₁ Some puzzling findings regarding the acquisition of verbs

₂ Joshua K. Hartshorne¹, Yujing Huang¹, & Lauren Skorb¹

₃ ¹ Department of Psychology and Neuroscience, Boston College

₄ Abstract

On the whole, children acquire frequent words earlier than less frequent words. However, there are other factors at play, such as an early “noun bias” (relative to input frequency, toddlers learn nouns faster than verbs) and a “content-word bias” (content words are acquired disproportionately to function words). This paper follows up reports of a puzzling phenomenon within verb-learning, where “experiencer-object” emotion verbs (A frightened/angered/delighted B) are lower frequency but learned earlier than “experiencer-subject” emotion verbs (A feared/hated/loved B). In addition to the possibility that the aforementioned results are a fluke or due to some confound, prior work has suggested several possible explanations: experiencer-object (“frighten-type”) verbs have higher type frequency, encode a causal agent as the sentential subject, and perhaps describe a more salient perspective on the described event. In three experiments, we cast doubt on all three possible explanations. The first experiment replicates and extends the prior findings regarding emotion verbs, ruling out several possible confounds and concerns. The second and third experiments investigate acquisition of chase/flee verbs and give/get verbs, which reveal surprising findings that are not explained by the aforementioned hypotheses. We conclude that these findings indicate a significant hole in our theories of language learning, and that the path forward likely requires a great deal more empirical investigation of the order of acquisition of verbs.

Keywords: language acquisition, perspective pairs, psych verbs, verb-learning

Introduction

Learning a word is more than just a function of having heard it. On the whole, children do acquire frequent words earlier than less frequent words, but that is hardly the end of the story (Goodman, Dale, & Li, 2008; Hansen, 2017). Controlling for frequency, nouns are learned earlier than verbs, which are in turn learned earlier than closed-class

words (Gentner, 1982; Goodman et al., 2008; Hansen, 2017). Controlling for word type, highly-imageable words are learned earlier than less-imageable words, perhaps because it is easier to identify the intended referent during conversation (Hansen, 2017; McDonough, Song, Hirsh-Pasek, Golinkoff, & Lannon, 2011). Indeed, raw frequency may be the wrong measure: rather than word knowledge slowly accumulating with each exposure, word-learning may be disproportionately driven by highly-informative learning opportunities, which give rise to “eureka” moments (Medina, Snedeker, Trueswell, & Gleitman, 2011).

However, none of these factors obviously explain a puzzling finding reported by Hartshorne, Pogue, & Snedeker (2015): emotion verbs that realize the experiencer as the direct object (*A frightened/pleased/angered B*; “frighten-type verbs”) are acquired substantially earlier than those that realize the experiencer as the subject (*A feared/liked/hated B*; fear-type verbs), despite being much lower frequency. Specifically, while English-speaking children have already acquired a handful of frighten-type verbs by the age of 4, they do not start acquiring fear-type verbs until about a year later. In particular, although four year-olds use words like *like* and *love* in spontaneous speech, they struggle to distinguish who did what to whom, treating *A loves B* as equivalent to *B loves A* (and similarly for other fear-type verbs). Interestingly, this had actually been observed in prior studies but either not remarked upon or overtly dismissed as the result of confounds (Bowerman, 1990; Braine, Brooks, Cowan, Samuels, & Tamis-LeMonda, 1993; Messenger, Branigan, McLean, & Sorace, 2012; Tinker, Beckwith, & Dougherty, 1988); by addressing those confounds and reporting multiple targeted studies using different methods, Hartshorne, Pogue, and Snedeker (2015)

The authors thank Jesse Snedeker, Tim O’Donnell, and Laura Lakusta for helpful feedback and suggestions, and Miguel Mejia, Zach Barker, Kayley Okst, Caitlin Garcia, Nicola Roux, Marissa Russell, Camille Phaneuf, Ernesto Gutierrez, Juliani Vidal, Hayley Greenough, Casey Nicastrì, Rudmila Rashid, Taylor Martinz, Kate Roberts, Madeleine McCanne, Sarah Al-Mayahi, Alice Lim, Everett Kim, Lily Feinberg, Zeen Naeem and the staff at the Acton Discovery Museum and Boston Children’s Museum for assistance with data-collection. Funding was provided by NSF 1551834, awarded to JKH. Data, code, and materials are available at <https://osf.io/k5xud/>. Author Contributions: JKH conceived of the study. JKH and LS designed the experiments. JKH, YH, and LS performed analyses. JKH and YH wrote the manuscript.

Correspondence concerning this article should be addressed to Joshua K. Hartshorne, McGuinn Hall 522, 140 Commonwealth Ave., Chestnut Hill, MA 02467. E-mail: joshua.hartshorne@bc.edu

32 established the finding as robust and something in need of explanation.

33 Hartshorne and colleagues interpret this finding as evidence for a “privileged link”
 34 between causality and sentential subjects. Specifically, a long line of research suggests both a
 35 cross-linguistic tendency for the agents of caused events to be mapped onto subject position
 36 and a corresponding learning bias on the part of young children to expect agents of caused
 37 events to be subjects of corresponding sentences (Braine, 1992; Dowty, 1991; Fisher, Gertner,
 38 Scott, & Yuan, 2010; Levin & Hovav, 2005; Lidz, Gleitman, & Gleitman, 2003; MacWhinney,
 39 1977; Marantz, 1982, 1982; Pinker, 1984; Strickland, Fisher, Keil, & Knobe, 2014). This
 40 set of findings has prompted a number of theorists – both Empiricist and Nativist – to
 41 argue for an innate bias for agents of caused events to be realized as sentential subjects
 42 (or whatever the analog of SUBJECT is in the account adopted by the theorist), and that
 43 this bias is a key part of what makes language learnable in the first place (for discussion,
 44 see Hartshorne et al., 2016). Critically, across a variety of languages, frighten-type verbs
 45 do encode causality (*A frightened B* means, roughly, “A caused B to feel fear”) whereas
 46 fear-type verbs do not (*A feared B* means, roughly, “A was disposed to feel fear about B”)
 47 (Hartshorne et al., 2016). (Note that while these findings overturned some earlier proposals
 48 that treated *fear* and *frighten* as synonymous, differing only in the syntax (Belletti & Rizzi,
 49 1988; Dowty, 1989; Grimshaw, 1990), those had never been tested empirically.) Thus, the
 50 “privileged link” would give children a leg up in learning frighten-type verbs, and might even
 51 impede the acquisition of fear-type verbs.

52 The data reviewed above are certainly consistent with the privileged link hypothesis,
 53 but you can draw an infinite number of lines through a single point; there are a number
 54 of reasonable alternative explanations. One alternative explanation is **type frequency**:
 55 While fear-type verbs have high token-frequency, they are relatively rare as types. English
 56 admits around 251 frighten-type verbs but only 49 fear-type verbs (Kipper, Korhonen,
 57 Ryant, & Palmer, 2008). The difference becomes even more stark when one considers that
 58 frighten-type verbs are exemplars of an extremely robust pattern (caused changes of state

Figure 1. Taking frequency in child-directed speech into account, four year-olds were substantially more likely to successfully interpret who did what to whom for frighten-type verbs relative to fear-type verbs (Hartshorne et al., 2015). Figure used with permission.

being realized in transitive syntax with the AGENT as the subject and the PATIENT as the direct object), whereas fear-type verbs are quite unusual (only a handful of other verbs have EXPERIENCERS as subjects). According to a variety of learning theories – including both Nativist and Empiricist – this high type frequency should benefit frighten-type verbs, as evidence about any one verb in the class ultimately applies to the others as well (Goldberg, 2006; Pinker, 1989).

A second alternative is **differential salience**: Most utterances are contextually ambiguous: there are many things someone could be talking about (Gleitman, 1990; Quine, 1964). For instance, it is often the case that when there is fearing going on, there is also some frightening; the speaker can choose which to comment on. If children find frightening more salient than fearing, they may be more disposed to successfully identify labeling of frightening than labeling of fearing. Unfortunately, no data currently speak to whether there is a salience asymmetry between frighten-type events and fear-type events.

There are a number of other, less easily testable options. For instance, it is possible that children’s exposures to fear-type verbs are relatively short on highly informative “eureka” moments, making them effectively lower-frequency. While there are methods for quantifying the availability of eureka moments, they are time-consuming and difficult to scale (cf. Medina et al., 2011). It may be that the thoughts encoded by fear-type verbs (a habitual disposition towards a particular emotion) are more complex, harder to represent, or emerge later in cognitive development than those encoded by frighten-type verbs (an externally-caused change of emotional state). It is not clear at the moment how this possibility would be tested. It may be that acquisition of frighten-type verbs is aided by the acquisition of other, related argument realization patterns, or that the acquisition of fear-type verbs is impeded by interference from other, contrasting argument realization patterns. This is currently difficult to test because the nature of argument realization patterns remains highly controversial (Levin & Hovav, 2005), and because very little is known about their acquisition outside of a handful of patterns and languages (for comprehensive reviews, see Ambridge et

86 al., 2018, 2020). There are no doubt other possibilities as well.

