

# The Anonymity of a Murmur: Internet (and Other) Memes

Simon J. Evnine

*Memes, of the kind found often (though not exclusively) on the internet, are an increasingly significant medium of expressive activity. I develop a theory of their ontological nature and, in parallel, an analysis of the concept meme. On my view, memes are abstract artifacts made out of norms for production of instances. The norms say things like ‘use a certain image; add text of a certain kind; the text should be delivered in two chunks, one at the top of the image, one at the bottom, etc.’ Instances of these memes are created when users follow (or flout) these norms and publish the results. The concept meme is analyzed in terms of the notion of memographic practice, a historically situated form of activity within which memes are created and their instances produced and made public.*

Memes, of the kind paradigmatically (though not exclusively) encountered on the internet, are now widespread.<sup>1</sup> To most of us, they are images captioned and re-captioned for humorous, political, and satirical purposes, sometimes made for the clearly aesthetic goals of exhibiting beauty, wit, and pathos. On its own, each is a mere whisper of creative expression, ephemeral and anonymous, too insignificant to be treated as a work of art. But together these whispers create a vast susurrant that restlessly adapts itself to new technologies and new modes of expression and communication. Nothing could be more aptly described by these prescient words of Michel Foucault:

I think that, as our society changes ... the author function will disappear, and in such a manner that fiction and its polysemous texts will once again function according to another mode ... All discourses ... would then develop in the anonymity of a murmur.<sup>2</sup>

But although scholars in communications and media studies have begun to explore memes, focusing on the issues that are of most interest to them—the forms in which memes are made and dispersed, the nature of the meaning-making they allow, and so on—philosophers have not yet got to work addressing the questions that particularly interest us and that, I hope, will be able to complement the work done in other disciplines. In this paper, I aim to make a start on a philosophical consideration of memes. I will develop both an account of the ontology of memes and an analysis of the concept *meme*. But a preliminary

- 
- 1 The term ‘meme’ was introduced by Richard Dawkins, in *The Selfish Gene* (Oxford: OUP, 1976), as part of a speculative theory of cultural evolution. In that theory, memes are the cultural counterparts of genes in biological evolution. There is some dispute over the relevance of Dawkins to the internet phenomenon of memes. My own view is that Dawkins’ very interesting theory is not useful in understanding the topic of this paper and I set it aside without further ado.
  - 2 Michel Foucault, ‘What is an Author?’, in Paul Rabinow (ed.), *The Foucault Reader* (New York: Pantheon, 1984), 101–120, at 119.

note of caution is in order. Philosophy, in its desire for clear and consistent accounts of ordinary phenomena, always runs up against the untrammelled nature of ordinary uses of significant terms. But that danger is even more pressing in the present case given how recent, and hence unsettled, the use of the term ‘meme’ is and given how the internet culture within which memes and talk of memes have flourished is itself so anarchic.<sup>3</sup> Any attempt at a philosophical treatment of such topics is bound to have a slight air of an exasperated parent with too many unruly children. I apologize in advance!

## 1. What is a Meme?

What is a meme? It is of the greatest philosophical importance to be clear that the form of words ‘what is an *F*?’ is capable of expressing two kinds of question. One kind of question is conceptual. Questions of this kind call for a definition, or analysis, of the concept *F* and will be answered by whatever one thinks is the appropriate way of defining or analysing or explicating a concept. Traditionally, this has been by giving illuminating, non-question-begging necessary and sufficient conditions (or something approximating these) for being an *F* and that is the approach I adopt here. But all sorts of other approaches to answering the conceptual question are available. The other kind of question is ontological. These are questions about the ontological nature of the things that fall under the concept *F*: of things that are *F*s, what are they? So, for example, if our question is ‘what is a teenager?’, the answer to the first kind of question might be something like ‘a person between the ages of 13 and 19’. That tells us about the concept *teenager*, about what counts as being a teenager. It provides a necessary and sufficient condition for the application of that concept. The answer to the second kind of question, however, will be simply ‘a person’. Of something that is a teenager, the answer to the question *what is it?* is that it is a person. It is not part of its nature that it is between the ages of 13 and 19. In this example, the answer to the conceptual question includes the answer to the ontological question. I am uncertain whether this must always be the case but in what follows I will assume that the relation between the answers to the two questions ‘what is a meme?’, the conceptual and the ontological, is like that between the answers to the two questions expressed by ‘what is a teenager?’: the ontological answer will specify a broader kind and the conceptual answer will take the form of giving the conditions under which something of that broader kind will fall under the kind *meme*.<sup>4</sup>

I shall initially (and mostly) focus on image macros. Image macros are memes that involve images that are copied from person to person, but customized by users through

3 Patrick Davison, ‘The Language of Internet Memes’, in Michael Mandiberg (ed.), *The Social Media Reader* (New York: NYU Press, 2012), 120–134, at 120, characterizes as the ‘unrestricted web’ that part of the internet in which creativity is valued over security. Although he is not thinking of the ‘security’ of established definitions and correct usage, his point extends to that as well. Memes, and talk about memes, have largely flourished in the unrestricted web (e.g. on 4chan, reddit, and so on). Their migration to what Davison calls the restricted web, including Facebook, usually happens later and is often seen as a kind of kiss of death by the passionate and energetic agents of the unrestricted web.

