results suggest that children’s over-regularization errors 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. The idea that the structure<sup>6</sup> and even semantics<sup>7</sup> of languages are innate hinges on arguments that language is otherwise unlearnable. These arguments ignore the pervasive role of expectation in learning. Learning in response to prediction error is evident in a wide range of animal,<sup>5 8 9</sup> human<sup>10 11 12</sup> and neurobiological<sup>13 14</sup> effects, and appears to be a driving force in language acquisition. Incorporating a more substantial view of learning into analyses of language and its development<sup>15 16 17</sup> can only broaden our understanding of them.

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.”<sup>18</sup>

A racehorse vanishes, its trainer murdered. Sherlock Holmes lights upon a critical piece of evidence: a dog remained silent throughout.<sup>19</sup> The fact that an expected event did not occur – the dog did not bark – provides a critical clue as to the identity of the culprit. Holmes’ deductions illustrate how much can be learned from the discrepancy between what is expected and what actually occurs. In what follows, we show how children use these discrepancies as an important source of information in learning, and that often, as in the curious incident of the dog in the nighttime, the *non*-occurrence of expected events provides a rich source of evidence.

Such evidence has often been ignored or marginalized in discussions of language learning,<sup>23 20 21</sup> dismissed as weak ‘indirect’ negative evidence that offers little to no assistance in the complex process of language acquisition.<sup>2</sup> However, the idea that children might make use of this kind of negative evidence need not be surprising: in animal learning, both positive evidence (successful predictions) and negative evidence (prediction-errors) are fundamental and uncontroversial aspects of learning.<sup>4 5 8-12</sup> Here we show that prediction-error provides a similar, critical source of negative evidence in human language learning.

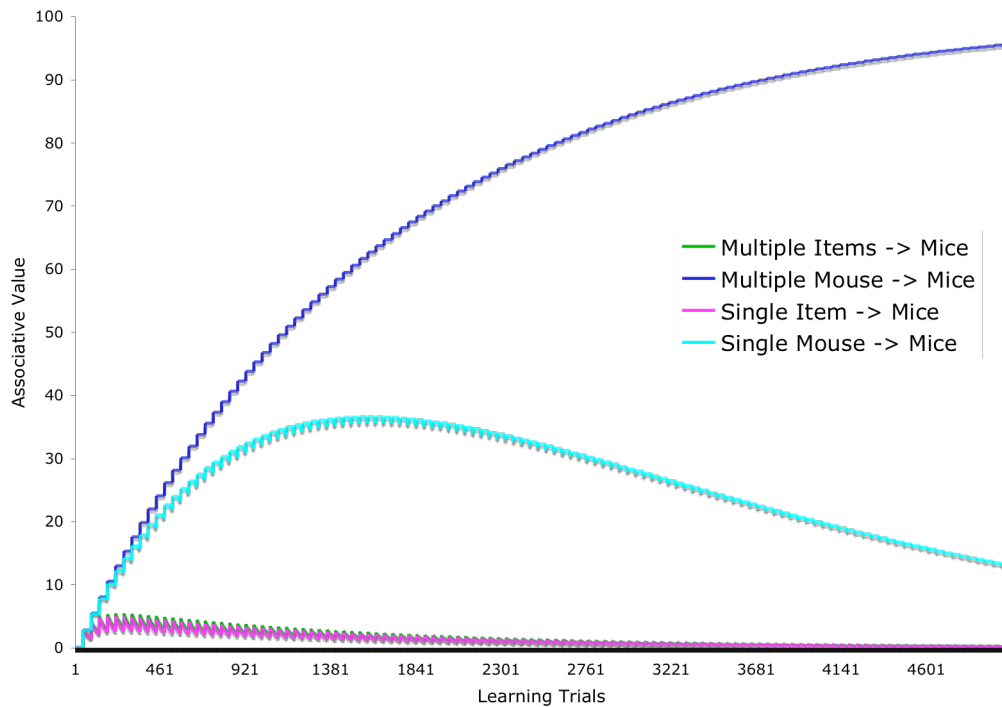


**Figure 1.** Four cues provided by mice that are relevant to plural and singular naming. All four cues always co-occur with the word “mice.” However, their covariance with other singular and plural nouns – and thus the distribution of error associated with them – differs in favor of the mice  $\Rightarrow$  “mice” mapping.

