

Running down the clock: The role of expectation in our understanding of time and motion

Michael Ramscar

Department of Psychology, Stanford University, Stanford, CA, USA

Teenie Matlock

Cognitive Science Program, University of California, Merced, CA, USA

Melody Dye

Department of Psychology, Stanford University, Stanford, CA, USA

Time is often talked about in terms of motion. People talk about themselves ‘moving’ through time, or about time ‘moving’ relative to them. Previous research has shown that attending to actual motion can influence judgements about time. Further, fictive motion language – figurative attributions of motion to static objects in space – has been shown to have much the same effect, suggesting that thought about space influences thought about time. However, evidence to date on fictive motion comes from experiments that included some degree of actual motion, such as drawing. In a series of four experiments, we tease apart the influence of actual motion and fictive motion language on people’s understanding of time. The results suggest that the similar ways in which people talk about motion through space and motion through time play an important part in their common underlying conceptualisation. This has important implications for our understanding of what comprises literal and metaphorical uses of language, and for the relationship between language, language use, and thought.

Correspondence should be addressed to Michael Ramscar, Department of Psychology, Stanford, CA 94305, USA. E-mail: ramscar@stanford.edu

This material is based upon work supported by the National Science Foundation under Grant Nos. 0547775 and 0624345 to Michael Ramscar. The authors thank Lera Boroditsky, Herbert Clark, Caitlin Fausey, Raymond Gibbs, Jr., Roddy Lindsay, Art Markman, Matt McGlone, Paul Maglio, and two anonymous reviewers for providing helpful comments on earlier drafts. The authors also thank Danielle Birkley, Nicole Gitcho, Michael Frank, Krysta Hays, Nilofaur Tahery, and Elsie Wang for assistance with data collection, and Justin Matthews and Ronald Yockey for assistance with data analysis.

© 2009 Psychology Press, an imprint of the Taylor & Francis Group, an Informa business

<http://www.psypress.com/lcp>

DOI: 10.1080/01690960903381166

Keywords: Concepts; Language and thought; Meaning; Metaphor; Prediction; Thinking for speaking.

People often talk about time in terms of motion. English speakers may describe themselves as *moving* through time toward or past events with statements such as ‘we’re entering the holidays’ or ‘we slipped past the due date’. They may also talk about events as *moving* toward or past themselves with statements such as ‘tough times are approaching us’ or ‘summer vacation has passed’. While this use of ‘borrowed’ language – talking about one thing in terms of another – might be dismissed as nothing more than an interesting curiosity, a growing body of evidence indicates that this metaphorical use of language may reflect an important process through which people form their understandings of the world (Boroditsky, 2000; Boroditsky & Ramscar, 2002; Clark, 1973; Gibbs, 1994; Lakoff & Johnson, 1980).

Take motion language, for example. In its literal uses, it is descriptive of paths and trajectories of objects, e.g., ‘people crossing roads’, ‘cars going through towns’, and so on. In its metaphoric applications (which are pervasive in everyday speech), motion language can also be descriptive of abstract domains, including emotion, thought, time, and notably, space itself. When motion through space is applied metaphorically, it is called *fictive motion* (see Matlock, 2004a, 2004b; Talmy, 2000). In fictive motion sentences, a motion verb is applied to a subject that is not literally capable of movement in the physical world, e.g., ‘the path swings along the cliff’. Fictive motion is so called because it is attributed to material states, objects, or abstract concepts that cannot (sensibly) be said to move through physical space.

Interestingly, not only has it been shown that thinking about *actual* motion influences people’s judgements about time (e.g., Boroditsky & Ramscar, 2002), but it also appears that thinking about fictive motion has the same effect (Matlock, Ramscar, & Boroditsky, 2005), suggesting that thinking about one abstract domain may influence people’s understanding of another. However, evidence to date about the influence of fictive motion on judgements of time comes only from tasks that include some degree of actual motion, such as drawing. This leaves open many possibilities regarding the nature of the relationship between actual motion, fictive motion, and time. Is the influence of fictive motion on people’s understanding of time rooted in a concrete, embodied conception of motion, such that both time and fictive motion are ultimately understood in terms of simulations of concrete experience (e.g., Langacker, 1999)? Or are the effects of fictive motion a product of the way that language itself influences the way we think (e.g., Slobin, 1996)? The results from four experiments exploring these questions support the idea that time can be independently influenced by fictive motion,

providing further evidence that thought about time can be influenced by thought about motion and space. We propose that the similar ways in which people talk about motion through space and motion through time is an important part of their underlying conceptualisation. This has many implications for our understanding of literal and metaphorical uses of language, for the relationship between language and thought, and for the way we think about the way we conceptualise the world.

CONCEPTIONS OF TIME IN EGO-MOVING AND TIME-MOVING METAPHORS

Metaphorical talk about time is often described in terms that ‘borrow’ from two different perspectives for conceptualising motion. In the *ego-moving* metaphor, time progresses along a time-line toward the future, while in the *time-moving* metaphor, time is a conveyor belt upon which events move from the future to the past like packages (Clark, 1973; Gentner, 2001; Lakoff, 1987; Lakoff & Johnson, 1980; McTaggart, 1908). Similar space–time metaphors exist in many cultures, and they are realised in different ways (see Boroditsky, 2001; Emanation, 1992; Evans, 2004; Moore, 2004; Núñez & Sweetser, 2006; Radden, 1996).

There is evidence to support the idea that the ego- and time-moving metaphors may embody distinct, systematic frameworks for conceptualising time. McGlone and Harding (1998), asked participants to read sentences phrased either in terms of an ego-moving metaphor (e.g., ‘we passed the deadline two days ago’), or a time-moving metaphor (e.g., ‘the deadline passed two days ago’). Participants then indicated the day of the week on which an event had occurred or would occur. Some of the sentences were ambiguous, for example: ‘*The meeting originally scheduled for next Wednesday has been moved forward two days*’, yields a different answer depending on whether an ego- or time-moving perspective is adopted towards ‘forward’. Participants in the ego-moving condition tended to interpret ‘moved forward’ accordingly, thinking the meeting had moved to Friday, whereas participants in the time-moving condition tended to think the meeting had moved to Monday (see also McTaggart, 1908).

Extending on this, Boroditsky (2000) asked participants to imagine either moving toward an object, or an object moving towards them. They then read and answered the question, ‘*Next Wednesday’s meeting has been moved forward two days. What day is the meeting now that it has been re-scheduled?*’ Boroditsky’s results suggested that participants’ conceptions of motion were significantly influenced by the scenario they had engaged with. If they had imagined themselves moving toward the designated object, they were more likely to adopt an ego-moving perspective and answer ‘Friday’. Conversely, if

they had imagined the designated object moving toward them, they were more likely to take a time-moving perspective and answer ‘Monday’.

Matlock et al. (2005) examined whether the influence of motion on the understanding of time would generalise to *fictive motion*, the figurative attribution of motion to non-moving objects that is common in everyday speech (see Matlock, 2004a, 2004b; Talmy, 2000). The experiments in Matlock et al. were designed to examine how figurative attributions of motion might shape how people think about the things that motion is attributed to – and in particular, how this might affect the way people think about time. Participants read sentences employing fictive motion language (hereafter, FM sentences), such as ‘the tattoo runs along his spine’, and sentences absent of fictive motion language (hereafter, NFM sentences), such as ‘the tattoo is next to his spine.’ They then drew a picture of what they had imagined, to ensure they paid attention and reflected on the meaning of the sentence. When participants were then asked to answer the question, ‘*Next Wednesday’s meeting has been moved forward two days. What day is the meeting now that it has been re-scheduled?*’ they gave significantly more Friday than Monday responses when they drew FM sentences, but gave about equal Friday and Monday responses when they drew NFM sentences. Indeed, the NFM results were similar to those produced when the ambiguous question is given with no context to a control group (previous studies have shown that participants are about equally likely to respond ‘Monday’ or ‘Friday’ in the absence of any priming; Boroditsky, 2000).

