The British Journal for the Philosophy of Science Advance Access published August 25, 2014

Brit. J. Phil. Sci. **0** (2014), 1–27

Analogue Magnitude Representations: A Philosophical Introduction

Jacob Beck

ABSTRACT

Empirical discussions of mental representation appeal to a wide variety of representational kinds. Some of these kinds, such as the sentential representations underlying language use and the pictorial representations of visual imagery, are thoroughly familiar to philosophers. Others have received almost no philosophical attention at all. Included in this latter category are analogue magnitude representations, which enable a wide range of organisms to primitively represent spatial, temporal, numerical, and related magnitudes. This article aims to introduce analogue magnitude representations to a philosophical audience by rehearsing empirical evidence for their existence and analysing their format, their content, and the computations they support.

doi:10.1093/bjps/axu014

Empirical discussions of mental representation posit a wide variety of representational kinds. Some of these kinds, such as the sentential representations underlying language use and the pictorial representations of visual imagery, are thoroughly familiar to philosophers. Others, such as cognitive maps, are somewhat less familiar. Still others have received almost no attention at all. Included in this latter category are 'analogue magnitude representations' (AMRs).[1]

AMRs are primitive representations of spatial, temporal, numerical, and related magnitudes. They are primitive because they represent magnitudes without presupposing the ability to represent any units of measurement or mathematically defined system of numbers. AMRs are also primitive ontogenetically and phylogenetically. They are present in six-month-old human infants, a wide variety of mammals, many birds, and at least some fish. In human adults with a formal education, AMRs exist alongside culturally acquired representations of space, time, and number. AMRs are analogue in a special sense to be explained in Section 2.

The nature of a representational kind is determined by its format, its content, and the computations it supports. Specifying the nature of any given representational kind is always controversial, but we have at least a rough grasp of the natures of some representational kinds. Thus, to a first approximation: the representations underlying language use have a sentence-like format, topic-neutral contents, and support syntactic and logical computations; the representations underlying visual imagery have a picture-like format, represent visible properties as laid out in egocentric space, and support computations such as rotation and scanning; and cognitive maps have a geometric format, represent the locations of entities in allocentric space, and support computations related to navigation such as localization and path-planning.[2] An analysis of the nature of AMRs should likewise provide insight into their format, content, and computations.

AMRs are overdue for analysis. There is now considerable evidence that they exist and play an important role in the cognitive lives of a wide range of animals, including humans. Given their relevance to the representation of categories such as space, time, and number which have long featured centrally

[1] While AMRs are rarely mentioned in the philosophical literature, there are a few notable exceptions. Laurence and Margolis ([2005]) discuss AMRs in the context of assessing the influence of natural language on the acquisition of natural number concepts; Pietroski *et al.* ([2009]) appeal to AMRs to analyse the semantics of the word 'most'; Burge ([2010]) discusses how AMRs of numerosity are represented in perception; and Beck ([2012]) argues that AMRs serve as an example of cognitive representations that are not systematically recombinable and thus have non-conceptual content. So far as I know, however, this article marks the first attempt to provide a general, philosophically oriented analysis of AMRs.

[2] Fodor ([1975]) and Block ([1983]) contain classic philosophical discussions of the nature of linguistic and imagistic representations, respectively. Rescorla ([2009]) contains an excellent philosophical analysis of cognitive maps.

in philosophical discussions of cognition, AMRs ought to be of intrinsic interest to philosophers. They should also interest philosophers because they bear upon a range of foundational issues in the philosophy of mind, including how mental representations are realized in the brain, the proper analysis of analogue representation, the requirements on representational content, the existence of non-conceptual content, and the nature of animal cognition. This article aims to introduce AMRs to a philosophical audience by reviewing evidence for their existence (Section 1) and carefully analysing their format (Section 2), their content (Section 3), and the computations they support (Section 4).

1 Background

1.1 Evidence of analogue magnitude representations

One interesting lesson to emerge from the study of cognition over the past few decades is that educated human adults are not alone in representing magnitudes such as distance, area, duration, numerosity,[3] and rate. A wide range of organisms, including fish, birds, 'lower' mammals such as rats, non-human primates, human children including pre-linguistic infants, and human adults from all cultures and educational backgrounds represent such magnitudes as well (for reviews see (Gallistel [1990]; Walsh [2003]; Dehaene [2011])). I begin by rehearsing a few highlights from this literature.

When Godin and Keenleyside ([1984]) fed a school of cichlid fish at differing rates from three separate tubes, the fish quickly apportioned themselves in ratios that matched the feeding rates, even before many of them had had a chance to be 'rewarded' by tasting the food. Thus, merely seeing one tube release food morsels at twice the rate of another was sufficient to lead twice as many fish to congregate in front of the first tube. A natural explanation is that the fish represent the rate at which food is dispersed from each tube and adjust their positions accordingly. From a biological perspective, this 'probability matching' behaviour makes sense since it is evolutionarily stable. If fish generally fed only at the source with the greatest payoff, then a non-conformist fish that instead fed at a source with a slightly lower pay out but no competitors present would recover more food. Natural selection thus encourages group foragers to adopt a probability-matching strategy.

Probability matching is widely observed throughout the animal kingdom. In a similar experiment that was performed on ducks by Harper ([1982]), two experimenters tossed morsels of bread to the ducks at two different rates. Within a minute, the ducks divided themselves in proportion to those rates.

[3] Numerosity is a number-like property. I discuss the difference between number and numerosity in Section 3.2.

But when one experimenter tossed morsels that were twice the size of the morsels tossed by the other experimenter, the ducks altered their strategy, and within five minutes repositioned themselves in proportion to the product of the morsel size and feeding rate. Gallistel ([1990], p. 358) summarizes this finding as follows:

> This result suggests that birds accurately represent rates, that they accurately represent morsel magnitudes, and that they can multiply the representation of morsels per unit time by the representation of morsel magnitude to compute the internal variables that determine the relative likelihood of their choosing one foraging patch over the other.

In other words, Gallistel interprets these results as showing that ducks not only represent magnitudes such as rate and physical size, but also perform computations such as multiplication over those representations.

Laboratory studies of rats paint a similar picture. Church and Meck ([1984]) trained rats to press one lever in response to a two-second sequence of two tones and a second lever in response to an eight-second sequence of eight tones. They then varied either the duration of the tones while holding number constant, or the number of the tones while holding duration constant. In both cases, the rats generalized from the training experiment. For example, when number was held constant at four tones, they pressed the first lever in response to a two- or three-second sequence and the second lever in response to a five-, six-, or eight-second sequence, thus showing that they represented the total duration of each tone sequence. But when duration was held constant at four seconds, they pressed the first lever in response to two or three tones and the second lever in response to five, six, or eight tones, thus showing that they also represented how many tones were in each sequence. The rats also immediately generalized their learning to new stimuli presented to new sensory modalities. For example, when presented with a two- or eight-second light sequence, they pressed the same lever they were trained to press in response to a two- or eight-second tone; and when presented with two or eight flashes of light, the rats pressed the same lever they were trained to press in response to two or eight tones. These cross-modal transfer experiments provide evidence that the rats represent abstract durations and numerical information and not just modality-specific stimuli.

While it will hardly be news to the reader that human adults keep track of durations, distances, numerosities, and other magnitudes, it may come as a surprise that human infants do the same, suggesting that such abilities are innate. Xu and Spelke ([2000]) presented six-month-old infants with a display of eight or sixteen dots until they habituated to the display. The infants were then presented with a test display of either eight or sixteen dots. Infants who were presented with a test display that had a novel number of dots looked significantly longer than infants who were presented with a test display that had the same number of dots as the habituation display. As infants are known to apportion their attention to stimuli they deem novel, and as other variables were controlled for (such as total dot area and dot density), this finding suggests that infants can make discriminations based on numerical information. Other dishabituation studies suggest that infants are likewise sensitive to duration (VanMarle and Wynn [2006]) and area (Brannon *et al.* [2006]).

