UC San Diego
UC San Diego Previously Published Works

Title

The Amelia Bedelia effect: World knowledge and the goal bias in language acquisition

Permalink

https://escholarship.org/uc/item/4rs630rk

Journal

Cognition, 128(3)

ISSN

00100277

Authors

Srinivasan, Mahesh
Barner, David

Publication Date

2013-09-01

DOI

10.1016/j.cognition.2013.05.005

Copyright Information

Peer reviewed
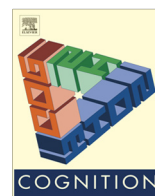
The Amelia Bedelia effect: World knowledge and the goal bias in language acquisition

Mahesh Srinivasan *, David Barner

University of California, San Diego, United States

ARTICLE INFO

ABSTRACT

How does world knowledge interact with syntax to constrain linguistic interpretation? We explored this question by testing children's acquisition of verbs like *weed* and *water*, which have opposite meanings despite occurring in the same syntactic frames. Whereas "weed the garden" treats "the garden" as a source, "water the garden" treats it as a goal. In five experiments, we asked how children learn these verbs. Previous theories predict that verbs which describe the transfer of an object with respect to its natural origin (e.g., "weed the garden") should receive source interpretations, whereas verbs that describe the transfer of an object with respect to something it is functionally related to (e.g., "water the garden") should receive goal interpretations. Therefore, acquiring world knowledge – about the natural origins and functional uses of objects – should be sufficient for differentiating between source and goal meanings. Experiments 1 and 2 casted doubt on this hypothesis, as 4- and 5-year-olds failed to use their world knowledge when interpreting these verbs and instead overextended goal interpretations. For example, children interpreted "weed the garden" to mean "put weeds onto a garden", even when they knew the natural origin of weeds. Experiment 3 tested children's interpretation of novel verbs and directly manipulated their access to relevant world knowledge. While younger children continued to exhibit a goal bias and failed to use world knowledge, older children generalized goal and source interpretations to novel verbs according to world knowledge. In Experiments 4 and 5, we confirmed that adults use world knowledge to guide their interpretation of novel verbs, but also showed that even adults prefer goal interpretations when they are made contextually plausible. We argue that children ultimately overcome a goal bias by learning to use their world knowledge to weigh the plausibility of events (e.g., of putting weeds into a garden).

© 2013 Elsevier B.V. All rights reserved.

1. Introduction

> "First," said Miss Emma, "weed the garden." [...].
> They [Amelia Bedelia and Effie Lou] went to the garden.
> "It does have a lot of weeds," said Effie Lou. She started to pull one.
> "Stop!" said Amelia Bedelia. "What are you doing?"

> "Trying to get the weeds out of the garden," said Effie Lou.
> "Get them out!" said Amelia Bedelia. "She said to weed the garden, not unweed it."
> "Oh," said Effie Lou. "I wonder why she wants more weeds."
> – Peggy Parish, Amelia Bedelia Helps Out (1979)

The meaning of a sentence is, to a large degree, a function of the meanings of its parts and how they are put together syntactically (Frege, 1892/1948). However, as shown by Peggy Parish's story about Amelia Bedelia, syntactic constructions can be ambiguous: although expressions like "weed the garden" and "milk the cow" denote events in which things are taken from sources, other

* Corresponding author. Address: Department of Psychology, UCSD, 9500 Gilman Drive, La Jolla, CA 92093-0109, United States. Tel.: +1 650 823 9488.
E-mail address: maheshs@ucsd.edu (M. Srinivasan).

syntactically identical expressions like "paint the house" and "water the plants" denote events in which things are put onto goals. In fact, some verbs admit both source and goal meanings – e.g., "feather the goose" vs. "feather the arrow", "seed the watermelon" vs. "seed the lawn", etc. In these cases, syntax alone cannot determine whether the verb denotes transfer from a source or to a goal. Instead, by some accounts, differentiating between sources and goals in these constructions requires that speakers draw on their knowledge of the world – e.g., the typical origins and uses of things like weeds, paint, and feathers (Barner & Bale, 2005; Buck, 1997; Clark & Clark, 1979; Harley, 1999; Kiparsky, 1997). In this paper, we explore whether such knowledge is used productively in the acquisition and use of verbs, to assign thematic roles to verb arguments. We present new evidence for a goal bias – which we call the Amelia Bedelia effect – and investigate how children overcome this bias to acquire verbs like *weed* and *milk.*

In English, the path of a motion event is typically denoted by satellite expressions that contain "into", and "out of", rather than by the verb itself (e.g., Talmy, 1985). Thus, in the expression "Amelia Bedelia tossed weeds into the garden", the preposition "into" indicates the direction of transfer of the weeds, and the noun phrase "the garden" denotes a goal. As a result, replacing "into" with "out of" in the same sentence changes "the garden" from a goal to a source.[1] Across the two sentences, the verb itself does not specify the path of transfer, but only a manner of motion – i.e., that Amelia Bedelia *tossed* the weeds, but did not *hurl* or *fling* them. In contrast, so-called "locatum verbs" like *weed* and *water* can specify opposite directions of transfer without employing words like "into" or "out of": in "weed the garden", "the garden" acts as a source, but in "water the garden", it acts as a goal (for additional examples of locatum verbs, see Table 1; see also Clark & Clark, 1979). This ambiguity raises the question of how thematic roles are assigned to such verbs and what information children draw on when learning them.

By some accounts, differences between how locatum verbs like *weed* and *water* express event structure could be attributed to differences between the things that these words name when they are used as nouns (Barner & Bale, 2005; Buck, 1997; Clark & Clark, 1979; Harley, 1999; Kiparsky, 1997). For example, according to Buck (1997), we interpret "weed the garden" as "take weeds out of the garden" because we know that weeds often originate in gardens, whereas we interpret "water the garden" to mean "put water into the garden" because water does not originate in gardens, but is often added to them. These criteria also help explain why the meanings of individual verbs like *feather* and *seed* are malleable, and can encode both transfer from a source and to a goal. When

feather describes the transfer of feathers in the context of their natural origin (e.g., "feather the goose"), it is given a source interpretation, but when it describes the transfer of feathers with respect to something feathers are functionally related to (e.g., "feather the arrow"), it is given a goal interpretation (see also Clark & Clark, 1979).

By these accounts, world knowledge is used to furnish default interpretations of locatum constructions, which would otherwise be ambiguous.[2] Thus, if learners lack relevant world knowledge – e.g., regarding where things originate and how they are typically used – they may be uncertain whether these verbs denote transfer from a source or to a goal. Faced with such uncertainty, learners could adopt a *conservative* strategy: they could wait to assign source or goal interpretations to verb constructions until they have direct, item-specific evidence regarding their meanings (see Tomasello, 2000, 2003, for a related proposal). Alternatively, learners could adopt a *productive* strategy: they could assign interpretations to locatum constructions in the absence of direct evidence regarding their meanings (see, e.g., Pinker, 1984, 1989; see also Baker (1979), for discussion).

Interestingly, an extensive literature suggests that, under uncertainty, language users may expect ambiguous expressions to denote transfer toward goals (see Papafragou (2010) for a review). First, both children and adults prefer to describe motion events (e.g., in which a bird flies from a tree to a flower) in terms of goals (e.g., "it flew onto the flower"), rather than sources (e.g., "it flew off of the tree"; for evidence from English, see Arnold, 2001; Lakusta & Landau, 2005, 2012; Papafragou, 2010; Stefanowitsch & Rohlde, 2004). Second, a goal bias is apparent not only in what language users choose to describe, but also in the

Table 1
Examples of locatum verbs in English.

Name	Mapping between argument and event structure	Attested examples
Goal	Noun form of verb ⇔ moved object	Paint the house, water the plants, salt the food
	Direct object ⇔ goal (endpoint of transfer)	
Source	Noun form of verb ⇔ moved object	Milk the cow, Dust the table, Weed the garden
	Direct object ⇔ source (starting point of transfer)	
Flexible (both source and goal)	Noun form of verb ⇔ moved object	Feather the goose/ arrow, Seed the watermelon/lawn, Skin the rabbit/drum
	Direct object ⇔ source or goal	

---

[1] By Talmy's analysis, the paths of motion events are characterized by the relation between a *figure* (i.e., an object that moves) and a *ground* (i.e., a reference object that the figure moves in relation to). Thus, "Amelia Bedelia tossed the weeds into the garden" specifies a goal path because the figure (i.e., "weeds") is moved toward the ground ("the garden"). Replacing "into" with "out of" makes the expression denote a source path, in which the figure is moved away from the ground (see Talmy, 1985, 2000).

[2] The idea that world knowledge constrains the assignment of thematic roles has been invoked to explain not only how words with noun and verb forms are interpreted, but also to explain how other forms of flexible language are interpreted, like nominal polysemy (Nunberg, 1979) and nominal compounds (Gagne & Shoben, 1997; Murphy, 1988).

structure of languages. Languages tend to make subtler distinctions involving goals than sources (Regier & Zheng, 2007). Further, when languages feature a source term (e.g., fell *out* of a box), they are also likely to have a goal expression (e.g., fell *into* a box), but not the opposite (Regier, 1997). Finally, a goal bias is present not just in language, but also in non-linguistic thought. When processing motion events involving intentional agents, pre-linguistic infants, children, and adults are better at remembering objects that serve as goals, rather than sources (Lakusta & Landau, 2005, 2012; Lakusta, Wagner, O'Hearn, & Landau, 2007; Papafragou, 2010), and are also better at remembering spatial relations involving goals, compared to those involving sources (Papafragou, 2010; Regier & Zheng, 2007).

These findings suggest that learners may adopt a productive strategy, and expect locatum verbs to denote transfer toward goals (rather than transfer from sources), in the absence of direct evidence that they do. Thus, if introduced to a novel substance called *dax* and told that a character has "daxed a tree", learners may assume that *dax* has been transferred *onto* the tree (rather than taken from it). Consistent with this, Fisher, Hall, Rakowitz, and Gleitman (1994) found that, when children were shown an event in which an elephant handed a ball to a bunny and heard "Look! *Blicking*!", they expected the verb *blick* to denote transfer toward a goal, as in the verb *give*, rather than transfer from a source, as in *take* (see Papafragou, 2010, for related evidence). Applied to the case of locatum verbs, these results suggest that children may be predisposed to assign goal interpretations to constructions like "milk the cow" and "weed the garden", which conventionally encode transfer from sources. Thus, if children lack world knowledge of where weeds originate, or are unable to deploy this knowledge when interpreting verbs, they may assume, like Amelia Bedelia, that "weed the garden" refers to putting weeds into a garden, rather than taking them out.

Previous studies indicate that children generally adopt productive learning strategies when learning words with noun and verb forms, such that they creatively assign interpretations to forms they could not have had experience with – e.g., as in a child's novel use of *cracker* in *I'm crackering my soup* (see Clark, 1982; see also Berman, 1999; Bushnell & Maratsos, 1984). However, no previous work has explored whether children over-extend goal interpretations to conventional source verbs like *weed* and *milk*, or more generally, whether children use their knowledge of the world to differentiate between source and goal interpretations. In Experiments 1 and 2, we tested this question, and found that children initially adopt a productive goal bias toward interpreting locatum verbs, resulting in striking misinterpretations of conventional source verbs like *milk* and *weed*. We call this finding the "Amelia Bedelia" effect. For example, we found that children often misinterpreted "milk the cow" to mean "put milk onto a cow" rather than "take milk out of a cow." Based on these findings, in our remaining experiments we asked how children might overcome a goal bias, to acquire adult-like meanings of conventional source verbs.