ANTICIPATION AND EXPECTATION: A FRAMEWORK FOR UNDERSTANDING FICTIVE MOTION

While it has been shown that experiences with motion, both actual and figurative, can influence people’s judgements about time, the nature of the relationship between actual motion, fictive motion, and time remains an open question. It has been suggested that the way fictive motion is conceptualised may be similar to that of actual motion (Langacker, 1987; Matlock & Richardson, 2004; Matsumoto, 1996; Richardson & Matlock, 2007; see Talmy, 2000 for discussion of various types of fictive motion; see also Matlock, 2004a, 2004b for a discussion of the semantics of fictive motion from a linguistic perspective). Some have proposed, for example, that motion understanding involves a sort of *simulation*, such that parsing metaphorical attributions of motion involves a mental scanning process. (Mental scanning accounts propose that, at some level, speakers *literally* adopt ego-moving or time-moving perspectives by simulating motion through those perspectives; e.g., Langacker, 1999; Matlock et al., 2005.) There is much to be said, intuitively, for explaining the relationship between

concrete domains and their analogical extensions in terms of deeper shared conceptual structures, as mental scanning accounts suggest. However, from a theoretical perspective, explaining people's understanding of metaphorical forms of motion by appealing to their understanding of actual motion simply transfers the problem; it is far from clear how actual motion is conceptualised or how our 'linguistic' understanding of the words relating to actual motion is realised (see Murphy, 2002 for a review).

Is there an alternative explanation for the similarities between people's understanding of actual and metaphorical motion, and the way that thinking (or talking) about one form of abstract motion influences the understanding of another? One obvious alternative account is suggested in the way in which talk about literal, fictive and temporal motion employs many of the *same* words and phrases (e.g., we can just as easily say 'the dog *runs by*' as 'the road *runs by*' as 'time *runs by*'). When people use the same words and phrases to express their understanding of literal and fictive motion, much of the same background knowledge and many of the same processes will be employed. Accordingly, it may be that some of the apparent conceptual overlap in these domains is a reflection of the cognitive consequences of using the same words to talk about different things. Similar ways of talking about different domains will involve similar patterns of 'thinking for speaking' (the various cognitive steps involved in processing those ways of speaking; see Slobin, 1996) and these shared patterns may in turn increase the underlying conceptual similarity between domains.

This idea might be formulated in two ways. The first, 'words as ephemera', formulation would assume that people's conceptions of literal, fictive, and temporal motion are independent, and that the priming results reviewed above are a relatively uninteresting consequence of the way that people's talk about these independent conceptions makes use of the same words. The second, 'words as culture' formulation, would note that attempts to cash out 'independent conceptions' and 'deep conceptual structures' (embodied or otherwise) have proven notoriously difficult (Murphy, 2002), and that there are many persuasive reasons to believe that the meanings of words are intimately bound up with the ways in which they are used (Wittgenstein, 1953; see also Tomasello, 1999). Under this formulation, much of 'conceptual structure' – especially that of abstract domains – is *inseparable* from the language in which it is expressed. This suggests that at least some aspects of people's conception of literal and abstract domains of motion are bound up with the ways in which motion words are used. By this account, rather than simply serving to express ideas, words may themselves contribute to the ways that cultures structure and encode abstract knowledge.

An obvious way in which words (and from a broader perspective, languages as cultural artifacts) might serve to structure knowledge in this

manner is by systematically shaping the expectations of speakers and listeners. Studies in a variety of research paradigms have revealed that when people are listening to or reading a sentence they build up *linguistic expectations*, anticipating upcoming words based on the structure and semantics of the prior discourse (e.g., Altmann & Mirković, 2009; Kutas & Federmeier, 2007; Tanenhaus & Brown-Schmid, 2008). While this may seem obvious in idiomatic phrases, such as ‘cross my heart and hope to ___’ or ‘hit the nail on the ___’, such studies have revealed just how extensively expectation pervades linguistic processes. Although studies to date have often focused on anticipation of a specific word, object, or event based on prior context (e.g., Altmann & Steedman, 1988; Tanenhaus, Spivey-Knowlton, Eberhard, & Sedivy, 1995; Otten & Van Berkum, 2007, 2008), it is clear that in natural speech, listeners are anticipating (probabilistically) a number of different possible words that might follow in a given speech stream (Altmann & Mirković, 2009; Tanenhaus & Brown-Schmidt, 2008).

One way these kinds of anticipatory inferences can be built up, over time, is by attending to the co-occurrence patterns of different words in speech. For example, ‘juggle balls’ is a much higher frequency pairing than ‘juggle chairs’ – over a hundred times more frequent – even though *balls* and *chairs* are similarly frequent words. If that mismatch seems obvious, it should; words do not co-occur with each other with equal frequency. Indeed, the distribution of words in languages is highly systematic (Baayen, 2001) and much of everyday discourse is made up of repetitive, highly predictable speech, containing highly frequent co-occurrence patterns (Tannen, 2007). Thus, in addition to making predictions about what comes next based on the content of the conversation, listeners are also making predictive inferences based on their knowledge of *how language works*; that is, how words co-occur in sensible, and less sensible, ways (e.g., Wicha, Bates, Moreno, & Kutas, 2003).

The kinds of expectations that people build up about words in listening and reading may have an important part to play in their conceptualisation of those words. Words are often thought of as being abstractions of objects and events in the world, but defining a simple relation between the thing being represented and the label that represents it is problematic (Murphy, 2002). Indeed, it has been argued that the meanings of words are better understood in relation to their *patterns of use*, rather than to the things in the world they appear to represent (Wittgenstein, 1953). When we talk about names, for example, we say things like ‘did you *catch* her name’, ‘his name is *mud*’, ‘they were *called* by name’, ‘she *made* a name for herself’, and so on. From this perspective, a ‘name’ is not only a word ‘by which something is called or known’, as the dictionary designates, but also a thing to be had, caught, muddied, cleared, called, and made. People ‘go by names’, they ‘throw names around’, they hope to see their ‘name in lights’, and on this view, the

meaning of ‘name’ is inextricable from its patterns of use: from the words it co-occurs with and the words that modify it, and the effect that these have on the way that people think about ‘name’.

These kinds of co-occurrence patterns offer a rich and readily available source of information for anyone learning to understand the world and the way that language relates to it, and there is considerable evidence to support the idea that people are sensitive to this information. Our suggestion is that people’s understanding of the patterns of use associated with motion words actually plays an important part in shaping their understanding of them. For instance, saying that time can ‘run out’ or ‘fly by’ influences what we understand time to *be* in the first place, because thinking about time in this way involves processes shared with other things that ‘run out’, ‘fly by’, or ‘stand still’ (Slobin, 1996). In a sense, the mind works metaphorically, associating words with other words that are used in similar ways.

If understanding results (at least in part) from predictive processes and the expectations produced by patterns of co-occurrence, then when we *use* words in similar ways they ought to become more closely aligned in meaning. This would suggest that saying literally ‘the man *runs* by’, fictively ‘the road *runs* along the river’, and figuratively ‘time *runs* out’, should, as a result of this common pattern of usage, more closely align our notions of how space and time operate (in this case, aligning them with an agentive – acting – verb: *run*). Accordingly, we suggest that the similar ways in which people talk about motion through space and motion through time is an important part of their common underlying conceptualisation.

LANGUAGE, MOTION AND EXPERIENCE

The experiments reported in Matlock et al. (2005) and reviewed above provide some reasons to believe that talk about figurative motion may influence people’s understanding of time. However, the question of whether fictive motion talk alone is sufficient to influence temporal understanding remains open. Participants in Matlock et al.’s studies had to generate a drawing before answering the ‘move forward’ question. There is reason to believe that this may have contributed to the temporal priming effect reported. For example, in Experiment 1 in Matlock et al. (2005), participants who provided Friday responses drew more extended depictions than those who provided Monday responses (e.g., sword for ‘the tattoo runs along his spine’ versus apple for ‘the tattoo is next to his spine’). In addition, participants who depicted fictive motion sentences included more motion elements in their drawings, such as bicycles and cars (see Matlock, Ramscar, & Boroditsky, 2003). It is possible that these differences in drawing could

have produced the differences in participants' responses about when the meeting was held, rather than the language used in the sentence primes.

