Casasanto, D. (2017). The hierarchical structure of mental metaphors. In B. Hampe (Ed.), *Metaphor: Embodied cognition and discourse* (pp. 46-61). Cambridge: Cambridge University Press.

3 The Hierarchical Structure of Mental Metaphors

Daniel Casasanto*

Chapter Preview

People think about abstract domains like TIME and GOODNESS metaphorically. This tendency may be universal. Yet the particular mental metaphors that people use can differ dramatically between individuals and groups, and may change from one moment to the next. Where do our mental metaphors come from, and how can they change so quickly in response to new patterns of experience? If mental metaphors are grounded in universal patterns of body–world interaction, how can different people have contrasting (and sometimes contradictory) mental metaphors? Hierarchical Mental Metaphors Theory (HMMT) provides an account of: (a) how mental metaphors are formed and changed, (b) how they can be fundamental to our understanding of abstract domains, yet at the same time remarkably flexible, and (c) how distinctive language-specific, culture-specific, and body-specific mental metaphors can arise from universal patterns of interaction with the physical and social world.

Introduction

According to theories of metaphorical mental representation, metaphors in language are more than just ways of talking, they are clues to a pervasive way of thinking (Lakoff & Johnson 1980). On this view, when people use expressions like a "*long* vacation," a "*high* price," or a "*close* resemblance," they are using mental representations of space (i.e. 'length,' 'height,' 'proximity') to scaffold mental representations in non-spatial conceptual domains (i.e. 'time,' 'value,' 'similarity'). Although initial evidence for metaphor theory was based on descriptive analyses of how people talk, there is now abundant experimental evidence that people also *think* metaphorically – even when they are not using any metaphorical language (or using language, at all)

1

(for a review, see Casasanto & Bottini 2014a). That is, people often think in "mental metaphors" (Casasanto 2008a, b): point-to-point mappings between nonlinguistic representations in a "source domain" (e.g. SPACE) and a "target domain" (e.g. TIME) that is typically more abstract (i.e. hard to perceive) or abstruse (i.e. hard to understand; Lakoff & Johnson 1980), which support inferences in the target domain.

The term "mental metaphor" is used contrastively with "linguistic metaphor" here, the former designating a mapping between non-linguistic mental representations, and the latter between linguistic representations.¹ Distinguishing mental metaphors from linguistic metaphors becomes particularly important in contexts like the present chapter where I raise questions like: "Do people who use different linguistic metaphors think in correspondingly different mental metaphors?" and "Do people sometimes think in mental metaphors that are absent from language?"

Where do our mental metaphors come from? Does everyone use the same mental metaphors, at least when they are thinking about universal experiences like the passage of time? Once established, do our basic mental metaphors ever change? Some of the answers to these questions offered by metaphor theory's founders (Lakoff 1993; Lakoff & Johnson 1999) are at odds with a growing body of experimental results. In this chapter I first illustrate the tension between some core tenets of Conceptual Metaphor Theory (Lakoff & Johnson 1999) and experimental tests of metaphorical thinking. I then describe a proposed resolution, Hierarchical Mental Metaphors Theory (Casasanto 2008b; Casasanto & Bottini 2014a, b), and review three sets of studies that serve as testbeds for this proposal.

2 Origin and Universality of Mental Metaphors: Puzzles and Paradoxes

According to Lakoff and Johnson (1999), our most basic mental metaphors are universal because we learn them from universal correlations between source and target domains in the natural environment:

We acquire a large system of primary metaphors automatically and unconsciously simply by functioning in the most ordinary of ways in the everyday world from our earliest years. We have no choice in this [...]

When the embodied experiences in the world are universal, then the corresponding primary metaphors are universally acquired [...] Universal conceptual metaphors are learned; they are universals that are not innate. (ibid.: 47, 56-57)

Once acquired, these basic mental metaphors are said to constitute "fixed conceptual mappings" (ibid.: 149), implemented in "permanent neural

connections ... across the neural networks that define conceptual domains" (ibid.: 46).

To summarize these points, our basic mental metaphors are proposed to be: (a) *learned* early in life on the basis of source–target correlations in the world; (b) *universal*, so long as the source–target correlations in the world are universal, and (c) *fixed*, by virtue of their implementation in permanent neural connections.