87 **Goals of the present study**

88 The goals of the present study are to a) address some remaining concerns about the
89 methods used by Hartshorne, Pogue, & Snedeker (2015), and b) provide more points in
90 order to constrain our lines.

91 Thus, first we replicate and extend Hartshorne, Pogue, & Snedeker (2015). In order
92 to obtain more power and precision for our results, we expanded the number of verbs tested
93 from 12 to 16, tested a wider age range (3-6 vs. 4-5), and addressed some potential confounds
94 in the original study.

95 Second, fear/frighten verbs are only one of a number of examples of “perspective
96 pairs”, where two classes of verbs describe the same sorts of events but with different
97 argument structures. Thus, we extended the investigation to two additional case studies of
98 “perspective pair” verbs, allowing us to deconfound causality, token frequency, type frequency,
99 and salience. In Experiment 2, we focus on chase-type and flee-type verbs, which differ in
100 terms of whether the pursuer is the subject (*A chased/pursued/followed B*) or an oblique
101 object (*A fled/escaped/ran from B*).¹ In Experiment 3, we focus on give-type and get-type
102 verbs, which differ in whether the SOURCE is the subject and the GOAL is an oblique (*A*
103 *gave/passed/sold B to C*) or *vice versa* (*A got/grabbed/bought B from C*). (Note that, like
104 nearly all verbs, fear/frighten, chase/flee and give/get verbs can appear in other sentence
105 frames as well; here, we focus on the ones that show the contrast most cleanly.)

106 Critically, causal semantics, frequency, and salience play out differently across the
107 three verb class pairs, allowing us to disentangle their effects (Table 1).

¹While chase-type and flee-type verbs can describe the same events, there is evidence that they – like fear-type and frighten-type verbs – are semantically distinct. Gleitman (1990) reports a personal communication from Steven Pinker arguing that intentional participation in the event is entailed only for the subject of chase/flee verbs: one can chase something that is not fleeing (e.g., a storm) or flee something that is not chasing (e.g., a tsunami). So far as we know, there has not been any systematic study of the the verb classes to determine whether this generalizes, though initial inspection suggests that it does. In any case, the exact semantics of these verbs will not be critical for the present study.

Table 1

For each of several factors, if that factor was determinative, which verbs would be acquired first.

	Exp1	Exp2	Exp3
pair	fear/frighten	chase/flee	give/get
token frequency	fear	chase	neither
causality	frighten	neither	neither
type frequency	frighten	chase	give
salience	??	?chase	get

Causal Semantics. As reviewed above, the “privileged link” for causality hypothesis predicts earlier learning for frighten-type (which have a causal semantics) relative to fear-type (which do not). However, semantic analysis typically ascribes causality to the subject of both give-type and get-type verbs, and to *neither* chase-type nor flee-type (Kipper et al., 2008; Pinker, 1989). Thus if the “privileged link” explains the early learning of frighten-type verbs, we should not expect any differences for chase/flee or give/get.

Type Frequency. We assessed type frequency using VerbNet, the largest compendium of verb classes in English (Kipper et al., 2008). In terms of type-frequency, as reviewed above, frighten-type verbs ($N = 251$) are far more numerous than fear-type ($N = 49$).

Chase-type verbs ($N = 22$) are more common than flee-type, of which there appear to only be around 4. Counting is complicated in that while VerbNet records three groups of chase-type verbs (classes 35.1, 35.3, and 51.6), it does not index flee-type verbs, perhaps because there are so few. *Flee from*, *escape from*, and *retreat from* are included in class 51.1, and *run from* in class 51.3.2, but only in the sense of escaping from a *place*, not an entity. Because flee-type verbs are not indexed in VerbNet, it is possible there are more that we have overlooked, though probably not enough to place the class on par with chase-type. Interestingly, there is a class of verbs (class 52; $N = 11$) in which the flier is the subject and

the pursuer is the direct object (*A dodged/eluded/evaded B*). Because they take direct objects, they belong in a different syntactic category from chase/flee verbs. We did not investigate them in addition to or instead of chase/flee verbs because they are all low-frequency and unlikely to be encountered by children.

Finally, give-type verbs ($N = 82$; classes 13.1, 13.2, 13.4.1, 13.4.2) are more numerous than get-type verbs ($N = 57$; classes 13.5.1, 13.5.2, 57), but the imbalance is not as stark as for the other perspective pair classes.²

Salience. As already noted, there is no evidence bearing on the question of whether frighten-type event construals are more or less salient than fear-type construals, and we do not investigate it here (it is currently clear how to do so). With regards to chase/flee, it has been argued that children and adults are biased to encode ambiguous events in terms of chasing rather than fleeing (e.g., Landau & Gleitman, 2015). However, the evidence is thin and mixed. Fisher, Hall, Rakowitz, and Gleitman (1994) indeed found that three year-olds and adults describing one ambiguous chase/flee scene in terms of chasing rather than fleeing. However, Gleitman, January, Nappa, and Trueswell (2007) ran a similar study with adults using two stimuli, finding a chase-bias for one and a flee-bias for the other. There do not appear to be any other empirical studies.

There is a more robust literature indicating that when presented with an ambiguous give/get event, both children and adults focus the getting over the giving. Children and adults are more likely to remark on and remember GOALS than SOURCES (Freeman, Sinha, & Stedmon, 1981; Fujita, 2000; Lakusta & Landau, 2005, 2012; Papafragou, 2010; Regier & Zheng, 2007). This suggests that get-type verbs should be privileged for two reasons. First, the GOAL is obligatory for get-type verbs and optional for many give-type verbs (*A sent/sold/passed B [to C]*). Second, it is widely argued that more salient entities are preferentially mapped onto sentential subjects.

²We exclude verbs of future having (*A awarded/bequeathed/owed B to C*) from give-type verbs. The semantics are critically different in that there is no caused change of possession, just a promised change of possession. If these are included in the count, the advantage in token-frequency for give-type verbs becomes more imbalanced at 110 to 57.

In terms of language acquisition, there is only limited evidence as to an advantage for get-type (or give-type) verbs. In studies looking at spontaneous description, subjects often avoid both give-type and get-type verbs, instead preferring THEME-subject verbs (*A walked/ran/rolled to B*). Lakusta and Landau (2005) similarly found that children were more likely to use known give-type verbs than get-type verbs when describing ambiguous scenes (Lakusta & Landau, 2005).³ A smaller novel-word learning study showed a similar bias towards give-type verbs (Fisher et al., 1994). However, Bowerman (1990) reports in a diary study of two children that the first get-type verb (*get*) emerges in spontaneous speech only just prior (1;8) to the first give-type verbs (*give/gimme, tell, and read me*, all of which emerge at 1;9). Since get-type verbs appear first but give-type appear in larger numbers, this seems to be a draw.

Summary of Predictions

The predictions for the three experiments are summarized in Table 1.

- If token frequency is the key factor, we would expect earlier acquisition of fear-type verbs compared to frighten-type; chase-type compared to flee-type; and roughly equal learning for give-type and get-type. This outcome seems unlikely given the results of Hartshorne, Pogue, and Snedeker (2015), but perhaps our replication-and-extension will show different results.
- If the "privileged link" hypothesis is correct, we should see earlier acquisition of frighten-type than fear-type verbs, with no systematic differences among chase/flee verbs or give/get verbs.
- If token frequency is the key factor, we should see earlier acquisition of frighten-type than fear-type verbs; chase-type relative to flee-type; and give-type relative to get-type.

³Lakusta and Landau (2005) argue that the salience hypothesis predicts a preference for give-type verbs. In particular, they posit that children should prioritize the mapping to the prepositional phrase and thus find it easier to map the more salient GOAL onto the prepositional phrase, rather than the SOURCE. This is an intriguing notion, but it is incompatible with most prominent theories of language acquisition, and it remains to be seen whether a new theory can be constructed around it. It is also inconsistent with the data we present below. Thus, we do not consider this hypothesis further in this paper.

- If salience is the key factor, we have no strong predictions for fear/frighten or chase/flee verbs, but a strong prediction of earlier acquisition of get-type relative to give-type verbs.

Thus, by utilizing three case studies, we are able disentangle the predictions of the four main hypotheses under consideration.