In designing the following studies, we had three main questions in mind: Does drawing influence temporal understanding in the absence of fictive motion language? Can fictive motion language affect people's conceptions of time in the absence of drawing? And: Can people's understanding of temporally ambiguous questions be manipulated by changing the pattern of their underlying linguistic expectations? We conducted four experiments to explore these questions, priming participants with either drawing *or* fictive motion language, and studying their responses to the ambiguous 'move forward' question (adapted from McClone & Harding, 1998). Experiment 1 investigated the independent impact of drawing on temporal understanding. Experiments 2, 3 and 4 then investigated whether language could shape participants' temporal understanding in the absence of drawing and whether it might do this by differentially priming their linguistic expectations. Experiment 2 was specifically designed to test the linguistic expectations hypothesis by examining whether it was possible to manipulate people's understanding of the 'move forward' question with lexical primes. Experiments 3 and 4 then sought to establish the degree to which fictive motion sentences alone could evoke temporal priming, and to explore the extensibility of the lexical prediction hypothesis in accounting for these phenomena.

EXPERIMENT 1

In seeking to better understand the relationship between actual and fictive motion, we first sought to examine whether the act of drawing would influence people's understanding of time in the absence of any motion language. As we noted above, since all of the experiments in Matlock et al. (2005) required participants to both draw and respond to fictive motion primes, it was unclear which of these factors produced the bias in their participants' answers to the ambiguous time question. In this experiment, participants were asked to draw short, medium, or long lines before answering the ambiguous 'move forward' question. Given that drawing lines requires participants to think both about drawing and about lines (which have beginnings and ends) and given that conceptions of time can be influenced by thought about movement, it seemed that drawing lines might influence the way people think about time. In particular, since drawing in this task requires the participant to move their pencil-holding hand across the page until it reaches an end point, it seems that drawing might implicitly force the participant to adopt an ego-moving perspective. Accordingly, we expect that this might bias participants toward a Friday response.

Method

Fifty-three participants saw a horizontal line at the top of a questionnaire page that was short (1 inch), medium (2.5 inches), or long (5 inches), and read the instructions, 'Use a pen or a pencil and copy the line three times'. In the middle, participants read and answered the 'move forward' question, *Next Wednesday's meeting has been moved forward two days. What day is the meeting now that it has been rescheduled?* At the bottom, they specified whether they had drawn their lines to left to right or right to left. Data from three participants whose lines went right to left were discarded so the analysis would include responses with uniform drawing direction.

Results and discussion

Overall, participants were more inclined to give a Friday response (64%) than a Monday response (36%). As shown in Figure 1 Friday responses increased proportionally to line length, specifically, from 43% (short length) to 59% (medium) to 88% (long). Percentages are based on 16, 17, and 17 participants, respectively. A Chi-square analysis showed that the overall effect of a bias towards responding 'Friday' after line drawing was reliable, $\chi^2(1) = 6.99, p = .008$.

The results are consistent with how people depicted fictive motion scenes in Matlock et al. (2005). Overall, the drawings in the FM condition contained longer paths and other figures than did drawings in the NFM condition. For example, highways in drawings of *The highway runs along the*

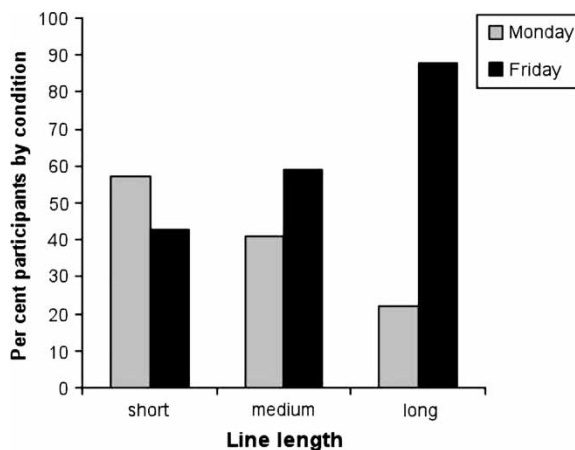


Figure 1. Participants were more likely to provide a Friday response as the length of their lines increased.

coast were reliably longer than highways in drawings of *The highway is next to the coast* (see Matlock et al., 2003, for discussion).

Why might drawing short lines have no affect on our participants' understanding, while drawing longer lines did bias people toward more Friday responses? One answer is suggested by Boroditsky and Ramscar (2002), who found that participants on trains were more likely to provide Friday answers to the ambiguous 'move forward' question at the beginnings and ends of journeys, but not in the middle of them. Boroditsky and Ramscar suggest that this is because participants were more likely to be thinking about motion *towards* a destination at the beginnings and ending of journeys, and it was this, rather than any actual movement associated with being on a train, that biased their participants' responses. Similarly, it is likely that being required to draw a longer line required our participants to focus more on the act of drawing the line (the longer lines could be less easily drawn in a single stroke) and correspondingly, to attend more to the end point of the drawing process. We suggest that it was this that affected their understanding of time.

This finding is consistent with Boroditsky and Ramscar (2002), and lends support to the notion that metaphorical motion understanding is related to literal motion understanding (since the act of drawing – which involves thinking about moving towards a goal – appeared to affect participants' temporal understanding). It is also somewhat consistent with the 'thinking for speaking' hypothesis, insofar as the finding that participants' temporal understanding became biased the more that they had to think about drawing towards an end point is consistent with the idea that the habitual thought processes associated with one kind of behaviour (in this case drawing) might influence another (thinking about time). Further, this finding suggests an alternative interpretation to the findings reported by Matlock et al. (2005): specifically, if thinking about drawing alone can influence participants' temporal thinking, then given that participants in all of Matlock et al.'s experiments drew fictive motion stimuli, it is possible that differences between the drawings participants produced for fictive and non-fictive motion primes – and *not* fictive motion language itself – was responsible for different patterns of data reported.

Accordingly, in Experiment 2, we sought to examine whether the patterns of priming for fictive motion language reported by Matlock et al. could be reproduced by having participants answer the temporally ambiguous '*next Wednesday's meeting ...*' question after reading fictive motion sentences, but *without* engaging in drawing. At the same time, we sought to establish whether our lexical prediction hypothesis could offer a plausible account of fictive motion priming by testing whether people's understanding of the 'move forward' question would be sensitive to changes in the expectations the FM sentences might be expected to evoke.

EXPERIMENT 2

How does fictive motion influence people's conceptualisation of time? Do different 'strengths' or 'magnitudes' of fictive motion affect underlying representations of actual motion, inducing corresponding differences in the way people imagine movement through time? Or do different number words prime conceptualisations of time in different ways?

One way in which it has been suggested that people might use literal motion to understand fictive motion is via imagistic parsing, in which they *imagine* a road running along a stream as they might recall a freight train racing down a track, thus engaging in some form of 'mental scanning' (e.g., Langacker, 1999). From this perspective, a fictive motion sentence may provide 'scan points' – discrete points that can be individuated and attended to one by one, such as the trees in 'four trees run along the driveway'. If this is so, understanding a fictive motion sentence involves mentally scanning along these points. It also suggests that reading a sentence about a series of *many* objects ought to lead to correspondingly more Friday responses than a sentence about a series of *few* objects (e.g., imagining *many trees* versus *a few trees* along a driveway), since it involves relatively more motion priming.

Matlock et al. (2005; Experiment 2) explored this idea by having participants read one of four fictive motion sentences that varied according to the number of scan points along a path, specifically, 4, 8, 10, or over 80 in *Four, Eight, Twenty, or Over Eighty pine trees run along the edge of a driveway*. The results from this 'scanning magnitude' experiment did not match up with the hypothesis. Participants could indeed be primed to provide more Friday responses by reading about (and drawing) trees running along driveways, but while a priming effect was produced by both 'eight trees' and 'twenty trees,' reading about 'four trees' or 'over eighty trees' produced no discernable change in participants' thinking about time.