1.2 Weber's law

A peculiar feature of magnitude discriminations deserves emphasis: they obey Weber's law, which holds that the ability to discriminate two magnitudes is determined by their ratio. As the ratio of two magnitudes approaches 1:1 they become harder to discriminate and beyond a certain threshold determined by the subject's 'Weber constant', they cannot be discriminated at all. As an illustration of this ratio sensitivity, consider again the experiments of Church and Meck ([1984]) in which the rats were trained to press the first lever in response to a two-second sequence of two tones and the second lever in response to an eight-second sequence of eight tones. When presented with a five-second sequence of five tones, the rats tended to press the second lever, for although five is equidistant between two and eight, the ratio of 5:2 is greater than the ratio of 8:5. When they were presented with a four-second sequence of four tones, however, the rats favoured neither lever. Since 8:4 = 4:2, four was the subjective midpoint for the rats between eight and two.

When explicit counting is not possible, the magnitude discriminations of human adults also exhibit the ratio sensitivity associated with Weber's law. For example, if two displays of dots are flashed too quickly for you to serially count the dots on each, your probability of correctly guessing which display has more dots will decrease as the ratio approaches 1:1, and beyond a certain ratio (roughly 7:8), your guesses will be at chance (Barth *et al.* [2003]). Similarly, in cultures where explicit counting is absent and there is no culturally acquired integer list, magnitude discriminations are almost always subject to Weber's law (Gordon [2005]; Pica *et al.* [2005]).[4] These results are consistent with the hypothesis that most educated human adults have two systems for representing magnitudes: a primitive 'analogue magnitude' system that is shared with a wide range of other organisms; and a uniquely human

[4] Why almost always? Two reasons: (1) discrimination errors are less common than Weber's law predicts for sets with four or fewer members (this is one finding in support of the existence of a separate 'small number' or 'object file' system); and (2) there is some evidence that when items can be lined up and put in one-to-one correspondence, participants can be induced to rely on Hume's principle to determine whether two sets are equinumerous.

magnitude system that depends upon the acquisition of concepts of natural numbers and units of measurement.

1.3 Scepticism about analogue magnitude representations

It is natural to wonder whether there is a simpler explanation of the results reviewed in this section that does not appeal to AMRs—an explanation that appeals to less sophisticated representations or perhaps to no representations at all. After all, researchers have been led down the garden path before. In the early twentieth century, many were convinced that the horse Clever Hans could perform arithmetic since he would stomp his hoof the appropriate number of times in response to problems posed by onlookers. But the psychologist Oskar Pfungst ([1965]) demonstrated that there was a simpler explanation: Hans was picking up on small unintentional movements in his audience that indicated when to start and stop stomping. Similarly, while the waggle dance of honeybees is often taken to represent the distance to a food source, there is evidence that the dance correlates less with distance than with optic flow—the amount an image moves over the retina—suggesting that bees don't represent distance as such after all (Esch *et al.* [2001]). These examples show that care needs to be taken before representations of magnitudes are posited. Alternative hypotheses need to be considered and ruled out.

In the case of the contemporary study of AMRs, experimenters have taken great care to avoid observer expectancy effects. Many of the techniques pioneered by Pfungst to test Hans are now routinely applied in comparative psychology. For example, experimental assistants are often kept blind about crucial aspects of the experiment to prevent them from inadvertently signalling the animals with which they interact. Studies of AMRs also routinely control for other alternative hypotheses. For instance, Church and Meck's cross-modal transfer experiments were in part designed to test whether animals represent simple sensory properties (such as retinal area or total illumination) rather than abstract properties such as numerosity and duration. The fact that rats are able to track magnitudes across disparate types of stimuli presented to disparate sensory modalities suggests that rats represent more than sensory properties. The overall body of empirical evidence thus strongly suggests—even if it does not apodictically prove—that magnitudes are represented. Moreover, while there are comprehensive and mathematically rigorous explanations of the findings discussed in this section that appeal to AMRs (for example, Gallistel [1990]), to the best of my knowledge there are no corresponding explanations of these results in non-representational terms, or in terms of representations of non-magnitudes. I thus take the existence of AMRs as a reasonable, yet defeasible, empirical hypothesis.[5]

2 Format

2.1 Carey's analogy

Most researchers agree that the ratio sensitivity associated with Weber's law is evidence that magnitude representations have a specific format: they involve some neural entity that is a direct analogue of the magnitude it represents (hence 'AMR'). As the ratio of two magnitudes approaches 1:1, the neural entities become increasingly similar and thus, assuming there is some noise in the system, harder to discriminate. To explain why researchers believe AMRs have an analogue format, Susan Carey ([2009], p. 118) provides a helpful analogy to the following external system of analogue number representations.

number	symbol
1:	__
2:	____
3:	______
4:	________
7:	______________
8:	________________

In this system, line length is a direct analogue of number. The greater the number represented, the longer the line. Suppose our brains deploy magnitude representations that are likewise analogue. Then just as it is harder for us to visually discriminate______________ from ________________ than __ from ____, we would expect our brains to find it harder to discriminate two

[5] Some philosophers may wonder how theories that posit AMRs relate to theories of the mind that aren't squarely in the symbolic mould, such as behaviourism, connectionism, and dynamical systems theory. Theories that posit AMRs are incompatible with behaviourism since they posit internal representations that are not governed by principles of association that characterize classical or instrumental conditioning. In fact, Gallistel ([1990]) and Gallistel and Gibbon ([2002]) argue that many phenomena that have traditionally been explained through conditioning are better explained by arithmetic computations over AMRs. By contrast, because representations can be realized in connectionist networks (even if they are often distributed across those networks), there is no conflict between connectionism and the existence of AMRs. Some researchers have even developed connectionist models for the implementation of AMRs (Church and Broadbent [1990]; Dehaene and Changeux [1993]). In principle, AMRs are also compatible with dynamical systems theory since a system's being describable dynamically need not preclude its being describable in terms of representations and computations. To the best of my knowledge, however, no one has developed dynamical models of the sorts of behaviours that AMRs are posited to explain.

analogue neural representations as the ratio of the magnitudes they represent approaches one. Carey ([2009], p. 458) remarks:

> We do not know how these analog representations are actually instantiated in the brain—larger quantities could be represented by more neurons firing or by faster firing of a fixed population of neurons, for example [...But however] analog magnitude representations are instantiated in the brain, their psychophysical signatures strongly suggest this type of representational scheme.

In other words, even if it remains controversial whether the brain uses a rate, population, timing, or other code to implement AMRs, the fact that magnitude discriminations are ratio sensitive suggests that AMRs are like Carey's external system of analogue number representations in being a direct analogue of the magnitudes they represent.

As I understand it, Carey's argument has the form of an inference to the best explanation. Magnitude discriminations display a peculiar property, ratio sensitivity, which calls out for explanation. The best explanation anyone can think of appeals to analogue representations. Thus, magnitude discriminations are probably supported by analogue representations. While I endorse this argument, there are three aspects of Carey's analysis that have the potential to mislead.

2.2 Neural realization

First, our ignorance about the neural realization of AMRs ought not to be overstated. One reason that AMRs are interesting is that we know much more about how they are realized in the brain than we know about how, say, a sentence in the language of thought is realized in the brain. The last fifteen years in particular have witnessed rapid progress. Experiments with functional magnetic resonance imaging have isolated an area in both parietal lobes inside the intraparietal sulcus that correlates with magnitude discriminations in humans and other primates regardless of the modality through which the magnitude is perceived or represented (Dehaene *et al*. [2003]). Moreover, the activation in this area exhibits its own type of ratio sensitivity: as the ratio of two magnitudes being compared approaches 1:1, the activation increases (Pinel *et al*. [2001]). In fact, researchers have found neurons in the intraparietal sulcus of monkeys that are tuned to specific magnitudes, with bell-shaped activation functions of fixed variance when plotted on a logarithmic axis (Nieder and Miller [2003]; Nieder and Miller [2004]). These activation functions are exactly what one would predict given Weber's law; as the ratio of two magnitudes approaches 1:1, the activation patterns of the neurons corresponding to those magnitudes become harder to tell apart. The activation patterns of these neurons thus mirror the behavioural patterns associated with magnitude discriminations. While many questions still exist about how these neurons are coordinated within larger populations, functional accounts of AMRs are increasingly being integrated with neuroscience. AMRs should thus serve as welcome targets for philosophers interesting in theorizing about the neural realization of mental representations.