Overview of analyses

In all three experiments, we submit the data to a mixed effects logistic regression with centered main effects of verb type, log frequency, and subject age in months.⁴ In order to improve convergence and avoid issues with singularity, we fit the model with partially-Bayesian regression with Wishart priors on the covariance matrix for random effects, using the blme package with bobyqa optimization (Chung, Rabe-Hesketh, Dorie, Gelman, & Liu, 2013).

We included random intercepts by subject and verb as well as a random slope of verb type by subject. We chose not to use a maximal random effects structure for three reasons. First, there is some debate about whether doing so is even desirable (Barr, Levy, Scheepers, & Tily, 2013; Bates, Kliegl, Vasishth, & Baayen, 2015; Matuschek, Kliegl, Vasishth, Baayen, & Bates, 2017). In particular, specifying overly complex models may substantially increase Type II error. Indeed, we found that for several of our analyses, a maximal random effects structure resulted in *no* significant effects, not even effects of age or log frequency, even though these are extremely well attested in the literature and clearly visible in our data. Second, the small number of verbs (ranging from 6 to 16 per study) means that random slopes are estimated based on very few data points. For instance, estimating a random slope for subjects for the interaction of verb type and log frequency involves calculations of slopes based on as few as 3 and at most 8 boolean responses. Third, empirically we found that in most cases random slopes led to worse fits (as assessed by BIC and AIC) – when fitting could even be done successfully. Analysis of the random effects structure for Exp.

⁴Not centering these variables frequency led to failures to converge.

3 did justify random slopes of verb type and log frequency by subject, but we felt using a different random effects structure for different experiments would impede comparison across experiments. Thus, we settled on a relatively simple random effects structure with only one random slope (verb type by subject), which is a fairly simple slope based on a relatively large amount of data.

While this is our primary analysis, we considered two other analysis methods. First, we fit the primary models using Bayesian regression as implemented in *brms* (Bürkner, 2017, 2018) and calculated Bayesian p-values with the `p_map` function from *bayestestR* (Makowski, Ben-Shachar, & Lüdtke, 2019).⁵ Unless data sets are very large, Bayesian analyses tend to be biased in favor of the null hypothesis, so it is not surprising that some of the effects that were significant in our main analyses are not significant in the Bayesian regressions, though for the most part the key results remained the same. The results of these models are reported in footnotes (in general, we use footnotes in lieu of supplementary materials). Second, we used model comparison to prune fixed effects from the primary model. Unfortunately, we obtained wildly different results using AIC, which generally favored retaining most or all of the fixed effects, and BIC, which is a very conservative method – particularly when data sets are not very large – and which generally favored eliminating most or even all of the fixed effects. Because the effects of log frequency and age are well-established and quite clear in the data, this suggests BIC is overly conservative. We do not describe these results in detail, but they are memorialized in our reproducible RMarkdown document, available at osf.io/k5xud. Because the Bayesian regression and BIC method both appear to be overly conservative, we do not consider their results as strongly indicative of the null, but it certainly does mean that the evidence we present below is not incontrovertibly strong.

We also report follow-up analyses for each experiment, focused each individual age group (3 year-olds, 4 year-olds, 5 year-olds, and 6 year-olds).

⁵For the intercept, fixed effects, and random effects, we used the relatively informative prior of a half-normal distribution centered at 0 and with a standard deviation of 1. We ran six chains of 8000 samples each, including 1000 warmup samples. For two of the models, we raised `alpha_adapt` to 0.9 in order to address a small number of divergent transitions.

R packages used to prepare this reproducible document include *papaja* (Aust & Barth, 2020), *knitr* (Xie, 2015), *ggplot2* (Wickham, 2016), *ggeffects* (Lüdtke, 2018), and *stargazer* (Hlavac, 2018).

Experiment 1: Fear/Frighten (Psych) Verbs

This experiment replicated the method of Exp. 1 of Hartshorne, Pogue, and Snedeker (2015), with the following changes: a) we added two new fear-type verbs and two new frighten-type verbs, b) we expanded the age range to 3 to 6, c) in order to accommodate the attention spans of the younger subjects, we split the verb lists so that each subject saw only half of the stimuli, and d) we presented half the subjects with critical verbs in the past tense, half in present tense.

The manipulation of tense requires some explanation. Tense affects fear-type and frighten-type verbs differently. While fear-type verbs refer to a habitual state whether used the present tense (*A fears B [always / *just the once]*) or past tense (*A feared B [always / *just the once]*), frighten-type verbs most naturally refer to a habitual or repeated event in the present tense (*A frightens B [always / *just the once]*) and a single event in the past tense (*A frightened B [?always / just the once]*). Unfortunately, Hartshorne, Pogue, and Snedeker (2015) did not explicitly control the tense used, so it is unclear whether the same tense was used for all verbs. It is unlikely this could explain poorer performance on fear-type verbs, which are unaffected by tense, but out of an abundance of caution, Exp. 1 included a tense manipulation in order to test the question directly.

Method

All research reported in this paper was approved by the Boston College and Harvard University Institutional Review Boards.

Materials. A total of 16 verbs were tested (8 fear-type and 8 frighten-type), including the 12 verbs from Hartshorne, Pogue, and Snedeker (2015) and 4 additional verbs. Part of the goal was to include three relatively high-frequency verbs (*enjoy*, *dislike*, *bug*) that

251 were overlooked in construction of the earlier experiments (Hartshorne, Pogue, and Snedeker
 252 (2015) did not have access to as complete a list of fear/frighten verbs as we do). The final
 253 verb (*anger*) is relatively low-frequency but was commonly used by children during a pilot
 254 verb-elicitation study and so was included.

255 We estimated frequencies for fear/frighten verbs in speech directed to children ages 36
 256 to 84 months (the age range in the present studies) by using childsr (Sanchez et al., 2019)
 257 to aggregate all speech (other than that by the target child) in North American English
 258 corpora in CHILDES (MacWhinney, 2000) where the target child was in the age range
 259 and where part-of-speech tags were available. This aggregate corpus consisted of 3,008,582
 260 tokens. Frequencies for our stimuli are shown in Table 2. A disproportionate number of
 261 high-frequency words were fear-type, though the difference between types across the 16 verbs
 262 did not reach significance ($\Delta M = 0.17$, 95% CI $[-2.15, 2.50]$, $t(11.44) = 0.16$, $p = .872$).⁶

Table 2

Frequency of verbs used in Experiment 1, in parts per million

type	verb	Frequency	LogFrequency
frighten	scare	216.7	5.4
frighten	surprise	57.2	4.1
frighten	bore	28.3	3.4
frighten	amaze	20.3	3.1
frighten	confuse	16.6	2.9
frighten	frighten	16.0	2.8
frighten	bug	9.0	2.3
frighten	anger	0.0	0.0
fear	like	2271.5	7.7
fear	love	272.2	5.6
fear	hate	55.8	4.0
fear	enjoy	49.5	3.9

⁶For these and other statistics, we used log-transformed frequencies.

type	verb	Frequency	LogFrequency
fear	trust	4.7	1.7
fear	fear	2.3	1.2
fear	admire	1.7	1.0
fear	dislike	0.3	0.3

For each of the 16 verbs, we created four scenarios, counterbalancing the pair of characters involved and which member of the pair experienced the emotion. We created two stimulus orders as follows: We did a mid-line split on token frequency for both fear-type and frighten-type verbs and placed the 4 highest-frequency of each type in the first half the list. We then created a second list reversing the orders *within* each half, so that the high-frequency verbs remained in the first half. The purpose of this was to allow participants to continue on to the second half if they have sufficient interest, but this rarely happened. Ultimately, we began randomly assigning participants to one half or the other. (For those participants from the initial phase of testing who had completed both halves, we excluded the second half.) The results was 64 lists: 2 (item set) by 2 (present vs past) by 2 (animal pair) by 2 (animal roles) by 2 (item order) by 2 (target response) design.

During testing, we discovered that an error in the lists such that for four of the low-frequency verbs (*bug*, *anger*, *enjoy*, *dislike*) only one of the character pairs was used. We fixed the lists and began replacing those subjects (N=27) who had not seen the intended character pairs. However, because we were unable to finish replacing them due to the pandemic, we included both the original and “replacement” subjects in order to improve power. Nonetheless, excluding the to-be-replaced subjects has no effect on the qualitative pattern of results.

Participants. We recruited 290 native English-speaking children between the age of 3 and 6 from the Greater Boston Area from parks, museums, and preschools. Of these, we excluded 47 children due to coding error or failure to complete more than 50% of the

experiment or both. We had intended to have 64 subjects per age group (one per list; see below), but delays due to the pandemic made it impossible to reach the full complement in some cases. In others, we ended up with more subjects than intended for the reasons described in “Procedure”. The final numbers included 54 3 year-olds, 73 4 year-olds, 82 5 year-olds, and 34 6 year-olds.