However, the mixed pattern of data reported by Matlock et al. (2005) does not necessarily discount the scanning hypothesis. It may have been the case – given the results we report in Experiment 1 – that drawing had an impact on these results. Participants in Matlock et al.'s study drew less extended pathways in response to '*four trees ...*' and often drew a block of trees to represent '*over eighty ...*' and it may be that this was either responsible for the bias in participants' answers, or that it affected the pattern of data produced. In the following experiment, we attempt to disambiguate the results by conducting a similar experiment without drawing.

If the scanning hypothesis is indeed correct, we might expect the priming effects that result from scanning to bear some relation to the amount of scanning that understanding a sentence involves; thus, the more scanning required, the more priming expected. On the other hand, if people's latent understanding of the co-occurrence patterns in language has a part to play in

shaping temporal understanding, then we would expect a markedly different pattern of results. Indeed, depending on the number words involved (and their co-occurrence patterns), the magnitude of priming should be discernable *independently* of the physical magnitudes described.

This is because the distributions of the words we use to describe magnitude (number words) and the words we associate with temporal understanding (time words) are closely entwined. An examination of the 385 million word *Corpus of Contemporary American English* (COCA; Davies, 2009) reveals that, in contemporary American English, the word most likely to follow *four*, *eight* and *twenty* is the same: *years* (for *eighty*, it is the second most likely, after *percent*). People seem to care about time, and they talk about time a great deal. Further, when people talk about time, they seem to care more about the future than the past. For all of the major time units that people talk about (seconds, minutes, hours, days, weeks, months, and years) the word *later* is one of the 5 most likely words to follow it. In contrast, *earlier* is not even in the top 10 most likely words to follow any of the major time unit words, and only in about half the cases is *earlier* even among the top 20 most likely following words (similarly, the number and time word sequence ‘in ___ years’ occurs with considerably greater frequency than ‘___ years ago’). This indicates that number words may prime a future-looking bias.

However, the detailed picture is slightly more complicated than this suggests. Because the frequency distributions of words obey Zipf’s Law (Zipf, 1935, 1949), the frequency with which one word follows another is heavily skewed, such that a few words have an extremely high probability of following say, the number *four*, while many other words have a much lower probability. Thus, relatively small differences in co-occurrence rank order (top 5 versus top 20) represent huge differences in frequency. In COCA, 5 of the 10 most likely words to be encountered after *four* are time words (by descending rank frequency: *years*, *months*, *days*, *hours* and *weeks*). Given that time words lead to the expectation of future biased words, it follows that the *more* a number word primes time words, the more likely it is to correspondingly prime future biased words, such as *later*. This means that unless all number words share the same distributional pattern with respect to time words, different number words will likely produce different priming effects.

Indeed, while many small integers (such as *four*) are highly likely to be used with – and thus to co-occur with – different time units, this is not necessarily the case for larger numbers (such as *eighty*). Only 2 of the 20 words most likely to follow *eighty* are time words. These distributional differences suggest an alternative account of the pattern of data reported in Matlock et al. (2005; Experiment 3): namely, that the different amounts of priming produced by the FM primes was a consequence of the different

degrees to which integers lead to the expectation of time words and thereby biased future understandings.

One benefit of this suggestion is that linguistic expectations are amenable to relatively detailed modelling, allowing this idea to be tested by assessing the degree to which a given number word should be expected to affect people's understanding of the ambiguous 'move forward' question. We would expect that if linguistic expectations are influencing people's answering of the ambiguous temporal question, then given the degree to which it leads to the expectation of time words, *four* ought to produce significant priming in the absence of drawing.

Accordingly, we constructed two models in order to estimate the probability that a time word would follow a given number word. Since time words are biased toward priming 'later', we hypothesised that the more strongly each number primed time words, the more strongly it would prime future words and thus, the Friday response. We thus expected that the degree to which number words are predictive of time words would be predictive of the variance in the degree to which different numbers of trees primed participants' responses to the ambiguous 'move forward' question.

Models

To provide for a larger range of values to model, we added five new integers (*ten*, *eleven*, *twelve*, *nineteen* and *a hundred*) to the four original number words tested in Matlock et al. (2005), Experiment 2 (*four*, *eight*, *twenty*, and *eighty*). In English, the distribution of integers has a Zipfian distribution in which frequency largely decreases in *inverse* proportion to the magnitude of an integer (subject to a degree of bias caused by factors such as the base-ten and imperial number systems). Thus of the words we modelled, *four* is the most frequent (occurring 145,699 times in the 385 million words in COCA, i.e., 378 times per million words; Davies, 2009), and *eight* the second most frequent (127 occurrences per million words in COCA), followed by *ten* (101), *a hundred* (83), *twenty* (44), *twelve* (24), *eleven* (16), *nineteen* (6), and *eighty* (6). The five new numbers were specifically chosen to reduce the strong negative correlation between word frequency and magnitude in English integers (the correlation between the magnitude of the nine numbers and their frequencies was $r = -0.5$).

We then used COCA to examine the 20 words that occurred with the highest frequency following each of the numbers mentioned in the priming stimuli (i.e., the words in the distribution that follow *four*, *eight*, *ten*, *eleven*, and so on). It is important to note that because of Zipf's Law (Zipf, 1935, 1949), the frequency distributions of even the most frequently co-occurring words is heavily skewed; in our experiment, the probability that a given time word would follow each number word was inversely proportional to its

rank in the frequency table. For example, the first most frequent word following ‘eight’ (years) is three times as likely as the second most frequent word (months) and almost five times as likely as the third most frequent word (hours). To avoid overestimating the influence of the highest frequency co-occurrences in our model, the raw frequency counts of co-occurrences were log transformed before their probabilities were calculated.

In Model 1, we first identified the time words that were amongst the twenty most frequent words in the distribution immediately following each of target integers (the full set included *seconds*, *minutes*, *hours*, *days*, *weeks*, *months*, *years*, and *decades*). For each time word, we then analysed the 20 most frequent words in the distribution immediately following each time word, and then estimated the future/past bias of each time word by measuring the log probability with which *earlier* and *later* were likely to be encountered in this distribution (as expected, in all cases, *later* occurred at a higher frequency than *earlier* in this distribution). These biasing probabilities were then used to modify the raw probabilities of each time word in the distribution following each integer (i.e., the log frequency with which *seconds* follows *ten* was multiplied by the log of probability of a future or past bias to return two values reflecting the degree to which *earlier* and *later* might be expected to follow). The future/past bias of each integer was then calculated by subtracting the log likelihood of *earlier* (given the set of time words in the distribution following each integer) from the log likelihood of *later* (given the same words/distribution). This allowed the model to take into account the way that different time words tend to be used with more or less future bias (for example, *seconds* tends to be much more future biased than *years*).

Model 2 was a simplified variation on Model 1, in which the biasing effect of each number was simply modelled as the log summed frequency of the 20 words in the distribution succeeding that integer multiplied by the log of the percentage of the 20 words that were time words. We then tested the predictions of these models empirically using the methods described by Matlock et al. (2005), with the notable change that participants answered the ambiguous ‘move forward’ question *without* drawing.

Method

A total of 399 participants read fictive motion sentences that varied only on the number of trees specified by the subject noun phrase. In addition to the four FM sentences that appeared in Matlock et al. (2005): *Four pine trees run along the edge of the driveway*, *Eight pine trees run along the edge of the driveway*, *Twenty pine trees run along the edge of the driveway*, or *Over eighty pine trees run along the edge of the driveway*, we also examined the participants’ responses to *Ten*, *Eleven*, *Twelve*, *Nineteen* and *A hundred pine*

trees run along the edge of the driveway. The new numbers were added both to allow for a more fine-grained analysis of the scanning and expectation hypotheses, and to control for the possibility that priming might simply be an effect of word frequency. While *ten* is considerably more frequent than *eleven* (more than 10 times more frequent in COCA; Davies, 2009), both the scanning hypothesis and our expectation-based models predict that they should produce roughly the same degree of priming.