2.3 Analogue representation

Second, the line segments (such as '______') that Carey analogizes to AMRs are continuous, and thus analogue in Goodman's ([1976]) sense of being 'dense' (between any two represented values, there is always a third represented value). But it is an open question whether AMRs themselves are continuous or dense since discrete symbols could likewise give rise to ratio sensitivity in magnitude discriminations. To see this, notice that the following external magnitude system would have served just as well to illustrate Carey's point.

number	symbol
1:	|
2:	||
3:	|||
4:	||||
7:	|||||||
8:	||||||||

Although these representations are discrete, they too are a direct analogue of the numbers they represent, and thus become harder to discriminate as their ratio approaches 1:1. Indeed, as Carey herself observes, AMRs could be represented by the total number of neurons firing within a given population. Since action potentials are 'all or none'—neurons either fire or they don't—there would then be no intermediate representation between a network of n neurons firing and a network of $n+1$ neurons firing.

How, then, should the relevant sense of analogue representation be specified if not in terms of continuity? Carey offers the following suggestion:

> Iconic representations are analog; roughly, the parts of the representation correspond to the parts of the entities represented. A picture of a tiger is an iconic representation; the word 'tiger' is not. The head in the picture represents the head of the tiger; the tail in the picture represents the tail. The 't' in 'tiger' does not represent any part of the tiger. ([2009], p. 458)

Carey's suggestion, in other words, is that iconic representations such as pictures are analogue because they have parts that correspond to the parts of the entity that is represented, and that representations such as the word 'tiger' are not analogue because they do not have parts that correspond to the parts of the entity represented. Carey thus seems to endorse what I will call the 'part analogue thesis' (PAT): *R* is an analogue representation of *X* if and only if the parts of *R* represent the parts of *X*. Applying PAT to AMRs, Carey ([2009], p. 458) illustrates why she holds that AMRs are analogue:

> Analog representations of number represent as would a number line—the representation of two (____) is a quantity that is smaller than and is contained in the representation for three (______).

Thus, Carey maintains that AMRs are analogue because they obey PAT: the parts of an AMR represent part of what the AMR as a whole represents.

Notice, however, that PAT is missing a quantifier. Must every part, or only some parts, of *R* represent part of *X*? If we choose the existential quantifier, PAT becomes too permissive. The sentence 'Bill is tall' represents an individual, Bill, and the property of being tall. So 'Bill' represents part of what that sentence as a whole represents, and PAT fails to distinguish AMRs from sentences. Using a universal quantifier doesn't make PAT any more plausible. If every part of *R* has to represent part of *X*, PAT becomes too restrictive. Recall Carey's suggestion that AMRs might be realized by a population of neurons firing. Certainly, a part of one of those neurons doesn't represent anything all on its own. Only the addition of each neuron (or perhaps the addition of a family or network of neurons) indicates a new magnitude. Moreover, there is nothing incoherent in the idea that greater magnitudes might be represented by smaller populations of neurons or lower neural firing rates. In that case, however, a part of a representation of a magnitude of (say) twenty would not represent part of twenty; it would represent more than twenty. PAT thus fails to capture the sense in which AMRs are analogue.

Maley ([2011]) provides a more promising account of analogue representation that can be applied to AMRs. Modifying his analysis only slightly, we can say that a representation, *R*, of a represented magnitude, *M*, is analogue if and only if: (1) there is some property, *P*, of *R* such that the quantity or amount of *P* determines *M*; and (2) as *M* increases or decreases by an amount *d*, *P* increases or decreases as a monotonic function of $M + d$ or $M - d$.[6] Thus, in order for an AMR that represents a magnitude, *M*, to be analogue, there has to be some property, *P*, of the AMR that increases or decreases monotonically with *M*. For example, consider three AMRs that represent four, five, and six

[6] This analysis differs from Maley's ([2011], p. 123) in two significant respects. First, Maley's account is couched in terms of representations of numbers, whereas the present account is couched in terms of representations of magnitudes. Obviously this alteration is necessary if we want to apply the account to AMRs since they represent non-numerical magnitudes such as duration and distance. Second, Maley's account appeals to a linear function rather than a monotonic function. Since AMRs are standardly interpreted as logarithmically compressed (Dehaene [2003]), this amendment is crucial. To his credit, Maley ([2011], p. 123, n. 2) anticipates the amendment.

flashes of light—call them AMR_4, AMR_5, and AMR_6, respectively. According to Maley's account, if these representations are to count as genuinely analogue, then there needs to be some property—presumably some neural property—that belongs to each and that monotonically increases or decreases from AMR_4 to AMR_5 to AMR_6. That property could be the rate at which a population of neurons fires, the number of neurons that fire, or something else entirely; but some such property has to exist that determines which magnitude each of the AMRs represent.

Notice that this account does not require analogue representations to be continuous or dense. Even if *P* turned out to be a non-continuous property, such as the number of neurons firing, the magnitude represented by *R* would be a monotonic function of the quantity or amount of *P* and would thus count as analogue. Additionally, because Maley's account avoids the problematic notion of a part, it avoids the difficulties that befall PAT.

The analogue nature of AMRs makes them unlike the symbols familiar from modern digital computers since there is no property of such symbols that monotonically increases or decreases with the magnitudes they determine. For example, although the binary symbols '00', '01', '10', and '11' represent magnitudes that can be ordered from least to greatest, there is no property of the symbols themselves that determines those magnitudes and increases or decreases monotonically with them.

The relationship between the analogue nature of AMRs and the language of thought hypothesis is less clear, primarily because proponents of the language of thought hypothesis are rarely explicit about the precise format that representations in the language of thought must have. While everyone agrees that representations in the language of thought must be 'language-like', paradigmatic languages have a range of properties, and it is unclear which of these properties representations in the language of thought need to possess. If language-like representations merely need to be compositional and support accuracy conditions, there is nothing stopping them from being analogue. However, it is also a familiar property of prototypical languages that their representations are not analogue. Like the binary symbols of digital computers, natural language words such as 'three' and 'four' do not have some property that monotonically increases or decreases with what they represent. Whether the analogue format of AMRs puts them at odds with the language of thought hypothesis thus ultimately depends on the vexed question of how 'language-like' representations in the language of thought must be.[7]

[7] In other work (Beck [2012], [forthcoming]), I have argued that there is a further respect in which AMRs are unlike paradigmatic linguistic representations: they lack the systematic recombinability of such representations. I return to this issue in Section 3.3.

2.4 Analogue magnitude representation components

There is a third way in which Carey's analysis may be misleading. There is nothing in the line lengths (such as '______') that she analogizes to AMRs to indicate that they correspond to number as opposed to distance, duration, rate, or any other magnitude. But the brain clearly has some means of distinguishing an AMR of seven individuals from an AMR of seven metres or seven seconds. Additionally, AMRs differ in their objects. An AMR of a seven-second light differs from an AMR of a seven-second tone. Thus, a more perspicuous depiction of AMRs would include a representation not just of the magnitude's size, but also what we might call its 'mode' and 'object'. We can thus depict each AMR as an ordered triplet that includes a size, mode, and object component. For example, we could depict an AMR of four fish as {________, NUMEROSITY, FISH}, an AMR of a three-second tone as {______, DURATION, TONE}, and so on. Notice that although the ratio sensitivity associated with Weber's law is evidence that the size component of AMRs is analogue, it does not provide evidence that the mode or object components are analogue. In fact, given that the analysis of analogue representation inspired by Maley is defined to apply only to representations of magnitudes, the mode and object components considered on their own cannot be analogue according to that analysis. Magnitudes are things that can be quantitatively related; yet modes and objects as such cannot be quantitatively related. The questions 'Is duration more or less than distance?' and 'Are light flashes greater or fewer than tones?' do not make sense. It is only the size component of AMRs, or AMRs considered as a whole, that can be quantitatively related, and thus analogue according to the analysis inspired by Maley.