One possibility, consistent with the observations of Buck (1997), is that children acquire adult-like meanings of locatum verbs once they learn about the natural origins and functional uses of relevant objects and substances. For example, to learn that "milk the cow" means "take milk out of the cow," children may need to learn only that milk typically comes from cows, and thus, that "the cow" is a source, rather than a goal, of the transfer of milk. On this view, children should not require direct, item-by-item evidence regarding how particular locatum verbs like *milk* and *weed* are used. Rather, children should interpret such verbs like adults once they have acquired their related noun forms and gained relevant world knowledge regarding the nouns' referents. Further, children may deploy their knowledge productively when interpreting novel verbs. For example, when taught that a novel substance called *dax* grows inside of trees, children may expect that "daxing a tree" involves taking *dax* out of a tree. In contrast, when taught that *dax* makes things shiny, learners may expect that "daxing a car" involves putting *dax* onto a car. In sum, world knowledge may directly, and productively, specify how locatum verbs encode event structure from early in life.

Alternatively, acquiring relevant world knowledge may not be sufficient for children to overcome a goal bias: instead, to learn individual verbs like *weed* and *milk*, children may require direct, item-specific evidence. On this account, although world knowledge may play a role in how locatum verbs have been derived from nouns historically – and thus explain why source verbs typically describe the transfer of objects in relation to their natural origins – it may not play a direct role in how these verbs are acquired. Thus, even if children know that milk comes from cows, they may not construct an adult-like meaning of "milk the cow" until they receive direct evidence of its meaning in the input. In support of this idea, some studies have suggested that children's early knowledge of verbs is limited to how they have heard those verbs used (see e.g., Tomasello, 2000, 2003). These studies indicate that children initially struggle to use verbs in new ways – e.g., to use an intransitive verb as a transitive verb – if novel uses of verbs have not been previously modeled (see Tomasello & Brooks, 1998; but see Fisher, 2002). Such findings also predict that children may require direct, item-specific evidence to acquire adult-like meanings of source verbs like *weed* and *milk*.

Finally, a third possibility is that children do not require direct item-specific evidence to overcome a goal bias, but also do not acquire source meanings of verbs like *weed* and *milk* simply by learning about the natural origins of weeds and milk. Instead, acquiring source meanings may require additional learning about how world knowledge relates to syntactic structures. This could take two forms. First, children might need to construct a rule – e.g., that whenever a construction involves a natural origin relation, as "milk the cow" does, it should receive a source interpretation. Alternatively, children might need to learn to use their knowledge pragmatically, to make inferences about the plausibility of source and goal meanings. For example, children could use their knowledge that milk comes from cows to recognize that milk is unlikely to be placed onto a cow, and thus, that an interlocutor is unlikely to intend

this when they say "milk the cow." Critically, whether acquiring adult-like meanings depends on using world knowledge to mediate rules or to support pragmatic reasoning, it may initially pose challenges for young children. Thus, some children could think that "milk the cow" means "put milk onto the cow," despite knowing that milk comes from cows.

The present study explored these possibilities in a series of five experiments. Experiments 1 and 2 explored the relationship between 4- and 5-year-old children's developing knowledge of specific objects and substances (e.g., milk, paint, etc.) and their interpretation of the related locatum verbs in English. These experiments reveal that, in the absence of relevant world knowledge, children default to a goal bias when interpreting locatum verbs. They also show that, even when children possess adult-like world knowledge, they initially struggle to use it to overcome a goal bias. Experiment 3 extended this approach to novel verbs, and explicitly manipulated children's access to relevant world knowledge, revealing a goal bias and failure to use world knowledge in younger children, but not in older children. Since all children received the same input, this suggests that the success of older children depends not on item-based learning alone, but also on a developing ability to draw productively on world knowledge. Experiments 4 and 5 explored how world knowledge is ultimately used to constrain verb interpretation – and thus, why younger children fail to use it – by testing adults. Using a similar method to Experiment 3, Experiment 4 showed that adults systematically draw on world knowledge to generalize goal and source interpretations to novel constructions. However, Experiment 5 showed that when goal interpretations have been made plausible, adults generally prefer them, even in contexts involving natural origins. This suggests that source interpretations are not adopted via application of a rule, but are instead inferred when goal interpretations are deemed implausible – e.g., because it is unusual to place a substance in a location where it naturally originates. Thus, even adults exhibit a goal bias, and may differ from children only in their ability to use world knowledge to evaluate the plausibility of events.

2. Experiment 1

As noted in the introduction, previous accounts propose that language users adopt default interpretations of locatum verbs like *weed* and *water* by drawing on their knowledge of the natural origins and functional uses of things in the world (Barner & Bale, 2005; Buck, 1997; Clark & Clark, 1979; Harley, 1999; Kiparsky, 1997). Early in life, however, children may lack sophisticated knowledge of many objects and substances, and may fail to recognize that such knowledge is relevant to verb interpretation.

In Experiment 1, we asked two specific questions to explore this idea. First, we asked how children might interpret these words in the absence of world knowledge. One possibility is that children only assign systematic interpretations to locatum verbs when they are provided with direct evidence regarding their meanings (see, e.g., Tomasello, 2000, 2003). Such a strategy would predict that,

when forced to choose, children should be equally likely to guess that an unfamiliar verb encodes transfer toward a goal or transfer away from a source. A second possibility, as noted in the introduction, is that children are initially biased to focus on goals rather than sources when interpreting language (Lakusta & Landau, 2005, 2012; Lakusta, Wagner, O'Hearn, & Landau, 2007; Papafragou, 2010; Regier, 1997). A goal bias would predict systematic misinterpretations of attested source verbs as goal verbs – e.g., such that children interpret the verb *milk* as "put milk onto something."

Second, we asked how children interpret locatum verbs once they have acquired relevant world knowledge. If world knowledge fully mediates interpretation, then once children know that milk comes from cows, they should think that "milking a cow" involves taking milk out of a cow. Another possibility, however, is that children must learn to use their world knowledge constrain verb interpretation. On this hypothesis, some children might know the origins of things like milk, but nonetheless fail to assign a source interpretation to the verb *milk*. To explore these questions, Experiment 1 tested the relationship between 4- and 5-year-old children's interpretation of attested locatum verbs (e.g., *milk* and *paint*) and their knowledge of the relevant objects and substances (e.g., milk, paint).

2.1. Method

2.1.1. Participants

The participants were 19 native English-speaking children (11 girls) between the ages of 4;1 and 6;0 ($M = 5;1$). An additional three children participated but were excluded due to failing the initial trials that gauged their understanding of the task (2), or for not cooperating (1). All children in this experiment and in Experiments 2 and 3 were either brought into the lab or recruited from daycares in the San Diego or Vancouver Island areas. Additionally, we tested 16 native English-speaking students (7 women) from the University of California, San Diego, between the ages of 18 and 29 ($M = 21$).

2.1.2. Materials and procedure

Each child participated in a pretest before proceeding to the critical judgment task. Adults participated only in the judgment task.

2.1.2.1. Pretest. Prior to the judgment task, children were pre-tested for their knowledge of the ten nouns that corresponded to the critical locatum verbs of the judgment task (described below). On each trial, children were shown four pictures separated into quadrants and were asked to point to the picture corresponding to the critical word – e.g., "Point to the milk." If the children were incorrect, they were given feedback. They then responded to questions about how the critical object/substance is used and where it originates from – e.g., "What do you do with milk?" "Where does milk come from?" Children's responses were compared to their responses on the judgment task.

2.1.2.2. Judgment task. On each critical trial, participants saw a video of either Big Bird or Zoe (from Sesame Street)

moving an object or substance from a source location to a goal location (e.g., moving milk from a cow to a jar). After the video, two additional characters, Elmo and Monkey, described what had happened, and participants were asked who provided a better description. Each description used a locatum verb that incorporated the name of the moved object/substance, but differed according to whether it labeled the source or goal location with the verb's direct object (e.g., "milk the cow" vs. "milk the jar"). For five of the critical items, the participants were presented with verbs that had conventional goal meanings ("Goal Items"; e.g., *salt*, *paint*), and for the other five items, they heard verbs with conventional source meanings ("Source Items" e.g., *milk*, *weed*; see Table 2).

For example, on one critical source trial, Big Bird took milk from a cow and poured it into a jar. Each component of the event was accompanied by a verbal description – e.g., "Big Bird's going to take the milk out of the cow. See? Big Bird took the milk from the cow. Now, let's see what he's going to do. Big Bird's going to put the milk into the jar. See? Big Bird put the milk into the jar." After the event, Monkey used *milk* as a source verb, "I know what happened! Big Bird *milked* the cow!" and Elmo used *milk* as a goal verb, "I know what happened! Big Bird *milked* the jar!" Thus, Monkey was correct if *milk* was interpreted as a source verb, because the direct object encoded the source (i.e., the cow). However, Elmo was correct if *milk* was interpreted as a goal verb, because the direct object encoded the goal (i.e., the jar). The critical goal items were structured similarly to the source items (see Table 2).

Table 2
Critical items of Experiment 1.

Event	Descriptions of events
Source items	
Big Bird takes milk from a cow;	Goal: "Big Bird milked the jar"
Puts it into a jar	Source: "Big Bird milked the cow"
Zoe takes a peel off of a banana;	Goal: "Zoe peeled the bowl"
Puts it into a bowl	Source: "Zoe peeled the banana"
Zoe takes dust off of a table;	Goal: "Zoe dusted the trash"
Puts it into the trash	Source: "Zoe dusted the table"
Big Bird takes weeds out of flowers;	Goal: "Big Bird weeded the trash"
Puts them into the trash	Source: "Big Bird weeded the flowers"
Big Bird takes a peel off of an orange;	Goal: "Big Bird peeled the box"
Puts it into a box	Source: "Big Bird peeled the orange"
Goal items	
Zoe takes water out of a bowl;	Goal: "Zoe watered the plant"
Puts it into a plant	Source: "Zoe watered the bowl"
Big Bird takes paint out of a cup;	Goal: "Big Bird painted the house"
Puts it into a house	Source: "Big Bird painted the cup"
Zoe takes salt out of a cup;	Goal: "Zoe salted the soup"
Puts it into soup	Source: "Zoe salted the cup"
Zoe takes butter out of a box;	Goal: "Zoe buttered the bread"
Puts it onto bread	Source: "Zoe buttered the box"
Big Bird takes a stamp off of a table;	Goal: "Big Bird stamped the letter"
Puts it onto a latter	Source: "Big Bird stamped the table"

Although the videos depicted the actions as realistically as possible, in some cases this was difficult: e.g., Big Bird did not "take milk out of a cow" by squeezing its udders (because the cow was an action figure with fake udders), but instead by taking a bucket full of milk from beneath the cow. To verify that the depicted actions satisfied the intuitions of adults, and could be labeled by conventional locatum verbs, adults participated in the same task as children.

Prior to receiving the ten critical items, participants were given four warm-up trials. These trials ensured that participants could attend to the events. Two of the warm-up trials tested participants' ability to attend to the goals of the events. For example, on one such trial, Big Bird took a book off of a table and put it into a box. Elmo then said, "Big Bird put the book on the table" and Monkey said, "Big Bird put the book in the box!" The other two warm-up trials were similarly structured but tested participants' ability to attend to the sources of events. Two children were excluded from analyses because they could not accurately respond to three out of four of the warm-up trials without feedback.

The videos corresponding to Elmo and Monkey's statements were presented side-by-side, sequentially, on a single screen. We created two item orders, counterbalancing the side on which Elmo and Monkey appeared. The character that appeared to the left of the participant always spoke first. The ten critical items were administered in a fixed order.

2.2. Results and discussion

Adult participants always preferred goal uses of goal verbs (e.g., "Zoe watered the flowers"; $M = 1.0$) and always preferred source uses of source verbs (e.g., "Big Bird milked the cow"; $M = 1.0$). Thus, for native English-speaking adults, conventional goal and source verbs reliably elicit goal and source interpretations, respectively. These results also indicate that the actions depicted in our videos were suitable for evaluating the acquisition of these verbs.