Each participant read one of the sentences, and then indicated whether the sentence made sense by checking ‘yes’ or ‘no’, and went on to answer the ‘move forward’ question, ‘*Next Wednesday’s meeting has been moved forward two days. What day is the meeting now that it has been re-scheduled?*’ Prior to the analysis, responses from seven participants were discarded: three because of an incorrect response (e.g., ‘Wednesday’), and four because the sentence was judged as not sensible.

Results and discussion

With regard to the numbers tested by Matlock et al. (2005), *four, eight, twenty* and *over eighty trees*, participants were overall more likely to respond ‘Friday’ (68%) than ‘Monday’ (32%). As the number of trees mentioned in the stimuli increased, the number of Friday responses decreased, specifically 83% (4 trees), 76% (8 trees), 64% (20 trees), and 48% (over 80 trees). Percentages were based on 29, 25, 28, and 29 participants, respectively. (Remaining responses were Monday.) A Chi-square test (linear-by-linear association) showed the effect was reliable, $\chi^2(1) = 8.58$, $p < .05$, and number of trees and Friday bias was negatively correlated ($r = -0.94$).

Regarding the new numbers of trees tested, *ten, eleven, twelve, and nineteen trees* all biased participants towards providing a Friday response: 70% (10 trees), 72% (11 trees), 71% (12 trees), and 66% (19 trees), respectively. These percentages were based on 53, 53, 42, and 44 participants, respectively. However ‘*a hundred trees*’, produced no bias, with 50% of participants responding Friday and 48% Monday (based on 96 responses, one of which was, ‘Wednesday’). A Chi-square test showed the biasing effect was reliable, $\chi^2(1) = 10.585$, $p < .03$, and number of trees and Friday bias were again negatively correlated, $r = -0.98$; analysing the data for the items tested in Matlock et al. and the new items together, $\chi^2(1) = 19.89$, $p < .03$, $r = -0.91$.

How do these results fit with the theoretical questions we raised above, and how do they compare with earlier results? First, these results overwhelmingly support the idea that reading fictive motion sentences can influence people’s understanding of time, even in the absence of drawing. In the current study, 68% of the participants who saw Matlock et al.’s prime sentences provided a Friday response (versus 61% in Matlock et al.). Friday

responses for most conditions were similar to those in Matlock et al. (2005): for 8 trees, 76 % (versus 80% in Matlock et al., 2005); for 20 trees, 64% (versus 61 %); and for over 80 trees, 48 % (versus 50 %). The results were quite different for 4 trees, however: 83% current (versus 55 %). With regard to the new materials tested here, participants were about equally biased to provide a Friday response for *ten* through *nineteen* (70 %, ten trees, 72 %, eleven trees, 71 %, twelve trees, 66 % nineteen trees, respectively) but showed no Friday bias for *a hundred trees* (50 %). This pattern of data rules out the possibility of a frequency effect (e.g., that ten would cause significantly more priming than eleven because of its much higher usage frequency, and that ten and a hundred would produce the same priming given their very similar frequencies).

Taken alongside the mixed pattern of data in Experiment 2 of Matlock et al. (2005), the strong priming effects demonstrated here suggest that fictive motion language is perfectly capable of influencing temporal thinking, even when the effects of drawing are controlled for. While these results are not altogether inconsistent with the notion of mental scanning, the overall negative correlation between number of scan points and amount of priming seems to run in the opposite direction of what might be expected if participants' temporal priming was the result of mentally simulating fictive motion, as one might reasonably expect more simulation to produce correspondingly more priming, not less. This, along with the lack of any detailed mechanistic account of mental scanning, suggests that something else might be responsible for the priming differences in the stimuli.

To examine whether lexical prediction might provide a plausible alternative explanation, we compared the responses provided by our participants given the various number primes with our models of the linguistic expectations the numbers could be expected to produce. As we expected, there was a good fit between the predicted priming and the degree of bias exhibited in the empirical data. The predictions of Model 1, in which we sought to account for the way that different time words might be expected to produce different degrees of future bias, correlated well with the pattern of data produced by our participants ($r = .76$, $t = 3.09$, $p < .01$). Perhaps surprisingly, however, the simpler model (2) performed even better ($r = .92$, $t = 5.46$, $p < .001$). Indeed, considered alone, the percentage of time words in the set of 20 words most likely to follow the numbers was a good predictor of our participants' bias ($r = .86$, $t = 4.46$, $p < .005$). Moreover, if *only* the proportion of time words amongst the 10 most frequent words in the distribution following each number word is considered, this correlation increases ($r = .96$, $t = 8.05$, $p < .0001$).

As we expected, the raw frequency of the numbers did not correlate significantly with bias in our participants' responses. An examination of the

partial correlations between the data, the frequency of the number words, and the proportion of time words at the top of the distribution following each number word, revealed a significant correlation between the distributional measure and the data when partialling for frequency ($r^2 = .879$, $t = 6.6$, $p < .001$), whereas the correlation between the data and frequency when the likelihood of encountering a time word in the distribution was partialled for was negligible ($r^2 = .091$, $t = 0.78$, $p > .45$).

With regard to magnitude, the results are less clear. An examination of the partial correlations between the data, the magnitude of the number words, and the proportion of time words at the top of the distribution following each number word, revealed a significant correlation between the distributional measure and the data when partialling for magnitude ($r^2 = .683$, $t = 3.6$, $p < .02$), whereas the correlation between the data and magnitude when the likelihood of encountering a time word in the distribution was partialled for was marginally significant ($r^2 = -.421$, $t = -2.9$, $p > .08$).¹

The results of these data are thus consistent with the suggestion that our participants' understanding of the ambiguous temporal question may be being shaped, at least in part, by their linguistic expectations. Moreover, as in many other paradigms in which people's sensitivity to the distributional properties of their languages have been explored, our data suggest that our participants are sensitive to the patterns of distribution for individual words. This is not to say that people do not make more general predictions (clearly people do have broader expectations). Rather, the lack of a correlation between simple frequency and priming in our data suggests that participant behaviour is sensitive to the fine-grained distributional properties of English. Further, as the *differences* in priming produced by individual number words attest, these prediction processes operate at a level far more detailed than simple, coarse-grained expectations such as 'integers predict time words, and these predict words that are future biased'.

To further explore these ideas, in Experiments 3 and 4 we sought to establish whether the patterns of priming for fictive motion language in the other experiments reported by Matlock et al. could be reproduced in the absence of drawing, and if so, whether these patterns would support our lexical prediction hypothesis.

¹ To get a measure of how well our corpus-based model of expectations captures real usage patterns, we correlated the log-transformed likelihood of a given time word succeeding a given number word in a model derived from COCA with the observed frequency with which each number and time word occurred in the frames 'Number Word + Time Word + *later*' versus 'Number Word + Time Word + *earlier*' in Google; i.e., we examined how well the predicted value of 'hours' given 'nineteen' in our model correlated with the summed number of hits returned by '*nineteen hours later*' and '*nineteen hours earlier*.' The fit between the predictive model and the observed data ($r = .95$, $p < .000001$) was encouraging.

EXPERIMENT 3

How will participants conceptualise time after reading a sentence that includes fictive motion? Will it differ from how they think about time following a sentence that does *not* include fictive motion?

Earlier studies have shown that without a priming context, participants are about equally likely to respond ‘Monday’ or ‘Friday’ to the ambiguous *move forward* question (see Boroditsky, 2000). In this experiment, we tried to bias participants in the FM condition towards a Friday response by having them read a series of sentences with the fictive motion word ‘runs’ (such as ‘the tattoo runs along his spine’). In the NFM condition, we used the phrase ‘next to’ in sentences such as ‘the highway is next to the coast’, which we did not expect to bias responses either way. Matlock et al. (2005; Experiment 1) found just this pattern of bias between the same FM and NFM conditions when drawing was included as part of the experiment.

Why might we expect to bias responses one way or the other in the absence of drawing? In light of our results from Experiment 2, we would expect that priming might result from linguistic expectations, with the distributional properties of the words in the priming sentences potentially affecting participants’ expectations and subsequent interpretations. To investigate this claim in Experiment 3, we decided to check the predicted priming effects against behavioural data. To first establish what linguistic priming effects would be predicted, we looked at time words that frequently follow ‘runs’ and ‘next to’ in COCA (Davies, 2009).