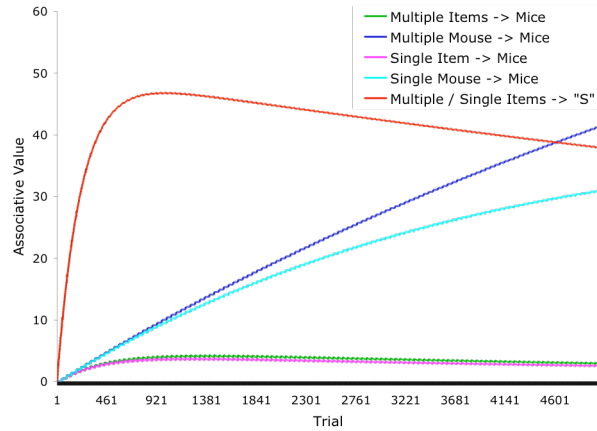
To explore how prediction-error might affect language learning, we constructed a model<sup>5</sup> of the way a child might learn to name singular and plural objects. English-speaking children tend to over-regularize irregular plurals (saying e.g., “mouses” instead of “mice”), a behavior that is rarely explicitly corrected. Because older children end up using only the adult form “mice,” and because there is no obvious reason for them to stop saying “mouses,” it has been argued that this presents a logical puzzle: how do they *learn* to do this?<sup>1 2 3</sup> The model was designed to have sufficient detail to answer this question, and allow predictions to be derived, while being simple enough for its mechanisms to remain transparent. It assumes that: 1) children are not formally taught their native languages;<sup>2</sup> 2) they (at least initially) learn words by hearing them in context;<sup>22 23 24</sup> and 3) the *error distribution* in the early linguistic environment – i.e., the combined value of both positive and negative evidence – favors the appropriate mappings (mice to “mice,” etc). For example, we assume that a child learning the word “mice,” will hear “mice” in relation to mice (or depictions of them), and must learn to associate the appropriate cues in the environment (mouse-things) with the word (“mice”).<sup>12 25</sup> These assumptions differ markedly from those of most other learning models, which treat the task facing the child as one of learning to transform a “word stem” (“cat”) into an “inflected form” (“cats”).<sup>26</sup>

The model simulates how cues to the irregular plural “mice” and a set of other singular and regular plural forms are learned. Figure 1 illustrates the four environmental cues that consistently co-vary with “mice,” and are thus most relevant to plural mouse-naming (they are separated out for explanatory convenience, but might be ranges of values on continuous perceptual dimensions as far as the model is concerned). Because all 4 of these cues—*thing*, *mouse-thing*, *things* and *mouse-things*—are present whenever mice are seen and “mice” is heard, a child cannot hope to isolate the cue(s) appropriate to naming “mice” on the basis of positive evidence alone, since each cue benefits from an identical amount of positive evidence. Crucially, however, the distribution of error associated with each cue differs. Since “mice” is usually heard when *mouse-things* are present (e.g., “look at those mice!”), there will be little error in the association between *mouse-things* and “mice.” However, given that most of the time that some *thing* is present in the child’s environment “mice” will *not* be heard (e.g., “cup,” or “daddy” might be heard instead), *thing* will result in a great deal of error as a cue to “mice.” Accordingly,

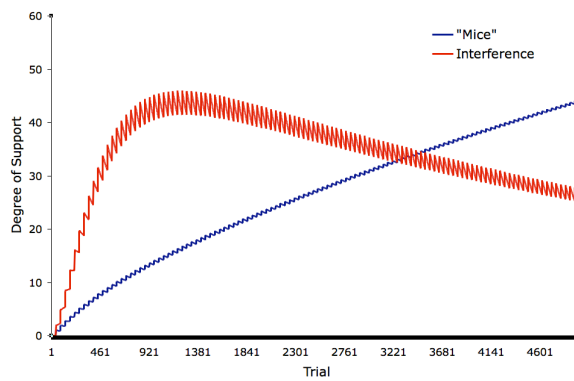
though *thing*, *mouse-thing*, *things* and *mouse-things* provide identical positive evidence for “mice,” the negative evidence each provides differs dramatically (Figure 2).