Fig. 1 displays children's performance for each of the critical items. When compared to chance (0.5), children reliably preferred goal uses of the conventional goal verbs (e.g., "Zoe watered the flowers"; $M = .97$, $SE = .02$; Wilcoxon $T = 190$, $n = 19$, $p < .001$) and source uses of the source verbs (e.g., "Big Bird milked the cow"; $M = .68$, $SE = .06$; $T = 155.5$, $n = 19$, $p < .05$).[3] However, interestingly, children also exhibited an "Amelia Bedelia effect": they more often adopted conventional interpretations of goal verbs than of source verbs ($T = 5.5$, $n = 19$, $p < .005$). Surprisingly, for three of the source verbs – i.e., *milk*, *dust*, and *weed* – children's performance did not differ from chance ($M = .53$, $SE = .08$; $T = 102.5$, $n = 19$, *ns*). This is striking because choosing goal interpretations in these cases required selecting expressions that are rarely, if ever uttered – e.g., "milk the jar" – over ones that are present in language – e.g., "milk the cow".

---

[3] Preliminary analyses for Experiments 1–3 did not find significant effects of gender, side of presentation, or order. We have thus excluded these factors from all analyses.

Fig. 1. The proportion of children's source and goal responses to the critical items of Experiment 1.

For the two source items involving *peel*, in contrast, children's performance was relatively adult-like. One possible reason for this is that the noun and verb forms of *peel* may not be related synchronically. Indeed, as a noun, *peel* tends to require an additional modifier – e.g., "*banana* peel", "*orange* peel". Further, it is possible to *peel* objects other than bananas and oranges, such as stickers (e.g., "peel the sticker off the page"). Thus, to interpret "peel a banana" or "peel an orange," children may not need to reason about the relation between a peel and a banana or orange. Below, we focus on children's errors on the other three source items.

First, we sought to rule out that children's errors might be due to the relative infrequency of source verbs relative to goal verbs. To assess this, we obtained frequency estimates of these verbs from the SUBTLEX-US corpus, which is tagged with part-of-speech information (Brysbaert & New, 2009; Brysbaert, New, & Keuleers, 2012).[4] Frequency estimates for each verb were computed by adding together the frequency of a verb in its uninflected form (e.g., *milk*) and in its various inflected forms (e.g., *milks, milking, milked*). Fig. 2 plots frequency estimates (per million words) for each of the critical verbs tested in Experiments 1 and 2. Disregarding *paint*, which was much more frequent than all of the other verbs, the source verbs we tested were not less frequent than the goal verbs (Average frequency of source verbs = 3.72 (*SD* = 3.26); Average frequency of goal verbs other than paint = 2.43 (*SD* = 1.83)). This suggests that children's errors for source verbs were unlikely to be an effect of those verbs' frequency. Indeed, verb frequency was not significantly correlated with children's performance in the judgment task (*r* = .30, *n* = 9, *ns*).

Another reason that children may have made "Amelia Bedelia" errors on *milk*, *weed*, and *dust*, is that they did not know the origin of these substances or how they are normally acted upon. Lacking such knowledge, children may have resorted to a goal bias, assuming that these verbs have goal rather than source meanings. To test this idea, we compared children's responses to pretest questions probing this knowledge – e.g., "Where does milk come from?" – to their performance on trials of the judgment task. For pre-test questions, we coded the proportion of times in which children referenced a take-off action or an appropriate source location in their responses – e.g., "you pull *weeds* out," "milk comes from *cows*," "*dust* comes from the floor." Chi-square tests revealed a reliable relationship between children's pre-test responses and their performance on the judgment task for the *dust* item, $\chi^2(1, N = 19) = 19.0$, $p < .001$, and a non-significant trend for the *weed* item, $\chi^2(1, N = 19) = 3.69$, $p = .055$.[5,6] Thus, some of children's errors may be explained by a failure to recognize the origin of things like weeds and dust.

However, it is unlikely that all of children's errors can be explained by a lack of such knowledge. For example, children expressed that milk comes from cows (*M* = .95) reliably more often than they preferred a source use of *milk* (*M* = .58; *t*(18) = 3.24, *p* < .005), and that weeds originate in the ground or should be removed (*M* = .79) reliably more often than they preferred a source use of *weed* (*M* = .42; *t*(18) = 3.24, *p* < .005). Children also expressed that dust should be swept up and thrown away (*M* = .79) more often than they preferred a source use of *dust* (*M* = .58), although this did not reach significance, *t*(18) = 1.46, *ns*. Thus, strikingly, children preferred to use conventional source verbs

---

[4] SUBTLEX-US is comprised of subtitles from American television and films. SUBTLEX-US frequency measures have been shown to better predictors of performance in word recognition tasks than frequency measures derived from books and newspapers (Brysbaert & New, 2009; Brysbaert et al., 2012).

[5] Due to our small sample sizes, we also conducted a parallel set of Fisher's exact tests for the *dust* and *weed* items, which revealed the same pattern of results (*dust*: *p* < .001; *weed*: *p* = .10).

[6] These analyses could not be conducted for the *milk* item and the two *peel* items due to insufficient data: no children failed to express the origin of banana and orange peels, and only one child failed to do so for milk.
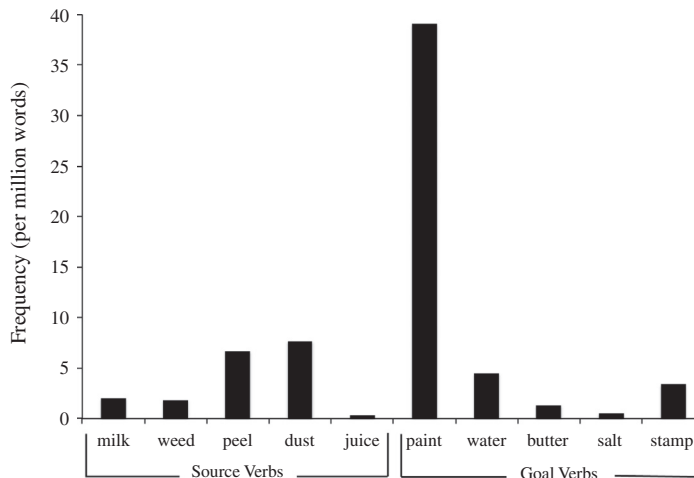
Fig. 2. SUBTLEX-US frequency estimates of the verbs tested in Experiments 1 and 2.

to denote transfer toward goals, despite having a relatively mature understanding of the relevant objects and substances.

Ultimately, English speakers systematically assign goal and source interpretations to conventional locatum verbs, as shown by the adult data. We found evidence that 4- and 5-year-olds are progressing toward this end-state: errors on *milk*, *weed*, and *dust*, declined as a function of age, $r(17) = .43$, $p < .05$ (one-tailed). This suggests that children initially over-extend goal interpretations to conventional source verbs – consistent with a goal bias – but scale back on these overextensions as they get older. One possibility is that with increasing age, children accumulate direct evidence in the language input regarding the meanings of attested source verbs, and correct their misinterpretations item-by-item. A second possibility is that as children get older, they get better at using world knowledge to guide interpretation, in the absence of direct, ostensive, evidence. We return to these possibilities in Experiment 3.

The results of Experiment 1 are consistent with the idea that children initially adopt a goal bias when interpreting locatum verbs – and fail to draw on world knowledge to overcome this bias – resulting in misinterpretations of conventional source verbs like *milk*. Another possibility, however, is that children have adult-like meanings for conventional source verbs but prefer descriptions involving goals relative to descriptions involving sources, because they tend to focus on the goals of events (see, Arnold, 2001; Lakusta & Landau, 2005, 2012; Papafragou, 2010; Stefanowitsch & Rohlde, 2004). Thus, after viewing Big Bird take milk from a cow and pour it into a jar, children may have understood that "Big Bird milked the cow" meant that Big Bird took milk from the cow. However, because they focused on the goal of the event, children may have judged that "Big Bird milked the jar" was the better description. Critically, this account predicts that when an expression like "milk the cow" is presented in isolation, children will interpret it without error. We tested this prediction in Experiment 2.

3. Experiment 2

Experiment 1 presented children with a single ambiguous event (e.g., in which Big Bird took milk from a cow and poured it into a jar) and asked them to choose between two descriptions using the critical verb (i.e., "Big Bird milked the cow" vs. "Big Bird milked the jar"). In contrast, Experiment 2 presented children with a single description using the verb (e.g., "Big Bird is going to milk a cow") and asked them to choose how it was best paraphrased (e.g., "Big Bird is going to put milk onto a cow" vs. "Big Bird is going to take milk out of a cow"). Thus, it tested whether children's errors in Experiment 1 were due to a genuine misinterpretation of conventional source verbs, or instead to a tendency to focus on the goals of events that include both sources and goals. Parallel items tested children's interpretations of conventional goal constructions, like "salt the food" (see Table 3). To examine the relationship between children's world knowledge and their interpretation of these constructions, we again probed children's knowledge of the relevant objects and substances in a pretest.

3.1. Method

3.1.1. Participants

The participants were 20 children (11 girls) between the ages of 3;10 and 5;11 ($M = 4;10$). An additional two children participated but were excluded due to failing the initial trials that gauged their understanding of the task (1), or for not cooperating (1).

3.1.2. Materials and procedure

Each child participated in the pretest before proceeding to the critical paraphrase task.

3.1.2.1. Pretest. Prior to the paraphrase task, children were pre-tested for their knowledge of the ten noun meanings that corresponded to the critical verbs that appeared in

Table 3
Critical items of Experiment 2.

Story	Paraphrases
Source items	
"Big Bird is going to milk a cow"	Goal: "Big Bird is going to put milk onto a cow"
	Source: "Big Bird is going to take milk out of a cow"
"Zoe is going to peel a banana"	Goal: "Zoe is going to put a peel onto a banana"
	Source: "Zoe is going to take a peel off of a banana"
"Zoe is going to dust a table"	Goal: "Zoe is going to put dust onto a table"
	Source: "Zoe is going to take dust off of a table"
"Big Bird is going to weed some flowers"	Goal: "Big Bird is going to put weeds onto some flowers"
	Source: "Big Bird is going to take weeds out of some flowers"
"Big Bird is going to juice an orange"	Goal: "Big Bird is going to put juice onto an orange"
	Source: "Big Bird is going to take juice out of an orange"
Goal items	
"Zoe is going to water a plant"	Goal: "Zoe is going to put water onto a plant"
	Source: "Zoe is going to take water off of a plant"
"Big Bird is going to paint a house"	Goal: "Big Bird is going to put paint onto a house"
	Source: "Big Bird is going to take paint off of a house"
"Zoe is going to salt some soup"	Goal: "Zoe is going to put salt onto some soup"
	Source: "Zoe is going to take salt out of some soup"
"Zoe is going to butter some bread"	Goal: "Zoe is going to put butter onto some bread"
	Source: "Zoe is going to take butter off of some bread"
"Big Bird is going to stamp a letter"	Goal: "Big Bird is going to put a stamp onto a letter"
	Source: "Big Bird is going to take a stamp off of a letter"

the paraphrase task. All aspects of the pretest were the same as in Experiment 1, except that we replaced the item that tested children's knowledge of orange peels with an item testing their knowledge of orange juice (for reasons described below).

3.1.2.2. Paraphrase task. Children were told that they would watch videos of Big Bird and Zoe to "find out what they are doing today". As before, videos were watched in the company of two other characters – Elmo and Monkey. After each video Elmo and Monkey each tried to describe what happened, and the child was asked to judge who was right.

On each trial children were told that Big Bird or Zoe was going to perform an action, which was described using either a conventional goal or source verb construction. For example, on one critical source trial, children heard, "Today, Big Bird is walking around a farm. Now, Big Bird's going to *milk a cow*. So what's Big Bird going to do?" The

children were then encouraged to repeat the critical phrase – e.g., "He's going to *milk a cow* – and were given assistance if needed. After this, children were shown another screen, split between two videos. In one video, Elmo paraphrased what Big Bird was going to do – e.g., "I know, Big Bird is going to put milk onto a cow" – while in the other, Monkey paraphrased what Big Bird was going to do – e.g., "I know, Big Bird is going to take milk out of a cow." The child was then asked to decide whether Elmo or Monkey was right.[7] We created two orders, counterbalancing the side in which Elmo and Monkey appeared.