Yet the claims that basic mental metaphors are *learned*, *universal*, and *fixed* are all challenged by experimental data that have accumulated over the past decade. The claim that primary metaphors are learned early in life on the basis of observed source–target correlations was called into question by the most direct test to date. De Hevia and colleagues (2014) showed that neonates between 0 and 3 days old are already sensitive to relationships between spatial, temporal, and numerical magnitudes that are encoded in linguistic expressions like "a *long* time" and "a *large* number." At 0 to 3 days, these infants presumably had no understanding of linguistic and cultural conventions linking these domains, and had very little experience with correlations between them in the natural world. The data presented by de Hevia and colleagues thus suggest that cross-domain mappings between space, time, and number may be innate, not learned.

The claim that mental metaphors are universal is challenged by numerous experiments demonstrating crosslinguistic, cross-cultural, and cross-individual variation in spatial mappings for basic domains of experience. For example, some languages talk about musical pitch in terms of one-dimensional spatial height, whereas other languages metaphorize pitch in terms of multidimensional spatial thickness; speakers' mental metaphors for pitch differ accordingly (Dolscheid et al. 2013). Some cultures depict temporal sequences as unfolding rightward across calendars, graphs, and written timelines, whereas other cultures depict them as unfolding leftward; people's mental timelines follow the directions of these culture-specific space-time mappings (Casasanto & Bottini 2014b; Fuhrman & Boroditsky 2010; Ouellet et al. 2010; Tversky et al. 1991). Right-handers implicitly associate positive ideas and emotions with the right side of space and negative ideas with the left, but left-handers show the opposite associations between space and emotional valence (Casasanto 2009a). These studies showing that mental metaphors can be language-specific, culture-specific, or body-specific also challenge the claim that cross-domain mappings in our minds are *fixed*; on the contrary, mental metaphors that are deeply entrenched and highly automatic can be changed - even reversed - after as little as 5 minutes of exposure to different cross-domain relationships. How can mental metaphors be grounded in universals of experience if they vary across people? How can they be fundamental to our conceptualizations of target domains if they can change in a matter of minutes?

3 A Proposed Solution: Hierarchical Mental Metaphors Theory

A solution to these apparent paradoxes emerges if we consider that even our most basic mental metaphors are constructed over multiple timescales, on the basis of multiple kinds of experience. According to Hierarchical Mental Metaphors Theory (HMMT; Casasanto & Bottini 2014a,b; see also Casasanto 2008b, 2014b; Dolscheid et al. 2013), the cross-domain mappings that people use at any moment are members of a superordinate family of mappings. The superordinate family is typically constructed on the basis of source-target relationships in the natural world. These relationships could be learned from early experiences with source and target domains, as Lakoff and Johnson (1999) suggest, or they could be part of infants' innate "core knowledge" (Srinivasan & Carey 2010). Cross-domain relationships like MORE TIME IS MORE DISTANCE are survival-relevant and could plausibly become encoded in the human genome. Whether learned or innate, each superordinate family of mental metaphors constitutes a set of mappings that can be used for scaffolding target-domain thinking and can be encoded in linguistic and cultural conventions. For a given target domain (e.g. PITCH), children acquire (or manifest) a superordinate family of source-domain mappings early in life. To the extent that source-target relationships in the natural world are found universally, superordinate families of mental metaphors should be universal.

Once learners are exposed to relevant conventions in language and culture, or to regularities in the way they use their particular bodies, a second process begins which can continue throughout the lifetime and gives rise to the diversity of mental metaphors found across individuals and groups. One of the mappings from a superordinate family (or a subset of the mappings) becomes strengthened through a process of competitive associative learning. For example, each time someone uses a linguistic metaphor, the corresponding non-linguistic source– target mapping is activated. Activating a mapping *strengthens* this source–target mappings in the same family, also *weakens* the competing source–target mappings in the same family, as a consequence. The process of strengthening the frequently activated mappings within a family, and of weakening their "sibling" mappings, follows naturally from the dynamics of long-term memory networks for families of associations (e.g. Anderson et al. 2000).

The process of strengthening and weakening specific mappings within superordinate families can account for several otherwise mysterious properties of mental metaphors. First, this hierarchical model can explain how mental metaphors can be grounded in universal source–target relationships in the natural world and yet be variable across people. Even if superordinate families are universal, the specific mappings that get used most frequently or automatically can vary across individuals and groups. On this view, there is no single answer to the question, "Are our mental metaphors universal?"