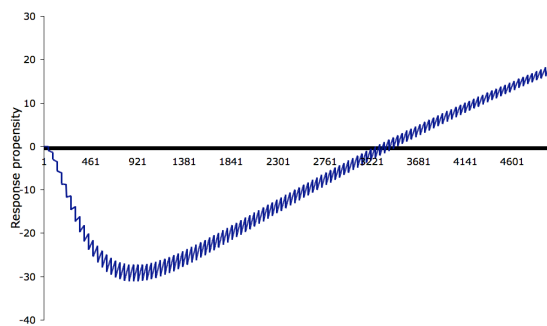
**Figure 2.** A model of the effect of error distribution on learning the four cues relevant to naming “mice” across time. The errors generated by *mouse-thing*, *thing* and *things* on singular and regular plural trials cause them to lose value to *mouse-things* as a result of negative evidence as expectations are violated. Notably, the *mouse-thing*  $\Rightarrow$  “mice” association only accrues negative evidence on mouse trials (when “mouse” is heard, not “mice”), and is unlearned more slowly than *thing* and *things*. This pattern is consistent with the results of tests of children’s irregular plural production,<sup>24</sup> which show that the likelihood of children producing a singular rather than over-regularized form (“mouse” rather than “mouses”) actually *increases* with age.



A



B

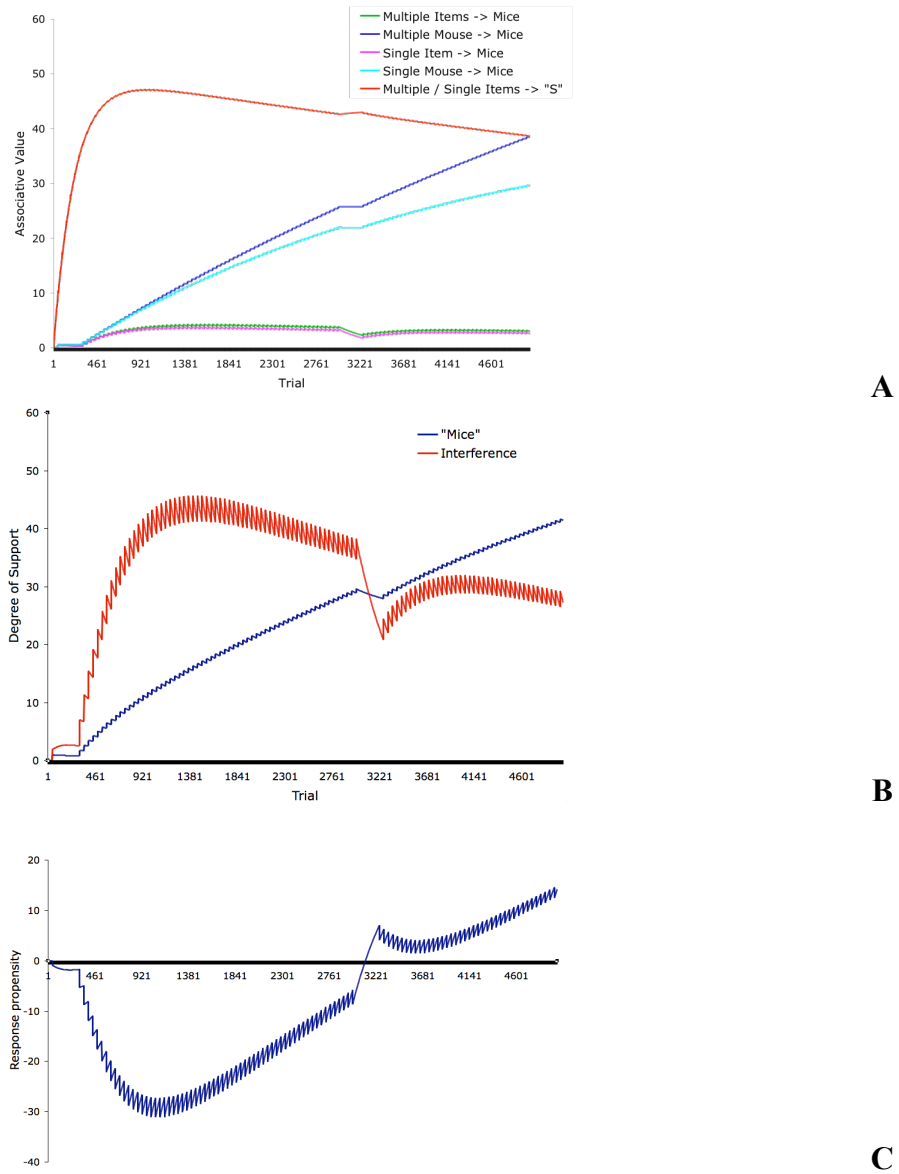


C

**Figure 3.** Panel A shows the values of the four cues to “mice” and the value of *thing* and *things* as cues to a common-form+/s/ (combined as a single value) developing over time. B shows the development of support for a correct or over-regularized form. C shows the response propensity of the model (negative values favor over-regularized responses; positive values favor correct irregular plural responses).