Procedure. Participants were run one at a time. An experimenter read stories accompanied by pictures. After each story, a second experimenter used a Mickey Mouse puppet to say what Mickey thought happened in the story. Sometimes what Mickey said was correct and sometimes it was incorrect. If the participant thought that what Mickey said was correct, they gave Mickey a cookie. If they thought that what Mickey said was incorrect, they gave him coal. The experimenter began with two practice stories involving action verbs (*hug* and *kiss*), followed by the eight test trials. Participants received explicit feedback on the two practice trials but only general affirmative reactions to the critical trials.

Results

Data for this and all experiments are available at osf.io/k5xud.

We conducted a preliminary analysis to assess whether tense systematically affected the results. We submitted the results to a partially-Bayesian mixed effects logistic regression with Wishart priors on the covariance matrix for random effects, using the *blme* package with *bobyqa* optimization (Chung et al., 2013). We included main effects of verb type and tense, along with their interaction, and random slopes of verb type by subject and tense by verb.⁷ The main effect of tense was not significant ($B = -0.13$, $SE = 0.15$, Wald’s $z = -0.88$, $p=0.38$), nor was the interaction of tense and type ($B = 0.17$, $SE = 0.29$, Wald’s $z = 0.59$, $p=0.56$). Thus, all subsequent analyses collapsed across tense.

We then conducted our main analyses, as described above in “Overview of Analyses.”

⁷This is actually a maximal design. The random slope of verb type by subject is included in all analyses, so its inclusion requires no additional justification. The random slope of tense by verb is a fairly simple slope (the effect is categorical) and Exp. 1 has a relatively large number of verbs (16). In any case, excluding this random slope does not change the pattern of results. Note that the random slope of tense is not included in other analyses because it is not relevant to those analyses.

Table 3
Regression results for psych verbs

	Estimate	Std. Error	z value	Pr(> z)
(Intercept)	0.352	0.074	4.761	0.00000
Age	0.289	0.056	5.156	0.00000
VerbType	0.246	0.144	1.700	0.089
LogFrequency	0.245	0.082	3.000	0.003
Age:VerbType	0.379	0.108	3.517	0.0004
Age:LogFrequency	0.112	0.060	1.869	0.062
VerbType:LogFrequency	0.343	0.162	2.118	0.034
Age:VerbType:LogFrequency	0.193	0.119	1.631	0.103

Results are shown in Table 3. The interaction of verb type by log frequency was significant, reflecting a larger effect of frequency on frighten-type verbs, as was the interaction of verb type by age, reflecting a larger advantage of frighten-type verbs in older subjects.⁸

As explained in “Overview of Analyses,” we followed up this analysis by considering each age group separately. For 3 year-olds, neither the main effects nor their interaction were significant ($ps \geq 0.12$). For 4 year-olds, the only effect to approach significance was that of log frequency ($B = 0.26$, $SE = 0.14$, Wald’s $z = 1.90$, $p=0.061$). By 5, there was a significant effect of log frequency ($B = 0.49$, $SE = 0.13$, Wald’s $z = 3.80$, $p=0.00014$), a marginal effect of verb type ($B = 0.41$, $SE = 0.23$, Wald’s $z = 1.80$, $p=0.076$), and a significant interaction ($B = 0.70$, $SE = 0.25$, Wald’s $z = 2.80$, $p=0.0056$). By the age of 6, the only significant effect was that of verb type ($B = 1.20$, $SE = 0.55$, Wald’s $z = 2.10$, $p=0.033$). Note that the sample sizes for 3 year-olds and 6 year-olds were smaller, perhaps explaining the fewer number of significant effects.

Fig. 2 visualizes these results along three dimensions. There is a clear overall advantage for frighten-type verbs beginning at around 50 months (Fig. 2 Top). Interestingly, while there is a clear upward trajectory for frighten-type verbs beginning at around 48 months, there is no clear pattern for fear-type verbs, on which performance at 84 months is similar

⁸The Bayesian regression revealed significant effects of age ($p < 0.00000001$) and log frequency ($p=0.034$), and a significant interaction of age and verb type ($p = 0.0023$)

to that at 36 months.

Fig. 2 Middle provides another window onto the pattern of results: children exhibit higher performance on frighten-type verbs, controlling for frequency, by 4 years old, with the effect growing substantially by 6 years old. This figure is consistent with our age-group-specific analyses: 3 and 4 year-olds appear to be close to chance on most items, though there is a hint of better performance for high-frequency verbs by the age of 4. 5 year-olds show a clear pattern of success largely restricted to high-frequency frighten-type verbs. By the age of 6, children are doing well on all the frighten-type verbs but performance remains low for most fear-type verbs.

The apparent lack of improvement with age on fear-type verbs may be due to learners misclassifying low-frequency fear-type verbs as frighten-type (Fig. 2 Bottom). The highest-frequency fear-type verbs (*like, love, hate*) may not have reached the performance levels of frighten-type verbs, but they did show gradual improvement across the age range. In contrast, four of the lower-frequency verbs (*enjoy, trust, dislike, fear*) actually showed a *decline* from 3 to 4 years old, with all four verbs actually declining to below-chance levels. This matches the results of Hartshorne, Pogue, and Snedeker (2015), who reported significantly below-chance performance on *trust* and *fear* (*enjoy* and *dislike* were not tested).

Discussion

The results of Exp. 1 confirmed the results of Hartshorne, Pogue, and Snedeker (2015) with more verbs and across a larger age range: children began learning frighten-type verbs by the age of 4-5, whereas fear-type verbs remained largely at near-chance levels even at the age of 6. As suggested by Hartshorne, Pogue, and Snedeker (2015), this appears to reflect a tendency to misanalyze fear-type verbs as being frighten-type verbs.

These results are consistent with both the privileged link hypothesis and the type frequency hypothesis. It is unknown how it matches the salience hypothesis, because we do not know whether fear-type or frighten-type event representations are more salient.

Figure 2. Results from Exp. 1. **Top:** Averaging across verbs within type, with LOESS smoothing over age. **Middle:** Accuracy for each verb against log frequency, split into four age groups, with linear regressions shown. **Bottom:** Performance on each verb across ages, aggregated into four age groups, with LOESS smoothing.

Experiment 2: Chase/Flee

Method

Materials. We selected the 3 highest-frequency chase-type verbs (*chase*, *follow*, *pursue*) and 3 highest-frequency flee-type verbs (*flee*, *escape*, *run*) (Table 4). Frequencies were determined as in Exp. 1, with the caveat that we restricted frequency counts for *run* to specifically the bigram *run from*. Two puppet subjects, Giraffe and Tiger, were used for each sentence. We created 6 stories which participants needed to act out with the two puppets. Each story was one sentence long (e.g., *Giraffe chased Tiger*) to minimize demands on working memory and clearly isolate the verb of interest. Giraffe and Tiger were counterbalanced between acting as the subject and the object for each verb. Four lists were created by counter-balancing across lists whether Giraffe or Tiger was the subject of each verb and by creating two item orders, one of which was the reverse of the other.

Table 4

Frequency of verbs used in Experiment 2, in parts per million

Type	Verb	Frequency	LogFrequency
chase	follow	91.4	4.5
chase	chase	76.8	4.4
chase	pursue	0.0	0.0
flee	run from	8.6	2.3
flee	escape from	4.7	1.7
flee	flee from	0.3	0.3

Participants. We recruited 208 children ages 3 through 6 from the Greater Boston Area. All participants were native English speakers. We excluded 41 children for failing to complete more than half of the items or experimenter error. Although we aimed for 40 subjects per age (10 per list), we ultimately obtained 43 3 year-olds, 43 4 year-olds, 40 5 year-olds, and 41 6 year-olds.

Procedure. The researchers began by familiarizing the participants with the puppets. Participants were then informed that they would be listening to stories about Giraffe and Tiger, and that they should act out those stories using the puppets. The stories were read aloud, and the researchers recorded whether or not the child correctly demonstrated the meaning of the verb using the puppets. This was operationalized by the child moving the puppets in the correct direction (i.e. the Giraffe towards the Tiger for “Giraffe chased Tiger”) or by reorienting the chaser to be facing the correct direction (i.e. the Giraffe faced the Tiger for “Giraffe chased Tiger”). The on-site record was later checked by another researcher using a video recording of the study.