Children received ten critical items, probing their knowledge of five goal constructions (e.g., "water a plant", "paint a house") and five source constructions (e.g., "milk a cow", "dust a table"; see Table 3). These constructions included the same critical verbs that were tested in Experiment 1, except that we replaced "peel an orange" with "juice an orange", on the grounds that the noun and verb forms of "peel" may not be related synchronically. The critical items were administered in a pseudo-random order. Prior to receiving these items, children were also tested on four warm-up trials to ensure that they understood the task. Two of these were put-on trials: on these trials, children were told that Big Bird or Zoe was going to put something somewhere – e.g., put a book in a box – and then Elmo and Monkey tried to say what happened. One character repeated the critical phrase verbatim (e.g., "Big Bird is going to put a book in a box"), and the other substituted "take-out" for "put-in" (e.g., "Big Bird is going to take a book out of a box"). The two take-out warm-up trials had an analogous structure and purpose. We excluded from our analyses those children who could not accurately respond to three out of four of the warm-up trials without feedback ($n = 1$).

3.2. Results and discussion

The dependent measure was the proportion of times children correctly chose goal paraphrases of the conventional goal constructions and source paraphrases of the conventional source constructions. Fig. 3 displays children's performance for each of the critical items. Children were reliably above chance (.5) at choosing goal paraphrases of the goal constructions ($M = .93$, $SE = .04$; $T = 208.5$, $n = 20$, $p < .001$) and at choosing source paraphrases of the source constructions ($M = .78$, $SE = .05$; $T = 199.5$, $n = 20$, $p < .001$). However, as in Experiment 1, children made significantly more errors on the source items, compared to the goal items, $T = 64.5$, $n = 20$, $p < .05$. Thus, consistent with a goal bias, children again exhibited an Amelia Bedelia effect, misinterpreting conventional source constructions as having goal meanings more often than the reverse.

---

[7] We are not claiming that these paraphrases fully define these verbs (e.g., that "take milk out of a cow" defines "milk the cow"). Our claim is only that one of the two paraphrases offered in each trial is part of the conventional interpretation of the verb construction, while the other is not: e.g., "milking a cow" involves taking milk out of a cow, but not putting milk onto a cow.
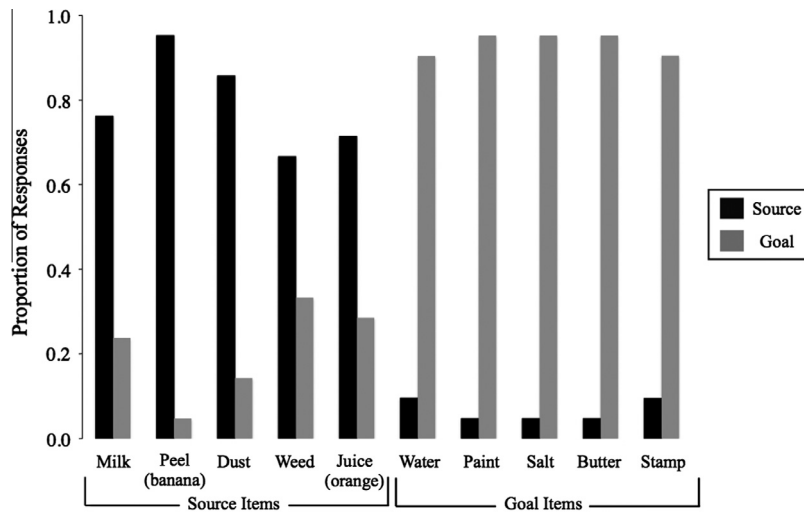
Fig. 3. The proportion of children's source and goal responses to the critical items of Experiment 2.

Despite continuing to exhibit an Amelia Bedelia effect, children in Experiment 2 appeared to perform better on source items than children in Experiment 1. To assess this, we compared performance on the four source items that children in each experiment received (i.e., "milk the cow", "dust the table", "weed the flowers", and "peel the banana"). There was a trend for children in Experiment 2 to respond more accurately to these items ($M = .80$, $SE = .05$) than children in Experiment 1 ($M = .63$, $SE = .07$; $U = 251$, $n = 39$, $p = .07$). This suggests that children's behavior in Experiment 1 – e.g., using "milk" to describe putting milk in a jar rather than taking it from a cow – cannot solely be explained by a misinterpretation of the semantics of source verbs like *milk*. Children may also prefer to use verbs to describe transfers toward goals, as opposed to transfers from sources, when presented with events that include both (see Arnold, 2001; Lakusta & Landau, 2005, 2012; Papafragou, 2010; Stefanowitsch & Rohlde, 2004).

Children's errors on source items could not be explained by a lack of world knowledge regarding the origin of things like milk, weeds, and dust. When children failed to express such knowledge in the pretest, they were no more likely to make errors on the paraphrase task: chi-square tests did not find reliable relationships between accuracy on the pre-test and on the paraphrase task for any of the source items (*milk*: $\chi^2(1, N = 20) = 0.74$, *ns*; *dust*: $\chi^2(1, N = 19) = 0.12$, *ns*; *weed*: $\chi^2(1, N = 20) = 1.83$, *ns*; *juice*: $\chi^2(1, N = 19) = 0.15$, *ns*).[8,9] Further, as in Experiment 1, children actually exhibited a relatively mature understanding of these substances and objects. For example, 90% of children expressed that milk comes from cows, 95% said that dust originates in the house and/or should be swept up and

thrown away, and 75% said that weeds originate in the ground and/or should be thrown away. Thus, it appears that children do not readily use their world knowledge to adopt goal and source interpretations. Strikingly, for children, the semantics of these constructions appears to override their world knowledge: although they know that milk comes from cows, and are unlikely to have ever witnessed milk being put onto a cow, some children think that this is the meaning of "milk a cow."

As in Experiment 1, we found evidence that children get better at adopting source interpretations between the ages of 4 and 6: children's "Amelia Bedelia" errors declined significantly as a function of age, $r(18) = .44$, $p < .05$ (one-tailed). Thus, although children initially over-extend goal interpretations, they correct their misinterpretations with age. One possibility is that children do this item-by-item, by accumulating direct evidence regarding the meanings of specific source constructions. For example, children could work out the meaning of "weed the garden" by observing that it labels an event in which weeds are taken out of a garden. A second possibility is that with age, children become increasingly able to deploy their world knowledge to guide interpretation – e.g., such that, by default, they adopt source interpretations under relations of natural origin. This would predict that, as children get older, they should become more adult-like not only in their interpretation of familiar verbs, but also in their interpretation of novel verbs, for which they have had no previous item-based experience. We explored this possibility in Experiment 3 by probing children's interpretations of verbs formed from novel nouns – for which they could not have item-based experience – while manipulating their knowledge of the natural origin and typical use of the nouns' referents.

---

[8] These analyses could not be conducted for the *peel* item because all children expressed that banana peels come from bananas or should be thrown away. Due to experimenter error, pre-test responses were only collected from 19 of the 20 children for the *dust* and *juice* items.

[9] Due to our small sample sizes, we also conducted a parallel set of Fisher's exact tests, which also yielded non-significant results.

4. Experiment 3

Experiment 3 explored the role of world knowledge and item-based learning in children's interpretations of

locatum constructions. To do so, we taught children novel nouns that referred to unfamiliar objects or substances that – we told them – exist only on Sesame Street. Children then learned the relationship between these objects/substances and either (1) the places on Sesame Street from which they naturally originate ("natural origin items"; e.g., that *dax* grows inside of a tree) or (2) the places to which they contribute some function ("functional relation items"; e.g., that *lorp* makes cars shiny). Then, in a paraphrase task like that in Experiment 2, children interpreted a novel verb construction in which the verb had been formed from the object or substance name, and the direct object had been formed from the place name (e.g., "dax a tree", "lorp a car"). According to previous theories (and as we will demonstrate in Experiment 4), mature language users draw on relational information to guide how thematic roles are mapped to verb arguments, allowing them to differentiate between source and goal interpretations (Barner & Bale, 2005; Buck, 1997; Clark & Clark, 1979; Harley, 1999; Kiparsky, 1997). For example, after learning that dax originates in a tree and being asked to interpret "dax a tree," the noun form of "dax" can be mapped to the substance that is transferred, and "a tree" to the source of this movement. In contrast, "a car" can be mapped to the goal of the movement of "lorp" in "lorp a car," because lorp adds a function to cars.

Experiment 3 had two objectives. A first objective was to test for the presence of a goal bias in novel verb learning. If children are biased toward goal interpretations, this should be evident not only in their misinterpretations of conventional source verbs (Experiments 1 and 2), but also in how they learn new verbs. Thus, children may generally prefer goal interpretations of novel verb constructions, regardless of whether they involve relations of natural origin or functional use. A second objective was to test whether, given equivalent experience, younger and older children differ in their ability to use world knowledge to differentiate between source and goal meanings. We reasoned that if the ability of older children to adopt source interpretations of attested source verbs (see Experiments 1 and 2) is due to item-based learning, their advantage should not extend to novel verbs. If, on the other hand, this ability stems from a developing capacity to use world knowledge to interpret verbs, older children may also

adopt source interpretations of novel verbs more often than younger children.

4.1. Method

4.1.1. Participants

The participants were 19 children (9 girls) between the ages of 3;10 and 5;11 (mean age = 4;10). An additional two children participated but were excluded for failing the initial trials that gauged their understanding of the task.

4.1.2. Materials and procedure

To begin each trial, we taught children a novel noun referring to a novel object or substance. On the natural origin trials, we then told the children that this object/substance was naturally derived from a source location – creating a relation to license a source interpretation. On the functional relation trials, we instead told the children that the critical object/substance added some function to a goal location – creating a relation to license a goal interpretation.

There were two natural origin trials and two functional relation trials (see Table 4). For example, on one natural origin trial, we introduced the children to a novel pink and orange substance called *dax* that is only found on Sesame Street. We told the children that *dax* grows inside of a special tree (which we also showed to the children), and that the muppets from Sesame Street like to eat it. Then, the children saw an animation in which they were told that Big Bird was walking outside and was going to "dax a tree." The children were then asked to repeat this critical phrase, and then watched as Elmo and Monkey tried to say what Big Bird was going to do – i.e., "put dax onto a tree" vs. "take dax out of a tree." Each item was formed on analogy with attested source or goal verbs: e.g., *dax* was analogous to verbs like *milk* and *juice*, and *lorp* was analogous to *wax* and *oil* (see Table 4). After the children judged who was right, we probed their knowledge regarding the critical object/substance. This was to ensure that the children had retained the information that we had provided when introducing the items. Following the natural origin trials, we asked children where the objects/substances come from, and what the muppets do with them. Following the functional relation trials, we asked children what the muppets do with the objects/substances and why.

As before, children also responded to four warm-up trials – two put-on trials and two take-off trials. For these items, children learned a novel word that referred to a novel object or substance, but the items were otherwise structured identically to the warm-ups in Experiment 2. We excluded from our analyses those children who could not accurately respond to three out of four of the warm-up trials without feedback ($n = 2$).