In the model, over-regularization occurs on “mice” trials because the cues *thing* and *things* produce expectation not only for “mice,” but for other nouns as well, producing a bias toward over-regularization errors. Figure 3 shows how over-regularization first arises out of the frequency of different word forms and the frequency and distribution of the cues to them, and then resolves itself as a result of the distribution of error amongst these same cues. This

developmental pattern exhibits the “U-shaped” learning curve often noted in the development and resolution of children’s tendency to over-regularize.<sup>27</sup>



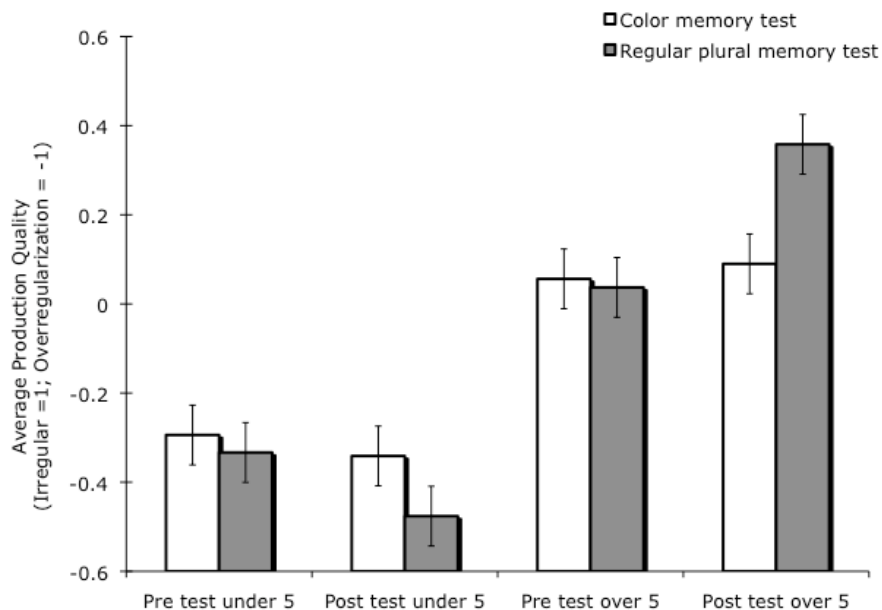
**Figure 4.** In this simulation, the model is not exposed to irregular forms between trials 100-330 and trials 2048-2278 (each corresponding to 5 “epochs” of training). Panel A shows the learning of the values of the four cues to “mice” and the value of *thing* and *things* as cues to a common-form+/s/. B shows the development of values supporting a correct or over-regularized form. C shows the model’s response propensity over time.

The formal properties of our model allow detailed predictions to be made about the circumstances that can lead to an increase or decrease in the rate of over-regularization in young children, depending on how much learning they've done already. Early in learning (trials 100-330), the model's limited experience means that irregulars are only weakly expected on regular learning trials. As a result, when irregulars are withheld at this stage, they are only weakly *unlearned* as possible outcomes (Figure 4A). This slow unlearning of the general cues *thing* and *things* as cues to "mice" (as a result of negative evidence on regular plural trials) when coupled with the rapid learning of *thing* and *things* as cues to an over-regularized form (as a result of positive evidence in regular plural trials), produces an interaction that leads to a gradual *increase* in over-regularization (Figure 4C).

When irregulars are withheld later on in learning (trials 2048-2278) however, when the model has more experience with irregular forms, *thing* and *things* will lead to greater (erroneous) expectation of irregular forms on regular trials, meaning that *thing* and *things* will be more strongly unlearned as cues to irregulars. Meanwhile, the value of *thing* and *things* as cues to an over-regularized form will have asymptoted, and will actually be in the process of losing value to specific cues to regulars (e.g., *dog / cat things*). Thus, at this point in learning, *thing* and *things* will be rapidly unlearned as cues to irregular forms as well as slowly unlearned as cues to over-regularized forms. The resulting interaction produces a sharp *decrease* in over-regularization (Figure 4C, trials 2048-2278). As a result, the model predicts that identical interventions— withholding irregulars—will produce opposite effects at different stages of learning.