Is the size component of AMRs truly independent of the mode and object components? Or is it an abstraction, imposed by us as theorists, on what in reality is an undifferentiated, homogenous representation? Although these questions have not been definitively resolved, they can be, and have been, subjected to empirical investigation. The basic strategy is to look for evidence of a common neural mechanism across AMRs that differ in their mode and object components. Any such common mechanism would then presumably be attributable to the size component, suggesting that the size component truly is independent of the mode and object components. There are several findings that have been cited in support of a common mechanism:

- Although ratio sensitivity varies across species and ages, there is evidence that the degree of sensitivity is the same for all modes within a given species at a given age. Thus, the ratios at which mature rats successfully discriminate durations are the same as the ratios at which they successfully discriminate numerosities (Meck and Church [1983]), and six-month-old infants show the same ratio sensitivity for area, numerosity, and duration (Feigenson [2007]).

- Attempts to intervene on the representation of size in one mode tend to have similar effects on the representation of size in all modes. For example, when methamphetamine is administered to rats, there is an identical accelerated shift in the representation of both numerosity and duration (Meck and Church [1983]).
- Behavioural tasks that tap into one mode tend to interfere with performance on tasks that tap into other modes. For example, adults are faster at comparing two digits when the greater digit is in a larger font compared to the lesser digit and slower when the greater digit is in a smaller font compared to the lesser digit (Henik and Tzelgov [1982]). Likewise, adults judge the duration of presentation of one digit to be longer than the (equal) duration of presentation of a second digit when the first digit is greater than the second (Oliveri *et al*. [2008]).[8]
- Associations within one mode tend to generalize to other modes, even in young infants. For example, if nine-month-old infants are shown that black striped rectangles are larger than white dotted rectangles, they expect the black striped rectangles to be more numerous and to last for a longer duration as well (Lourenco and Longo [2010]).
- Evidence from both brain deficits and brain imaging point to overlapping loci for magnitude representations of various modes in the inferior parietal cortex of humans and other primates (Walsh [2003]; Jacob *et al*. [2012]).

To be sure, these various lines of evidence are hardly conclusive. Other explanations of the results are possible. Perhaps the ratio sensitivity of AMRs of different modes all improve at the same rate across ontogeny not because they share a common size mechanism, but because organisms tend to use their AMRs of distance, duration, numerosity, and so on equally often, leading them all to be independently honed at the same rate. Similarly, other results might be picking up not on a common size mechanism, but on homogenous AMRs that are instantiated in overlapping or adjacent neural networks. For example, if AMRs of duration and numerosity were embedded in sufficiently proximal neural networks, tasks that activate one might spread activation to the other, leading to interference effects or generalizations from one domain to the other. These possibilities show that the strategy of searching for an

[8] Interestingly, not every magnitude shows the same interference effects. For example, although loudness discriminations are subject to Weber's law, loudness does not interfere with distance discriminations in the way that duration does (Srinivasan and Carey [2010]), suggesting that representations of loudness do not recruit the same size component that AMRs of distance, duration, and numerosity seemingly draw upon.

autonomous size component needs to be carried out with care, but they do not undercut the prospect that thoughtful and patient experimentation will eventually allow us to determine whether AMRs are, in fact, structured from a dissociable size component.

3 Content

Having discussed the format of AMRs, I now want to focus on three questions about their contents. Do AMRs have representational content? If so, what do they represent? And what type of contents do they have?

3.1 Do analogue magnitude representations have representational content?

Following Burge ([2010]), we can distinguish three clusters of theories concerning representational content. At one end sit theories that impose very high demands on the possession of mental representations, such as the ability to speak a language (Davidson [1975]), the ability to engage in conceptual thought (Evans [1982]), the ability to participate in the 'space of reasons' (McDowell [1994]), or the ability to represent objectivity as such (Strawson [1959]).[9] While these demanding theories can acknowledge the existence of AMRs, they will insist that AMRs are only genuinely representational when they are present in organisms that speak a language, engage in conceptual thought, or exhibit other sophisticated cognitive feats. Demanding theories are thus likely to deny that fish, birds, 'lower' mammals, and even human infants represent magnitudes. Consequently, demanding theories sit uneasily with the empirical evidence, reviewed in Section 1, that AMRs do real work explaining the behaviours of these animals, such as why a rat will press lever A rather than lever B following the presentation of a particular stimulus. Although I cannot defend the opinion here, I also agree with Burge ([2010]) that the arguments advanced in favour of demanding theories of representational content are not compelling. I am thus inclined to join Burge in rejecting such theories as unmotivated and empirically dubious.[10]

[9] While Davidson, Evans, McDowell, and Strawson each say things that suggest that they place high demands on the possession of mental representations (Burge [2010], pp. 137–288), interpreting them is not entirely straightforward. For present purposes, I bracket questions of interpretation. The important point is that there is a way of reading each of them that suggests a demanding theory of representation.

[10] Some proponents of demanding theories of representation might allow that AMRs are representational even in 'lower' animals provided that they are treated as sub-individual states. Some subsystem of the rat's represents magnitudes but the rat does not represent magnitudes. Yet given that it is the surprisingly intelligent behaviour of the rat that needs to be explained, including the rat's ability to learn, this reply strikes me as implausible. Such sophisticated behaviours, it seems to me, are best explained in terms of mental states that belong to rats.

At the other extreme sit 'deflationary' theories of representation, which seek to reduce the notion of representational content to other notions, such as information and learning (Dretske [1981]), biological function (Millikan [1989]), or functioning isomorphism (Gallistel [1990]). Because these reductive analyses are relatively easy to satisfy, representations in the deflationary sense are relatively easy to come by. All of these theories are thus likely to count AMRs as having representational content, even in relatively simple organisms. Gallistel ([1990]) is the one deflationary theorist I am aware of who explicitly discusses AMRs. According to Gallistel ([1990], p. 15)

> The brain is said to represent an aspect of the environment when there is a functioning isomorphism between some aspect of the environment and a brain process that adapts the animal's behaviour to it.

Thus, Gallistel counts AMRs as representational because there is a functioning isomorphism between the brain states in which AMRs are realized and magnitudes in the environment. For example, an organism's AMR of a seven-second duration is isomorphic to that duration and is used to help the organism adapt its behaviour to that duration.[11]

Burge ([2010]) criticizes deflationary accounts of representation for obscuring a more robust, non-reductive notion of representation that is rooted in the perceptual constancies. According to Burge, perceptual constancies enable individuals to distinguish what is happening at their surfaces from what is happening in the world, thereby giving rise to non-trivial veridicality conditions. To mark this distinction, Burge reserves the term 'representation' in the first instance for states that are dependent on perceptual constancies. Since many types of sensory states do not involve constancies, Burge does not count them as representational—even if they register information or involve a functioning isomorphism. Burge ([2010], pp. 476, 494–6) thus rejects Gallistel's reasons for counting AMRs as representational; a mere functioning isomorphism is not sufficient for representation. Nevertheless, Burge ([2010], p. 476) does count AMRs as representational because he presumes that AMRs are (at least much of the time) based on perceptual inputs, such as representations of bodies, to which perceptual constancies have already been applied. In fact, I think Burge could go further than he actually does. Because animals represent magnitudes as constant across considerable changes in sensory stimulation (as when rats represent eight tones as numerically equivalent to eight light flashes), AMRs themselves exhibit a type of perceptual

[11] I set aside general difficulties with Gallistel's analysis. Isomorphism is surely too strong a requirement on representation since it eliminates the possibility of error. It should be possible for an AMR to misrepresent a worldly magnitude, though in that case the AMR will not be isomorphic to the magnitude it represents. A more plausible suggestion is that isomorphism serves as a standard of representational accuracy for AMRs in the way that truth serves as a standard of representational accuracy for sentences.

constancy. Information about magnitudes is treated as invariant across diverse inputs. Proponents of Burge's constancy-based approach can thus stake out a middle ground between demanding and deflationary theories of representation that still counts AMRs as representational.