Finally, we grouped the two functional relation items and the two natural origin items into separate blocks, and counterbalanced the order in which they were presented. Because we also counterbalanced the side on which Elmo and Monkey appeared, this resulted in four task versions. Pictures of novel objects and substances were either chosen from the Novel Object and Unusual

Table 4
Critical items of Experiment 3.

| Construction	Description of conceptual relation	Similar to:
Natural origin items		
"*Dax* a tree"	*Dax* is a substance that grows inside of a tree. The muppets like to eat the *dax*	*Milk, juice*
"*Moop* a fruit"	*Moops* grow inside of a kind of fruit. The aliens like to eat the fruit but not the *moops*	*Seed, pit, core*
Functional relation items		
"*Tiv* some soup"	*Tiv* is a special substance. The muppets think *tiv* makes soup taste really good	*Salt, pepper, spice*
"*Lorp* a car"	*Lorp* is a special substance. The muppets think *lorp* makes their cars look shiny	*Polish, wax, oil*
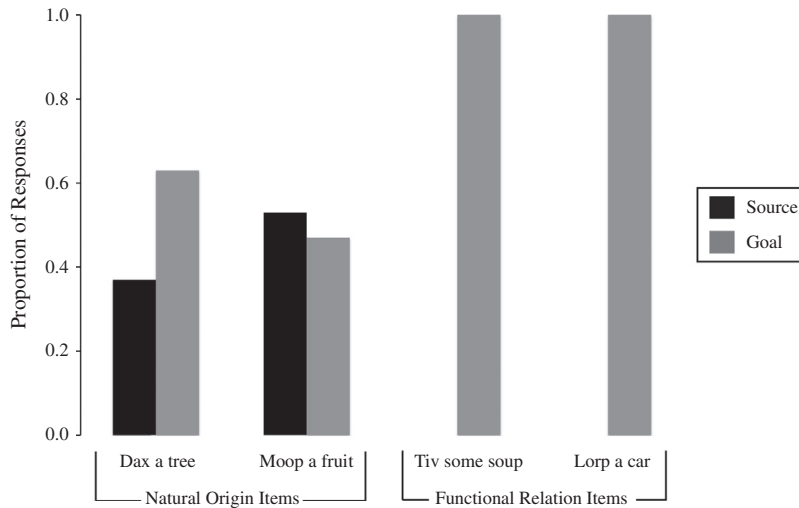
Fig. 4. The proportion of children's take-off and put-on responses to the critical items of Experiment 4.

Name (NOUN) Database ([Horst, 2009](#)) or located via Google Image Search.

4.2. Results and discussion

Our dependent measure was the proportion of times children "correctly" chose the source interpretation on the natural origin trials, and the goal interpretation on the functional relation trials. [Fig. 4](#) displays children's performance for each of the critical items. Children reliably adopted goal interpretations on the functional relation trials (e.g., after learning that *lorp* makes cars shiny and judging "lorp a car"; *M* = 1.0). Meanwhile, although children adopted goal interpretations significantly less often on the natural origin trials (e.g., after learning that *dax* comes from a tree and judging "dax a tree"; *T* = 0, *n* = 12, *p* < .01), they did not adopt source interpretations on these trials more often than chance (*M* = .45, *SE* = .09). Thus, children again exhibited an Amelia Bedelia effect, making reliably more "errors" on the natural origin trials than on the functional relation trials, *T* = 0, *n* = 19, *p* < .005. This finding converges with that of Experiments 1 and 2, and suggests that children rely on a default goal bias to learn new locatum verbs, even when conceptual criteria dictate, according to previous accounts (and the data of Experiment 4 below), that those verbs should denote transfer from sources.

Children's errors on the natural origin trials could not be attributed to ignorance of the origins of the critical objects and substances, as they were consistently able to indicate where these things come from: 95% of the children said that *dax* comes from a tree, and 95% said that *moops* come from fruits. Children, however, were not as good at indicating the functions of these things: only 68% correctly said that the muppets take out the *dax* and eat it, and only 42% said that the muppets remove the *moops* but do not eat them. Interestingly, some children confused the functions of these things by saying that they should be applied to the places from which they originate – e.g., that the function of *dax* is to put it onto a tree (11%) or that *moops*

should be put into fruits (11%). This is unlikely to be due to poor memory for functions, given that children were readily able to remember functions on the functional relation trials: 84% said that *tiv* makes soup taste good, and 89% said that *lorp* makes cars shiny. Instead, children may have misattributed put-on functions to *dax* and *moops* because they had assigned goal-directed interpretations to their constructions – e.g., by interpreting "dax a tree" as "put dax onto a tree".

Children adopted more source interpretations of the natural origin items as a function of age, *r*(17) = .41, *p* < .05 (one-tailed). This suggests that children initially adopt a goal bias when learning novel verb constructions, such that they misinterpret conventional source constructions as having goal meanings. It also suggests that the decline of goal overextensions (Experiments 2 and 3) is not purely a result of item-by-item learning of attested source constructions, because it extends to novel constructions. It would appear that, as children get older, they get better at using world knowledge productively, to adopt source and goal interpretations of novel locatum constructions.

However, these results leave open *how* world knowledge ultimately constrains interpretation. Why do younger children fail to use their knowledge to differentiate between source and goal interpretations, and what changes as children age? If simply acquiring relevant world knowledge is insufficient for constructing adult-like meanings, what additional learning is needed? In the Introduction, we contrasted two possible ways in which children might learn to use world knowledge when interpreting verbs. First, children might need to construct rule-like mappings that are mediated by world knowledge: e.g., that (1) whenever a construction involves a natural origin relation, as "milk the cow" and "moop the fruit" do, it receives a source interpretation, and that (2) whenever a construction involves a functional relation, as "butter the bread" and "lorp the spaceship" do, it receives a goal interpretation. If children have failed to make these generalizations, they may think, for example, that "milk the cow" means

"put milk onto a cow", despite knowing that milk originates in cows (as observed in Experiments 1 and 2).

Second, children might need to learn to use their world knowledge pragmatically, to make inferences about the likely meanings of constructions. By this account, source interpretations could be selected when – on the basis of natural origin information – goal interpretations are deemed implausible. For example, children could use their knowledge that milk comes from cows to recognize that (1) milk is unlikely to be put onto cows, and thus that (2) an interlocutor is unlikely to intend this meaning when they say "milk the cow." This could allow children to infer that the interlocutor instead intends "take milk out of the cow." Thus, if children have difficulty using world knowledge to reason pragmatically about the plausibility of events, they may fail to adopt adult-like interpretations, despite possessing relevant world knowledge.

Consistent with the second proposal, Experiments 1 and 2 showed that, when given evidence that a goal interpretation of a conventional source verb was plausible, children were more likely to prefer it. For example, in Experiment 1, when children were shown an event in which milk was taken from a cow and poured into a jar, they often chose a description in which *milk* was used as a goal verb ("Big Bird milked the jar"). In contrast, in Experiment 2, in which children were asked to choose between goal and source paraphrases of "milk the cow" in absence of an event involving the transfer of milk, children were more likely to prefer the source interpretation. This pattern of results suggests that when the context makes a goal interpretation of a verb plausible – e.g., when milk has been put into a goal location – that interpretation is generally preferred. However, in absence of such a context, a goal interpretation can be deemed implausible on the basis of world knowledge alone, allowing a source interpretation to be inferred.

To more directly differentiate between the rule-based and pragmatic accounts, the final two experiments tested how adults use world knowledge when interpreting novel locatum verbs. We reasoned that if acquiring source meanings requires using world knowledge to mediate rules, adults should adopt source interpretations of constructions whenever they involve a relation of natural origin. Alternatively, if acquiring source meanings requires using world knowledge to evaluate the plausibility of goal interpretations, adults should be less likely to adopt source interpretations when goal interpretations have been made plausible, even when relations of natural origin are present. In Experiments 4 and 5, we tested these predictions by examining adults' interpretations of novel locatum verbs, while manipulating the plausibility of source and goal interpretations. Experiment 4 explored adults' interpretations of novel verbs in absence of contexts making source and goal interpretations plausible, using a design similar to Experiment 3. We expected that adults would generalize source and goal interpretations according to information about an object's natural origin and functional use, just as older children did in Experiment 3. Finally, Experiment 5 tested whether adults would differentiate between source and goal interpretations when these interpretations were made equally plausible by the context.

5. Experiment 4

Here, we explored whether adults use world knowledge to differentiate between source and goal interpretations of novel verbs, in the absence of contexts making these interpretations plausible. To do so, we first taught participants names for novel objects and substances from a fictitious planet called Samsara. As in Experiment 3, participants were taught the relationship between these things and either (1) the places on the planet from which they originated ("natural origin trials"; e.g., that *moop* grows around a fruit) or (2) the places to which they contributed some function ("functional relation trials"; e.g., that *lorp* makes spaceships shiny). Participants were then presented with verb constructions formed from the object/substance name and the place name (e.g., "moop a fruit") and asked what they meant (e.g., they could choose "take moop off of a fruit," "put moop onto a fruit" or "neither of the above."). Critically, as in Experiment 3, participants were given no evidence that either of these interpretations were plausible – e.g., they were not presented with scenarios in which moop was taken from or placed onto a fruit. Thus, we probed whether adults can use world knowledge alone to constrain their interpretation of locatum verbs.

The developmental trend observed in Experiment 3 – the sensitivity of older but not younger children to world knowledge – suggests that adults may use world knowledge productively to interpret locatum verbs. This would predict a preference for source interpretations of the natural origin items (e.g., "moop a fruit"), and for goal interpretations of the functional relation items (e.g., "lorp a spaceship"). If, alternatively, adults fail to draw productively on world knowledge when interpreting locatum constructions, they should not differentiate between the natural origin and functional relation items, and should either respond randomly, or adopt a default interpretation, like a goal bias.

5.1. Method

5.1.1. Participants

The participants were 32 native English-speaking adults (17 women) between the ages of 19 and 57. They were recruited and tested via Amazon's MTurk platform (www.mturk.com). Participants were given monetary compensation in exchange for participation. We used MTurk's exclusion capabilities to ensure that our participants were only logging-in within the United States, and also tracked IP addresses to ensure that participants did not participate multiple times. Finally, to ensure that participants were attending to the task and competent speakers of English, they completed three morphological control trials, as described below.

5.1.2. Materials and procedure

Participants were told that they would read stories about the things that exist on a planet named Samsara, and answer questions about the meanings of phrases. Each of the critical stories first used a novel word to describe a novel object or substance and then described the relation

between this object or substance and another place. For example, in one of the critical natural origin items, *moop* was described as a substance that covers a special kind of fruit that grows on Samsara. The critical functional relation items, in contrast, did not describe a place of origin, but instead described the function that a thing contributes to a place – e.g., *lorp* was described as a substance that makes spaceships shiny. Each story was formed on analogy with attested source or goal verbs: e.g., *moop* was analogous to verbs like *peel* and *shell*, and *lorp* was analogous to *wax* and *oil* (see Table 5).

After each story, participants were presented with a verb construction that used the novel word as a verb – e.g., "moop a fruit" or "lorp a spaceship". They were asked what the phrase meant, and were asked to choose a source interpretation (e.g., "Take moop off of a fruit"), a goal interpretation (e.g., "Put moop onto a fruit"), or "Neither of the above". There were eight critical items in total, including four natural origin items and four functional relation items (see Table 5). Stories were presented in a quasi-random order, and intermixed with the seven control items described below.

We included four control items in which the described relations were not directly relevant to interpreting the target constructions. This was to test whether participants were basing their interpretations of the constructions in the critical items according to the specific relations provided in the stories. For example, in one natural origin control item, participants learned that *blick* labels an object that grows inside a vegetable – e.g., a relation that could license a source interpretation of "blick a vegetable". However, instead of judging the meaning of "blick a *vegetable*", participants were asked to judge the meaning of "blick an *animal*". Thus, although a relation of natural origin had been described in the story, it was not relevant to the construction participants had to judge – i.e., the relation between the *blick* and animal had not been described. Consequently, if participants base their interpretations only on relevant relations provided in the stories, they should not show a clear preference between source and goal readings of "blick an animal." In addition to judging two natural origin control items, participants also judged two functional relation control items, which had an analogous structure and purpose.

Table 5
Critical items of Experiment 4.