To see whether these predictions held, we asked 38 four- and 40 six-year old children to name a series of six irregular plural nouns and six regular plural pairings to establish a baseline rate of over-regularization for each child. The children were then randomly assigned to two

groups. In the experimental condition children were required to exercise their knowledge of plural nouns by telling a cookie monster whether or not depictions of regular plural noun-objects had the same name as items they had previously named in the pre-test. In the control condition, children performed a similar old/new name task with colors. Children in each condition were then given a post-test identical to the pre-test. Children's performance in these tests overwhelmingly supported the model's predictions. A 2 (pre- to post- test) x 2 (age) x 2 (condition) repeated measures ANOVA of the children's plural production revealed a significant interaction between age, training type and pre- to post-test performance ( $F(1,348) = 6.627$ ,  $p < .01$ ), and a significant interaction between age and pre- to post-test performance ( $F(1,348) = 12.584$ ,  $p < 0.0001$ ). Older children in the experimental condition improved their irregular production between pre- and post- test ( $t(88) = 3.562$ ,  $p < 0.0005$ ), while the performance of the younger children declined ( $t(83) = 1.927$ ,  $p < 0.05$ ). There was little change in the performance of either age group in the control condition (see Figure 5).



**Figure 5.** Pre and post-test performance by age and condition. The data is plotted as the number of correct forms minus over-regularized forms averaged across each pair of trials. Error bars are SEM.

Testing memory for regular plural noun-objects significantly reduced plural over-regularization in six-year-olds. Although our particular testing and training regime may have affected the strength of these results, it is clear that children in the experimental condition learned about irregular plural labels even though none were present in the training trials. To the extent that this result is surprising, it may be due to common misunderstandings of the way learning works,<sup>28</sup> and particularly, how prediction-error provides a source of negative evidence to learners.

Overwhelmingly, research into language learning has pre-occupied itself with the observable: with what a child hears or sees. Because many researchers have assumed that children must learn from “positive evidence” alone,<sup>2 3 21</sup> linguistic theory has been heavily influenced by Gold's<sup>29</sup> demonstration of the limitations of such learning. This is regrettable, because as Gold himself noted, his proof applied to a highly formal conception of language,<sup>12 30</sup> and it showed *either* that only the most trivial class of languages is learnable *or else* that children gather negative evidence, “in a way we do not recognize.”<sup>29</sup> Since Gold’s time, it has become clear that language processing involves prediction at every conceivable level,<sup>12 17 31</sup> and that processes responsive to prediction error are ubiquitous in learning.<sup>10 - 13</sup> Our results suggest that these processes enable children to correct their own mistakes in learning language, and thus, that there simply is no logical problem in the way that children who say “mouses” can, without explicit correction, become adults who say “mice.”

We would not claim that it follows from these results that all aspects language are learned. Rather, we have shown how learning theory can give a detailed, principled account of the specific pattern of data associated with children’s over-regularization errors. A hundred and

fifty years ago, Darwin's account of the evolution of the eye demonstrated how evolutionary theory could similarly provide a detailed, principled account of the specific pattern of characteristics of eyes across evolutionary time. Just as evolutionary theory is preferable to teleological arguments for "design," we believe that learning theory is preferable to the claim that language must result from an unspecified "universal grammar," simply because it isn't immediately obvious how children learn it.<sup>2 3 6 7 21</sup> Incorporating the evidence provided in linguistic distributions<sup>32</sup> into analyses of language learning will increase our understanding of its cultural and biological basis, and may reap many useful social and scientific benefits.