Without taking sides between them, I have been arguing that both deflationary theories of representation such as Gallistel's and constancy-based theories of representation such as Burge's will count AMRs as representational. Moreover, unlike demanding theories of representation, they will count AMRs as representational regardless of whether they occur in mature humans, infants, or non-human animals. It is one thing, however, to say that AMRs are representational and quite another matter to say what they represent. I now want to consider this further question.

3.2 What do analogue magnitude representations represent?

Thus far, I have been speaking loosely of AMRs representing seven flashes of light or a duration of seven seconds, but the appeal to precise integers such as seven to characterize the contents of the size component of AMRs is misleading given that the ability to respond to a given value is imprecise. For example, if we trained a rat to press a lever seven times in exchange for food and then plotted its responses, we would get a bell curve whose mean and mode centred on seven, but with the vast majority of the rat's responses falling slightly above or below seven. Relatedly, AMRs fail to respect the successor relation, which defines the integers, by failing to treat the distance between any two successive integers as exactly one (Carey [2009], p. 295). Because of Weber's law, the difference between two and three is represented as greater than the difference between three and four. Thus, whatever the size component of AMRs represent, it is not the positive integers.

These difficulties also spell trouble for Gallistel and Gelman's ([2000]) proposal that the contents of AMRs are best characterized by real numbers. As Burge ([2010], p. 481) observes, the integers are a subset of the real numbers, so if AMRs cannot represent integers nor can they represent the real numbers. Gallistel and Gelman reason that insofar as the format of the size component of AMRs is analogue, it must be continuous, and thus that the size component of AMRs must represent real numbers, which are likewise continuous. But this line of reasoning is problematic. As we saw in Section 2.3, it is a mistake to suppose that AMRs must be continuous just because they are analogue. Weber's law can be explained by discrete symbols that are a direct analogue of the magnitudes represented. Moreover, as Laurence and Margolis ([2005]) emphasize, Gallistel and Gelman seem to be seduced by a spurious format-content conflation. Just as discrete representational vehicles such as 'π' or '$\sqrt{2}$' can stand for real numbers, continuous representational vehicles can surely stand for something other than real numbers. Thus, even if AMRs were to have a continuous format, it simply wouldn't follow that they must represent the real numbers.

Rather than interpreting AMRs in terms of integers or real numbers, Carey ([2009], p. 135) proposes to interpret them as approximate cardinal values. For example, an AMR might represent 'approximately seven light flashes' or 'approximately seven seconds'. But of course one cannot represent 'approximately seven' without being able to represent 'seven', and so Carey's suggestion fares no better than Gallistel and Gelman's. Moreover, Carey's suggestion fails to specify how approximate AMRs are. While one might reach for greater specificity by appealing to Weber's law, Weber's law is itself defined in terms of ratios among integers, which we were trying to avoid.

Burge ([2010], pp. 482–3) recommends looking to Eudoxus' theory of pure magnitudes to articulate the contents of AMRs.[12] Eudoxus developed his theory of pure magnitudes in order to handle ratios that bedevilled the Pythagoreans. Because the Pythagoreans maintained that whole numbers are the basis of all ratios, they could not express incommensurable ratios such as that between the side of a square and its diagonal. (In our terminology, they could not express irrational numbers such as $\sqrt{2}$.) Eudoxus responded by defining ratios and proportions without appealing to numbers. Instead, he defined ratios in terms of size relations among homogenous magnitudes themselves, and proportions in terms of comparative size relations of ratios.[13] Thus, two line lengths can enter into one ratio; two weights can enter into another ratio; and then those two ratios can be compared and found to be proportional or not. Discrete magnitudes such as whole numbers can enter into ratios and proportions as well, and are understood in the same way as ratios and proportions that involve continuous magnitudes such as line length or weight. Thus, the key feature of Eudoxus' theory, according to Burge, is that its concept of pure magnitude does not differentiate between continuous and discrete magnitudes. Pure magnitudes are like the size component of AMRs in that they cannot be expressed using numbers. This leads Burge to hypothesize that AMRs refer to pure magnitudes.

Burge's hypothesis is suggestive, and seems like a step in the right direction. It is plausible that the size components of AMRs do not refer to numbers, which makes pure magnitudes inviting candidates to serve as their contents. Nevertheless, Eudoxean pure magnitudes are imperfect fits for AMRs in at least one respect: they fail to capture the imprecision that is inherent in AMRs.

[12] Eudoxus' theory of pure magnitudes is described by Euclid ([1956], Book 5). I am indebted to (Sutherland [2006]) and (Stein [1990]) for exposition.

[13] Eudoxus' principal insight was that such comparisons could be spelled out in terms of equimultiples. See (Euclid [1956], Book 5, Definition 5; Sutherland [2006], pp. 536–7; Stein [1990], pp. 166–9).

A ratio of two pure magnitudes can be of unlimited precision—hence their ability to capture the ratio between two incommensurable magnitudes such as the side of a square and its diagonal. But a ratio of two AMRs will be inherently imprecise, as Weber's law exemplifies.

The difficulty of characterizing the size component of AMRs transfers to the mode component. Although it is tempting to say that 'number' is one value that the mode component of AMRs can take, this claim is questionable given that the size component doesn't correspond to integers, real numbers, or any other system of numbers that mathematicians have rigorously defined.

How, then, can we capture the numerical and size contents of AMRs? The empirical literature contains the hint of a suggestion. When researchers discuss AMRs, they often use the term 'numerosity' in place of 'number'. Burge complains, ' "numerosity" is a hedge term meant to apply to number-like properties' ([2010], p. 472). I want to suggest that it is more charitably interpreted as a neologism for the approximate number-like properties that AMRs track. Just as physicists will introduce a new term to refer to a hypothesized particle whose nature is only partially understood, psychologists have introduced the term 'numerosity' as a placeholder for the incompletely understood number-like property that AMRs represent. 'Numerosity' is thus no more of a hedge term than 'Higgs boson'.

In the course of elucidating what numerosities are, we can then appeal to some of the insights we have already encountered. Drawing on Burge's discussion, we can note that numerosities are like Eudoxean pure magnitudes in that they do not refer to numbers. At the same time, we can draw on Carey's suggestion to insist that numerosities are approximate in a way that pure magnitudes are not. Thus, we can use a phrase such as 'approximately seven' to describe the content of an AMR while simultaneously resisting the idea that any such phrase expresses, or is synonymous with, that content. The phrase helps us as theorists to refer to and understand that content even if it doesn't express it. A fuller description of that content would also characterize how approximate it is. In practice, this is often done mathematically by specifying the Weber constant (a measure of the ratio sensitivity) or variance (a measure of the variability around the mean response) associated with an AMR. While researchers use rational numbers to specify these measures, they do not thereby impute possession of rational number concepts to the organism whose AMRs they are characterizing. They use rational numbers as tools to describe the content of AMRs, not to express those contents.

There are difficulties that attach to characterizing the non-numerical modes of AMRs as well. In particular, the contents 'duration', 'distance', 'area', 'rate', and so on might seem too generic since there are different ways that they can be represented, corresponding to the units used to represent them. For example, durations can be represented in seconds, galactic years, or light-miles, and distances can be represented in yards, metres, or light-years. The question thus arises of how AMRs represent these modes.