4 This is actually not exactly true in the fourth and final version of the answers I give.

the addition of text to each new copy. For many people, they are paradigms of what a meme is. In what follows, I develop answers to our two ‘what is a meme?’ questions that I think work well for image macros. I shall then indicate how to generalize these answers to apply to memes that are not image macros as we gradually confront just how wide the domain of memes is. Some well-known examples of image macros are *Socially Awkward Penguin*, *Batman Slapping Robin*, and *Archaic Rap*. The image for *Socially Awkward Penguin* is a photograph of a penguin against an abstract, blue background. That for *Batman Slapping Robin* is a panel from a 1965 Batman comic in which, as the name suggests, Batman slaps Robin. In *Archaic Rap*, a self-portrait by the eighteenth-century painter Joseph Ducreux is used.<sup>5</sup> In each of these cases, users add text of a particular kind. In *Socially Awkward Penguin*, the text is about situations in which one feels socially awkward. In *Archaic Rap*, the text should be an archaic rendering of lyrics from a rap song. The text associated with *Batman Slapping Robin* should involve a situation in which (youthful) eagerness or annoyance is slapped down by the hand of cynicism or experience.<sup>6</sup> The addition of text to these images by users takes place within a certain context. Users have seen the efforts of others, on Facebook, 4chan, or elsewhere. They have intuited, more or less well, what one is supposed to do and, in the light of their understanding, they produce their own efforts (often through specially designed websites) that they then post on the internet. This situation is what makes it proper to think of memes as involving copying in some sense and what accounts for the use of Dawkins’ neologism for the phenomenon in question. The image is copied by one user from another, but undergoes mutation, with respect to its textual element, with each copy. (And, of course, people often copy without mutation when they simply repost someone else’s effort. But this kind of copying is not what lies at the heart of the phenomenon of memes.) In addition, there is a meta-level of activity in which examples of these image macros are discussed, commented on, up-voted, down-voted, criticized, collected, replied to in kind, and so on. Let us call all this *memographic practice*. Memographic practice includes many memographic practices around particular image macros, though of course there is more to memographic practice in general than merely the sum of the individual memographic practices, since some memographic practice may deal with memes in general, or memes of some sub-kind, without being specifically about any particular image macro. I will return to the notion of memographic practice in due course. For now, I have tried to describe it in as neutral a way as possible. Memographic practice is not just an adventitious feature of the production and enjoyment of memes. As almost all writers on the subject recognize, this practice is an essential part of the point of memes. No attempt to understand the phenomenon could possibly succeed if it did not in some way or other accord a significant role to it. Where different theorists diverge is over how to incorporate it into an account of memes and what the consequences are of incorporating it in one way rather than another.

5 These images are not associated with stable URLs but an internet search of the names I have given will lead to examples of them.

6 I have a lot more to say about *Batman Slapping Robin* in my *A Certain Gesture: Evnine’s Batman Meme Project and its Parerga!* (in preparation) and its associated blog *A Certain Gesture: Evnine’s Batman Meme Project and its Parerga! A blog about a book about a meme*, <<https://evninebatman.wordpress.com>> accessed 25 June 2018.

In the following, I will develop an account of memes that identifies them with something that is common to the many things that people create as part of memographic practice. These common contents of the many things produced are what we refer to with such terms as ‘*Socially Awkward Penguin*’, ‘*Batman Slapping Robin*’, and ‘*Archaic Rap*’. What, exactly, is the common content that a meme is, on my suggestion? Since, focusing on image macros, producing the instances involves adding idiosyncratic, user-generated text to a common image, it is a very attractive and natural thought to identify the common content with the relevant image itself, without words. Jonathan Carter seems to adopt this view when he writes:

Each of these images [e.g. the image of Batman slapping Robin] is its own meme.<sup>7</sup>

But, as attractive as it is, I think the identification of memes with images should be resisted. We noted above how absolutely crucial to understanding memes is memographic practice. Some of the role of memographic practice will appear when we turn our attention to the conceptual question expressed by ‘what is a meme?’. But we can give a better answer to the conceptual question by building in more to the ontological answer of what memes are than simply an image. To hold that a meme itself is just an image leaves no connection, in the ontology of memes, to memographic practice. (Another reason for resisting the identification of memes with images is that this will make it much harder to generalize our account to other kinds of memes that do not involve images.)

We get an important clue about what the common content is in the continuation of the sentence I quoted above from Carter:

Each of these images is its own meme, with unique affective dimensions and implied narratives.<sup>8</sup>

Let us think about that ‘with’. If one considers *Batman Slapping Robin*, it is clear that the wordless image associated with it already has an affective dimension and implied narrative. Not only is there implied narrative (implications of things that might have preceded the depicted slap or that might follow it), there is even explicit narrative, since by the conventions of reading such images, Robin’s speech act occurs before, and indeed occasions, the slap and Batman’s verbal response. It is also highly affect-laden since it depicts a violent act in which a child is being beaten. It is easy to see how the various images users create by adding text to the wordless image are constrained in terms of affect and narrative by the affective nature and implied narrative of the wordless image itself. But many image macros are not like that at all. The image involved in *Socially Awkward Penguin*, for example, does not have or express much affect or imply any narrative. The portrait of Joseph Ducreux associated with *Archaic Rap* may express affect and imply narrative (it is, after all, a striking painting), but if it does, it is in ways that bear very complex relations to the kinds of text associated with it. These complex ways cannot simply be ‘read off’ the image. The relevant expressions of

7 Jonathan Carter, ‘Modal Ethos: Scumbag Steve and the Establishing of Ethos in Memetic Agents’, in Moe Folk and Shawn Apostel (eds), *Establishing and Evaluating Digital Ethos and Online Credibility* (Hershey, PA: IGI Global, 2017), 291–308, at 293.