Construction	Description of conceptual relation	Similar to:
Natural-origin items		
"*Dax* an animal"	*Dax* is a special substance that grows on an animal. The aliens eat the animal, but not the *dax*	*Skin, feather, scale*
"*Moop* a fruit"	*Moop* is a fuzzy, greenish substance that covers a fruit. The aliens eat the fruit but not the *moop*	*Peel, shell, husk*
"*Jop* a fruit"	*Jops* are small, soft objects that are found inside of fruits. The aliens like to eat the fruits, but not the *jops*	*Seed, pit, core*
"*Tig* an animal"	*Tig* is a liquid that is produced by an animal. The aliens like to use *tig* for cooking	*Milk, juice*
Functional relation items		
"*Tiv* some food"	*Tiv* is a dark-orange powder. The aliens think *tiv* adds taste to food	*Salt, pepper, spice*
"*Lorp* a spaceship"	*Lorp* is a whitish, gel-like liquid. The aliens think *lorp* makes their spaceships look shiny	*Polish, wax, oil*
"*Wug* a tail"	A *wug* is a reddish, soft article of clothing. *Wugs* protect the aliens' tails from rain	*Dress, glove, mask*
"*Buck* a field"	*Bucks* are special structures. They protect the aliens' vegetable fields from thieves	*Fence, partition, barricade*
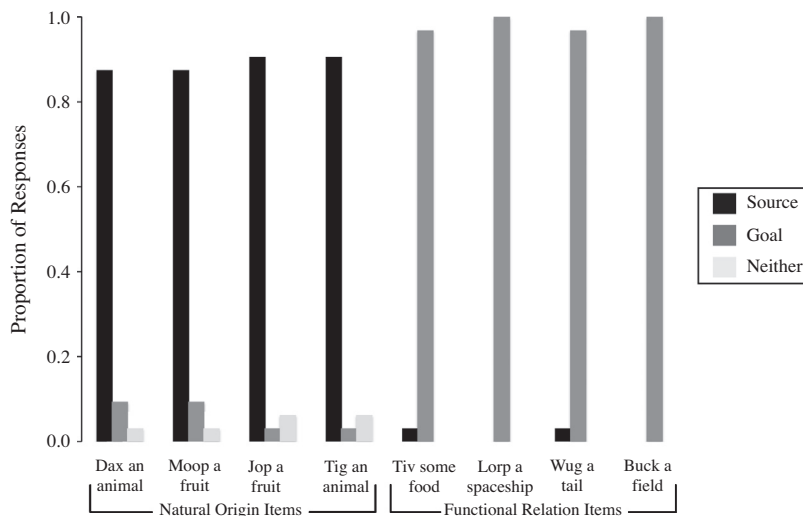
Fig. 5. The proportion of source, goal, and "neither of the above" responses to the critical items of Experiment 4.

Finally, we also included three morphological control items to ensure that participants were attending to the task, and were competent English speakers. In these items, participants learned a novel word and then judged a phrase in which the word had undergone a morphological change. For example, in one item, participants learned that when an alien on Samsara curls up into a ball, it is said that the alien *gazzers*. Participants then judged the meaning of "An alien gazzered for three hours", and were asked to decide whether the phrase described a past action, a future action, or neither.

5.2. Results and discussion

Our dependent measure on the critical trials was the proportion of trials on which participants chose source, goal, or "neither of the above" interpretations of the critical constructions. We defined chance responding as 0.33, as there were three possible answers for each trial.

Fig. 5 displays adults' responses for each of the critical items. On the critical functional relation trials – e.g., when participants learned that *lorp* makes spaceships shiny and then judged "lorp a spaceship" – participants chose goal interpretations ($M = .98$, $SE = .02$) reliably more often than chance ($z = 41.9$, $p < .001$) and more often than they chose source interpretations ($M = .02$, $SE = .02$; $z = 31$, $p < .001$) or "neither of the above" ($M = 0$; $z = 63$, $p < .001$). On the other hand, on the critical natural origin trials – e.g., when participants learned that *moop* comes from a fruit and then judged "moop a fruit" – participants chose source interpretations ($M = .85$, $SE = .04$) reliably more often than chance ($z = 12.5$, $p < .001$) and more often than they chose put-on interpretations ($M = .09$, $SE = .03$; $z = 10.88$, $p < .001$) or "neither of the above" ($M = .05$, $SE = .02$; $z = 13.08$, $p < .001$).

Participants' responses to the control items indicate that they based their interpretation of constructions on the specific relations that had been described in the stories. In response to the natural origin control trials – e.g., in which participants learned that *blick* comes from a vegetable but then had to judge "blick an *animal*" rather than "blick a vegetable" – participants chose "neither of the above" ($M = .56$, $SE = .08$) more often than chance ($z = 2.90$, $p < .01$) and more often than they chose goal interpretations ($M = .14$, $SE = .06$, $z = 3.48$, $p < .005$) or source interpretations ($M = .30$, $SE = .07$, $z = 1.87$, $p = .07$). Participants responded similarly to the functional relation control trials, replying "neither of the above" ($M = .67$, $SE = .08$) more often than chance ($z = 4.47$, $p < .001$) and more often than they chose goal interpretations ($M = .31$, $SE = .02$, $z = 2.35$, $p < .05$) or source interpretations ($M = .02$, $SE = .02$, $z = 8.29$, $p < .001$). Finally, participants performed well on the morphological change controls (e.g., past-tense or progressive-tense inflection), indicating that they were attending to the task and competent speakers of English. They responded correctly on almost all trials, and significantly more often than expected by chance ($M = .93$, $SE = .03$, $z = 20.6$, $p < .001$).

Consistent with the sensitivity of older but not younger children to world knowledge in Experiment 3, the results of Experiment 4 suggest that mature language users draw

productively on world knowledge to assign thematic roles to the arguments of novel verb constructions. Thus, to interpret "moop a fruit", adults may map the noun form of *moop* to the substance that is transferred and "the fruit" to the source of this movement, because they know that *moop* comes from a fruit. An alternative possibility, however, is that participants in Experiment 4 used world knowledge, but not to assign thematic roles to verb arguments. Specifically, they may have paid little attention to the syntactic frames of the constructions when responding to items, and may have based their judgments exclusively on world knowledge. For example, having learned that *moop* grows around a fruit, participants may have expected that the most likely event would involve taking moop off of a fruit, and may have selected "take moop off of a fruit" without attending to the syntax of "moop a fruit."

If this alternative account is correct, and participants ignore the syntax of the constructions, they should respond similarly to constructions with inverted word orders – e.g., they should also interpret "fruit a moop" to mean "take moop off of a fruit" after learning that moop grows around a fruit. To test this, we conducted a follow-up study using the same methods, but in which adult participants were asked to judge constructions that had either *standard* word orders (e.g., "moop a fruit" and "lorp a spaceship"), or *inverted* word orders (e.g., "fruit a moop" and "spaceship a lorp"). There were 35 participants (22 women), ranging in age between 19 and 65 (Mean age = 33). The order in which participants judged standard and inverted constructions was randomized across trials. The results of this study confirmed that participants attended to the specific syntactic frames in which constructions had been embedded. As before, when judging constructions with standard word orders, participants readily adopted goal interpretations for critical functional relation trials (e.g., for constructions like "lorp the spaceship"; $M = .89$, $SE = .03$; $z = 10.94$, $p < .001$) and source interpretations for natural origin trials (e.g., for constructions like "moop the fruit"; $M = .81$, $SE = .05$; $z = 5.98$, $p < .001$). However, participants did not adopt goal and source interpretations of constructions with inverted word orders; they most often chose "neither of the above" in response to inverted constructions on the critical goal trials (e.g., "spaceship the lorp"; $M = .69$; $SE = .07$; $z = 2.74$, $p < .01$) and in response to inverted constructions on the critical source trials (e.g., "fruit the moop"; $M = .72$, $SE = .07$; $z = 3.06$, $p < .005$).

Together, the results of Experiment 4 indicate that mature language users productively use world knowledge to adopt source and goal interpretations of novel locatum verbs, in the absence of any other evidence regarding the meanings of these verbs. However, as discussed in Experiment 3, these findings leave open exactly how world knowledge is used to constrain interpretation. On one hand, world knowledge could be used to mediate rule-like mappings: e.g., whenever a construction involves a natural origin relation (e.g., "milk the cow"), it could be given a source interpretation. On the other hand, world knowledge could be used to support pragmatic reasoning about the likely meanings of constructions: e.g., constructions involving natural origin relations (e.g., "milk the cow") could be given source interpretations because goal

interpretations are deemed implausible in those contexts (e.g., because milk is unlikely to be put onto cows). Thus, according to the pragmatic account – but *not* the rule-based account – manipulating the plausibility of a goal interpretation should affect whether a natural origin construction receives a source interpretation. Experiment 5 tests this prediction.

6. Experiment 5

In Experiment 4, adults restricted source interpretations of novel constructions to relations of natural origin, and goal interpretations to cases in which the relevant object/substance added a function to a place. Here, we examined whether adults continue to respect these constraints when the goal and source interpretations of a construction are each made plausible. Participants were taught about novel objects or substances and their relations to places just as they were in Experiment 4. But after learning about these relations, they were told that one alien on the planet had put an instance of the object or substance in a place (e.g., put *moop* onto a fruit), and another alien had taken a different instance of this object or substance away from another place of the same kind (e.g., taken *moop* off of another fruit). They were then asked which action could be labeled by the novel locatum construction (e.g., which alien "mooped a fruit"). Thus, both goal and source interpretations were made plausible, in that a transfer toward a goal and a transfer away from a source had each taken place.

We reasoned that if source interpretations are inferred when goal interpretations are deemed implausible on the basis of world knowledge (as proposed by the pragmatic account), they may not arise when those interpretations have been made plausible. Thus, adults, like children, may exhibit a goal bias when both interpretations are plausible. If, in contrast, source interpretations arise whenever a relation of natural origin exists (as proposed by the rule-based account), they should still arise even when goal interpretations have been made plausible. This would predict that adults should prefer source interpretations of natural origin constructions, just as in Experiment 4.

6.1. Method

6.1.1. Participants

The participants were 31 native English-speaking adults (17 women) between the ages of 20 and 64 (*M* = 34). Participants were recruited and tested via Amazon's MTurk.

6.1.2. Materials and procedure

Each of the critical stories presented participants with a name for a novel object or substance and then described the relation between this object/substance and another place (e.g., that *moop* is a substance that covers a kind of fruit). These descriptions were the same as those used in Experiment 4 (see Table 5), and either described relations that had led to source interpretations (e.g., involving natural origin) or goal interpretations (e.g., involving functional use) in Experiment 4. After participants read these descriptions, they learned that one alien had put the object/sub-

stance in the place, and that another alien had taken a different instance of that object/substance away from another place of the same kind (the order of these actions varied across trials). Participants were then asked what happened in the story, and responded by selecting between phrases in which the aliens' actions were labeled by novel verb constructions.

For example, in one critical natural origin item, participants learned that *moop* grows around a kind of fruit (see Table 5) – a relation that led to source interpretations in Experiment 4. They then learned that an alien named Gamma "went up to one of these fruits and took moop off of the fruit" and that another alien named Alpha, who was carrying some moop with him "went up to another fruit and put moop onto the fruit." Participants were then asked what happened in the story, and could choose "Gamma mooped a fruit" (a transfer from a source), "Alpha mooped a fruit" (a transfer to a goal), or "Neither of the above." The critical functional relation items were similarly adapted from Experiment 4 (see Table 5). There were eight critical items, consisting of four natural origin items and four functional relation items.

We also included two functional relation and two natural origin control items to assess whether participants expect that constructions can label transfers from sources or toward goals even when the relevant objects/substances and places encoded by the construction are unrelated. These stories taught participants the name of an object or substance and described its relationship to a particular place – e.g., that a *blick* is an object that grows inside of a vegetable (these were the same descriptions as in Experiment 4). Participants then read that one alien took the object/substance (e.g., the *blick*) away from a *different* place than had been described (e.g., that Theta took a *blick* from an *animal*, rather than from a *vegetable*), and that another alien put another instance of the object/substance on another place of that kind (e.g., that Mu put a *blick* onto another one of those animals). The participants then judged which of these actions were best labeled by the novel verb construction that incorporated the unrelated place name: e.g., whether Theta or Mu "blicked an animal".

Finally, we also included an additional three morphological control items to ensure that speakers were attending to the task and capable of understanding morphological changes to novel words. As in Experiment 4, on these trials, participants first learned a novel word and its meaning – e.g., that when an alien curls up into a ball, the alien *gazzers*. They were then told about what two different aliens did, and judged which of those actions were better described by a phrase incorporating a morphologically different version of the word. For example, in the *gazzer* item, participants read that "Yesterday, an alien named Sigma curled up into a ball and rested for three hours" and that "Tomorrow, an alien named Beta will curl up into a ball and rest for three hours". Participants were then asked to choose whether "Sigma gazzered for three hours", "Beta gazzered for three hours", or "Neither of the above".