Early models of AMRs were iterative and appealed to pulses of energy generated by a pacemaker and stored in an accumulator. According to the accumulator model of Meck and Church ([1983]), the rat's pacemaker releases a pulse every 200 milliseconds. Thus, according to their model duration is measured in units that are extensionally equivalent to 200 milliseconds, such that a three-second duration (say) would be equivalent to fifteen of these units. The accumulator model has fallen out of favour, however, since it falsely predicts that when the model is running in its numerosity mode, the time it takes to represent a set is proportional to the set's size (Carey [2009], pp. 132–3). More recent models, which operate in parallel, are not obviously committed to any units. For example, according to Church and Broadbent's ([1990]) model, durations are associated with the phases of a set of neural oscillators of differing periods. Thus, a given duration will typically be associated with several oscillators whose individual periods total to the duration being measured (ignoring error). As a result, there is no unit in this model by which a total duration is measured. Of course, as theorists we might use seconds, minutes, or some other unit of time to describe a total duration, but such units do not play a functional role in the model itself. To borrow a helpful phrase from Christopher Peacocke ([1986]), AMRs might thus be characterized as 'unit free'.[14]

Turning finally to the object component of AMRs, it is notable that AMRs are not restricted to representing objects from only one modality. For example, AMRs can represent visual objects such as light flashes or dots on a screen as well as auditory objects such as tones. Nor are AMRs restricted to attributing magnitudes to spatially well-defined objects. AMRs can also assign magnitudes to abstract events, such as a sequence of rabbit jumps (Wood and Spelke [2005]). In fact, there is no evidence I am aware of suggesting that the objects of AMRs have any restrictions beyond those imposed by the broader representational capacities of the subject.[15]

[14] Peacocke introduces this term to characterize the content of perceptual experiences of magnitudes, as when you visually experience the length of a piano without experiencing it in feet or metres. I hypothesize that the phenomenon Peacocke isolates is grounded in AMRs, i.e. that our conscious experiences have a unit-free character because those experiences are generated by AMRs which are themselves unit free.

[15] Burge ([2010], p. 472) suggests that when AMRs are used in perception, they are limited to representing perceivable entities. But he does not consider whether AMRs are so limited when they are used outside of perception—for example, when a person considers the abstract question whether the union of two sets with thirty-five and twenty-four members, respectively, would have greater or fewer than ninety members in total. There is evidence, however, that AMRs are involved in such abstract comparisons when the objects of comparison cannot be perceived (Dehaene [2011], p. 239ff.). I thus see no reason to deny that AMRs can represent non-perceivable entities outside of perception.

3.3 What content types do analogue magnitude representations have?

We have just been considering what AMRs represent. There is a further question, however, concerning how AMRs represent what they do. What type of contents do AMRs have?

As I understand them, mental contents have the function of capturing how things are from the perspective of the thinker. They fulfil this function, in part, by marking a thinker's mental abilities. Where a mental state is the result of two separate abilities, many philosophers have wanted to use a content that is structured from two elements to capture it. For example, if a person who can think that Amy died rich and that Amy died poor would exercise a single ability—the ability to think about Amy—in the course of thinking each thought, we can mark that fact by attributing a common element—the concept 'Amy'—to the content of each of those thoughts (Evans [1982], pp. 100–5). Since the conceptual thoughts of human beings seem to be structured from discrete abilities in this sense, many philosophers have wanted to use structured contents to capture them.

We have seen reason to think that AMRs are also structured from discrete abilities. For example, the ability to represent seven flashes of light seems to share something in common with the ability to represent a sequence of seven tones. More generally, the ability to deploy AMRs seems to decompose into abilities to represent a size, a mode, and an object. If that is right, then we have reason to view the contents of AMRs as structured.

Another important dimension along which mental contents can differ is their fineness of grain. An argument often advanced in favour of fine-grained contents is that they can accommodate distinct modes of presentation of the same entity. For example, a person might have the ability to think about Venus under the Hesperus mode of presentation while lacking the ability to think about Venus under the Phosphorus mode of presentation. This consideration applies to AMRs. The ability to think about a given magnitude using numbers and units differs in kind from the ability to think about it using AMRs. For example, the ability to represent a ten-second duration as ten seconds is quite different from the ability to represent it in the unit-free manor associated with AMRs. There is thus reason to view AMR contents as composed from modes of presentation.

The idea that contents are structured from modes of presentation is likely to bring to mind 'Fregean thoughts', which are structured from senses. In fact, by claiming that AMR contents are structured from modes of presentation it may seem that I have identified the contents of AMRs with Fregean thoughts. It would be too hasty to draw that conclusion, however, and for two reasons.

First, if properties are individuated finely enough, it may be possible to explain the difference between representing a ten-second duration as ten seconds and representing it in the unit-free manor associated with AMRs in terms of the representation of distinct properties. Thus, just as some argue that the difference between the Hesperus and Phosphorus modes of presentation boils down to the difference between representing the properties 'visible in the evening' and 'visible in the morning', it is open to someone to argue that the difference between unit-laden and unit-free modes of presentation boils down to whether integer and unit properties are represented. In that case, however, there would be no need to appeal to Fregean senses, which are supposed to be distinct from represented properties.

Second, at least as typically conceived, Fregean thoughts and their components—senses—are meant to mark a particular type of ability that I'll call 'conceptual'. One feature of conceptual abilities is that they are systematically recombinable. The conceptual ability to think that *a* is *F* and that *b* is *G* entails the conceptual ability to think that *a* is *G* and that *b* is *F*. In other words, conceptual abilities obey what Evans ([1982]) calls the 'generality constraint', which holds that the conceptual thoughts one can think are closed under all meaningful recombinations of the constituents of the sentences that best express them. It is questionable, however, whether AMRs obey the generality constraint. I have argued elsewhere (Beck [2012], [forthcoming]) that because of Weber's law, a thinker using AMRs can have the ability to represent that a magnitude of nine is less than a magnitude of eighteen, and that a magnitude of ten is less than a magnitude of twenty, but not that a magnitude of nine is less than a magnitude of ten nor that a magnitude of eighteen is less than a magnitude of twenty. Moreover, because Weber's law is traceable to the analogue format of AMRs themselves, I argued that this failure of recombinability is a failure of representational competence and not merely a failure of discriminative performance. It is in the nature of AMRs that they are unable to represent that one magnitude is less than another when their ratio exceeds a certain threshold. If that is right, the sorts of representational abilities that underlie AMRs differ in kind from conceptual abilities. They do not exhibit the same unfettered recombinability. We thus have reason to distinguish the contents of AMRs from Fregean thoughts even if both types of content are structured from modes of presentation. We should conclude that AMRs have a *sui generis* type of non-conceptual content instead.

4 Computations

Supposing that AMRs exist, the question of what organisms can do with them arises. What sorts of computations do they support?

4.1 Arithmetic computation

Above all, AMRs are associated with arithmetic computations, including comparison, addition, subtraction, multiplication, and division. Recall, for example, the ability of ducks to position themselves in proportion to the rate at which experimenters toss morsels of bread. This ability embodies a capacity for division since the ability to represent rate plausibly derives from the ability to divide representations of numerosities by representations of durations (Gallistel [1990], pp. 351–83). It also embodies an ability to compare two rates and calculate how much greater one is than the other. Recall as well that the ducks altered their strategy when one experimenter tossed morsels that were twice the size of those tossed by the other, thereby exemplifying an ability to multiply morsel size by feeding rate.