Second, HMMT can explain how mappings can change rapidly in response to new patterns of experience. In the examples that will be reviewed below, participants are implicitly using "new" mappings that differ from– and in some cases directly contradict – the mappings they ordinarily use, after only brief experimental interventions. This rapid change is possible because the "new" mappings introduced during the experiment are not really new; rather, they are members of the same superordinate family as the mappings participants normally use and can become strengthened through repeated use to the point that they are (at least temporarily) stronger than the mappings that participants ordinarily use.²

Here I will focus on the three sets of mental metaphors mentioned above, whose use is conditioned by different streams of physical and social experience. Spatial representations of a particular dimensionality or directionality serve as the metaphorical source domain that structures people's mental representations in the target domains of MUSICAL PITCH, TIME, and EMO-TIONAL VALENCE. These target domains are more "abstract" than the domain of space insomuch as they are impossible to see or touch; in the cases of time and valence, they are impossible to experience through any of the five senses. That is, we can see the spatial length of a rope or the height of a ladder, but we can never *see* the length of a vacation (i.e. its duration) or the height of a musical note (i.e. its auditory frequency).³

Spatializing these non-spatial domains in our minds may make our experiences of pitch, time, and valence easier to imagine, compare, or remember. It may be a human universal to conceptualize these domains in terms of space (cf. Eitan & Timmers 2010; Whorf 1956), but the particulars of these spatial representations vary across groups of people, according to the particulars of their linguistic, cultural, or bodily experiences. The mechanism underlying all of these effects of experience, I will suggest, is the same: Habitual experiences cause a certain mental metaphor to be activated frequently, strengthening this source–target association in memory, at the expense of competing associations within the same family of mappings.

4

Spatial Representations of Musical Pitch: Universals and Language-Specificity

In many languages, pitch is metaphorized in terms of vertical space: Highfrequency pitches are "high" and low-frequency pitches "low." But this is not the only possible spatial metaphor for pitch. In other languages like Farsi, Turkish, and Zapotec, high-frequency pitches are "thin" and low-frequency pitches are "thick" (Shayan et al. 2011). Beyond talking about pitch using spatial words, do people think about pitch using spatial representations? Several studies suggest that speakers of "height languages" like English

activate vertical space–pitch mappings when judging pitches (e.g. Pratt 1930; Roffler & Butler 1968). In one set of experiments, Dolscheid et al. (2013) investigated (a) whether people still activate space–pitch associations even when they are not using language, and (b) whether speakers of "height languages" and "thickness languages" tend to use the same non-linguistic space–pitch associations, or whether their mental metaphors for pitch are shaped by their experience of using linguistic metaphors.

Like English, Dutch describes pitches as hoog ('high') or lag ('low'), but in Farsi, high pitches are 'thin' (nāzok) and low pitches are 'thick' (koloft), as already noted. Dolscheid et al. (2013) tested Dutch and Farsi speakers on a pair of non-linguistic pitch reproduction tasks in which participants were asked to reproduce the pitches of tones that they heard in the presence of irrelevant spatial information: either lines that varied in their height (height interference task) or their thickness (thickness interference task). Dutch speakers' pitch estimates were strongly affected by irrelevant spatial height information. On average, a given tone was sung back higher when it had been accompanied by a line that was high on the computer screen, and lower when it had been accompanied by a line that appeared low on the screen. By contrast, lines of various thicknesses had no measurable effect on Dutch participants' pitch estimates. Farsi speakers showed the opposite pattern of results. Lines of varying heights had no measurable effect on Farsi speakers' pitch estimates, but tones accompanied by thin lines were sung back higher, and tones accompanied by thick lines were sung back lower.

4.1 Differences in Nonlinguistic Pitch Representations Not Due to Verbal Labeling during the Task

The pattern of spatial interference on people's pitch reproduction performance reflected the space–pitch metaphors in their native languages: Dutch speakers could not help incorporating irrelevant height information into their mental representations of pitch (but could ignore thickness), whereas Farsi speakers could not help incorporating irrelevant thickness information into their mental representations of pitch (but could ignore height). This pattern cannot be explained by differences in overall accuracy of pitch reproduction, or in differences in musical training between groups.

Importantly, this pattern also cannot be explained by participants using language covertly during the task: labeling the pitches they needed to reproduce as "high/low" or "thick/thin." This explanation was ruled out by the experimental design, in which nine different pitches were paired with each of nine different spatial heights or thicknesses. This "crossing" of all of the levels of pitch and of space meant that variation in each domain was orthogonal to variation in the other: There was no correlation between space and pitch in the

stimuli. As such, covertly labeling high pitches as "high" (or "thin") and labeling low pitches as "low" (or "thick") could not produce the observed effects of space on pitch reproduction; on the contrary, labeling pitches using the spatial metaphors in one's native language during the task could only work *against* the effects we predicted and found.