To model any potential differences precisely, we looked at the likelihood that either future or past-looking spatial and temporal words occurred after ‘runs’ and ‘next to’ respectively (e.g., *before*, *sooner*, versus *later*, *after*). As we noted above, the Zipfian distribution of words means that relatively small differences in rank order can represent huge differences in frequency. Accordingly, we examined just the hundred most frequent words to follow ‘runs’ and ‘next to’, which provided a good estimation of what someone exposed to a representative sample of English might expect given each of these words. Further, to prevent the few words that co-occurred with ‘runs’ and ‘next to’ with extremely high frequency from overwhelming lower frequency co-occurrences in our analysis (which might overestimate the effect of the high frequency co-occurring words on participants’ expectations), we analysed the logarithm of their frequencies. For each word, we added the logs for all future-looking words and all past-looking words and then compared them to assess the degree to which ‘runs’ and ‘next to’ were more likely to cue future- or past-looking words.

In the FM condition, ‘runs’ biases future-looking words (like *after* or *later*) over past-looking words (like *before* or *earlier*) by a log ratio of about 2:1. In contrast, in the NFM condition, ‘next to’ co-occurs with

future and past-looking words at virtually the same rate (log ratio of 1:1.09). Thus, if linguistic expectations have a part to play, seeing ‘runs’ should prime participants to expect words like *later*, moving them toward the ‘Friday’ response, while seeing ‘next to’ should have no discernable effect. In Experiment 3, we tested these predicted priming effects against participant data.

Method

A total of 144 Stanford University students volunteered for partial credit in an introductory psychology course. The instructions and stimuli appeared on a single page in a booklet that contained unrelated materials. Each participant volunteered in only one experiment. Participants were asked to, ‘Please read the sentence below. Does it make sense?’ Next they read either an FM sentence (e.g., ‘the tattoo runs along his spine’) or a NFM sentence (e.g., ‘the tattoo is next to this spine’), and indicated whether the sentence was sensible by circling ‘yes’ or ‘no’. Then they answered the *move forward* question: ‘Next Wednesday’s meeting has been moved forward two days. What day is the meeting now that it has been rescheduled?’ There were five sentences in each condition, as shown in Table 1. All were from the first experiment in Matlock et al. (2005). Responses from six participants were discarded: one because of an incorrect response, and five because the sentence was judged to be non-sensible.

Results and discussion

Consistent with our predictions, fictive motion language influenced how participants answered the ‘move forward’ question. Participants were more

TABLE 1
Experiment 1 stimuli

FM sentences
The bike path runs alongside the creek
The tattoo runs along his spine
The bookcase runs from the fireplace to the door
The highway runs along the coast
The county line runs along the river
NFM sentences
The bike path is next to the creek
The tattoo is next to his spine
The bookcase is between the fireplace and the door
The highway is next to the coast
The county line is next to the river

likely to provide a Friday response after reading fictive motion sentences than they were after reading non-fictive motion sentences, $\chi^2(3) = 4.75$, $p < .05$. Of the 77 participants who read a FM sentence, such as ‘the tattoo runs along his spine’, 32% went on to provide a Monday response and 68% went on to provide a Friday response. Of the 61 participants who read a NFM sentence, such as ‘the tattoo is next to his spine’, 51% gave a Monday response and 49% gave a Friday response.

The pattern of participants’ responses to the ‘move forward’ question replicates the findings of Matlock et al. (2005; Experiment 1) in the absence of drawing, and is consistent with the lexical prediction hypothesis. Our analysis of the materials employed in Matlock et al. (2005; Experiment 1) revealed that the FM verb ‘runs’ primes future-oriented words, and so could be expected to produce a Friday bias, while ‘next to’ did not have these distributional properties, and thus was neither expected to, nor did, produce this effect.

EXPERIMENT 4

In all the experiments discussed thus far, FM sentences evoked a strong ‘Friday’ response in participants. In Experiment 4 we wished to see whether the ‘direction’ of fictive motion could be used to differentially influence participants’ understanding of time. Matlock et al. (2005; Experiment 3) examined whether FM sentences could be used to prime both future *and* past responses by asking participants to read and depict an FM sentence about a road *going to* New York, or a FM sentence about a road *coming from* New York (Stanford was the implied starting point or end point because all participants completed the task there). The results suggested that FM sentences could indeed prime both responses. Experiment 4 examined whether the findings of Matlock et al. would replicate when participants simply read and comprehended the FM primes, without drawing them.

The verbs ‘comes’ and ‘goes’ in Experiment 4 also offered another opportunity to test the lexical prediction hypothesis. An analysis of the co-occurrence patterns of ‘comes’ and ‘goes’ in COCA reveals that the distribution of future- and past-looking time words (i.e., *before*, *sooner*, versus *later*, *after*) is opposite for the two verbs. ‘Goes’ is followed by future-looking words twice as often as past-looking words, whereas ‘comes’ shows precisely the opposite pattern. (Using the same methods elaborated in Experiment 3, the log ratio of future to past-looking words for ‘goes’ is 2.17:1; for ‘comes’, the trend is reversed, with the log ratio of past to future 1.83:1.) In other words, seeing the verb ‘comes’ ought to prime participants to expect words like *earlier* and *sooner*, while seeing the verb ‘goes’ should have the opposite priming effect. We thus expected ‘comes’ to move

participants toward a ‘Monday’ response and ‘goes’ to move participants toward a ‘Friday’ response.

Method

A total of 106 participants read one of two FM sentences, *The road goes all the way to New York* or *The road comes all the way from New York* before answering the ‘move forward’ question. After calculating the Monday and Friday responses in each condition, we discarded two participants’ incorrect responses (e.g., Tuesday).

Results and discussion

In Experiment 4, participants were more likely to provide a Friday response when primed with fictive motion going away from them versus fictive motion coming toward them. Of the 60 participants who read *goes to* FM sentences, 34% said Monday and 66% said Friday. Of the 44 participants who read *comes from* FM sentences, 61% said Monday and 39% said Friday, $\chi^2(1) = 5.57, p = .025$.

These results broadly replicate those of Matlock et al. (2005; Experiment 3) while controlling for the influence of drawing, and are consistent with other findings (Boroditsky, 2000; Boroditsky & Ramscar, 2002; see also Núñez, Motz, & Teuscher, 2006). Further, as our analysis of ‘comes’ and ‘goes’ reveals, the pattern of participant responses to the ambiguous ‘move forward’ question is once again consistent with the idea that our participants’ understanding may have been shaped by linguistic expectations raised by the phrasing of the prime sentences.

GENERAL DISCUSSION

Four experiments examined the influence of drawing and fictive motion language on people’s understanding of time. Each examined the way that motion (literal or metaphorical) influenced participants’ responses to an ambiguous ‘move forward’ question (see Boroditsky, 2000; McGlone & Harding, 1998).

In Experiment 1, we found that participants showed a bias towards a Friday response after responding to a request to copy a physical line; however, this bias was only evident as the lines became longer. Experiments 2, 3, and 4 then examined participants’ responses to the ambiguous ‘move forward’ question after they had read sentences involving fictive motion, a figurative way of using motion words to describe static arrangements of objects in space. Experiment 2 showed that using different number words to describe fictional trees ‘running’ along a driveway differentially primed

Friday bias. Further, this variance correlated well with the degree to which each integer led to the expectation of a time word (corpus analyses indicated that this would in turn lead to the expectation of future oriented words, such as *later*), while correlating only weakly with the magnitude of the integers, and hardly at all with their frequency. Experiment 3 then revealed that while the language used in fictive motion sentences biased participants towards Friday responses, non-fictive motion sentences expressing similar content did not. Finally, Experiment 4 revealed that the priming effect that fictive motion language has on people's temporal understanding was not limited to pushing them to future-looking interpretations of temporally ambiguous sentences. While participants reading about a road 'going to' New York did show a future bias, those reading about a road 'coming from' New York were biased in the opposite direction, towards an earlier (past) interpretation.