Results

The mixed effects logistic regression revealed that every main effect and interaction was significant except the interaction of verb type and log frequency, which trended towards significance (Table 5).⁹ To aid interpretation of the three-way interaction between verb type, age, and log frequency (Table 5), we plotted marginal effects (Fig. 3). These show that, when frequency is held constant, the flee-type verbs are learned faster.

As explained in “Overview of Analyses,” we followed up this analysis by considering each age group separately. For 3 year-olds, neither the main effects nor their interaction were significant ($ps \geq 0.12$). For 4 year-olds, there were significant main effects of log frequency ($B = 1.30$, $SE = 0.43$, Wald’s $z = 3$, $p=0.0025$) and verb type ($B = -1.40$, $SE = 0.69$, Wald’s $z = -2$, $p=0.043$), though their interaction was not significant ($B = -0.52$, $SE = 0.87$, Wald’s $z = -0.60$, $p=0.55$). At age 5, there remained significant main effects of log frequency ($B = 2.90$, $SE = 0.65$, Wald’s $z = 4.50$, $p=0.0000078$) and verb type ($B = -2.70$, $SE = 1.10$, Wald’s $z = -2.50$, $p=0.014$), and the interaction approached significance ($B = -2.30$, $SE = 1.30$, Wald’s $z = -1.80$, $p=0.073$). By the age of 6, all effects were significant: log frequency ($B = 15$, $SE = 3.50$, Wald’s $z = 4.20$, $p=0.000028$), verb type ($B = -13$, $SE = 4.20$, Wald’s z

⁹The results of the Bayesian regression were reasonably similar, with a significant main effect of age ($p < 0.00000001$) and a trend towards an effect of log frequency ($p = 0.074$), as well as significant interactions of age and log frequency ($p < 0.00000001$) and of age x log frequency x verb type ($p = 0.016$).

Figure 3. Marginal effects for the interaction of verb type, (z-scored) age, and frequency in chase/flee verbs (Exp. 2). Note that the curves plotted are inferred by the model; they do not represent direct observations. For instance, there was no flee-type verb with a log frequency of 0.

Table 5

Regression results for chase/flee verbs, no interactions with Frequency

	Estimate	Std. Error	z value	Pr(> z)
(Intercept)	1.364	0.279	4.890	0.00000
Age	1.008	0.186	5.423	0.00000
VerbType	-1.732	0.506	-3.421	0.001
LogFrequency	1.675	0.320	5.237	0.00000
Age:VerbType	-0.676	0.286	-2.359	0.018
Age:LogFrequency	0.918	0.167	5.507	0.00000
VerbType:LogFrequency	-1.104	0.643	-1.717	0.086
Age:VerbType:LogFrequency	-1.038	0.326	-3.189	0.001

393 = -3, $p=0.0025$), and their interaction ($B = -24$, $SE = 5.80$, Wald's $z = -4.10$, $p=0.00004$).

394 Fig. 4 visualizes these results along three dimensions. Across ages, there is no obvious
395 difference between chase-type and flee-type verbs (Fig. 4 Top). However, at all ages there is a
396 clear numerical advantage for flee-type verbs once frequency is taken into account, something
397 visible in all four age groups (Fig. 4 Middle). This is consistent with the significant three-way
398 interaction observed in the statistical results. It is also visible in the by-verb plots (Fig. 4
399 Bottom): accuracy increases with age nearly as fast for *escape from* and *run from* as for
400 *chase* and *follow*, despite being much lower frequency (Table 4) .

401 Discussion

402 The results of Exp. 2 indicate an early advantage of flee-type verbs relative to
403 chase-type verbs, holding frequency constant, albeit perhaps not as pronounced as that for
404 frighten-type verbs relative to fear-type verbs. This militates against a key role for token
405 frequency or (possibly) salience. The results are also inconsistent with the privileged link
406 hypothesis, in that no privileged link is expected to apply to either chase-type or flee-type
407 verbs, and thus they should have been learned equally quickly, modulo token frequency.

408 Interestingly, one chase-type verb and one flee-type verb each exhibited stubbornly
409 sub-chance levels: Children responded as if *pursue* was flee-type and *flee from* was chase-type
410 – something that had not resolved even by the age of 6. This is puzzling, and we do not

Figure 4. Results from Exp. 2. **Top:** Averaging across verbs within type, with LOESS smoothing over age. **Middle:** Accuracy for each verb against log frequency, split into four age groups, with linear regressions shown. **Bottom:** Performance on each verb across ages, aggregated into four age groups, with LOESS smoothing.

411 have much to say about it at the moment other than that it is further evidence against a
 412 privileged link applying to chase/flee verbs, since that would entail misanalysis of *either*
 413 chase-type *or* flee-type verbs, not both.

414 It should be noted that the observed differences between chase-type and flee-type verbs
 415 are due primarily to four verbs (the differences in learning outcomes for *pursue* and *flee from*
 416 are fairly similar). Thus, different results for a single verb would have substantially changed
 417 the statistical outcomes. Unfortunately, because flee-type verbs are so rare type-wise, not
 418 much can be done about this other than to consider other perspective pairs as well.

419 Experiment 3

420 Method

421 **Materials.** The four give-type verbs and four get-type verbs were selected from the
 422 Verbnet Unified Verb Index. These were the four highest-frequency verbs that we were able
 423 to identify for each type (Table 6). The two groups did not differ significantly in terms of
 424 frequency ($\Delta M = 0.17$, 95% CI $[-4.38, 4.73]$, $t(4.28) = 0.10$, $p = .922$).

Table 6
Frequency of verbs used in Experiment 3, in parts per million

type	stem	Frequency	LogFrequency
get	get	7026.2	8.9
get	buy	428.4	6.1
get	grab	66.1	4.2
get	receive	4.7	1.7
give	give	1333.9	7.2
give	pass	86.8	4.5
give	send	76.1	4.3
give	sell	62.8	4.2

We constructed videos depicting one male and one female experimenter exchanging objects with one another. Fig. 5 shows an example: the man gives an apple to the woman, who then reciprocates by giving a hammer to the man. The participant would then be asked either *what did the boy give?* or *what did the girl give?* The fact that we could query either character allowed us to counter-balance SOURCES and GOALS within the same stimuli. We counter-balanced another way as well: we made pairs of videos that differed in which character moved first, which should wash out any bias towards the first- (or last-) mover. Similarly, the videos were designed to allow querying one give-type verb and one get-type verb: the same video was used for either *give* or *get*, another for *send* or *receive*, and yet another for *buy* and *sell*. An exception to this design was forced by *grab*, which was paired with *pass*. For obvious reasons, the same video could not be used for both. Moreover, we found videos with reciprocal grabbing to be confusing. Thus, rather than have the man and woman both grab from each other within the same video, we created two videos – one with grabbing by the man and one with grabbing by the woman – and placed them one above the other. Instead of counter-balancing which character acted first, we counter-balanced which video was on top. Because *pass* was matched with *grab* for purposes of counter-balancing, we constructed the *grab* verbs in a similar way.

Four lists were constructed by counterbalancing the order in which the verbs were queried and, for a given verb, which character was queried and which character moved first in the video. To achieve this, we used a single fixed order of the videos (which a caveat describe below), counter-balancing which character is asked about and which verb in each verb pair was queried. The caveat is that this was obviously not possible for *pass* and *grab*. Instead, for these we swapped the order of the videos (so on two lists, *pass* was presented before *grab*, and the reverse was true for the other two lists) and the vertical positioning of the videos (see above).

Finally, these four lists were triplicated by making two more sets of videos, each with a different pair of actors and different sets of objects, for a total of 12 lists.

Participants. We recruited 452 native English-speaking children ages 36 to 83 months old from the Greater Boston Area. We excluded 28 children due to experimenter error or failure to complete more than half the trials. While we had intended to recruit 9 subjects per age group per list, we fell 8 subjects short due to recruitment restrictions. Thus, we finished with 424: 107 3 year-olds, 105 4 year-olds, 109 5 year-olds, and 104 6 year-olds.

Procedure. Videos were presented on an iPad using Keynote. Prior to watching each video, participants viewed the opening frame and were asked to point to the target objects, with corrections provided as necessary. This helped ensure they could identify the objects well enough to answer the subsequent questions. The video was then played, and the participant was asked the question involving the give/get verb for that video.

Results

The three-way interaction of age, log frequency, and verb type was significant, as was the interaction of age and verb type and the main effects of frequency and age (Table 7).¹⁰ The data plots provide context. Averaging within verb-type, there is a clear early advantage for get-type verbs, which disappears by age 6 in part because performance improves with age for give-type but not get-type verbs (Fig. 6 Top). Plots of individual verbs (Fig. 6 Bottom) show steady improvement with age for all give-type verbs, with much less improvement for three of the get-type verbs and a substantial decline in performance for *receive*.