If AMRs support arithmetic computations and are part of the innate cognitive hardware of human beings, then human children should exhibit a primitive ability to engage in arithmetic prior to being trained in it. Barth *et al.* ([2006]) set out to test this prediction. Using animated displays in which sets of coloured dots move behind or emerge from an occluder, they tested the ability of pre-school children to compare, add, and subtract sets of dots. For example, the children might see a set of twenty-five blue dots move behind an occluder, then see twenty-five more blue dots join them behind the occluder, then see thirty red dots and be asked whether there are more (occluded) blue dots or (unoccluded) red dots. Children succeeded on comparison, addition, and subtraction tasks even after controlling for non-numerical variables such as dot circumference, area, and density. More recently, McCrink and Spelke ([2010]) used a similar paradigm to show that children could succeed on non-symbolic multiplication tasks before being schooled in multiplication or division. Both sets of results suggest that AMRs support a primitive type of arithmetic that does not depend upon formal mathematical training.[16]

4.2 Practical deliberation

To fully appreciate the computational power of AMRs, it is essential to notice that they can be used not only to represent how the world is, but also how the cognizer would like it to be. That is, AMRs can play a desire-like role in addition to a belief-like role. Consider the long-tailed hummingbird, which forages by recovering nectar from a variety of widely dispersed sources, often flying at least half a kilometre for any one feeding. Because the bird is too

[16] Arithmetic computations over AMRs are discussed at length in (Gallistel [1990]). See also (Brannon *et al.* [2001]; Flombaum *et al.* [2005]; McCrink and Wynn [2004]; Beran and Beran [2004]).

small to store much energy, and consumes energy rather quickly, there is considerable pressure for it to optimize its frequent foraging runs—to find an interval that is long enough for the harvest to have replenished since the previous visit, but not so long that the bounty is likely to have been pilfered by a competitor. That in turn requires the bird to represent the rate at which various nectar sources replenish after depletion and the temporal intervals between visits. It can then compute an estimation of the amount of nectar that is currently at each source, and use that value as a proxy for the utility of visiting each source. In controlled environments using artificial flowers that are filled with sugar water at intervals set by the experimenter, it can be shown that birds will, in fact, optimize their visits (Gill [1988]).

Other examples of how AMRs might encode desires or utilities are not hard to dream up. A robin might use its AMR of the rate of return of worms in a given field as a proxy for the desirability of foraging there, a monkey might use its AMR of the numerosity of predators stalking it to estimate the utility of retreat, and a human child might use its AMR of the size of two pieces of cake as a measure of the desirability of each. The ability to use AMRs in this way is important, as it opens up the possibility of embedding an entire process of practical deliberation over AMRs within a formal decision-theoretic framework, such as expected utility theory, whereby an animal calculates the expected utility of each of a range of actions and then chooses the action with the maximum expected utility.

To see how this might work for a simplified case, suppose that a robin is deciding between two actions: foraging for worms in the field (A_1) or foraging for berries in the forest (A_2). And suppose that the world can be in one of two possible states, raining (S_1) or not raining (S_2), where rain increases the prevalence of worms in the field but has no immediate effect on how many berries are available in the forest. Assuming that robins value worms and berries equally, they could calculate the expected utility of each action as follows (where $u(A|S)$ is the desirability that the robin assigns to action A given that the world is in S, and $Pr(S|A)$ is the subjective probability that the world will be in S given that action A is performed):

$$EU(A_1) = [u(A_1|S_1) \times Pr(S_1|A_1)] + [u(A_1|S_2) \times Pr(S_2|A_1)],$$

$$EU(A_2) = [u(A_2|S_1) \times Pr(S_1|A_2)] + [u(A_2|S_2) \times Pr(S_2|A_2)].$$

As we've already seen, $u(A|S)$ can be calculated by the robin's AMR of the rate of return of worms in the field and berries in the forest during rain, and no rain, in the past. But how will the robin calculate the values of $Pr(S|A)$, i.e. the probability of rain? One possibility is that these values will be based upon the correlations in its past experiences between the amount of sky that is covered by clouds and rain. Since measures of area can be captured using AMRs, the robin's practical deliberations about where to forage might thus be based entirely on AMRs.

I do not mean to insist that cognizers use expected utility theory as opposed to any other formal decision-theoretic framework; its simplicity just makes it useful for illustration. The key point is that given some formal decision-theoretic framework or other, various probabilities and subjective utilities can be represented entirely through AMRs, allowing AMRs to percolate all the way through practical deliberation. We thus arrive at a model of one primitive form of non-conceptual practical deliberation. Using this model we can begin to see how a host of sophisticated behaviours could be explained by appealing only to AMRs, raising the prospect that AMRs can account for a significant portion of the behaviour of 'lower' organisms. Moreover, even in organisms where AMRs typically operate by interacting with a range of more sophisticated representations, they can still play a significant role in guiding behaviour. Whether you see a large piece of cake as preferable to a small piece will depend on whether you are dieting, but your AMR of the size of the cake can contribute to your decision in either case.

5 Conclusion

Over the past three decades, empirical work on AMRs has exploded as researchers have come to recognize the outsized role that these representations play in cognition. By contrast, AMRs are rarely discussed in the philosophical literature—in spite of philosophers' enduring interest in the representation of time, space, and number; and in spite of the relevance of AMRs to a host of central issues in the philosophy of mind concerning neural realization, analogue representation, representational content, non-conceptual content, and animal cognition. In this article I have sought to take some first steps towards redressing this imbalance by drawing attention to AMRs and analysing their format, their content, and the computations they support. Although my analysis of AMRs has only scratched the surface, I hope that it will assist philosophers who want to give AMRs a more central place when they theorize about the mind.

Acknowledgements

I presented material from this article at the annual meeting of the Southern Society for Philosophy and Psychology in Austin, Texas, in February 2013. Thanks to those who braved an 8 a.m. start to provide me with feedback. I am also indebted to my colleagues Brian Huss, Henry Jackman, Josh Mugg, and especially Kristin Andrews for valuable comments on drafts, and to two anonymous referees for this journal whose helpful suggestions spawned substantial improvements.

Department of Philosophy
York University
4700 Keele Street
Toronto, Ontario
Canada, M3J 1P3
jbeck@yorku.ca

References

Barth, H., Kanwisher, N. and Spelke, E. [2003]: 'The Construction of Large Number Representations in Adults', *Cognition*, **86**, pp. 201–21.

Barth, H., La Mont, K., Lipton, J., Dehaene, S., Kanwisher, N. and Spelke, E. [2006]: 'Non-symbolic Arithmetic in Adults and Young Children', *Cognition*, **98**, pp. 199–222.

Beck, J. [2012]: 'The Generality Constraint and the Structure of Thought', *Mind*, **121**, pp. 563–600.

Beck, J. [forthcoming]: 'Analogue Magnitudes, the Generality Constraint, and Nonconceptual Thought', *Mind*.

Beran, M. J. and Beran, M. M. [2004]: 'Chimpanzees Remember the Results of One-by-One Addition of Food Items to Sets over Extended Time Periods', *Psychological Science*, **15**, pp. 94–9.

Block, N. [1983]: 'Mental Pictures and Cognitive Science', *The Philosophical Review*, **92**, pp. 499–541.

Brannon, E. M., Lutz, D. and Cordes, S. [2006]: 'The Development of Area Discrimination and Its Implications for Number Representation in Infancy', *Developmental Science*, **9**, pp. F59–64.

Brannon, E., Wusthoff, C., Gallistel, C. R. and Gibbon, J. [2001]: 'Numerical Subtraction in the Pigeon: Evidence for a Linear Subjective Number Scale', *Psychological Science*, **12**, pp. 238–43.

Burge, T. [2010]: *Origins of Objectivity*, New York: Oxford University Press.

Carey, S. [2009]: *The Origin of Concepts*, New York: Oxford University Press.

Church, R. M. and Broadbent, H. A. [1990]: 'Alternative Representations of Time, Number, and Rate', *Cognition*, **37**, pp. 55–81.

Church, R. M. and Meck, W. H. [1984]: 'The Numerical Attribute of Stimuli', in H. L. Roitblat, T. G. Bever and H. S. Terrace (*eds*), *Animal Cognition*, Hillsdale, NJ: Lawrence Erlbaum, pp. 445–64.

Davidson, D. [1975]: 'Thought and Talk', in S. Guttenplan (*ed.*), *Mind and Language*, Oxford: Oxford University Press, pp. 7–23.