8 Ibid., my emphasis.

affect and narrative implications that go ‘with’ images, in their connection with memes, therefore, must be capable of being associated with their respective memes not only (or not at all, in some cases) through the imagistic properties of the images but through external imposition. That imposition is accomplished through memographic practice. Broadly speaking, people start using the images in ways that connect them to certain affects and/or narratives, others respond and imitate, there is consolidation and refinement through the early stages of a meme’s history, and so there comes to obtain the requisite association between images, on the one hand, and affective dimensions and implied narratives, on the other. In some cases, this is helped along by affective dimensions and implied narratives present in the images themselves but, even here, there may need to be external imposition of the relation between image and meme (whether the affects and narratives of the memes are consistent, or at odds, with the affects and narratives suggested by the image).<sup>9</sup>

The roles that the suggestions of affect and narrative play are as norms. They require, prohibit, allow, encourage, or discourage certain kinds of textual additions. Memographic practice creates an expectation that someone who wants to add text to a particular image will do it in a way that conforms to certain constraints on affect, narrative, subject matter, theme, and so on that have come to be associated with that image. Indeed, once we start thinking in terms of norms, we can see that the relevant constraints go well beyond what we might call material constraints concerning such things as affect and narrative and encompass formal constraints as well. For example, for many image macros, in adding text to an image, one is expected to have two chunks of text, one at the top and one at the bottom, often rhetorically structured as set-up and punch-line. *Socially Awkward Penguin* is subject to all those formal norms and is a paradigm of image macros. Its material constraints require that the text be about situations in which people have felt themselves to be socially awkward. *Archaic Rap*, by contrast, has no particular affective or narrative requirements, but its material requirement is that the text be an archaic rendering of lines from a rap song. Its formal requirements are those of the typical image macro (two text chunks, etc.) except that the rhetorical structure of set-up and punch-line drops out owing to the particular nature of its material constraints. *Batman Slapping Robin* is rather atypical of image macros. The placement of text is not at the top and bottom of the image but in the speech bubbles that are part of the image. Also, because the image depicts a dynamic scene with two interacting people, the rhetorical structures available are rather unconstrained, as long as the text chunks can function as two successive contributions to a conversation. Its material constraint, as noted above, is that the text represent (youthful) eagerness or annoyingness slapped down by experience or age.

Putting all this together, a natural account of what a meme is begins to emerge. We shall offer several refinements as we pursue this idea but, as a first pass, we can answer the ontological question as follows:

Meme<sub>Ont</sub> (first pass): A meme is an ordered pair <I, N> where I is an image and N is a set of norms.

9 See the notion of ‘centrifugal multimodality’ in Lillian Boxman-Shabtai and Limor Shifman, ‘Evasive Targets: Deciphering Polysemy in Mediated Humor’, *Journal of Communications* 65 (2014), 977–998 for intentional dissonance between text and image.

What *Archaic Rap* is, for example, is the ordered pair consisting of the image of the portrait of Ducreux and a set of norms that includes the following: there should be two chunks of text, at the top and bottom of the image, and the text should be an archaic rendition of some lines from a rap song.  $\text{Meme}_{\text{Ont}}$  is an answer to the ontological question ‘what is a meme?’. It is not a definition or analysis of the concept *meme*. One cannot glean from it much of what it is to be a meme just as one cannot glean much of what it is to be teenager from the ontological claim that teenagers are persons. The concept, as I indicated, will be defined in terms of restrictions that a pair consisting of an image and a set of norms must meet in order for that pair to fall under the concept *meme*. We may, as our first pass at defining the concept, say something like the following (there are many more or less equivalent, but perhaps slightly different, ways to frame this):

$\text{Meme}_{\text{Con}}$  (first pass): M is a meme if and only if M is an ordered pair  $\langle I, N \rangle$ , where I is an image and N a set of norms such that the norms in N have come to be associated with I through memographic practice and people produce images by modifying I in accordance with N as part of that memographic practice.

The notion of memographic practice is clearly a vital part of the definition of the concept *meme*, on my account. (I address worries about circularity in 2.C) below. But, unlike our rejected attempt to identify the common content simply with an image, memographic practice is also connected in an interesting way to the answer to the ontological question ‘what is a meme?’. This is because there is nothing about an image, as such, that suits it for being taken up by a memographic practice. Norms, by contrast, in their very nature, are the sorts of things to govern practice. According to  $\text{Meme}_{\text{Ont}}$ , the ontological nature of memes has them already primed and ready to be taken up in some kind of social practice.

## 2. Refinements to and Comments on our Answers

A) We may immediately simplify these answers in a way that also extends their reach in a satisfying way. Instead of seeing an image macro as a pair of an image and a set of norms, we may simply take as one of the norms that a certain image be used to add text to. So, instead of seeing *Archaic Rap* as the pair of the image of Ducreux and the set of norms that say ‘add two chunks of text...’ etc., we may see it simply as the set of norms that says ‘Use the image of Ducreux, add two chunks of text...’. This allows us to give a very simple answer indeed to the question about the ontology of memes:

$\text{Meme}_{\text{Ont}}$  (second pass): A meme is a set of norms.

And we can correspondingly simplify our conceptual answer to:

$\text{Meme}_{\text{Con}}$  (second pass): M is a meme if and only if M is a set of norms N for producing images such that the norms in N have come to be associated with each other through memographic practice and people produce images in accordance with N as part of that memographic practice.

The reformulations here, besides being simpler than the originals, have two notable advantages. First, although many image macros involve one particular image (as all of our examples

do), there are also image macros where any image of a certain kind may be used. For example, *Feminist Ryan Gosling* requires only that one use an image of Ryan Gosling in which he looks fetching. No specific such image is required. This is easily subsumed by the modified definition since the relevant norm may say either ‘use this particular image’ (as they do with our original examples) or ‘use any image of the following kind. . .’ (as they do with *Feminist Ryan Gosling*). Secondly, they also easily point the way to further refinement to encompass all kinds of memes, and not just image macros. The ontological answer can stand as is, since we have already dropped any requirement to include images in the second pass (only let us now refer to it as a ‘third pass’ to keep the levels of answers to the two questions aligned). All we need to do to effect the widening is to replace, in the conceptual answer, ‘images’ with some general term that can refer to any of the sorts of things that can be instances of a meme. I shall employ the term ‘things’ but other choices are available and I do not mean anything to ride on my decision:

Meme<sub>Con</sub> (third pass): M is a meme if and only if M is a set of norms N for producing things such that the norms in N have come to be associated with each other through memographic practice and people produce things in accordance with N as part of that memographic practice.