Fig. 6 provides some context for the three-way interaction between age, log frequency, and verb type: There is initially an advantage for get-type verbs – particularly low-frequency ones – but that dissipates by ages 4 and 5 and may even start to reverse by age 6. These qualitative observations were largely confirmed by quantitative analysis. For 3 year-olds, the main effect of log frequency was significant ($B = 0.92$, $SE = 0.31$, Wald's $z = 3$, $p=0.0028$) and the main effect of verb type trended towards significance ($B = -0.87$, $SE = 0.50$, Wald's $z = -1.70$, $p=0.081$), though their interaction was not significant ($B = 0.92$, $SE = 0.62$,

¹⁰The Bayesian regression showed a trend towards the three-way interaction ($p=0.08$), a significant interaction of age and verb type ($p<0.00000001$), and a significant main effect of age ($p<0.00000001$). The other effects did not come even close to significance.

Figure 5. Stills from one of the videos for *give*, used in Exp. 3.

Table 7
Regression results for *give/get* verbs

	Estimate	Std. Error	z value	Pr(> z)
(Intercept)	1.604	0.314	5.107	0.00000
Age	0.614	0.064	9.545	0
VerbType	0.097	0.631	0.154	0.877
LogFrequency	0.976	0.404	2.417	0.016
Age:VerbType	0.725	0.141	5.143	0.00000
Age:LogFrequency	0.036	0.078	0.464	0.643
VerbType:LogFrequency	0.281	0.807	0.348	0.728
Age:VerbType:LogFrequency	-0.358	0.156	-2.304	0.021

Wald's $z = 1.50$, $p=0.14$). At ages 4 and 5, the only effect to even approach significance was that of log frequency (4 year-olds: $B = 1.10$, $SE = 0.48$, Wald's $z = 2.30$, $p=0.021$; 5 year-olds: $B = 1.10$, $SE = 0.52$, Wald's $z = 2.10$, $p=0.035$). At the age of 6, no effects or interactions were significant ($ps \geq 0.16$, $SE = 0.95$, Wald's $z = 0.41$, $p=0.68$).

Discussion

As with psych verbs and chase/flee verbs, the give/get verbs showed a clear effect of verb type. However, it was in many ways the opposite of what was observed for the first two perspective pairs: rather than an advantage for one verb-type emerging between the ages of 3 and 6, we observed an early advantage for get-type verbs that dissipated.

It may not be necessary to make much of this difference. Three year-olds showed much higher accuracy on give-type and especially get-type verbs than they did on either type of psych verbs or chase/flee verbs – something that is consistent with their relatively early emergence in spontaneous speech (Bowerman, 1990) and extremely high token frequency. Get-type verbs must have diverged from give-type verbs at an earlier age, and had we tested two year-olds, our results may have looked (more) qualitatively similar to those of the other two perspective pairs, simply shifted earlier in development.

Regardless, the findings again contradict the type-frequency hypothesis, which anticipated, if anything, an advantage for give-type verbs. The results are, however, consistent with the salience hypothesis.

On the face of it, the results are inconsistent with the privileged link hypothesis. As reviewed above, both give-type and get-type verbs are analyzed as involving events caused by their subjects, so they should, on this hypothesis, be equally easy to learn.

However, there are the puzzling U-shaped results for *receive*. Given the very large number of subjects, statistical fluke seems unlikely. One possibility is that this reflects a misanalysis of the meaning of *receive*, much like what we observed for some fear-type verbs. While VerbNet analyzes all four of our get-type verbs as having causal semantics (roughly:

Figure 6. Results from Exp. 3. **Top:** Averaging across verbs within type, with LOESS smoothing over age. **Middle:** Accuracy for each verb against log frequency, split into four age groups, with linear regressions shown. **Bottom:** Performance on each verb across ages, aggregated into four age groups, with LOESS smoothing.

cause-self-to-have), this is far more intuitively true for *grab*, *get*, and *buy* than it is for *receive*. Critically, the VerbNet analysis is backed by linguistic analysis and not any systematic empirical data. If *receive* differs from the other give/get verbs in lacking a causal semantics – and if there is indeed a privileged link between causal agents and sentential subjects – then the same explanation of U-shaped learning we considered for fear-type verbs might apply. Specifically, children may initially misanalyze *receive* as having a causal agent as its subject. Note, however, that a separate explanation would still be needed for the early success on the other three get-type verbs relative to give-type verbs.

General discussion

We investigated the development of high-frequency verbs in three “perspective pair” classes: emotion (psych) verbs, chase/flee verbs, and give/get verbs. Taking token frequency into account, we found substantially earlier acquisition of frighten-type verbs relative to fear-type verbs, a smaller but significant advantage for flee-type verbs relative to chase-type verbs, and a much earlier emergence of get-type verbs as compared to give-type verbs (though by four years of age, this difference had dissipated).

In all three experiments, we saw some evidence of U-shaped learning. This was clearest for fear-type verbs, where fully half of them showed a decline to below-chance performance – replicating and extending earlier observations by Hartshorne, Pogue, and Snedeker (2015). However, we also saw apparent U-shaped learning and/or below-chance performance for two frighten-type verbs (*anger*, *amaze*), a chase-type verb (*pursue*), a flee-type verb (*flee*), and a get-type verb (*receive*). The large numbers of comparisons suggest some caution, but given the large samples and the prior observations of Hartshorne, Pogue, and Snedeker (2015), it is likely that at least some of these findings are reliable.

Methodological limitations

Before discussing the theoretical implications, we consider several limitations in the data and the strength of the evidence.

Most broadly, we only investigated the acquisition of 30 verbs across three perspective pair classes in a single language. This is certainly a substantial improvement on prior work, which focused on narrow age ranges and dealt primarily with fear/frighten verbs in English (Bowerman, 1990; Braine et al., 1993; Hartshorne et al., 2015; Messenger et al., 2012; Tinker et al., 1988), with the exception of one study of the emergence in the spontaneous speech of two children of four give/get verbs (Bowerman, 1990). A few additional studies looked at relative preference for one perspective pair or the other in elicited naming, with mixed results (Fisher et al., 1994; Gleitman et al., 2007; e.g., Lakusta & Landau, 2005). While our studies included most of the verbs in the three perspective pair classes that were sufficiently high frequency in child-directed speech to be plausibly known by young children, it is quite possible the model fits would be different if other verbs had been available. Moreover, the results of these three perspective pairs may be unrepresentative even within English, and may not generalize to other languages – particularly languages that organize argument structure differently, such as ergative or agglutinative languages. It would certainly not be the first time such generalization failed (Evans & Levinson, 2009; Yarkoni, 2019).

Related to that point, while the sample sizes are not small by the standards of language development research, they are nonetheless not large by the standards of statistical analysis, particularly when investigating three-way interactions (Hartshorne & Schachner, 2012; Vankov, Bowers, & Munafò, 2014). Some comfort is given by the fact that in all three experiments, the key interactions that were significant in our partially-Bayesian analyses were also significant or trended towards significance in the more conservative fully-Bayesian analyses. (The lone exception was the interaction of age and log frequency in Exp. 1.) However, interpretation of the p-values we report must take into account the fact that the small number of verbs precluded fitting maximal random effects, which may or may not be anti-conservative (Barr et al., 2013; Bates et al., 2015; Matuschek et al., 2017).

Similarly, while the estimates of frequency in child-directed speech are based on the largest dataset available (all part-of-speech-tagged corpora involving native English-speaking,

556 North American children in CHILDES-db in the focused age range) the resulting dataset is
557 not actually very large – around 3 million tokens, which is far less than what one child hears
558 in a single year (Hart & Risley, 1995) and is likely to be biased by the sampling strategies used
559 by the researchers. Moreover, many of the part-of-speech tags are automatically generated
560 and of variable quality. Additionally, with the exception of *run from*, we counted all uses of
561 the verbs, rather than uses in the critical syntactic frames. This was driven by a practical
562 consideration (automatically extracting syntactic frames is difficult, particularly for spoken
563 corpora), but determining whether it is reasonable theoretically will require more exact
564 theories of language acquisition than we currently have.

565 The three experiments use three different methods: Truth Value Judgment, act-out,
566 and question-answering. These differences were driven by the semantics of the verbs, which
567 rendered different methods more or less natural. Consider, for instance, how one would run
568 an act-out task for fear-type verbs, which describe a habitual disposition, not any particular
569 action. Nonetheless, this methodological variation limits direct cross-study comparison.