6.2. Results and discussion

Our dependent measures on the critical trials were the proportion of trials on which participants chose the phrase

that labeled a transfer from a source, the phrase that labeled a transfer toward a goal, or chose "neither of the above". We defined chance responding as 0.33, as there were three possible answers for each trial.

Fig. 6 displays adults' responses for each of the critical items. On the critical functional relation trials – which described relations that licensed goal interpretations in Experiment 4 – participants readily chose to label transfers toward goals (e.g., that "Alpha lorped a spaceship" when he put *lorp* onto a spaceship, and *lorp* makes spaceships shiny). They chose to label transfers toward goals reliably more often than chance ($M = .84$, $SE = .04$; $z = 12.39$, $p < .001$) and more often than they chose to label transfers from sources (e.g., that "Epsilon lorped a spaceship" when he took *lorp* off of a spaceship; $M = .06$, $SE = .03$; $z = 11.56$, $p < .001$) or said "neither of the above" ($M = .09$; $SE = .03$; $z = 11.23$, $p < .001$). In contrast, on the critical natural origin trials – which described relations that licensed source interpretations in Experiment 4 – participants did not reliably choose to label transfers from sources (e.g., that "Alpha daxed an animal", when he took *dax* off of an animal, and *dax* comes from an animal). They did not choose to label transfers from sources reliably more often than chance ($M = .42$, $SE = .07$; $z = 1.35$, *ns*) or more often than they chose to label transfers toward goals (e.g., that "Beta daxed an animal", when he put *dax* onto an animal; $M = .42$, $SE = .07$; $z = 0$, *ns*).

Recall that in Experiment 4, when participants learned that an object/substance originated in a place (e.g., that *dax* comes from an animal), they expected the construction formed from the object/substance and place name (e.g., "dax an animal") to encode transfer from a source. The results of Experiment 5 suggest such interpretations may have arisen because goal interpretations were deemed implausible, making source interpretations more natural. When participants in Experiment 5 were told that transfers toward goals had taken place (e.g., that an alien had put *dax* onto an animal) – making goal interpretations plausi-

ble – they no longer reliably adopted source interpretations of these constructions. Thus, although source interpretations can be inferred on the basis of world knowledge, these interpretations remain malleable, and can be overridden by information regarding the plausibility of events. This pattern of results mirrors that observed in Experiments 1 and 2 with children. Recall that children were less likely to adopt source interpretations of attested source verbs like *milk* and *weed* when goal interpretations of those verbs were contextually plausible (e.g., when Big Bird "milked a jar" by pouring milk into a jar; Experiment 1) compared to when no information was given about their plausibility (Experiment 2). Thus, for both children and adults, increasing the plausibility of a goal interpretation reduces the likelihood that a verb will receive a source interpretation.

In contrast, however, increasing the plausibility of source interpretations does not appear to reduce the tendency to adopt goal interpretations, providing evidence for a goal bias. In Experiment 4, when adult participants learned that an object/substance added some function to a place (e.g., that *lorp* makes spaceships shiny), they expected its construction (e.g., "lorp a spaceship") to encode transfer to a goal. Participants continued to prefer these interpretations in Experiment 5, even when they learned that a transfer from a source had been carried out (e.g., that an alien had taken *lorp* off of a spaceship), making a source interpretation plausible. Similarly, children were at ceiling at adopting goal interpretations of attested goal verbs (e.g., "butter"), whether source interpretations of those verbs were made contextually plausible (Experiment 1) or not (Experiment 2).

Interestingly, responses to the functional relation and natural origin control trials of Experiment 5 suggest that locatum constructions can be used to label transfers toward goals, even when the transferred objects/substances are not related to their goal locations. For example, in the natural origin control trials – e.g., in which participants
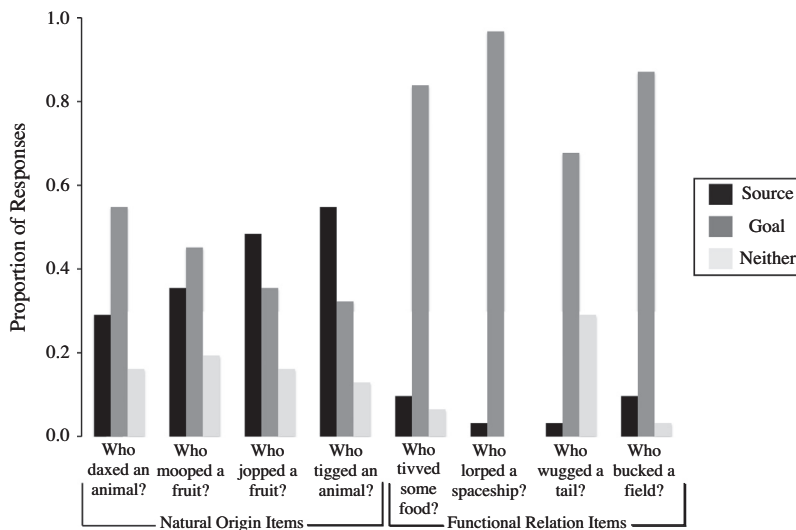

Fig. 6. The proportion of source, goal, and "neither of the above" responses to the critical items of Experiment 5.

learned that *blicks* come from vegetables but then had to judge which alien "blicked an animal" – participants chose to label transfers toward a goal reliably more often than chance (e.g., that "Mu blicked an animal" when he put a *blick* onto an animal; $M = .52$, $SE = .09$; $z = 2.19$, $p < .05$), and more often than they chose to label transfers away from sources (e.g., that "Theta blicked an animal" when he took a *blick* off of an animal; $M = .11$, $SE = .05$; $z = 3.59$, $p < .005$). Participants responded similarly to the functional relation control trials, choosing to label transfers toward goals reliably more often than chance ($M = .56$, $SE = .08$; $z = 2.95$, $p < .01$), and more often than they labeled transfers away from sources ($M = .10$, $SE = .05$, $z = 4.41$, $p < .001$). Recall that participants in Experiment 4 – who were not told that any transfers had taken place – did not adopt goal or source interpretations of these items (their modal response was "neither of the above"). This again suggests that when a transfer toward a goal has taken place – and therefore that a goal interpretation is plausible – it can be labeled by a locatum verb. Finally, as in Experiment 4, participants were reliably above chance when responding to the morphological control items ($M = .87$, $SE = .04$; $z = 11.9$, $p < .001$), indicating that they were attending to the task.

Together, these results suggest that the goal bias exhibited by young children (Experiments 1–3) may be present even in adulthood: when a goal interpretation of a construction is plausible, it is generally preferred, even for adults. Adults may differ from young children only in their ability to use their world knowledge to recognize that a particular goal interpretation is implausible (e.g., of "milk the cow"), and instead infer a source interpretation. This plausibility account may also help explain why source interpretations are limited to relations involving natural origin in the first place: it is pragmatically implausible to put milk onto a cow, weeds into a garden, or feathers onto a goose (see Clark & Clark, 1979).

7. General discussion

Although syntax often provides critical clues to meaning (Gleitman, 1990; Landau & Gleitman, 1985), it cannot alone differentiate the interpretation of constructions like "milk the cow" and "paint the house", which encode opposite meanings using the same syntactic frames. Our studies show that, to resolve whether such constructions denote transfer toward goals or from sources, adult language users draw productively on their world knowledge, and restrict their interpretations according to the plausibility of source and goal transfers. However, children initially struggle to use world knowledge to guide interpretation, and instead exhibit a productive goal bias, leading to misinterpretations of attested source verbs (an Amelia Bedelia effect). Together, our studies elucidate the developing interface between world knowledge and verb semantics, providing new evidence for a goal bias in language acquisition, and indicating how this bias is overcome.

One major conclusion of our experiments is that adults draw productively on world knowledge to mediate the assignment of thematic roles to locatum constructions.

Adults do not limit source and goal interpretations to attested verbs, but also generalize them to novel verbs, basing their interpretations on the origins and uses of objects and substances (as predicted by the proposals of Buck, 1997; Clark & Clark, 1979; Kiparsky, 1997). In support of this idea, when adults in Experiment 4 learned that a substance labeled something that naturally originated in a place (e.g., that *dax* comes from an animal), they expected the construction formed from the substance and place name (e.g., "dax an animal") to label transfer from a source. However, when they learned that the relevant object or substance did not originate in a place but instead added a function to it (e.g., that *lorp* makes spaceships shiny), adults expected the construction (e.g., "lorp a spaceship") to label transfer toward a goal. Thus, for adults, source and goal interpretations of verbs like *milk* and *paint* do not merely reflect item-based lexical knowledge, but instead indicate a generative semantics rooted in world knowledge.

A second conclusion of our studies is that, early in life, children adopt a productive learning strategy toward acquiring locatum verbs, assigning them goal interpretations in the absence of direct evidence regarding their meanings. This leads children to behave like Amelia Bedelia, and over-extend goal interpretations to attested source verbs. In both Experiments 1 and 2, children made errors when interpreting conventional source verbs – e.g., responding that Big Bird is going to "put milk onto a cow" when told that he will "milk a cow". In contrast, they rarely misinterpreted conventional goal verbs – e.g., "paint the house" – to denote transfer from sources. Experiment 3 confirmed that a goal bias guides children's acquisition of novel verbs. Under similar conditions to those that gave rise to goal and source interpretations for adults (Experiment 4), children reliably adopted goal interpretations of novel constructions but failed to adopt source interpretations. Thus, in absence of item-specific evidence in the input, children expect locatum verbs to have goal meanings. This is consistent with previous reports that children produce innovative verbs to denote transfer toward goals – e.g., "I'm *crackering* my soup" (see Clark, 1982; see also Berman, 1999; Bushnell & Maratsos, 1984).

Our findings present novel evidence of a strong goal bias in language acquisition. As noted in the introduction, previous work has suggested that goals are more salient than sources, in both language and in non-linguistic cognition (see, e.g., Lakusta & Landau, 2005; Lakusta, Wagner, O'Hearn, & Landau, 2007; Papafragou, 2010; Regier, 1997). A goal bias could lead language users to prefer that verbs describe transfer toward goals, relative to sources – as children did in Experiments 1–3, and as adults did in Experiment 5. Alternatively, language users could also prefer goal interpretations of locatum verbs due to properties of the English lexicon. In particular, because English includes a greater number of attested goal constructions than source constructions (see Clark & Clark, 1979), the majority of constructions children first learn may have goal meanings. This could lead learners to expect that novel verbs will generally denote transfer toward goals. However, as noted before, an asymmetry between descriptions of goals and sources is not specific to English, but is instead

robust cross-linguistically. Thus, even if properties of English are proximally responsible for the results reported here, a deeper goal bias in linguistic structure or in non-linguistic cognition may be needed to explain the presence of this asymmetry in English.

That children adopt a productive learning strategy toward acquiring locatum verbs converges with – and adds support to – previous reports of how children acquire generative linguistic structures (see, e.g., Pinker, 1984, 1989; Pinker, Lebeaux, & Frost, 1987; Rabagliati, Marcus, & Pylkkanen, 2010). For example, to acquire forms of nominal polysemy and verb alternations, children also adopt productive strategies that initially result in over-extension errors. For instance, for adults, words for physical objects can label the abstract content those objects contain (e.g., "The *DVD* was interesting"), but words for abstract content cannot label physical media (e.g., "The *movie* is round" is infelicitous). However, young children initially allow both licensed and unlicensed shifts – e.g., agreeing that "a movie can be round" (Rabagliati et al., 2010). Further, adults passivize some kinds of verbs (e.g., "The car was *owned* by Amelia Bedelia"), but not others (e.g., "The car was *had* by Amelia Bedelia" is infelicitous), whereas children initially violate these selectional restrictions and produce unattested passives – e.g., "I don't like being falled down on!" (Wasow, 1981; see also Pinker et al., 1987). Note that in each of these cases, children need to acquire a single productive mapping and restrict its application to admit exceptions. In principle, the case of locatum verbs should be more challenging, since children need to acquire *two* productive mappings that directly conflict with one another. This makes it all the more striking that children adopt a productive strategy, and suggests that such strategies are robust in language acquisition.