## References

- <sup>1</sup> Baker, C. Syntactic theory and the projection problem. *Linguistic Inquiry*, **10**, 533–581, (1979)
- <sup>2</sup> Pinker, S. *Language learnability and language development*. (Harvard University Press, 1984)
- <sup>3</sup> Pinker, S. Clarifying the logical problem of language acquisition. *Journal of Child Language*, **31**, 949-953 (2004)
- <sup>4</sup> Kamin L.J. in: Campbell B, Church R (eds). *Punishment and Aversive Behaviour*. (Appleton-Century-Crofts, 1969)
- <sup>5</sup> Rescorla R.A. and Wagner A.R. in Black & Prokasy (Eds.), *Classical Conditioning II: Current Research and Theory*. (Appleton-Century-Crofts, 1972)
- <sup>6</sup> Chomsky, N. Review of Verbal Behavior, by B.F. Skinner. *Language* **35**: 26-57 (1959).
- <sup>7</sup> Chomsky, N. *New Horizons in the Study of Language and Mind*. (Cambridge University Press, 2000)
- <sup>8</sup> Pearce, J. M., & Hall, G. A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, **87**, 532-552 (1980)
- <sup>9</sup> McLaren, I. P. L., & Mackintosh, N. J. An elemental model of associative learning: I. Latent inhibition and perceptual learning. *Animal Learning & Behavior*, **28**, 211-246, (2000)
- <sup>10</sup> Miller R.R., Barnet R.C. & Grahame N.J. Assessment of the Rescorla-Wagner Model, *Psychological Bulletin*, **117**(3), 363-386, (1995)
- <sup>11</sup> Siegel, S.G., & Allan, L.G. The widespread influence of the Rescorla-Wagner model, *Psychonomic Bulletin and Review*, **3**(3), 314-321, (1996)
- <sup>12</sup> Ramscar, M., Yarlett, D., Dye, M., Denny, K., & Thorpe, K. Feature-Label-Order effects and their implications for symbolic learning. *Cognitive Science* (in press.)
- <sup>13</sup> Hollerman J.R., & Schultz W. Dopamine neurons report an error in the temporal prediction of reward during learning. *Nature Neuroscience*, **1**: 304-309 (1998)
- <sup>14</sup> Waelti, P., Dickinson, A. and Schultz, W. Dopamine responses comply with basic assumptions of formal learning theory. *Nature* **412**: 43-48 (2001)
- <sup>15</sup> Seidenberg, M. S. Language acquisition and use: Learning and applying probabilistic constraints. *Science*, **275**, 1599-1604. (1997)
- <sup>16</sup> Rohde, D. L. T. and Plaut, D. C. Language acquisition in the absence of explicit negative evidence: How important is starting small? *Cognition* **72**, 67-109 (1999)
- <sup>17</sup> Altmann, G.T.M. & Mirković, J. Incrementality and prediction in human sentence processing. *Cognitive Science*, **33**: 1–27 (2009)

- <sup>18</sup> Conan Doyle, A. *The Memoirs of Sherlock Holmes: The Adventure of Silver Blaze* (Newnes, 1894)
- <sup>19</sup> Gregory, R. Book review: The Memoirs of Sherlock Holmes: The Adventures of Silver Blaze (by Conan Doyle, A), *Nature*, **445**, 152 (2007)
- <sup>20</sup> Brown, R., & Hanlon, C. in J. R. Hayes (Ed.), *Cognition and the development of language*. (Wiley, 1970)
- <sup>21</sup> Marcus, G. F. Negative evidence in language acquisition. *Cognition*, **46**, 53–85 (1993)
- <sup>22</sup> Tomasello, M. *Constructing a Language*. (Harvard University Press, 2003)
- <sup>23</sup> Smith, L.B., & Yu, C. Infants Rapidly Learn Word-Referent Mappings via cross-situational Statistics. *Cognition*, **106**, 333-338, (2008)
- <sup>24</sup> Ramscar, M. & Yarlett, D. Linguistic self-correction in the absence of feedback: A new approach to the logical problem of language acquisition. *Cognitive Science*, **31**, 927-960 (2007)
- <sup>25</sup> Wittgenstein, L. *Philosophical Investigations* (Blackwell, 1953)
- <sup>26</sup> Rumelhart D. E. & McClelland J. L. On learning past tenses of English verbs. In Rumelhart & McClelland (eds) *Parallel Distributed Processing: Vol 2* (MIT Press, 1986)
- <sup>27</sup> Brown, R. *A First Language: the Early Stages*. (Harvard University Press, 1973)
- <sup>28</sup> Rescorla R.A. Pavlovian Conditioning: It's Not What You Think It Is, *American Psychologist*, **43**(3), 151-160 (1988)
- <sup>29</sup> Gold, E. M. (1967) Language identification in the limit. *Information and Control*, **10**, 447-474
- <sup>30</sup> Johnson, K. Gold's Theorem and Cognitive Science. *Philosophy of Science*, **71**, 571-592, (2004)
- <sup>31</sup> Tanenhaus, M.K. & Brown-Schmidt, S. Language processing in the natural world. In Moore, B., Tyler, L. & Marslen-Wilson, W. (eds.) *The perception of speech: from sound to meaning*. *Philosophical Transactions of the Royal Society B*, **363**, 1105-1122 (2008)
- <sup>32</sup> Baayen R.H. *Word Frequency Distributions*. (Kluwer Academic Publishers, 2001)