The ‘things’ in question might be images, videos, songs, imaginary book titles, actions, or anything else. For example, consider the meme *Planking*. There is a set of norms that tells a person to extend herself horizontally in a rigid pose, preferably in an unlikely place or situation, have someone take a picture, and then post that picture on the internet. One might argue that the picture itself is the instance of the meme and hence we are still dealing with images, but it is at least plausible to hold that the instance of the meme is the particular action itself and the picture just a proof. There certainly *could* be something like *Planking* that did not require one to, or even required one not to, post a picture. In that case, it would certainly be an action that was the instance of the meme.<sup>10</sup> Or consider #scalebackabook. This Twitter hashtag is associated with norms that require one to produce (and post on Twitter with the hashtag) a scaled back version of a book title (for example, *The Country of Frozen Yogurt Star* instead of *The Country of Ice Cream Star*, or *Oliver Bend* instead of *Oliver Twist*). Here the instance of the meme is an imaginary book title or possibly a tweet but, in either case, not an image. These cases show how our third pass answers can be extended beyond the realm of image macros to take in other types of memes, including (as in the case of *Planking*) memes that are not necessarily internet-based.

B) I have outlined a view on which memes are sets of norms for the production of things (images, actions, tweets, book titles, etc.) and users produce instances of these memes by following the norms in the set. But it is undeniable that many people apply the term ‘meme’ to those instances. There is little point in arguing which things memes *really* are, common contents or instances. The term is commonly used for both and it is

---

10 In case you think I am here stretching the extension of *meme* beyond what is reasonable, I note that other established theorists of memes discuss pre-internet memes (such as *Kilroy was Here*) and take *Planking* to be of roughly the same nature. See Limor Shifman, *Memes in Digital Culture* (Cambridge, MA: MIT Press, 2014), who describes *Planking* as a meme which ‘basically involves people lying face down with their arms at their sides in public places’ (27).

understandable why that is so. We should therefore recognize a benign ambiguity in the term ‘meme’ between what we can think of as its common content sense and its instance sense. Where necessary, I will subscript the term with ‘CC’ or ‘I’ to disambiguate. We can answer the conceptual question ‘what is a  $meme_I$ ?’ as follows:

$Meme_{I-Con}$ : M is a  $meme_I$  if and only if M is an instance of a  $meme_{CC}$ .

This is not circular since our conceptual answer to what a meme, in the common content sense, is did not employ the concept  $meme_I$ . But the two definitions together do imply a priority relation. The priority in question is conceptual, not ontological: the concept  $meme_I$  is defined in terms of the concept  $meme_{CC}$  but not vice versa. This is as it should be. What is primary, from a conceptual point of view (i.e. from the point of view of understanding what memes are) is the potential for a certain kind of production.

With respect to the ontology of  $memes_I$ , there simply is no general account. Where the  $meme_{CC}$  in question is an image macro, its instances will be images; where it is something like *Planking*, its instances will be actions; where it is like the Twitter hashtag #scalebackabook, they will be book titles or tweets. This result is salutary. Although we do have a uniform account of the ontology of  $memes_{CC}$ , we should not expect there to be one for their instances, given how many different things can be produced by norms that are formulated by memographic practice. Nor are  $memes_{CC}$  ontologically prior to  $memes_I$ . Indeed, as we shall see in F) below,  $memes_I$  might turn out to be ontologically prior in some sense to  $memes_{CC}$ .

C) This brings us to the role of memographic practice in the answer to the conceptual question ‘what is a  $meme_{CC}$ ?’ I have said relatively little so far about memographic practice but it is clear that this concept is doing a lot of important work in the definition. There are many sets of norms that are not memes either because they fail to be norms for producing things at all, or they fail to be norms that have been put together by memographic practice for producing things as part of that practice. For example, there is a norm that applies to speakers of English to apply the word ‘dog’ to dogs, if one wishes to talk about dogs.<sup>11</sup> One might think of each such application as something produced by someone in accord with the norms. But the norm will only count as a  $meme_{CC}$  (and hence the applications as  $memes_I$ ) if memographic practice is involved in the consolidation and application of the norm. Intuitively, one wants to exclude something like this from counting as a meme; it certainly seems, on the face of it, not theoretically useful or appealing to class it along with *Planking* or *Batman Slapping Robin*. We therefore require a robust hold on the notion of memographic practice. One can imagine various ways of characterizing this notion, with knock-on effects on the extension of the concept *meme*, defined in terms of memographic practice. This latitude seems fine—there is no good reason to think the concept is, in the absence of some theoretical assistance, entirely determinate or clearly bounded.<sup>12</sup>

11 The norm would need to be formulated with greater care than I have done here.

12 See Shifman, *Memes in Digital Culture* and Ryan Milner, *The World Made Meme: Public Conversations and Participatory Media* (Cambridge, MA: MIT Press, 2016) for extended treatments of what I have called memographic practice. Both these works have helped me formulate my own thoughts about this topic. For reasons of space I cannot go into the similarities and differences between their accounts and mine.