Rather than an effect of using language "online" during the task, these experiments show an effect of people's previous experience using either one linguistic metaphor or the other, and thereby strengthening either one mental metaphor or the other (i.e. strengthening a non-linguistic HEIGHT-PITCH or THICKNESS-PITCH mapping in memory). To confirm that the observed effects did not depend on participants covertly labeling pitches during the task, Dolscheid et al. (2013) repeated the height interference task in Dutch speakers with the addition of a concurrent verbal suppression task. On each trial of the task, participants had to rehearse a novel string of digits while perceiving and reproducing the pitches. Secondary tasks like this have been used across many experiments to prevent participants from labeling the stimuli (e.g. Winawer et al. 2007). As predicted, verbal suppression had no effect on the results of the pitch reproduction task. Dutch speakers still showed strong height-pitch interference, consistent with an "offline" effect of participants' previous experience using language on their subsequent non-linguistic pitch representations (see also Casasanto 2008b).

4.2 Does Using Different Linguistic Metaphors Cause People to Use Different Mental Metaphors?

The results reviewed so far show a correlation between people's linguistic metaphors and their non-linguistic mental metaphors, but they do not provide any evidence that language *causes* Dutch and Farsi speakers to mentally represent pitch differently. Dolscheid et al. (2013) reasoned that if using THICKNESS-PITCH metaphors in language is what causes Farsi speakers to activate THICKNESS-PITCH mappings implicitly when reproducing pitches, then exposing Dutch speakers to similar THICKNESS-PITCH metaphors in language should cause them to reproduce pitches like Farsi speakers. A new sample of Dutch speakers were recruited and assigned to one of two training conditions: Participants in the "thickness training" group learned to describe pitches using Farsi-like metaphors (e.g. "a tuba sounds thicker than a flute"), whereas the other half in the "height training" group (i.e. the control group) described pitches using standard Dutch metaphors (e.g. "a tuba sounds lower than a flute"). After about 20 minutes of this linguistic training, participants in both groups performed the non-linguistic thickness interference task described above. Whereas "height-trained" participants showed no effect of irrelevant thickness information on their pitch estimates, "thickness-trained"

participants showed a thickness interference effect that was statistically indistinguishable from the effect found in native Farsi speakers. Even a brief (but concentrated) "dose" of thickness metaphors in language was sufficient to influence Dutch speakers' mental metaphors, demonstrating that linguistic experience can cause the differences in non-linguistic pitch representations found across natural language groups.

Notably, Dolscheid et al. (2013) also included a training condition in which Dutch-speaking participants were taught to use "thick" and "thin" to describe pitches in a way that contradicts both the linguistic mappings in languages like Farsi and the relationships between thickness and pitch in the natural world (e.g. they learned to use expressions like "a tuba sounds *thinner* than a flute"). This training intervention was identical to the "Farsi-training" condition described above in every way except for the pairing of high/low with thick/ thin. Although participants learned to use the "reverse-Farsi" mapping with high accuracy (95%), this linguistic training had no effect on their non-linguistic mental representations of pitch.

The contrast between the Farsi-like and reverse-Farsi training supports a prediction of HMMT: People should be able to adopt mappings that are included in a superordinate family of mappings (e.g. the SPACE–PITCH mappings that are evident in the natural world) more easily than they can adopt mappings that are not included in the superordinate family (i.e. mappings that run contrary to, or orthogonal to, the source–target mappings found in the natural world).

4.3 When Does Language Shape Space–Pitch Mappings?

The results reviewed up to this point leave open the question: Do space-pitch metaphors in language cause people to develop the corresponding nonlinguistic space-pitch mappings, or does using linguistic metaphors change how likely people are to use a preexisting mental metaphor? To evaluate these possibilities, Dolscheid et al. (2014) tested 4-month-old infants on a pair of space-pitch congruity tasks. Infants heard pitches alternately rising and falling while they saw a ball rising and falling on a screen (height congruity task) or a cylinder growing thicker and thinner (thickness congruity task). For half of the trials, changes in pitch and space were congruent with the HEIGHT-PITCH and THICKNESS-PITCH mappings encoded in Dutch and Farsi, respectively, and for the other half of the trials they were incongruent with these SPACE-PITCH mappings. The data showed that infants looked longer at congruent SPACE-PITCH displays than at incongruent displays. This was true both in the height congruity condition (consistent with an earlier experiment by Walker et al. 2010) and in the thickness congruity condition. There was no difference in the magnitude of the congruity effect between conditions,

suggesting that there was no difference in the strength of the HEIGHT–PITCH and THICKNESS–PITCH mappings in the infants' minds.