570 Relatedly, the chase/flee experiment has a potential confound in that flee-type verbs
571 have prepositions but chase-type verbs do not. In principle, children might have responded
572 by ignoring the verb and focusing on the preposition. This can, of course, cut both ways:
573 the particular sentence frame used by chase-type verbs is relatively rare compared to the
574 transitive, which may make it more difficult for children to comprehend. This consideration
575 just further highlights the need for data on more different kinds of perspective pair verbs.

576 Finally, as noted in the Introduction, the present studies test whether participants
577 understood who did to whom. They do not directly address *what* was done. It is possible
578 that subjects succeeded on these tasks without understanding the differences between *love*
579 and *hate* or between *grab* and *get*.

Theoretical implications

With the above caveats, the clearest finding is that it is not clear what drives successful verb learning. For all three perspective pairs, learning systematically departed from token frequency. While the emotion verb findings and give/get findings were consistent with the privileged link hypothesis, the chase/flee and give/get results largely contradicted it. The same was true of the type frequency hypothesis: The fear/frighten verb results support the hypothesis but the chase/flee and give/get results contradict it.

The salience hypothesis is more difficult to evaluate. It makes no predictions about fear/frighten verbs, since we currently do not know which perspective is more salient. It correctly predicts the earlier learning of get-type verbs – at least, if one adopts our interpretation of prior work on goal salience (see above). As we reviewed in the Introduction, there is currently some suspicion that the “chase” perspective is more salient than the “flee” perspective, which would be inconsistent with our finding of early learning of flee-type verbs. However, that suspicion is based on sparse data with inconsistent results. So one could reasonably argue that we have no predictions about chase/flee verbs, either.

In summary, the privileged link hypothesis is 1 for 3, the token frequency hypothesis is 1 for 3, and the salience hypothesis is at best 1 for 1 and at worst 0 for 2. Thus, salience might do a bit better, but perhaps only because its predictions are less clear. Thus, one obvious direction for future research is to obtain a clearer understanding of which perspective is more salient for different perspective pairs, in order to better test the hypothesis. This might be aided by developing a more nuanced, precise version of the hypothesis. Salience is a phenomenon that requires its own explanation. Perhaps a perspective is more salient because it is more simply represented (chasing involves a simple goal of being where the target is, whereas the goal of fleeing involves a negation: not being where the pursuer is), because it is more temporally concentrated (frightening happens at a distinct time and place, whereas fearing is an ongoing state of affairs), or because of a recency effect (*get* highlights the end of an event, whereas *give* highlights its beginning; cf. Regier and Zheng (2007)). A more

precise account would allow us to make predictions without necessarily having to obtain direct evidence of salience (i.e., figuring out how to measure the salience of hard-to-depict event perspectives like FEAR). For instance, researchers working on the psychophysics of action perception found it fairly straightforward to make artificial agents that chase (a simple heat-seeking policy works quite well), whereas designing artificial agents that could flee effectively required a much more complex policy (Tang et al., 2021). While this actually makes the wrong prediction in the present study (chase-type verbs were learned later, not earlier), a well-defined, quantitative simplicity-based theory of salience is potentially within reach.

Another direction would be to better characterize the quality of learning opportunities. Our findings are based on acquiring verbs earlier or later than might be expected based on input frequency. This implicitly assumes that all encounters with a verb are equally informative, which is not the case. As we reviewed above, Medina and colleagues (2011) argue that word learning is primarily driven by the rare highly informative encounter. While they present a method for identifying these ‘eureka’ moments, it requires time-intensive hand annotation. Currently-available annotated corpora are vanishingly small; developing one large enough to test whether frighten-type, flee-type, and get-type verbs have more than their fair share of eureka moments will either require an enormous amount of work or some mechanism for automating the annotation (e.g., through machine learning). Note, moreover, Hartshorne, Pogue, and Snedeker (2015) raise some reasons for being skeptical about this explanation, at least with respect to fear/frighten verbs. They suspect that because frighten-type events are ephemeral, speakers are unlikely to remark upon them as they happen, whereas because fear-type states are ongoing, the reverse may be true for them. As a result, it may be easier for children to connect a fear-type utterance with its co-temporal referent than a frighten-type utterance with its non-co-temporal referent.

Encounters with verbs can be more or less informative in other ways as well. Hartshorne, Pogue, and Snedeker (2015) reported that high-frequency fear-type verbs such as *like*, *love*,

634 and *hate* occur primarily with one of a small number of subjects – mostly *I* and *you*. They
635 note that this might induce children to treat these high-frequency bigrams as set phrases.
636 As a result, most uses of these verbs would fail to provide much information about their
637 argument structure. In that case, the frequency analyses above should be redone excluding
638 those high-frequency bigrams. This is not trivial – for one thing, it requires a principled
639 way of determining which bigrams are of sufficiently high frequency – and we leave it to
640 future work.

641 Conclusion

642 Hartshorne, Pogue, and Snedeker (2015) reported a puzzling finding: relatively old
643 children failed to understand extremely high-frequency fear-type verbs long after they had
644 acquired a number of lower-frequency frighten-type verbs. They proposed a number of
645 possible explanations based on current understanding of verb-learning. Of the ones that
646 are currently testable, none are clearly consistent with the present results. More broadly,
647 the current study revealed that this is not a funny fact about fear/frighten verbs, but may
648 in fact be a common phenomenon – one that has gone largely undetected and remains
649 essentially without explanation. This conclusion is based on only three case studies in a
650 single language: the empirical picture may be even more complicated than it appears so far.
651 All we can say at the moment is that something is going on, and we do not understand it.
652 This should concern us, because if we are missing a large part of the empirical description
653 of language acquisition, our theorizing may be entirely misdirected. At the very least, it is
654 incomplete. There is more in heaven and earth than is dreamt of in our philosophy, so some
655 new philosophy is needed.

- 656 Ambridge, B., Barak, L., Wonnacott, E., Bannard, C., Sala, G., Zwaan, R., & Ferreira,
657 F. (2018). Effects of both preemption and entrenchment in the retreat from verb
658 overgeneralization errors: Four reanalyses, an extended replication, and a meta-analytic
659 synthesis. *Collabra: Psychology*, 4(1).
- 660 Ambridge, B., Maitreyee, R., Tatsumi, T., Doherty, L., Zicherman, S., Pedro, P. M., ...
661 others. (2020). The crosslinguistic acquisition of sentence structure: Computational
662 modeling and grammaticality judgments from adult and child speakers of english, japanese,
663 hindi, hebrew and k'iche'. *Cognition*, 202, 104310.
- 664 Aust, F., & Barth, M. (2020). *papaja: Create APA manuscripts with R Markdown*. Retrieved
665 from <https://github.com/crsh/papaja>
- 666 Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for
667 confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*,
668 68(3), 255–278.
- 669 Bates, D., Kliegl, R., Vasishth, S., & Baayen, H. (2015). Parsimonious mixed models. *arXiv*
670 *Preprint arXiv:1506.04967*.
- 671 Belletti, A., & Rizzi, L. (1988). Psych-verbs and θ -theory. *Natural Language & Linguistic*
672 *Theory*, 291–352.
- 673 Bowerman, M. (1990). Mapping thematic roles onto syntactic functions: Are children helped
674 by innate linking rules? *Linguistics*, 28(6), 1253–1289.
- 675 Braine, M. D. (1992). What sort of innate structure is needed to “bootstrap” into syntax?
676 *Cognition*, 45(1), 77–100.
- 677 Braine, M. D., Brooks, P. J., Cowan, N., Samuels, M. C., & Tamis-LeMonda, C. (1993).
678 The development of categories at the semantics/syntax interface. *Cognitive Development*,
679 8(4), 465–494.
- 680 Bürkner, P.-C. (2017). brms: An R package for Bayesian multilevel models using Stan.
681 *Journal of Statistical Software*, 80(1), 1–28. <https://doi.org/10.18637/jss.v080.i01>
- 682 Bürkner, P.-C. (2018). Advanced Bayesian multilevel modeling with the R package brms.
683 *The R Journal*, 10(1), 395–411. <https://doi.org/10.32614/RJ-2018-017>