Without attempting an exhaustive characterization of memographic practice, here are some thoughts. 1) As I suggested at the opening of the paper, memes are like whispers of art. As with art in general, people make them to express their feelings, to entertain, to comment on various things, or simply to exercise their skill, virtuosity, and wit. Memographic practice, then, should be seen as a kind of artistic practice. It is, however, an artistic practice that, relative to historical norms of preceding centuries, reconfigures—in the manner suggested by the quotation from Foucault I gave at the outset—the relation between producer and product. Memes are, so to speak, art without artists. Seeing memographic practice as a form of artistic practice will serve to exclude, say, practices around greeting people, or the application of words, from counting as memographic. It is also the key to understanding a couple of other distinctive features of memographic practice. 2) Jerrold Levinson's historical definition of art emphasizes how art is an on-going practice in which there is, except in its earliest moments, a continuous awareness of its history.<sup>13</sup> For a practice to count as memographic, there must be a strong sense, within the practice, of making things that are responsive to other products of the practice.<sup>14</sup> The norm-governed practice of shaking hands, for example, will not count as memographic (and hence individual acts of hand-shaking will not be instances of a meme) because one does not, in general, take one's hand-shakes to be riffs on, responses to, or developments or transformations of, earlier hand-shakes. (Though some people may evolve such a self-aware practice and in that case their behaviour might be memographic and their hand-shakes memes<sub>1</sub>.) 3) Following from this self-awareness, memographic practice should be seen as a kind of play.<sup>15</sup> The practice provides a kind of framework for friendly competition. The meme<sub>CC</sub> provides a set of (more or less) agreed-upon rules and the practice provides an invitation for people to show what they can do within those rules. This is very clear in spontaneous and local memes, such as the Facebook example I discuss in section F) below.

4) We must pay some attention to how the description of the practice involves the concept *meme* itself. It would clearly be circular simply to characterize memographic practice as a practice around the production of memes since here, as theorists, we would be using the concept *meme* in our explanation of a component of the definition of the concept. But there is obviously nothing illicit in allowing that people engaged in the practice might themselves understand their practice in terms of the concept *meme*. They may often think of themselves as making memes. But there is something stronger we can allow, too, that is derived from approaches to art like Levinson's historical definition and institutional theories of people like Danto and Dickie. It may be not just an adventitious but a deep and central feature of the practice that it involves the application of the concept *meme* by those

13 Jerrold Levinson, 'Defining Art Historically', *BJA* 19 (1979), 21–33.

14 Shifman's definition of *meme* picks up on this: 'a group of digital items sharing common characteristics of content, form, and/or stance; that were created with awareness of each other; and were circulated, imitated, and/or transformed via the Internet by many users' (*Memes in Digital Culture*, 7; italics mine).

15 This point is made by several theorists of memes, e.g. Yuval Katz and Limor Shifman, 'Making Sense? The Structure and Meanings of Digital Memetic Nonsense', *Information, Communication & Society* 20 (2017), 825–842 and Jens Seiffert-Brockmann, Trevor Diehl, and Leonhard Dobusch, 'Memes as Games: The Evolution of a Digital Discourse Online', *New Media & Society* (2017), <<https://doi.org/10.1177/1461444817735334>> accessed 30 August 2018.

engaging in it. There is a kind of foundational part of the practice that does not involve application of the concept of the product. This can be described in terms of the concepts *copying*, *changing*, *publicizing*, *image*, *text*, and so on. This foundational level accretes, both in time and in its logical complexity, further layers that become conceptualized in terms of the practice itself and its products. Hence the memographic practice in terms of which *meme* is analysed comes to involve the application by practitioners of that very concept. There is nothing mysterious, or logically suspicious, about this phenomenon. It is pervasive in human affairs and has been studied in a variety of forms by philosophers and others.

D) In B) above, I pointed out what I called a benign ambiguity in the term ‘meme’ between its common content and its instance senses. These two senses are closely related, as we saw, and it is natural that the term has come to function in these ways. There is, however, another ambiguity lurking which is less benign, in the sense that it extends the use of the term outside of the circle of interrelated concepts and practices that unite both  $\text{memes}_{\text{CC}}$  and  $\text{memes}_i$  and hence threatens confusion.<sup>16</sup> It is now quite common for people to describe as a meme any internet image with words that have been added. It is clear how this linguistic practice might arise. For memes that are image macros, one produces an instance by adding words to an image. Hence, adding words to an image comes to be seen itself as sufficient for making a  $\text{meme}_i$ , even under circumstances where there is no set of norms that govern how words are to be added to that image—in other words, even where there is no relevant  $\text{meme}_{\text{CC}}$ , or memographic practice, in existence. Yet this use of the term is not just a matter of a semantic elision. There is, in fact, a whole alternative theoretical discourse that takes memes to be just the juxtaposition of words and images. Theoretical statements about what a meme is are made along these lines, though without the distinction between the conceptual and ontological questions, it is often hard to know how to evaluate such claims. The artist Barbara Kruger is often brought in as an early, pre-internet pioneer of the conjunction of image and text. For example, Maggie Williams and Lauren Razzore write:

An internet meme consists of a still (or sometimes animated) image with a short piece of text superimposed over the scene... [D]ebating the artistic value of memes may also have its place. For example, we might ask what differentiates an internet meme from the work of a modern artist like Barbara Kruger? Kruger’s feminist collages frequently rely upon meaningful juxtapositions between borrowed photographs and short graphic texts; memes are very similar, in that they make meaning by combining existing images with words inserted later by a secondary viewer/creator.<sup>17</sup>

Clearly this kind of approach is perfectly valid in its own terms, but it uses the word ‘meme’ in a significantly different way from that envisaged here although, given the use of the term to mean *meme*, it is often very difficult to know, in particular cases, whether an author describes something as a meme because it is an instance of an image macro or because it is simply an image with added text. In the alternative theoretical discourse I am

16 And recall that I have completely set aside the term’s original, Dawkins-related use.

17 Maggie Williams and Lauren Razzore, ‘Medieval Memes’, in Gail Ashton (ed.), *Medieval Afterlives in Contemporary Culture* (London: Bloomsbury, 2015), 322–332, at 322–323.