What do these results tell us about the relationship between actual, literal, and metaphorical motion? First, while the results of Experiment 1 are consistent with the idea that *thinking* about actual motion might affect metaphorical motion, they are less consistent with the idea that metaphorical motion understanding (in particular people's understanding of 'motion' in relation to time) is affected by movement alone. The act of simply drawing a line was not sufficient to bias our participants' temporal understanding (see also Boroditsky & Ramscar, 2002); rather, it seemed that a bias was only evident when the drawing task became sufficiently engaging. This is consistent with Boroditsky and Ramscar's (2002) suggestion that temporal understanding is subject to bias from people's *thinking* about motion, rather than their simply *being* in motion.

Similarly, the results of Experiments 2, 3, and 4 revealed that reading fictive motion sentences – which did not describe any actual motion – served to bias our participants' temporal understanding even in the absence of a drawing task. Taken together, these results allow us to begin to tease apart the relationship between actual, literal, and fictive motion. First, the findings of Experiment 1 (along with similar findings by Boroditsky & Ramscar, 2002) suggest that the experience of actually moving alone is insufficient to bias people's understanding of time. Given the popularity of the idea that human understanding, including people's understanding of words and language, is ultimately embodied (see for example Lakoff & Johnson, 1999; and the mental scanning proposals described above), it is interesting that embodied action appears to have no effect on people's understanding unless they are actively engaged in it. Although few embodiment proposals have been articulated with sufficient detail to evaluate whether our data is consistent with the overall idea of embodiment (or not), our findings do provide some interesting constraints on the ways in which ideas about embodiment and simulation are to be fleshed out.

Taken together with the lack of specification offered by alternative accounts, the pattern of data observed offers support to our proposal that linguistic expectation – thinking that is shaped by the demands of speaking (Slobin, 1996) – has a direct impact on the way people form their understanding of the world. On this view, the actual words and the actual patterns of words that people use to describe both literal and imaginative motion do not merely *describe* or point to some underlying conceptual structure, but rather, *are integral to it*.

The idea that thinking for speaking might have an impact on conceptual structure was first proposed as part of the debate over linguistic relativity (Slobin, 1996), and is easily illustrated in this context: If the meaning of words is derived – even in part – from their patterns of use, significant conceptual differences ought to be found between languages with highly distinct co-occurrence patterns. For example, Boroditsky, Phillips, and Schmidt (2001) asked a group of native Spanish speakers and a group of native German speakers to describe what came to mind when they were shown a picture of a bridge. The German speakers chose words like ‘beautiful’, ‘elegant’, and ‘fragile’, while the Spanish speakers chose words like ‘strong’, ‘sturdy’, and ‘towering’. The effect was observed repeatedly for 24 different objects with opposite grammatical genders in either language. In each case, both sets of speakers were more likely to generate ‘masculine’ adjectives for masculine gendered items and ‘feminine’ adjectives for feminine gendered items.² What causes these differences in Spanish and German speakers? Boroditsky et al. suggest that native speakers ‘imbue’ nouns with the characteristics of their grammatical gender – feminine in the case of German (*die brücke*) and masculine in the case of Spanish (*el puente*). Just how they do this, however, remains open to question.

One way of accounting for the data described by Boroditsky et al. is to assume that the set of associations each noun carries varies by language. If associations (or expectations; Rescorla & Wagner, 1972) covary with the ways in which masculine and feminine markers are used (and vice versa), we would expect disparate usage patterns to lead to the development of dissimilar associations (and thus dissimilar connotations and meanings). In this way, culture informs language, and language culture.

At one level, this suggestion is obvious (see Everett, 2005). For example, it seems evident that in a language in which the word *industrial* was often

² The masculinity or femininity of the adjectives was rated by a set of independent judges. Unsurprisingly, corpora seem to validate many of the intuitions (and suspicions) we have about how we associate adjectives by gender. For example, a search of COCA for adjectives preceding *man* came up with words like ‘big’, ‘rich’, ‘handsome’, ‘wise’, ‘leading’, ‘wild’, and ‘brave’, while a search for adjectives preceding *woman* came up with words like ‘beautiful’, ‘pregnant’, ‘pretty’, ‘single’, ‘intelligent’, and ‘slender’.

spoken of in glowing terms, this would be reflected in common daily usage (e.g., in association with words like ‘major’, ‘advanced’, ‘modern’, ‘leading’, ‘national’, and ‘great’). This would be similarly reflected in the co-occurrence patterns of a language in which *industrial* was often spoken of disparagingly (e.g., in association with words like ‘pollution’, ‘emissions’, ‘chemicals’, ‘espionage’, and ‘waste’). One view of these patterns of co-occurrence is that they are derivative, the product of some common underlying ‘concept’. We suggest something different. On the view put forward here, co-occurrences (*usage patterns*) are not merely derivative; rather, they are part of a cultural artifact – the language of a community – which shapes peoples’ linguistic expectations. As such, co-occurrences play a part in determining an individual’s understanding of *what words mean*. Thus, it is likely that the co-occurrence patterns for ‘woman’ in a culture with strict gender roles are decidedly different from those in a more liberal society. Both individually and collectively, the patterns of co-occurrence of ‘woman’ with other words, and the expectations these usage patterns produce, are an integral aspect of individual’s and societies’ concepts of *woman* itself.

While this view is a departure from most standard approaches to the relationship between words and their meanings, it satisfies many of the peculiar constraints that apply to theories of mental representation. There are many reasons to believe that human knowledge cannot be easily or neatly bounded into ‘domains’ that words can simply refer to (Lakoff & Johnson, 1999; Murphy, 2002; Wittgenstein, 1953). This in turn presents a massive obstacle to simple, unitary characterisations of the nature of human understanding. While at first blush it might seem that more abstract domains simply borrow conceptual structure from more ‘literal’ experiential domains, the reality is almost certainly more complex. It is likely that even the structure of experiential domains derives in part from the structure of language. As Wittgenstein (1953) famously pointed out, seemingly straightforward experiential concepts, such as the kinds of things we call ‘games’, do not exist independently of language and culture. Rather, it appears that the structure of ‘game’ comes, at least in part, from the way that the word *game* is used (Wittgenstein, 1953).

Our proposal, that patterns of usage are part of the structure of ‘concepts’, can be seen as fleshing out at least one mechanism through which *use* may influence meaning (Wittgenstein, 1953). From this point of view, literal motion is not a ‘concept’ from which metaphorical motion borrows structure, and the common language used to describe literal and metaphorical motion does not merely reflect those borrowings. Rather, peoples’ understandings of literal and metaphorical motion are the product of shared cognitive processes that generate an ‘understanding’ when making a plan, performing an action, or answering an ambiguous question in an experiment (for example). The shared patterns of co-occurrence in talk about

literal and metaphorical motion are not mere artifacts of some underlying conceptual structure, but rather make up a critical *part* of that structure: one that may influence the way that someone makes a plan, performs an action, or answers an ambiguous question in an experiment.

If we accept that even our understanding of experiential concepts has a linguistic component, and that the particular ways in which people use language may impact their conceptual understanding, it follows that *both* literal and abstract ideas will be shaped to some degree by the way that languages are used. From this perspective, language can be seen as more than a referential code; rather, language offers a medium for structuring, encoding, and transmitting cultural knowledge (Tomasello, 1999).

This view of conceptual structure is consistent with the Wittgensteinian conception of knowledge and language ‘as an ancient city: a maze of little streets and squares, of old and new houses, and of houses with additions from various periods; and this surrounded by a multitude of new boroughs with straight regular streets and uniform houses’ (Wittgenstein, 1953). Wittgenstein’s broader suggestion is that language itself might not be so neatly bounded; that there is no simple process that corresponds to ‘language’, but rather that communication is a ‘form of life’, involving a variety of related practices. While precisely what this larger suggestion entails is something that remains to be determined, our findings here suggest that this more radical view of the nature of human communication – and in particular of the relationship between ‘language’ and ‘other’ aspects of culture and cognition – may be worthy of further consideration.