Four-month-olds are completely unable to produce SPACE–PITCH metaphors in language and are also, presumably, unable to understand them. Yet, they are sensitive to two of the SPACE–PITCH metaphors that are found in languages like Dutch and Farsi, and in their speakers' non-linguistic pitch representations. These results suggest that people who use different linguistic metaphors for pitch come to think about pitch differently, not because language instills in them one spatial mapping instead of the other, but rather because language strengthens one of their preexisting space–pitch mappings, at the expense of the other.

4.4 Hierarchical Construction of Spatial Metaphors for Pitch

How could infants who are sensitive to both height-pitch and thickness-pitch mappings turn into adults who appear to activate only one of these mappings when they represent pitch? This process can be understood in terms of HMMT. First a superordinate "family" of mappings is established, which in the case of space and pitch includes both the height-pitch and thickness-pitch mappings. These mappings may be constructed, over either ontogenetic or phylogenetic time, on the basis of observable correlations between space and pitch in the natural world. The HEIGHT-PITCH mapping reflects the fact that people involuntarily raise their larynxes, chins, and sometimes other body parts (e.g. their eyebrows) when they produce higher pitches, and lower them when they produce lower pitches. It also reflects a statistical tendency for higher pitches to originate from higher locations, and lower pitches from lower locations (Parise et al. 2014). The thickness-pitch mapping reflects a pervasive correlation between pitches and the size of the objects or creatures that produce them: Consider the different pitches produced by, for example, strumming thin vs. thick strings on a guitar; banging on a large steel oil drum vs. a small steel can; barking by a small dog vs. a big dog. Although Dolscheid et al.'s (2014) data confirm that both the HEIGHT-PITCH and THICKNESS-PITCH mappings are present in infants' minds, they leave open the question of exactly how and when these mappings become established initially.

Whatever the ultimate origin of space–pitch mappings in pre-linguistic children may be, when children learn metaphors in language, a second process begins. The findings in adults reported by Dolscheid and colleagues (2013) suggest that each time people use a linguistic metaphor like "a *high* pitch" they activate the corresponding mental metaphor, strengthening this mapping at the expense of competing mappings in the same family of space–pitch associations. As a consequence, speakers of height languages like Dutch and English come to rely on vertical spatial schemas to scaffold their pitch representations

more strongly than on multidimensional spatial schemas, whereas speakers of thickness languages like Farsi come to rely on multidimensional spatial schemas, more strongly than vertical spatial schemas.

According to HMMT, the process of strengthening certain mental metaphors via the use of the corresponding linguistic metaphors results in the *weakening* of other members of the family of mappings – but this does not cause these dispreferred mappings to be extinguished. This aspect of the theory may explain the representational flexibility demonstrated in the training experiment by Dolscheid and colleagues (2013). Dutch speakers could be induced to use a non-linguistic thickness–pitch mapping (like Farsi speakers) after only a brief training intervention because no spatial mappings had to be created or destroyed; rather, the new pattern of language experience boosted the strength of the thickness–pitch mapping that had presumably been present in the Dutch speakers' minds since infancy, causing them to think about pitch in a way that was not new, just rarely used.

5 Spatial Representations of Temporal Sequences: Universals and Culture-Specificity

Spatial metaphors for time are common across languages (Alverson 1994). In English, time appears to flow along a sagittal (front-back) axis: the future is "ahead" and the past is "behind." No known spoken language uses the lateral (left-right) axis to talk about time conventionally: Monday comes before Tuesday, not to the left of Tuesday (Cienki 1998). Yet, despite the total absence of left-right metaphors in spoken language, there is strong evidence that people implicitly associate time with left-right space and that the direction in which events flow along people's imaginary lateral timelines varies systematically across cultures. In a seminal study (Tversky et al. 1991), children and adults were asked to place stickers on a page to indicate where breakfast and dinner should appear relative to the lunch sticker, in the middle of the page. Whereas English speakers placed breakfast on the left and dinner on the right of lunch, Arabic speakers preferred the opposite arrangement. This crosscultural reversal in the lateral space-time mapping has been corroborated by reaction time tasks (e.g. in English vs. Hebrew speakers; Fuhrman & Boroditsky 2010; Ouellet et al. 2010).