describing in this section, the role of common content and memographic practice is displaced from the heart of the account of what memes are (in both the ontological and conceptual sense) and becomes simply an interesting fact about the way memes, taken merely as images with added text, may be produced and circulated on the internet. Nor can the approach extend naturally to non-image based memes like *Planking* and *#scalebackabook*. But the biggest point of difference between an approach that sees any image with added text as a meme and my own approach is that, on my view, one cannot tell that something is a meme<sub>1</sub> merely by looking at it. Even if a person were to add two chunks of text to the image of *Socially Awkward Penguin* in a way that reflected a socially awkward situation, the result would not be a meme<sub>1</sub>, an instance of the meme<sub>CC</sub> *Socially Awkward Penguin*, unless it was produced as part of a memographic practice. At a minimum, this means produced by someone who knew that this image was supposed to be used in this way. Something is a meme<sub>1</sub>, on my view, only if it has the right kind of history. The historical requirements for something to be a meme, according to the alternative discourse, are much thinner. I say 'thinner' rather than 'non-existent' because even on this approach, it is possible to hold that something is a meme only if it is intentionally made, only if it involves adding text to an image that already exists, and so on. But setting aside these very general historical requirements, one could say that for the alternative discourse, *meme* is not a historical concept whereas on my own approach, it is.

Yet even this alternative discourse might be incorporated into my own theoretical framework. For it might be that the existence of memographic practice in general has been responsible for the coming into being of a very general meme<sub>CC</sub> the only norm of which is 'add some text to an existing image', or something like that. Call this *MEME*.<sup>18</sup> In that case, any time somebody adds text to an image, as part of this on-going practice, they are making an instance of *MEME*.<sup>19</sup> Any instance of an image macro of a more specific kind, such as *Batman Slapping Robin*, will also be an instance of *MEME*; but there might be instances of *MEME* that are not instances of narrower meme kinds. But even if this is so, we must remember that *MEME* is simply one meme among many and that to give an account of it is not to give an account of memes in general. I shall return to this issue in F) below.

E) I have associated memes with norms and accounted for the singularity of a meme with respect to the plurality of norms it may be associated with via the notion of a set. A meme, I said, is a set of norms. Philosophers almost always reach for sets to do this kind of work and I think that is often a mistake. My own position in this paper, so far, is such a mistake and my answers to the ontological and conceptual questions posed by 'what is a meme?', therefore, must be taken as approximations, to be replaced by something better.

18 I use all capitals here to distinguish the name of this particular meme<sub>CC</sub> from my use of the italicized '*meme*' to name a concept.

19 This idea was suggested to me by Rosa Vince, at the University of Sheffield, and further illuminated by a conversation with Haley Mathis of the University of Miami. I should stress the importance of 'as part of this on-going practice'. The creation of instances of *MEME* is a particular, situated practice that came into existence with the internet and does not include the addition of words to pictures in comics, advertising, etc., which are parts of their own historical practices.

Before coming to the replacement, let me describe what I take to be problematic about relying on sets in this instance (and, *mutatis mutandis*, in many other instances). One problem, perhaps less pressing than the one I shall come to next but still worthy of note, is that sets are not the kinds of things that can be created; they exist automatically as long as their members do. Something like a meme, however, is a cultural artefact, essentially the product of intentional creation. (I discuss the creation of memes in the following section.) But the more pressing problem is that sets cannot change their members, and have their members essentially. Memes, by contrast, should not be thought of as incapable of changing the norms they are associated with or as being associated with just those norms essentially. For example, it is, arguably, one of the norms for most image macros that they use Impact font. But surely, the very image macros for which that is a norm might have been associated with a different norm requiring some other font. It is a highly contingent matter that Impact became so associated with image macros.<sup>20</sup> If these memes were sets of norms that included the norm ‘use Impact,’ those very sets could not have existed with a different norm, for example ‘use Comic Sans’. Similarly, we might think that although it was originally one of the norms of *Archaic Rap* that the song lyrics put in archaic English come from a rap song, it is now no longer the case and lyrics may come from a wider variety of songs. Again, if the meme were a set of norms, we could not account for such changes over time.

Philosophers who favour the use of sets for ontological purposes such as the analysis of changeable entities with contingent features may resort to more complex constructions. For example, to deal with change over time, we could take a meme to be a function from times to sets of norms. I do not think this strategy is helpful but I will not pursue the issue here. Instead, I shall state my preferred way of thinking about all this.<sup>21</sup> A large class of the objects for which set-theoretic analyses are inappropriate are artefacts. An artefact is a hylomorphically complex entity. It is made out of its matter by a maker’s imposing a concept on that matter (not to be confused with predicating a concept of it) by an act of intentional making. The result is a *sui generis* kind of object, the identity conditions of which (over time, and modally) are determined by the concept imposed and the act of its imposition. The same account can hold for abstract artefacts as well as concrete ones. So a musical work involves the imposition of a concept (such as *symphony*) onto a sound structure;<sup>22</sup> a fictional character involves the imposition of the concept *fictional character* onto some

20 See Kate Brideau and Charles Berret, ‘A Brief Introduction to Impact: “The Meme Font”’, *Journal of Visual Culture* 13 (2014), 307–313 for the story. A referee for this journal has suggested that the use of Impact is not a norm for image macros but merely a common feature of them, stemming from the availability of the font in most meme generators. I take the point and agree with the distinction between normative and merely common or prevalent features. I do not have a settled view about whether the use of Impact is normative or not but, if it is, I think it provides a good example of an obviously contingent norm associated with various memes. Of course, I am not committed to thinking that *all* norms are inessential to their memes.

21 The approach is developed at length in my *Making Objects and Events: A Hylomorphic Theory of Artifacts, Actions, and Organisms* (Oxford: OUP, 2016).