## Supplementary Methods

### Modeling

The Rescorla-Wagner model<sup>5</sup> is a widely used learning rule that has been applied to numerous learning effects in animals and humans. While it cannot account for all 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 this analysis is consistent with a wide range of learning models, in which equivalent simulations could be implemented.<sup>12</sup>

The Rescorla-Wagner model simulates changes in the associative strengths ( $V$ ) between individual cues  $i$  and a response  $j$  as the result of discrete exposure trials, where  $n$  indexes the current trial:

$$\Delta V_{ij}^n = \alpha_i \beta_j (\lambda_j - V_{total}) \quad (1)$$

Learning is governed by a discrepancy function ( $\lambda_j - V_{TOTAL}$ ) where  $\lambda_j$  is the value of the predicted event (the maximum amount of associative strength that an event  $j$  can support) and  $V_{total}$  is the predictive value of a set of cues (the sum of the associative strengths between all cues  $i$  present on the current trial). The amount of learning on a given trial is determined by two factors: the individual saliency of cues, denoted by a parameter  $\alpha_i$  (where  $0 \leq \alpha_i \leq 1$ ), and the overall learning rate  $\beta_j$  (where  $0 \leq \beta_j \leq 1$ ). These determine the rate at which the discrepancy between  $\lambda_j$  and  $V_{TOTAL}$  reduces. Thus, in trials in which there is positive evidence – i.e., in which expected outcomes *do* occur – the rule produces a negatively accelerated learning curve (the

result of events being better predicted) and asymptotic learning over repeated trials (as events become fully predicted).

What is not widely appreciated about this kind of learning is what happens when a predicted event does *not* occur (i.e., when there is *negative evidence*). If a cue predicts something that doesn't follow,  $\lambda_j$  (the expected outcome) takes a value of zero because it didn't occur. In such cases any discrepancy ( $\lambda_j - V_{TOTAL}$ ) has a negative value, resulting in a reduction in the associative strength between the cues present on that trial and the absent outcome  $j$ . Because negative evidence diminishes cue values, and because the total amount of value a given outcome can support is finite, cues *compete* with one another for relevance, producing learning patterns that usually differ greatly from those that would arise by simply recording the correlations between cues and outcomes (a common misconstrual of learning<sup>28</sup>).

In the simulation in Figure 2,  $\lambda = 100\%$  for each word,  $\alpha_i=0.07$  for the dimension cues and  $\beta_j=0.4$ . In all other simulations,  $\lambda = 100\%$  for each word,  $\alpha_i=0.07$  for all of the relevant dimensions, and  $\beta_j=0.07$  (the slower learning rate allows early learning patterns to be more clearly seen, but has no impact on the shape of early learning, or any interactions between cues). At the outset of the simulation, the model assumes that children begin with no knowledge of nouns (trial 1). The model then learns 24 regular singular forms, 20 regular plurals, 1 irregular singular and 1 irregular plural, with each training “epoch” thus comprising 46 trials, such that at trial 4601, the model has been exposed to 100 irregular plurals. Crucially, exposure to irregulars is withheld at two time points (between trials 100-330 and trials 2048-2278) in order to examine the effect of “negative evidence” at different stages of learning.

The model assumes that when a child is asked to name a picture of mice, the child's memory activates the word *mice*, because this is the phonological form the child has learned to

associate with the semantic representation of mice (however, what the child actually says is contingent on both the strength of the representation of mice, and the degree to which other forms interfere with “mice” production). It further assumes that in a language such as English, where multiple instances of count-nouns are usually named using a different phonological form to single count-nouns, a child must learn to discriminate between single and multiple items when naming (meaning that set size will serve as a cue to whether forms are singular or plural).<sup>12</sup>

Importantly, the simulation also assumes that the phonological forms of regular singular and plural (+S) nouns are only distinguished temporally, by the occurrence (in plurals) or non-occurrence (in singulars) of a sibilant after a common form. While this smoothes over some of the differences between the single and plural forms of regular nouns (such as different sibilant allomorphs, co-articulation effects, etc), it allows the model to capture the idea that while the poor discrimination of individual plurals may not result in interference in regular plural production, it can and does interfere with irregular plural production, leading to over-regularization. This is because regular plurals resemble one another with respect to their key phonological indicator of plurality (the sibilant), whereas irregular plurals usually resemble neither regular plurals nor each other. For a young 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 interfere with irregular plural production (“mouses”).