Dehaene, S. [2003]: 'The Neural Basis of the Weber–Fechner Law: A Logarithmic Mental Number Line', *Trends in Cognitive Sciences*, **7**, pp. 145–7.

Dehaene, S. [2011]: *The Number Sense: How the Mind Creates Mathematics*, New York: Oxford University Press.

Dehaene, S. and Changeux, J. P. [1993]: 'Development of Elementary Numerical Abilities: A Neuronal Model', *Journal of Cognitive Neuroscience*, **5**, pp. 390–407.

Dehaene, S., Piazza, M., Pinel, P. and Cohen, L. [2003]: 'Three Parietal Circuits for Number Processing', *Cognitive Neuropsychology*, **20**, pp. 487–506.

Dretske, F. I. [1981]: *Knowledge and the Flow of Information*, Cambridge, MA: MIT Press.

Esch, H. E., Zhang, S., Srinivasan, M. V. and Tautz, J. [2001]: 'Honeybee Dances Communicate Distances Measured by Optic Flow', *Nature*, **411**, pp. 581–3.

Euclid [1956]: *The Thirteen Books of Euclid's Elements*, translated by T. L. Heath, New York: Dover.

Evans, G. [1982]: *The Varieties of Reference*, New York: Oxford University Press.

Feigenson, L. [2007]: 'The Equality of Quantity', *Trends in Cognitive Sciences*, **11**, pp. 185–7.

Flombaum, J. I., Junge, J. A. and Hauser, M. D. [2005]: 'Rhesus Monkeys (*Macaca mulatta*) Spontaneously Compute Addition Operations over Large Numbers', *Cognition*, **97**, pp. 315–25.

Fodor, J. [1975]: *The Language of Thought*, New York: Crowell.

Gallistel, C. R. [1990]: *The Organization of Learning*, Cambridge, MA: MIT Press.

Gallistel, C. R. and Gelman, R. [2000]: 'Non-verbal Numerical Cognition: From Reals to Integers', *Trends in Cognitive Sciences*, **4**, pp. 59–65.

Gallistel, C. R. and Gibbon, J. [2002]: *The Symbolic Foundations of Conditioned Behaviour*, Mahwah, NJ: Lawrence Erlbaum.

Gill, F. B. [1988]: 'Trapline Foraging by Hermit Hummingbirds: Competition for an Undefended Renewable Resource', *Ecology*, **69**, pp. 1933–42.

Godin, J. J. and Keenleyside, M. H. A. [1984]: 'Foraging on Patchily Distributed Prey by a Cichlid Fish (Teleosti, Cichlidae): A Test of the Ideal Free Distribution Theory', *Animal Behaviour*, **32**, pp. 120–31.

Goodman, N. [1976]: *Languages of Art: An Approach to a Theory of Symbols*, Indianapolis, IN: Hackett.

Gordon, P. [2005]: 'Numerical Cognition without Words: Evidence from Amazonia', *Science*, **306**, pp. 496–9.

Harper, D. G. C. [1982]: 'Competitive Foraging in Mallards: Ideal Free Ducks', *Animal Behaviour*, **30**, pp. 575–84.

Henik, A. and Tzelgov, J. [1982]: 'Is Three Greater than Five? The Relation between Physical and Semantic Size in Comparison Tasks', *Memory and Cognition*, **10**, pp. 389–95.

Jacob, S. N., Vallentin, D. and Nieder, A. [2012]: 'Relating Magnitudes: The Brain's Code for Proportions', *Trends in Cognitive Sciences*, **16**, pp. 157–66.

Laurence, S. and Margolis, E. [2005]: 'Number and Natural Language', in P. Carruthers, S. Laurence and S. P. Stich (*eds*), *The Innate Mind: Structure and Content*, New York: Oxford University Press, pp. 216–35.

Lourenco, S. and Longo, R. [2010]: 'General Magnitude Representation in Human Infants', *Psychological Science*, **21**, pp. 873–81.

Maley, C. [2011]: 'Analog and Digital, Continuous, and Discrete', *Philosophical Studies*, **155**, pp. 117–31.

McCrink, K. and Spelke, E. S. [2010]: 'Core Multiplication in Childhood', *Cognition*, **116**, pp. 204–16.

McDowell, J. [1994]: *Mind and World*, Cambridge, MA: Harvard University Press.

Meck, W. H. and Church, R. M. [1983]: 'A Mode Control Model of Counting and Timing Processes', *Journal of Experimental Psychology: Animal Behaviour Processes*, **9**, pp. 320–34.

McCrink, K. and Wynn, K. [2004]: 'Large-Number Addition and Subtraction in Infants', *Psychological Science*, **15**, pp. 776–81.

Millikan, R. [1989]: 'Biosemantics', *Journal of Philosophy*, **86**, pp. 281–97.

Nieder, A. and Miller, E. K. [2003]: 'Coding of Cognitive Magnitude: Compressed Scaling of Numerical Information in the Primate Prefrontal Cortex', *Neuron*, **37**, pp. 149–57.

Nieder, A. and Miller, E. K. [2004]: 'A Parieto-Frontal Network for Visual Numerical Information in the Monkey', *Proceedings of the National Academy of Sciences of the United States of America*, **101**, pp. 7457–62.

Oliveri, M., Vicario, C. M., Salerno, S., Koch, G., Turriziani, P., Mangano, R., Chillemi, G. and Caltagirone, C. [2008]: 'Perceiving Numbers Alters Time Perception', *Neuroscience Letters*, **438**, pp. 308–11.

Peacocke, C. [1986]: 'Analogue Content', *Proceedings of the Aristotelian Society, Supplementary Volume*, **60**, pp. 1–17.

Pfungst, O. [1965]: *Clever Hans: The Horse of Mr. von Osten*, New York: Henry Holt.

Pica, P., Lemer, C., Izard, V. and Dehaene, S. [2005]: 'Exact and Approximate Arithmetic in an Amazonian Indigene Group', *Science*, **306**, pp. 499–503.

Pietroski, P., Lidz, J., Hunter, T. and Halberda, J. [2009]: 'The Meaning of "Most": Semantics, Numerosity, and Psychology', *Mind and Language*, **24**, pp. 554–85.

Pinel, P., Dehaene, S., Riviere, D. and LeBihan, D. [2001]: 'Modulation of Parietal Activation by Semantic Distance in a Number Comparison Task', *Neuroimage*, **14**, pp. 1013–26.

Rescorla, M. [2009]: 'Cognitive Maps and the Language of Thought', *British Journal for the Philosophy of Science*, **60**, pp. 377–407.

Srinivasan, M. and Carey, S. [2010]: 'The Long and the Short of It: On the Nature and Origin of Functional Overlap between Representations of Space and Time', *Cognition*, **116**, pp. 217–41.

Stein, H. [1990]: 'Eudoxus and Dedekind: On the Ancient Greek Theory of Ratios and Its Relation to Modern Mathematics', *Synthese*, **84**, pp. 163–211.

Strawson, P. F. [1959]: *Individuals: An Essay in Descriptive Metaphysics*, New York: Methuen.

Sutherland, D. [2006]: 'Kant on Arithmetic, Algebra, and the Theory of Proportions', *Journal of the History of Philosophy*, **44**, pp. 533–58.

VanMarle, K. and Wynn, K. [2006]: 'Six-Month-Old Infants Use Analogue Magnitudes to Represent Duration', *Developmental Science*, **9**, pp. F41–9.

Walsh, V. [2003]: 'A Theory of Magnitude: Common Cortical Metrics of Time, Space, and Quantity', *Trends in Cognitive Sciences*, **7**, pp. 483–8.

Wood, J. and Spelke, E. [2005]: 'Infants' Enumeration of Actions: Numerical Discrimination and Its Signature Limits', *Developmental Science*, **8**, pp. 173–81.

Xu, F. and Spelke, E. [2000]: 'Large Number Discrimination in 6-Month-Old Infants', *Cognition*, **74**, pp. B1–11.