22 This account of musical works, which I develop in *Making Objects and Events* and in ‘Constitution and Qua Objects in the Ontology of Music’, *BJA* 49 (2009), 203–217 is very close to Levinson’s view of musical works as indicated sound structures, as set out in his ‘What a Musical Work is’, *Journal of Philosophy* 77 (1980), 5–28.

properties (being handsome, being clever, being rich, etc.). In these cases, the very same object could have existed had the sound structure or the properties on which the maker imposed the concept been different (though perhaps not entirely different). Furthermore, the matter of such abstract artefacts can change, as when a composer or author re-writes a work or, in the case of fictional characters, provides a sequel in which the properties out of which the character is made are added to. Memes are abstract artefacts. The norms the set of which I treated as being identical to a meme are in fact its matter. The meme-maker imposes the concept *meme* onto these norms in an act of creation but the very same meme could have been made out of somewhat different norms.

In the light of this account, I give a fourth (and final) pass to answering the ontological and conceptual questions expressed by ‘what is a meme?’:

Meme<sub>Ont</sub> (fourth pass): A meme is an abstract artefact made out of norms.

Meme<sub>Con</sub> (fourth pass): M is a meme if and only if M is made, as part of memographic practice, out of norms for producing things as parts of that memographic practice.

A couple of comments on these fourth pass answers will be in order. First, the making need not be by one person and its period may be imprecise. This, I believe, is the case for many cultural artefacts and I discuss this at greater length in the following section. Secondly, my account of the nature of artefacts requires the imposition of the concept under which the artefact essentially falls in its creation. We thus see here the significance of the point I made above—that, although as theorists we should not describe memographic practice using the concept *meme*, we can, and here we see we should, allow that the practice may involve the deployment of the concept by its practitioners, even though the concept is itself analysed in terms of that practice. This, as I noted, is part of the self-aware nature of meme-making, along with many other cultural and artistic practices.

F) In E), one of my objections to taking memes as sets of norms, rather than as comprising norms in some other way, had to do with creation. Sets exist automatically if their members exist. But memes, I argued, are the kinds of things that are created, that have to be made in order to exist. It behooves us now to consider the process of creation of memes a little more closely. It is memes<sub>CC</sub> that pose the problem; the nature of the creation of memes<sub>CC</sub> seems straightforward. If we are thinking of an image macro, an instance of it is created by a user deliberately and intentionally adding text to the (or an) image specified by a norm that the meme is associated with. As with the production of anything, there may be vagueness concerning exactly when it comes to exist but we can easily see that it comes to exist within a certain, narrowly-defined interval, through creative activity that is similar to that involved in, say, the creation of a limerick. Memes<sub>CC</sub>, on the other hand, raise quite different issues. The processes of their creation are much more diffuse and often spread over many people, and the deliberation and intention directed at the object of creation is likely to be highly attenuated.

Presumably (continuing to focus on image macros) memes<sub>CC</sub> come to exist as people produce images with text and a social practice around these images either comes to bring into being various norms or consolidates existing norms in packages so that further practice involving the image comes to be governed by (or answerable to) those norms. It is quite likely that in many cases a user first adds text to an image without having any very

specific ideas about how other users should respond. But even if a user does have a specific idea, those ideas may be quickly overwhelmed by how other users actually respond and which elements of the original text they take as guides to the norms that govern further practice.

Here is an example of the coming into being of a meme<sub>CC</sub> on a local and restricted level, not this time an image macro but something more like the case of the twitter hashtag #scalebackabook that we looked at above. On 13 March 2017, Shannon Dea (a philosopher) made the following post on Facebook:

Just saw a strawberry-filled doughnut... and now I'll be humming that Beatles song all day. #strawberryfilledforever

Her Facebook friends naturally (and probably expectedly) took up the challenge, seeing the post as invitation to ‘throw down’ themselves. But what exactly were they supposed to throw down? Funny versions of song titles? Funny versions of Beatles song titles? Food-related versions of song titles? Food-related versions of Beatles song titles? Would titles of things other than songs count? Dea did not explicitly say, though she might have done. Almost immediately, Mark Morton responded with four food-related variants of Beatles song-titles (including ‘Can’t Buy Me Lunch’ and ‘I Want to Hold Your Ham’) and this set the practice pretty firmly for following entries. A couple of cases, however, did not conform. After a posting of the Muppets singing ‘Hey Food’ Andrew Howat posted a Muppets version of Tom Waits songs. And Geoffrey Curry had a food-related line that simulated a line from a Beatles lyric and incorporated a clever reconceptualization of its title (‘That’s a lot of sugar dear. Prudence’).<sup>23</sup> The thread did not continue for very long, and the community involved was, no doubt, very tolerant and good-natured. But one could see, if there had been a longer period of time and a nastier community ethos, that standards for ‘doing it right’ would emerge. The meme<sub>CC</sub> in question would have been clarified—whether it would include a norm that specifically required Beatles song titles, etc.—and possibly enforced. (For example, *Archaic Rap* is associated with the norm that the text should be an archaic rendering of lines from a rap song. But I have somewhere seen, I no longer know where, people being criticized for making instances of it using lines from the non-rap song ‘Call Me Maybe’.) The meme’s creation is a collective effort, in this case, since Dea did not specify ‘how to go on’. Together, the people imposed the concept *meme* onto some norms, creating the meme in the process of performing it. (The example also illustrates the ludic nature of memographic practice that I mentioned in section C) above.