The sagittal mapping of time, enshrined in linguistic metaphors, has been proposed to arise from the canonical experience of moving forward through space (not backward or sideways) due to the construction of our feet, hands, and sensory organs all of which are directed toward the front of our bodies (Clark 1973). As we use these bodies to move forward through the world, objects that we will encounter in the future lie literally ahead of us, and objects we have already passed lie behind us. Thus, a correlation between anteriority and the

future and between posteriority and the past is reinforced nearly every time we walk (or run, bike, drive, fly, etc.). But where does the lateral mapping of time come from?

The left-right mapping of temporal sequences has been hypothesized to arise from our experience with the written word. As we read or write, we move our eyes and attention through both space and time, from left to right for some orthographies (e.g. Roman script) and from right to left for others (e.g. Arabic script). In English, for each line of text we read we begin on the left side of a page (at an earlier time) and arrive at the right side (at a later time). Thus, reading English entails a correlation between "progress" through space and time, from left to right. To find out whether experience using one orthography or another is sufficient to determine the direction of the mental timeline, Roberto Bottini and I asked Dutch participants to perform a space-time congruity task on stimuli written in standard (left-to-right) Dutch orthography, mirror-reversed orthography, or orthography that was rotated either 90 degrees upward or downward (Casasanto & Bottini 2014b). When participants judged temporal phrases written in standard orthography, their reaction times were consistent with a rightward-directed mental timeline: Past-related phrases (e.g. "a day earlier") were judged faster with the left hand, and future-related phrases (e.g. "a year later") with the right hand. After a few minutes of exposure to mirror-reversed orthography, however, participants showed the opposite pattern of reaction times; their implicit mental timelines were reversed. When standard orthography was rotated 90 degrees upward or downward, participants' mental timelines were rotated, accordingly.

5.1 Separating Effects of Language and Culture on Space–Time Mappings

These data show that experience of reading is sufficient to determine the direction of people's implicit mental timelines (though they do not rule out the possibility that other culture-specific practices, such as gesturing about time or using calendars, could influence people's lateral representations of time, as well). But why is this an effect of *cultural experience* as opposed to *linguistic experience*? Language can be considered an aspect of culture, and in many cases it is difficult to disentangle linguistic and non-linguistic practices. At first glance, reading might appear to be an ambiguous case, since it is a cultural overlay on natural language. Yet it is notable that reading is extremely recent and rare in human history; although reading may seem integral to language use in our culture, only a tiny fraction of all of the humans who have ever used language have been able to read. More to the point, in Casasanto and Bottini's experiments, *language was held constant across the orthography conditions*. The words and phrases in natural language were invariant; all that changed was the

direction and orientation of the orthography in which they appeared. Thus, changes in orthography determined the flow of time in people's minds independently of any changes in the structure or content of language.

5.2 Hierarchical Construction of Spatial Metaphors for Temporal Sequence

How could a few minutes of exposure to a new orthography completely reverse people's usual mental timeline, established over a lifetime of reading experience? As in the case of language, space, and pitch, HMMT may explain the flexibility of this culture-dependent mental metaphor. To elaborate, Casasanto and Bottini (2014b) proposed that people's implicit associations between source and target domains can be characterized, not just as families of mappings but alternatively as a set of nested intuitive hypotheses (Goodman 1983). At the top of the hierarchy is the overhypothesis, which comprises a family of specific hypotheses. In this case, the overhypothesis could be: "Progress through time corresponds to change in position along a linear spatial path." This correspondence could be learned as children observe the relationship between space and time in moving objects, or it could be innate (Casasanto 2010; de Hevia et al. 2014; Srinivasan & Carey 2010). Either way, the overhypothesized association between space and time is presumably universal across cultures, and it should be omnidirectional, since more time passes as moving objects travel farther in any direction.

Once children are exposed to cultural practices with consistent directionality, they accumulate a preponderance of evidence for one specific hypothesis. For Dutch children, reading and writing experience provides evidence for the specific hypothesis "Progress through time corresponds to rightward change in position along a linear spatial path," strengthening this hypothesis at the expense of its competitors and causing Dutch speakers to use a rightwarddirected mental timeline by default. Exposure to a different orthography in the experimental setting increased the weight of evidence for one of the participants' overhypothesized (but culturally dispreferred) space–time mappings, strengthening it to the point that it influenced behavior, transiently weakening their culturally preferred mapping as a consequence.