To simulate the conditions that might thus lead to over-regularization, the model assumes that on “mice” trials, the cues *thing* and *things* not only result in some expectation of “mice,” but also in some expectation of a form (roughly) common to singulars and plurals followed by a +s/

ending (i.e., an over-regularized form). The strength of this expectation arises out of two factors:

- 1) the degree to which *thing* and *things* are active as cues on mice trials (as Figure 3 shows, learners will come to ignore these cues over time as they better discriminate specific items) and
- 2) the degree to which *thing* and *things* have been learned as cues to an over-regularized form.

An interference function can thus be calculated by multiplying the current learned values of *thing* and *things* as cues to the regular +s/ form by both the active values of *thing* and *things* on mice trials and a spread parameter  $S$  ( $0 \leq S \leq 1$ ), which we assigned  $S=0.1$ . (While this value was chosen to fit the model in the diagrams, exploring a range of values of  $S$  revealed that the *shape* of this function is largely independent of the value of  $S$ .)

While this interference function simplifies the similarity relationships between items in the model (in reality, *things* like rats might be expected to produce more interference in “mice” production than would books<sup>14</sup>), it allows the model to capture the idea that children are learning to classify objects at the same time that they are learning to name them.<sup>12</sup> Importantly, though differences in class similarity will affect the individual trajectories with which items are learned, they will not affect the *shape* of that learning, which is determined by similarities between names rather than items (regular plurals have +s in common, whereas irregular plurals do not). This simplification thus allows the model to simulate the way that poor discrimination in early learning may result in uncertainty about which name is associated with a given class of objects, without adding unnecessary detail regarding the degrees of similarity between classes of objects.

## **Behavioral Experiment**

The children were recruited from a database of volunteers living in the vicinity of Palo Alto, California. The average ages were 4 years and 6 months for the four year olds, and 6 years

and 7 months for the six year olds. The children were randomly assigned to two groups, both of which were pre-tested on plural production. The irregular items in both the pre- and post- tests were MOUSE-MICE, CHILD-CHILDREN, SNOWMAN-SNOWMEN, GOOSE-GEESE, TOOTH-TEETH and FOOT-FEET; regular matches were RAT, DOLL, COW, DUCK, EAR and HAND. Although children at these ages over-regularize these irregular plurals, they have reliable knowledge of their correct forms.<sup>24</sup>

In the experimental condition the children performed an old/new task, telling a cookie monster whether or not depictions they were shown were things they had named in the pre-test. The absence of overt naming responses by children was intended to reduce the effect of perseverative biases on post-test performance, allowing a better measure of the representation of the items of interest. 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 or not. As words’ phonological representations are cued by their semantics,<sup>35</sup> these measures could be expected to produce prediction errors and latent learning.<sup>12</sup>

<sup>24</sup> The children were asked to tell the cookie monster “yes or no,” to indicate whether they had already seen these depictions or not. 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 didn’t 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 pre-test and half were foils. The children were thus tested on 12 new and 12 old items per block.

In the control condition, children were shown 6 color slides after the pre-test, 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 each condition was equal. Both sets of children were then post-tested on exactly the same set of depictions used in the pre-test.

### **Supplementary References**

- <sup>35</sup> Meyer, D.E. and Schvaneveldt, R.W. Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology*, **90**, 227-23 (1971)

### **Acknowledgements**

This material is based upon work supported by the National Science Foundation under Grant Nos. 0547775 and 0624345 to Michael Ramscar. We are grateful to Daniel Yarlett for many insightful discussions of learning theory and language, and to Dan, Bradley Love, Ewart Thomas and Gordon Bower for helpful feedback on these specific ideas. We thank Rowan and Ian Goddard for inspiring this project.

### **Author Correspondence**

Please contact MR with any questions or comments regarding this manuscript via email (ramscar@stanford.edu).