There is a potentially perplexing feature of this process, however. Are the early efforts in the process themselves memes<sub>i</sub>? My definition of the concept *meme<sub>i</sub>*, recall, was that something is a meme<sub>i</sub> if and only if it is an instance of a meme<sub>CC</sub>. But the early efforts are parts of the processes whereby those memes<sub>CC</sub> come to exist and hence cannot be instances of them. Is Dea’s own ‘Strawberry Filled Forever’ (the hashtag formatting was not picked up on in subsequent entries) a meme<sub>i</sub>? Were the earliest images with text that

23 Many thanks to Shannon Dea and all those named in this example for permission to make use of their Facebook postings.

helped to establish the existence of *Batman Slapping Robin* memes?<sup>7</sup> I suggest that these items do have a different status from genuine instances that are made after the practice has got going. Those first entries are not made *according to* a practice associated with a meme<sub>CC</sub> because no such practice yet exists. But if they are not instances of a meme<sub>CC</sub>, nonetheless, ontologically, they are exactly the same as later, genuine instances. In the case of image macros, for example, they are all just images.

Here we should remind ourselves of the way in which the alternative discourse around memes discussed in D) seemed to be ultimately subsumable by my own account. There exists a general practice of meme-making. A person may associate text with an image, or come up with a food-related variant of a Beatles song title, with or without the hope or intention that it will catch on and that others will produce images or song title variants that resemble the original one in specific ways. All these products are, given the existing social practice, instances of the very general meme<sub>CC</sub> I called *MEME*. Some do catch on, in ways that may, or may not, reflect the hopes and intentions (if there were any) of the original maker(s). If they do catch on, a new meme<sub>CC</sub> may come to exist, and subsequent images or song titles, that resemble the originals in certain ways (the text conforms to the same subject matter, rhetorical structure, etc.) will be instances of that new meme<sub>CC</sub>, though their inspirations were not. Perhaps the originals are given a retroactive and honorary status as instances of the new meme<sub>CC</sub>. They are still memes<sub>I</sub> because they are instances of *MEME*, but may be only honorary instances of *Batman Slapping Robin* or *Food-Related Beatles Song-Title Variants* (to give that meme a snappy name). Their epigones are genuine, and not honorary, instances of both *MEME* and *Batman Slapping Robin* or *Food-Related Beatles Song-Title Variants*.

A final point on the topic of creation. The way in which memes<sub>CC</sub> come to exist makes them existentially dependent on the original exemplars that establish, or assemble into packages, the norms that constitute them. Although, as we have been seeing, those original exemplars are not yet instances of the memes<sub>CC</sub> they help to bring into existence, they are nonetheless memes<sub>I</sub> since they are instances of *MEME*. Thus, memes<sub>CC</sub> are existentially dependent on memes<sub>I</sub> although, as we saw earlier, memes<sub>I</sub> are conceptually dependent on memes<sub>CC</sub>.

G) A final, important point is that for something to be an instance of a meme<sub>CC</sub> that is associated with a given set of norms, it is not necessary that it actually comply with all those norms. Where the failure to comply with norms is accidental, users may be subject to correction (as noted in the previous section). More interesting are cases where makers deliberately flout the norms associated with a meme. The mechanism seems parallel to the mechanism Paul Grice describes under the rubric of conversational implicature.<sup>24</sup> Conversants share acceptance of various conversational norms, according to Grice, but can flout these norms deliberately to implicate content that they do not explicitly say. A famous example involves an exchange in which A asks B whether Jones is a good philosopher and B replies that Jones has good hand-writing. The answer deliberately flouts the norm of relevance and allows B to implicate, without actually saying, that Jones is not a good philosopher. Likewise, meme-makers can flout commonly recognized norms to

24 Paul Grice, *Studies in the Way of Words* (Cambridge, MA: Harvard University Press, 1991).

create predictable effects and layers of meaning. Since producing a meme does not have semantic content in the way a linguistic utterance does (of course, the words in a meme<sub>1</sub> have semantic content, but that cannot be said to be semantic content of the meme<sub>1</sub> as a whole or the act of posting the meme), the notion of ‘memographic implicature’ will, generally, also lack semantic content. But memographic implicature and conversational implicature seem to be two varieties of a more general phenomenon of exploiting norms to convey meaning by flouting them.

Finally, we should note that, as indicated in E) above, the norms associated with a given meme<sub>CC</sub> might change over time. This will happen as people, either deliberately or haphazardly, introduce new norms or neglect existing ones without flouting them in memographic implicature. To recognize that the norms might change, however, does not imply that a given meme can continue to exist through *any* kind and amount of change. The conditions of identity of memes over time will be vague, often problematic, and not susceptible to rigorous treatment. In this respect, memes are like many other things that can undergo some but not any changes while persisting—languages, games, clubs, tables, and indeed just about anything!

### 3. Conclusion

I have offered answers to both the ontological and conceptual questions expressed by ‘what is a meme?’. I have further elaborated those answers and drawn out many of their consequences. The overall picture is somewhat complex. The initial distinction between the ontological and conceptual issues, the need to distinguish memes<sub>CC</sub> from memes<sub>1</sub>, the introduction of the meme<sub>CC</sub> *MEME*, the identification of several forms of ambiguity haunting the word ‘meme’—all of these make for a rather messy story in the end. But we are in the early days of the use of the term and the culture in which it is used is not all that amenable to constraint and precision. The account is unavoidably complex because the phenomena it seeks to account for are complex.<sup>25</sup>

Simon J. Evnine  
University of Miami  
[sevnine@miami.edu](mailto:sevnine@miami.edu)

---

25 Parts of this paper were presented at the University of Sheffield, Ben Gurion University of the Negev, Mumbai University, the Indian Institute of Technology Bombay, and the University of Maryland. The generous but critical reception in all these venues helped refine the work in progress. Detailed comments from two referees for this journal were enormously helpful. Since completing the paper, I discovered that the idea that a meme is made out of norms for producing instances had been anticipated by Anthony Cross in the blog post ‘The Curious Case of Pepe the Frog’ on the *Aesthetics for Birds* website (published online 26 January 2017, <<https://aestheticsforbirds.com/2017/01/26/the-curious-case-of-pepe-the-frog-by-anthony-cross>> accessed 24 May 2018).