5.3 Spatial Representations of Emotional Valence: Universals and Body-Specificity

Several years ago I proposed a theory of *bodily relativity*, by analogy to the theories of linguistic and cultural relativity described above (Casasanto 2009a). By hypothesis, the contents of our minds are constructed, in part, through our physical interactions with the environment. People with different kinds of

bodies interact with their environment in systematically different ways. Therefore, myriad aspects of their thinking should vary relative to the particulars of their bodies. The spatial mapping of emotional valence has provided one fruitful testbed for this proposal (for reviews, see Casasanto 2011, 2014a).

Across languages and cultures, good things are often associated with the right side of space and bad things with the left. This association is evident in positive and negative idioms like "my right-hand man" and "two left feet." Beyond language, people also conceptualize good and bad in terms of left–right space, but not always in the way linguistic and cultural conventions suggest. Rather, people's implicit associations between space and valence are "body specific." When asked to decide which of two products to buy, which of two job applicants to hire, or which of two alien creatures looks more trustworthy, right- and left-handers respond differently. Right-handers tend to prefer the product, person, or creature presented on their right side but left-handers tend to prefer the one on their left (Casasanto 2009a). This pattern persists even when people make judgments orally, without using their hands to respond. Children as young as 5 years old already make evaluations according to handedness and spatial location, judging animals shown on their dominant side to be nicer and smarter than animals on their non-dominant side (Casasanto & Henetz 2012).

Beyond the laboratory, the association of 'good' with the dominant side can be seen in left- and right-handers' spontaneous speech and gestures (Casasanto & Jasmin 2010): In the final debates of the 2004 and 2008 US presidential elections, positive speech was more strongly associated with right-hand gestures and negative speech with left-hand gestures in the two right-handed candidates (George W. Bush, John Kerry), but the opposite association was found in the two left-handed candidates (John McCain, Barack Obama).

In summary, a body-specific mental metaphor links lateral space and emotional valence, and influences the way people think and communicate about positive and negative ideas. The observed space–valence mappings cannot be explained by influences of language or culture since, in the case of the good-isleft mapping in left-handers, the implicit mental metaphor goes *against* the explicit good-is-right mapping enshrined in linguistic idioms and other cultural conventions (e.g. raising the right hand to swear to tell the truth).

5.4 Experiential Basis of Lateral Space–Valence Mappings

Where do body-specific space-valence mappings come from? All of the results reviewed so far demonstrate correlations, but Casasanto (2009a) proposed a *causal* relationship between the way people use their hands and the way they implicitly spatialize 'good' and 'bad.' In general, greater motor fluency leads to more positive feelings and evaluations: People like things better when they are easier to perceive and interact with (e.g. Ping et al. 2009). Bodies are

lopsided. Most of us have a dominant side and a non-dominant side and therefore interact with the physical environment more fluently on one side of space than on the other. As a consequence right-handers, who interact with their environment more fluently on the right and more clumsily on the left, come to implicitly associate 'good' with 'right' and 'bad' with 'left,' whereas left-handers form the opposite association (Casasanto 2009a).

To test this proposal, Evangelia Chrysikou and I studied how people think about 'good' and 'bad' after their dominant hand has been handicapped, either due to brain injury or to something much less extreme: wearing a bulky ski glove. In one experiment, right-handed university students performed a motor fluency task, arranging dominoes on a table while wearing a cumbersome glove on either their left hand (which preserved their natural right-handedness) or on their right hand (which turned them temporarily into left-handers, in the relevant regard). After about 12 minutes of lopsided motor experience, participants removed the glove and performed a test of space–valence associations, which they believed to be unrelated. Participants who had worn the left glove still thought RIGHT was GOOD, but participants who had worn the right glove showed the opposite LEFT-IS-GOOD bias, like natural lefties (Casasanto & Chrysikou 2011).