- 684 Chung, Y., Rabe-Hesketh, S., Dorie, V., Gelman, A., & Liu, J. (2013). A nondegenerate pe-
685 nalized likelihood estimator for variance parameters in multilevel models. *Psychometrika*,
686 78(4), 685–709. <https://doi.org/10.1007/s11336-013-9328-2>
- 687 Dowty, D. R. (1989). On the semantic content of the notion of “thematic role.” In *Properties*,
688 *types and meaning* (pp. 69–129). Springer.
- 689 Dowty, D. R. (1991). Thematic proto-roles and argument selection. *Language*, 67(3),
690 547–619.
- 691 Evans, N., & Levinson, S. C. (2009). The myth of language universals: Language diversity
692 and its importance for cognitive science. *Behavioral and Brain Sciences*, 32(5), 429–448.
- 693 Fisher, C., Gertner, Y., Scott, R. M., & Yuan, S. (2010). Syntactic bootstrapping. *Wiley*
694 *Interdisciplinary Reviews: Cognitive Science*, 1(2), 143–149.
- 695 Fisher, C., Hall, D. G., Rakowitz, S., & Gleitman, L. R. (1994). When it is better to receive
696 than to give: Syntactic and conceptual constraints on vocabulary growth. *Lingua*, 92,
697 333–375.
- 698 Freeman, N. H., Sinha, C. G., & Stedmon, J. A. (1981). The allative bias in three-year-olds
699 is almost proof against task naturalness. *Journal of Child Language*, 8(2), 283–296.
- 700 Fujita, H. I. I. (2000). A cognitive approach to errors in case marking in japanese agramma-
701 tism. In *Constructions in cognitive linguistics: Selected papers from the fifth international*
702 *cognitive linguistics conference, amsterdam, 1997* (Vol. 178, p. 123). John Benjamins
703 Publishing.
- 704 Gentner, D. (1982). Why nouns are learned before verbs: Linguistic relativity versus natural
705 partitioning. *Center for the Study of Reading Technical Report; No. 257*.
- 706 Gleitman, L. R. (1990). The structural sources of verb meanings. *Language Acquisition*,
707 1(1), 3–55.
- 708 Gleitman, L. R., January, D., Nappa, R., & Trueswell, J. C. (2007). On the give and
709 take between event apprehension and utterance formulation. *Journal of Memory and*
710 *Language*, 57(4), 544–569.
- 711 Goldberg, A. E. (2006). *Constructions at work: The nature of generalization in language*.

- 712 Oxford University Press on Demand.
- 713 Goodman, J. C., Dale, P. S., & Li, P. (2008). Does frequency count? Parental input and the
714 acquisition of vocabulary. *Journal of Child Language*, 35(3), 515–531.
- 715 Grimshaw, J. (1990). *Argument structure*. the MIT Press.
- 716 Hansen, P. (2017). What makes a word easy to acquire? The effects of word class, frequency,
717 imageability and phonological neighbourhood density on lexical development. *First*
718 *Language*, 37(2), 205–225.
- 719 Hart, B., & Risley, T. R. (1995). *Meaningful differences in the everyday experience of young*
720 *american children*. Paul H Brookes Publishing.
- 721 Hartshorne, J. K., O'Donnell, T. J., Sudo, Y., Uruwash, M., Lee, M., & Snedeker, J.
722 (2016). Psych verbs, the linking problem, and the acquisition of language. *Cognition*,
723 157, 268–288.
- 724 Hartshorne, J. K., Pogue, A., & Snedeker, J. (2015). Love is hard to understand: The
725 relationship between transitivity and caused events in the acquisition of emotion verbs.
726 *Journal of Child Language*, 42(3), 467.
- 727 Hartshorne, J. K., & Schachner, A. (2012). Tracking replicability as a method of post-
728 publication open evaluation. *Frontiers in Computational Neuroscience*, 6, 8.
- 729 Hlavac, M. (2018). *Stargazer: Well-formatted regression and summary statistics tables*.
730 Bratislava, Slovakia: Central European Labour Studies Institute (CELSI). Retrieved
731 from <https://CRAN.R-project.org/package=stargazer>
- 732 Kipper, K., Korhonen, A., Ryant, N., & Palmer, M. (2008). A large-scale classification of
733 english verbs. *Language Resources and Evaluation*, 42(1), 21–40.
- 734 Lakusta, L., & Landau, B. (2005). Starting at the end: The importance of goals in spatial
735 language. *Cognition*, 96(1), 1–33.
- 736 Lakusta, L., & Landau, B. (2012). Language and memory for motion events: Origins of the
737 asymmetry between source and goal paths. *Cognitive Science*, 36(3), 517–544.
- 738 Landau, B., & Gleitman, L. R. (2015). 10 height matters. *Structures in the Mind: Essays*
739 *on Language, Music, and Cognition in Honor of Ray Jackendoff*, 187.

- 740 Levin, B., & Hovav, M. R. (2005). Argument realization.
- 741 Lidz, J., Gleitman, H., & Gleitman, L. R. (2003). Understanding how input matters: Verb
742 learning and the footprint of universal grammar. *Cognition*, 87(3), 151–178.
- 743 Lüdecke, D. (2018). Ggeffects: Tidy data frames of marginal effects from regression models.
744 *Journal of Open Source Software*, 3(26), 772. <https://doi.org/10.21105/joss.00772>
- 745 MacWhinney, B. (1977). Starting points. *Language*, 152–168.
- 746 MacWhinney, B. (2000). *The CHILDES project: The database* (Vol. 2). Psychology Press.
- 747 Makowski, D., Ben-Shachar, M. S., & Lüdecke, D. (2019). bayestestR: Describing effects
748 and their uncertainty, existence and significance within the bayesian framework. *Journal*
749 *of Open Source Software*, 4(40), 1541. <https://doi.org/10.21105/joss.01541>
- 750 Marantz, A. (1982). On the acquisition of grammatical relations in linguistik als kognitive
751 wissenschaft. *Linguistische Berichte Braunschweig*, (80), 32–69.
- 752 Matuschek, H., Kliegl, R., Vasishth, S., Baayen, H., & Bates, D. (2017). Balancing type i
753 error and power in linear mixed models. *Journal of Memory and Language*, 94, 305–315.
- 754 McDonough, C., Song, L., Hirsh-Pasek, K., Golinkoff, R. M., & Lannon, R. (2011). An
755 image is worth a thousand words: Why nouns tend to dominate verbs in early word
756 learning. *Developmental Science*, 14(2), 181–189.
- 757 Medina, T. N., Snedeker, J., Trueswell, J. C., & Gleitman, L. R. (2011). How words can
758 and cannot be learned by observation. *Proceedings of the National Academy of Sciences*,
759 108(22), 9014–9019.
- 760 Messenger, K., Branigan, H. P., McLean, J. F., & Sorace, A. (2012). Is young children’s
761 passive syntax semantically constrained? Evidence from syntactic priming. *Journal of*
762 *Memory and Language*, 66(4), 568–587.
- 763 Papafragou, A. (2010). Source-goal asymmetries in motion representation: Implications for
764 language production and comprehension. *Cognitive Science*, 34(6), 1064–1092.
- 765 Pinker, S. (1984). *Language learnability and language development*. Cambridge, Mass.:
766 Harvard University Press.
- 767 Pinker, S. (1989). Learnability and cognition: The acquisition of argument structure.

- 768 Quine, W. (1964). Word and object.
- 769 Regier, T., & Zheng, M. (2007). Attention to endpoints: A cross-linguistic constraint on
770 spatial meaning. *Cognitive Science*, 31(4), 705–719.
- 771 Sanchez, A., Meylan, S. C., Braginsky, M., MacDonald, K. E., Yurovsky, D., & Frank, M.
772 C. (2019). Chldes-db: A flexible and reproducible interface to the child language data
773 exchange system. *Behavior Research Methods*, 51(4), 1928–1941.
- 774 Strickland, B., Fisher, M., Keil, F., & Knobe, J. (2014). Syntax and intentionality: An
775 automatic link between language and theory-of-mind. *Cognition*, 133(1), 249–261.
- 776 Tang, N., Gong, S., Liao, Z., Xu, H., Zhou, J., Shen, M., & Gao, T. (2021). Jointly perceiving
777 physics and mind: Motion, force and intention. In *Proceedings of the annual meeting of*
778 *the cognitive science society* (Vol. 43).
- 779 Tinker, E., Beckwith, R., & Dougherty, R. (1988). Markedness and the acquisition of
780 emotion verbs. *Eastern States Conference on Linguistics*.
- 781 Vankov, I., Bowers, J., & Munafò, M. R. (2014). Article commentary: On the persistence
782 of low power in psychological science. *Quarterly Journal of Experimental Psychology*,
783 67(5), 1037–1040.
- 784 Wickham, H. (2016). *ggplot2: Elegant graphics for data analysis*. Springer-Verlag New York.
785 Retrieved from <https://ggplot2.tidyverse.org>
- 786 Xie, Y. (2015). *Dynamic documents with R and knitr* (2nd ed.). Boca Raton, Florida:
787 Chapman; Hall/CRC. Retrieved from <https://yihui.org/knitr/>
- 788 Yarkoni, T. (2019). The generalizability crisis. *Behavioral and Brain Sciences*, 1–37.