Manuscript received March 2007

Revised manuscript accepted September 2009

First published online Month/year

REFERENCES

- Altmann, G. T. M., & Steedman, M. (1988). Interaction with context during human sentence processing. *Cognition*, 30, 191–238.
- Altmann, G. T. M., & Mirković, J. (2009). Incrementality and prediction in human sentence processing. *Cognitive Science*, 33, 1–27.
- Baayen, R. H. (2001). *Word frequency distributions*. Dordrecht, the Netherlands: Kluwer Academic.
- Boroditsky, L. (2000). Metaphoric structuring: Understanding time through spatial metaphors. *Cognition*, 75, 1–28.
- Boroditsky, L. (2001). Does language shape thought? English and Mandarin speakers’ conceptions of time. *Cognitive Psychology*, 43(1), 1–22.
- Boroditsky, L., Phillips, W., & Schmidt, L. A. (2001). Sex and similarity: How do quirks of language shape the way we think? *Proceedings of the 42nd Annual Meeting of the Psychonomic Society*, Orlando, FL.
- Boroditsky, L., & Ramscar, M. (2002). The roles of body and mind in abstract thought. *Psychological Science*, 13, 185–188.

- Clark, H. H. (1973). Space, time, semantics, and the child. In T. E. Moore (Ed.), *Cognitive development and the acquisition of language* (pp. 26–63). San Diego: Academic Press.
- Davies, M. (2009). The 385+ million word corpus of contemporary American English (1990–2008): Design, architecture and linguistic insights. *International Journal of Corpus Linguistics*, 14, 159–190.
- Emanation, M. (1992). Chagga ‘Come’ and ‘Go’: Metaphor and the development of tense-aspect. *Studies in Language*, 16, 1–33.
- Evans, V. (2004). *The structure of time: Language, meaning and temporal cognition*. Philadelphia, PA: John Benjamins.
- Everett, D. (2005). Cultural constraints on grammar and cognition in Pirahã: Another look at the design features of human language. *Current Anthropology*, 46, 621–646.
- Gentner, D. (2001). Spatial metaphors in temporal reasoning. In M. Gattis (Ed.), *Spatial schemas and abstract thought* (pp. 203–222). Cambridge, MA: MIT Press.
- Gibbs, R. W. (1994). *The poetics of mind: figurative thought, language, and understanding*. New York: Cambridge University Press.
- Kutas, M., & Federmeier, K. D. (2007). Event-related brain potential (ERP) studies of sentence processing. In G. Gaskell (Ed.), *Oxford handbook of psycholinguistics* (pp. 385–406). Oxford, UK: Oxford University Press.
- Lakoff, G. (1987). *Women, fire, and dangerous things: What categories reveal about the mind*. Chicago, IL: University of Chicago Press.
- Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. Chicago: University of Chicago Press.
- Lakoff, G., & Johnson, M. (1999). *Philosophy in the flesh: the embodied mind and its challenge to western thought*. New York: Basic Books.
- Langacker, R. W. (1987). *Foundations of cognitive grammar, Vol. 1: Theoretical prerequisites*. Stanford, CA: Stanford University Press.
- Langacker, R. W. (1999). *Grammar and conceptualization*. Berlin: Mouton de Gruyter.
- Matlock, T. (2004a). The conceptual motivation of fictive motion. In G. Radden & K. U. Panther (Eds.), *Studies in linguistic motivation* (pp. 221–248). Berlin: Mouton de Gruyter.
- Matlock, T. (2004b). Fictive motion as cognitive simulation. *Memory and Cognition*, 32, 1389–1400.
- Matlock, T., Ramscar, M., & Boroditsky, L. (2003). The experiential basis of meaning. In *Proceedings of the Twenty-fifth Annual Conference of the Cognitive Science Society* (pp. 792–797). Mahwah, NJ: Lawrence Erlbaum.
- Matlock, T., Ramscar, M., & Boroditsky, L. (2005). On the experiential link between spatial and temporal language. *Cognitive Science*, 29, 655–664.
- Matlock, T., & Richardson, D. C. (2004). Do eye movements go with fictive motion? *Proceedings of the Twenty-Sixth Annual Meeting of the Cognitive Science Society* (pp. 909–914). Mahwah, NJ: Lawrence Erlbaum.
- Matsumoto, Y. (1996). Subjective motion and English and Japanese verbs. *Cognitive Linguistics*, 7, 183–226.
- McGlone, M. S., & Harding, J. L. (1998). Back (or forward?) to the future: The role of perspective in temporal language comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24, 1211–1223.
- McTaggart, J. (1908). The unreality of time. *Mind*, 17, 457–474.
- Moore, K. E. (2004). Ego-based and field-based frames of reference in space to time metaphors. In M. Achard & S. Kemmer (Eds.), *Language, culture and mind* (pp. 151–166). Stanford, CA: CSLI Publications.
- Murphy, G. L. (2002). *The big book of concepts*. Cambridge, MA: MIT Press.
- Núñez, R., Motz, B., & Teuscher, U. (2006). Time after time: The psychological reality of the ego- and time-reference-point distinction in metaphorical construals of time. *Metaphor and Symbol*, 21, 133–146.

- Núñez, R., & Sweetser, E. (2006). With the future behind them: Convergent evidence from Aymara language and gesture in the crosslinguistic comparison of spatial construals of time. *Cognitive Science*, 30, 401–450.
- Otten, M., & Van Berkum, J. J. A. (2007). What makes a discourse constraining? Comparing the effects of discourse message and scenario fit on the discourse-dependent N400 effect. *Brain Research*, 1153, 166–177.
- Otten, M., & Van Berkum, J. J. A. (2008). Discourse-based word anticipation during language processing: Prediction of priming? *Discourse Processes*, 45, 464–496.
- Radden, G. (1996). Motion metaphorized: The case of ‘coming’ and ‘going’. In E. Casad (Ed.), *Cognitive linguistics in the redwoods: The expansion of a new paradigm in linguistics* (pp. 423–458). Berlin: Mouton de Gruyter.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning: Vol. II. Current research and theory* (pp. 64–99). New York: Appleton-Century-Crofts.
- Richardson, D. C., & Matlock, T. (2007). The integration of figurative language and static depictions: An eye movement study of fictive motion. *Cognition*, 102, 129–138.
- Slobin, D. I. (1996). From ‘thought and language’ to ‘thinking for speaking’. In J. J. Gumperz & S. C. Levinson (Eds.), *Rethinking linguistic relativity* (pp. 70–96). Cambridge, UK: Cambridge University Press.
- Talmy, L. (2000). *Toward a cognitive semantics. Vol. 1: Concept structuring systems*. Cambridge, MA: MIT Press.
- Tanenhaus, M. K. & Brown-Schmidt, S. (2008). Language processing in the natural world. In B. C. M. Moore, L. K. Tyler, & W. D. Marslen-Wilson, (Eds.), *The perception of speech: from sound to meaning. Philosophical Transactions of the Royal Society B: Biological Sciences*, 363, 1105–1122.
- Tanenhaus, M. K., Spivey-Knowlton, M. J., Eberhard, K. M., & Sedivy, J. E. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science*, 268, 1632–1634.
- Tannen, D. (2007). *Talking voices: repetition, dialogue, and imagery in conversational discourse*. Cambridge, UK: Cambridge University Press.
- Tomasello, M. (1999). *The cultural origins of human cognition*. Cambridge, MA: Harvard University Press.
- Wicha, N. Y., Bates, E. A., Moreno, E. M., & Kutas, M. (2003). Potato not pope: Human brain potentials to gender expectation and agreement in Spanish spoken sentences. *Neuroscience Letters*, 346, 165–168.
- Wittgenstein, L. (1953). *Philosophical investigations*. Oxford, UK: Blackwell.
- Zipf, G. K. (1935). *The psychobiology of language*. New York: Houghton-Mifflin.
- Zipf, G. K. (1949). *Human behavior and the principle of least-effort*. Cambridge, MA: Addison-Wesley.