5.5 Hierarchical Construction of Spatial Metaphors for Valence

Even a few minutes of altered motor experience can change people's implicit associations between space and emotional valence, causing a reversal of their usual judgments. HMMT provides a potential explanation of this representational flexibility. In the case of mental metaphors linking lateral space and valence, the overhypothesis may be: "The fluent region of space is good." For right-handers, who act more fluently on the side of their dominant hand, typical motor experience provides a preponderance of evidence for the specific hypothesis that "the right side of space is good," whereas typical motor experience for left-handers increases the strength of the evidence for the hypothesis that "the left side of space is good." In terms of memory networks, this means that either the association between 'right' and 'good' or the association between 'left' and 'good' is strengthened at the expense of the competing associations – which are weakened but not lost and can therefore be strengthened again by new patterns of motor experience.

6 Hierarchical Construction of Language-, Culture-, and Body-Specific Mental Metaphors

The case studies described in this chapter illustrate ways in which mental metaphors can be constructed via similar mnemonic processes being driven by

different kinds of linguistic, cultural, or bodily experiences. Early in ontogenetic time (or perhaps over phylogenetic time), *families of source-target mappings* are constructed, which reflect sets of observable relationships between source and target domains in the natural world. Specific members of these families are then strengthened according to an individual's languagespecific, culture-specific, or body-specific experiences. As a result, other mappings in these families are weakened. This process of competition among mappings in long-term memory explains why *all* of the mappings within a given family are not active at once: for example, why adult Dutch speakers do not typically conceptualize pitch using representations of both height *and* thickness, even though both mappings appear to be equally active in infants' minds (Dolscheid et al. 2013; Dolscheid et al. 2014).

The hierarchical structuring of mental metaphors in terms of source–target families and their constituent members may also explain how spatial source–target mappings can be fundamental to our conceptions of non-spatial domains, yet also remarkably flexible. It is notable, for example, that about 5 minutes of exposure to mirror-reversed writing did not simply *modulate* the direction of participants' mental timeline, it *completely reversed* its direction, as indicated by a reversal of reaction–time congruity effects (Casasanto & Bottini 2014b). Had space–time congruity effects been extinguished by exposure to the new orthography, this would have been consistent with the possibility that people cease to represent time spatially when their preferred mental metaphor is challenged. Instead, the reversal of these effects indicates that participants did not abandon a spatial mapping of time; rather they rapidly adopted a *different* mental timeline, consistent with their new orthographic experience.

How are the various members of an overhypothesized family of mappings preserved in long-term memory, even though some of the mappings may never be reinforced explicitly (e.g. by the use of corresponding linguistic metaphors or by cultural conventions)? Presumably they are maintained by the recurrence of the same sorts of physical experiences that are ultimately responsible for the family's construction. Even in a language group that exclusively uses height metaphors for pitch in language, physical thickness–pitch relationships can still be observed (e.g. in the sound of thick vs. thin guitar strings). Even in a left-to-right reading culture, it should still be possible to observe moving objects progressing through space and time from right to left (and in all other directions). Even right-handers occasionally experience greater fluency with their left hand, or on their left side of space. The preservation of dispreferred mappings explains why they can be adopted so quickly when people are given new patterns of experience.

By seeking to understand common mechanisms by which our mental metaphors are shaped by language-specific, culture-specific, or body-specific Comp. by: Kumar Stage: Proof Chapter No.: 3 Title Name: Hampe Date:23/2/17 Time:09:59:42 Page Number: 61

patterns of experience, we can better understand the origins of our thoughts, the extent of cognitive diversity, and the dynamism of our mental lives.

Notes and Acknowledgments

- * This research was supported in part by a James S. McDonnell Foundation Scholar Award (#220020236) and a National Science Foundation grant (BCS #125710). Thanks to Laura Staum Casasanto for helpful discussion.
- 1 This distinction is blurred by the use of the terms "conceptual metaphor" (Lakoff & Johnson 1980), "primary metaphor" (Lakoff & Johnson 1999), or just "metaphor" (Lakoff 1993) to refer to mental metaphors, since these terms are often used to refer to linguistic metaphors, as well.
- 2 HMMT concerns universals and variability within basic mental metaphors; this theory should not be confused with the "inheritance hierarchies" that Lakoff (1993) proposed to explain how complex metaphors (like LOVE IS A JOURNEY, A CAREER IS A JOURNEY, A LIFE IS A JOURNEY, etc.) are constructed and related.
- 3 A common objection to this claim is that we can see the length of a vacation as it stretches across the grid of our wall calendar, and the height of a musical pitch as it is notated on a musical staff. But in these cases, we are not seeing time or pitch, at all; rather, we are seeing a spatial representation of time or of pitch –consistent with the theory that we habitually use space to represent time and pitch metaphorically, both in the mind and on the page.