Finally, an important contribution of our studies is to shed light on children's developing ability to use world knowledge to constrain the assignment of thematic roles. Surprisingly, our data indicate that merely acquiring relevant world knowledge is not sufficient for children to overcome a goal bias and adopt source interpretations. In particular, children misinterpreted conventional source constructions and failed to generalize source interpretations even when they could express knowledge regarding the natural origin of the critical objects or substances – e.g., even when they could say that milk comes from a cow or that *dax* comes from a tree. Thus, children do not initially appear to use world knowledge to restrict their interpretations, as adults do. However, our experiments indicate that as children get older, they correct their goal overextensions and adopt source interpretations of attested verbs. In theory, it is possible that children acquire adult-like meanings for conventional source verbs by accumulating direct evidence in the input (see Tomasello, 2000, 2003). However, the fact that older children are also better able than younger children to adopt source interpretations of novel verbs (Experiment 3) suggests that item-based learning is not alone responsible for acquiring adult-like meanings. With age, children begin to use world knowledge productively – like adults – to constrain interpretation in the absence of direct, ostensive evidence.

One reason that children may initially struggle to use world knowledge is because doing so could require relatively sophisticated pragmatic reasoning. Our results suggest that adults adopt source interpretations not simply according to their knowledge of whether objects have natural origins, but because they use such knowledge to evaluate which of two possible meanings is more plausible and likely to be the intended meaning of an interlocutor. Consistent with this hypothesis, we found evidence that, for adults, the plausibility of events modulates the acceptability of goal interpretations. When adults in Experiment 5 learned that one alien had taken a substance away from its natural origin – e.g., by taking *moop* from a fruit – and another alien had put that substance onto that location – e.g., by putting *moop* onto a fruit – they did not prefer to use "moop a fruit" to label a transfer from a source, as they had in Experiment 4. Thus, when a transfer toward a goal has taken place – making a goal interpretation plausible – adults no longer reliably adopt source interpretations. This suggests that source interpretations may be inferred in contexts in which goal interpretations are implausible – e.g., as with milk and cows. This may explain, in part, why children initially overextend goal interpretations, but gradually scale back on their overextensions as they grow older. Children could initially have difficulty either with using world knowledge to reason about the plausibility of goal interpretations, or with inferring source interpretations when they deem goal interpretations implausible.

Before concluding, it is worth noting that, although we referred to information regarding the origin and use of objects as types of "world knowledge", such content could be represented and integrated with verb interpretation in a number of different ways (Clark & Clark, 1979; Kiparsky, 1997). One possibility, advanced by Pustejovsky (1995), is that this information is encoded within the structure of lexical items themselves (see also Kiparsky, 1997). For example, to account for systematic forms of polysemy, Pustejovsky argues that nouns specify what their referents are made of, what form they have, what their function is, and how they are created (see also Moravcsik, 1998). By his account, a verb like *bake* specifies an act of creation in "bake the cake", but only a change of state in "bake the potato", because of lexically-specified information regarding how cakes and potatoes come into being. That is, because cakes – but not potatoes – are created by baking, "bake the cake" receives a creative reading. Such lexical structures could also constrain the interpretation of locatum verbs – e.g., a source interpretation of "milk the cow" could arise compositionally because the representation of *milk* specifies that milk comes from animals like cows. The results of Experiment 5 suggest that such interpretations would be pragmatically defeasible, in contexts making goal interpretations plausible – e.g., when milk has been put onto a cow.

Alternatively, the interpretation of locatum verbs could depend not on lexically-specified knowledge, but instead on non-linguistic conceptual structure, context, and the ability to reason pragmatically about the beliefs and goals of interlocutors (see, e.g., Clark & Clark, 1979; Nunberg, 1979). For example, conventional interpreta-

tions of verbs like *milk* and *paint* could be supported by generic conceptual knowledge regarding the ontogeny of those substances and how they are typically used. Novel interpretations of verbs could arise via pragmatic reasoning. For example, in some cases, a verb's interpretation is not predicted by the canonical role of an object or substance, but instead by how that object or substance is used in a particular context. Thus, "bottle the policemen" can apply to an event in which bottles have been thrown at policemen, even though this is not a canonical use of bottles, and is unlikely to be lexically-specified (Clark & Clark, 1979). Verb meanings can also depend critically on specific knowledge shared between interlocutors. For example, "The man *Napoleoned* for the camera" could mean that a man posed for a photo by tucking his hand into his jacket (see Clark & Gerrig, 1983). Such an interpretation could arise if the way Napoleon posed in portraits is part of the common ground and can be easily recognized as relevant to the sentence. Future studies should extend the approach presented here to test the role of context in children's acquisition and interpretation of locatum verbs, to help differentiate between lexicalist and pragmatic accounts of creative language use.

Acknowledgments

References

Arnold, J. (2001). The effect of thematic roles on pronoun use and frequency of reference continuation. *Discourse Processes, 31*, 137–162.

Baker, C. L. (1979). Syntactic theory and the projection problem. *Linguistic Inquiry, 10*, 533–581.

Barner, D., & Bale, A. (2005). No nouns, no verbs? Rejoinder to panagiotidis. *Lingua, 115*, 1169–1179.

Berman, R. (1999). Children's innovative verbs versus nouns: Structured elicitations and spontaneous coinages. In L. Menn & N. Bernstein Ratner (Eds.), *Methods for studying language production* (pp. 69–93). Mahwah, NJ: Erlbaum.

Brysbaert, M., New, B., & Keuleers, E. (2012). Adding part of speech information to the SUBTLEX-US word frequencies. *Behavior Research Methods, 44*, 991–997.

Brysbaert, M., & New, B. (2009). Moving beyond Kucera and Rancis: A critical evaluation of current word frequency norms and the introduction of a new and improved word frequency measure for American English. *Behavior Research Methods, 41*, 977–990.

Buck, R. A. (1997). Words that are their opposites: Noun to verb conversion in English. *Word, 48*, 1–14.

Bushnell, E. W., & Maratsos, M. P. (1984). "Spooning" and "basketing": Children's dealing with accidental gaps in the lexicon. *Child Development, 55*, 893–902.

Clark, E. V., & Clark, H. H. (1979). When nouns surface as verbs. *Language, 55*(4), 767–811.

Clark, H. H., & Gerrig, R. J. (1983). Understanding old words with new meanings. *Journal of Verbal Learning and Verbal Behavior, 22*(5), 591–608.

Clark, E. V. (1982). The young word maker: A case study of innovation in the child's lexicon. In E. Wanner & L. Gleitman (Eds.), *Language acquisition* (pp. 390–428). Cambridge: Cambridge University Press.

Fisher, C. (2002). The role of abstract syntactic knowledge in language acquisition: A reply to Tomasello (2000). *Cognition, 82*, 259–278.

Fisher, C., Hall, D. G., Rakowitz, S., & Gleitman, L. (1994). When it is better to receive than to give: Syntactic and conceptual constraints on vocabulary growth. In L. Gleitman & B. Landau (Eds.), *Acquisition of the lexicon* (pp. 333–375). Cambridge, MA: MIT Press.

Frege, G. (1892/1948). On sense and reference. *The Philosophical Review*, 207–230 (Translated by M. Black).

Gagne, C. L., & Shoben, E. J. (1997). Influence of thematic relations on the comprehension of modifier–noun combinations. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23*(1), 71–87.

Gleitman, L. (1990). The structural sources of verb meaning. *Language Acquisition, 1*(1), 3–55.

Harley, H. (1999). Denominal verbs and aktionsart. In A. van Hout, L. Pylkkanen, & H. Harley (Eds.). *Papers from the Upenn/MIT roundtable on the lexicon* (Vol. 35, pp. 73–85). MIT Working Papers in Linguistics.

Horst, J. S. (2009). The Novel Object and Unusual Name (NOUN) database. <http://www.lifesci.sussex.ac.uk/research/wordlab/NOUN-database.pdf>.

Kiparsky, P. (1997). Remarks on denominal verbs. In A. Alsina, J. Bresnan, & P. Sells (Eds.), *Complex predicates* (pp. 473–499). Palo Alto, CA: CSLI Publications.

Lakusta, L., & Landau, B. (2005). Starting at the end: The importance of goals in spatial language. *Cognition, 96*, 1–33.

Lakusta, L., & Landau, B. (2012). Language and memory for motion events: Origins of the asymmetry between source and goal paths. *Cognitive Science*, 1–28.

Lakusta, L., Wagner, L., O'Hearn, K., & Landau, B. (2007). Conceptual foundations of spatial language: Evidence for a goal bias in infants. *Langauge Learning and Development, 3*, 179–197.

Landau, B., & Gleitman, L. (1985). *Language and experience: Evidence from the blind child*. Cambridge, MA: Harvard University Press.

Moravcsik, J. M. E. (1998). *Meaning, creativity, and the partial inscrutability of the uuman mind*. Stanford: CSLI.

Murphy, G. L. (1988). Comprehending complex concepts. *Cognitive Science, 12*, 529–562.

Nunberg, G. (1979). The non-uniqueness of semantic solutions: Polysemy. *Linguistics and Philosophy, 3*, 145–184.

Papafragou, A. (2010). Source-goal asymmetries in motion representation: Implications for language production and comprehension. *Cognitive Science, 34*, 1064–1092.

Parish, P. (1979). *Amelia Bedelia helps out*. New York: Greenwillow Books.

Pinker, S. (1984). *Language learnability and language development*. Cambridge, MA: Harvard University Press.

Pinker, S. (1989). *Learnability and cognition: The acquisition of argument structure*. Cambridge, MA: MIT Press.

Pinker, S., Lebeaux, D. S., & Frost, L. A. (1987). Productivity and constraints in the acquisition of the passive. *Cognition, 26*, 195–267.

Pustejovsky (1995). *The generative lexicon*. Cambridge, MA: MIT Press.

Rabagliati, H. A., Marcus, G. F., & Pylkkanen, L. (2010). Shifting senses in semantic development. *Cognition, 117*(1), 17–37.

Regier, T. (1997). Constraints on the learning of spatial terms: A computational investigation. In R. Goldstone, P. Schyns, & D. Medin (Eds.), *Psychology of learning and motivation. Mechanisms of perceptual learning* (Vol. 36, pp. 171–271). San Diego: Academic Press.

Regier, T., & Zheng, M. (2007). Attention to endpoints: A cross-linguistic constraint on spatial meaning. *Cognitive Science, 31*, 705–719.

Stefanowitsch, A., & Rohlde, A. (2004). The goal bias in the encoding of motion events. In G. Radden & K.-U. Panther (Eds.), *Studies in linguistic motivation*. Berlin: Mouton de Gruyter.

Talmy, L. (2000). *Towards a cognitive semantics*. Cambridge, MA: MIT Press.

Talmy, L. (1985). Lexicalization patterns: Semantic structure in lexical forms. In T. Shopen (Ed.), *Language typology and syntactic description* (pp. 57–149). New York: Cambridge University Press.

Tomasello, M. (2000). The item-based nature of children's early syntactic competence. *Trends in Cognitive Sciences, 4*(4), 156–163.

Tomasello, M. (2003). *Constructing a language: A usage-based theory of language acquisition*. Cambridge, MA: Harvard University Press.

Tomasello, M., & Brooks, P. (1998). Young children's earliest transitive and intransitive constructions. *Cognitive Linguistics, 9*, 379–395.

Wasow, T. (1981). Comments on the paper by Baker. In C. L. Baker & J. J. McCarthy (Eds.), *The logical problem of language acquisition* (pp. 324–329). Cambridge, MA: MIT Press.