



Taxonomising delusions: content or aetiology?

Peter Clutton, Stephen Gadsby & Colin Klein

To cite this article: Peter Clutton, Stephen Gadsby & Colin Klein (2017): Taxonomising delusions: content or aetiology?, *Cognitive Neuropsychiatry*, DOI: [10.1080/13546805.2017.1404975](https://doi.org/10.1080/13546805.2017.1404975)

To link to this article: <https://doi.org/10.1080/13546805.2017.1404975>



Published online: 20 Nov 2017.



Submit your article to this journal [↗](#)



View related articles [↗](#)



View Crossmark data [↗](#)



Taxonomising delusions: content or aetiology?

Peter Clutton ^a, Stephen Gadsby ^b and Colin Klein ^a

^aDepartment of Philosophy, Macquarie University, Sydney, NSW, Australia; ^bSchool of Philosophical, Historical and International Studies, Monash University, Melbourne, VIC, Australia

ABSTRACT

Introduction: Many theoretical treatments assume (often implicitly) that delusions ought to be taxonomised by the content of aberrant beliefs. A theoretically sound, and comparatively under-explored, alternative would split and combine delusions according to their underlying cognitive aetiology.

Methods: We give a theoretical review of several cases, focusing on monothematic delusions of misidentification and on somatoparaphrenia.

Results: We show that a purely content-based taxonomy is empirically problematic. It does not allow for projectability of discoveries across all members of delusions so delineated, and lumps together delusions that ought to be separated. We demonstrate that an aetiological approach is defensible, and further that insofar as content-based approaches are plausible, it is only to the extent that they implicitly link content to aetiology.

Conclusions: We recommend a more explicit focus on cognitive aetiology as the grounds for delusion taxonomy, even when that would undermine traditional content-based boundaries. We also highlight the iterative and complex nature of evidence about aetiologically grounded taxonomies.

ARTICLE HISTORY

Received 2 April 2017

Accepted 1 November 2017

KEYWORDS

Delusions; aetiology; taxonomy; belief; content

Introduction

The importance of taxonomy

Good theorising is aided by good taxonomy of the target domain. Mendeleev's periodic table laid the foundations for modern chemistry. Darwin's insights were not possible without the taxonomic work of generations of naturalists. Theorising and classification are not exclusive interests; scientific progress involves an iterative refinement of both (Wimsatt, 1976). Yet explicit taxonomic *principles* are valuable at each stage of that refinement.

So too in the neuropsychology of delusions. Theorising about delusions typically involves finding factors common to many cases. As Young writes,

The idea of psychiatric syndromes derives its power from the fact that a basic requirement is an agreed and reliable way of classifying the phenomena encountered. Psychiatrists therefore tend to look for syndromes which are defined by clusters of co-occurring symptoms. (2000, p. 55)

Two patients with Capgras syndrome will differ in innumerable ways. One might think it is his wife who has been replaced by an imposter, another her husband; one justifies this belief by pointing to the shape of the wife's eyes, the other the poor fit of the husband's clothes. Yet these differences are incidental to whether we group them together. Conversely, the two may have features in common with patients suffering from other delusions. Yet it is worth distinguishing patients suffering from different delusions, for they might have too little in common to make for fruitful theorising.

Much has been written about the principles which distinguish *delusions* from other, non-delusional states (Bayne, 2016; Bortolotti, 2010; Clutton & Gadsby, 2017; Miyazono, 2015), and also about the theories behind diagnoses of different mental disorders in line with clinical manuals like the DSM (Follette & Houts, 1996). Less attention has been paid to the general principles by which we might group and distinguish various delusions from one another, and assign particular patients to categories.¹ Much of what there is tends to be implicit and ad hoc. Consider, for example, the argument over whether Cotard's delusion (the belief that one is dead, rotting, etc.) and Capgras differ only in the patient's attributional style, wherein Cotard's patients internalise their feelings of unfamiliarity and Capgras patients externalise them onto others (Langdon & Coltheart, 2000; McKay & Cipolotti, 2007; Young, Leafhead, & Szulecka, 1994). This is (in part) a debate about whether Capgras and Cotard's patients should be lumped together as variants of the same delusion (or at least as species of the same genus). Yet it is hard to know what principles should guide an answer to this question.

Grounds and evidence

We will argue that cognitive neuropsychiatrists ought to taxonomise delusions by their *aetiology* rather than any other features. Most importantly, we claim, the *content* of a delusion should be ultimately irrelevant to its location in a finished classificatory system. Cotard's and Capgras delusions might be lumped together as the same disorder (or as species of the same genus) because of their shared aetiology. More importantly, delusions which involve otherwise identical contents might be split if they arise from different underlying causes: two patients who claim that their leg does not belong to them could have different disorders, despite having the same (delusional) belief.

We believe that this position follows from principles that are widely (though not universally) accepted. The consequences of these principles have not, however, been fully appreciated: many cognitive neuropsychiatrists working on delusions implicitly taxonomise delusions by a hybrid of content and aetiology. In part because of this, there are live debates over both taxonomy and theory that are difficult to solve. The goal of this paper is to encourage a shift towards a solely aetiological principle, and thereby make progress on specific issues.

Our argument is about the *grounds* for taxonomising delusions. That is, it is a story about what actually makes it the case that two patients ought to be grouped together as having the same delusional disorder (and, if there are superordinate groupings, how those go as well). It is important to distinguish this from a claim about the *evidence* that one has about a classification. Many things can furnish at least provisional evidence about a taxonomy of delusions, including facts about belief content. But content is only relevant insofar as it is a guide to aetiology: it is not grounds on its own.

This may seem like a recondite distinction. But it is mirrored by classificatory practice in other fields. A close analogy is perhaps with the cladistic approach to biological classification of species (O'Brien & Lyman, 2003; Sterelny & Griffiths, 1999).² Cladists admit only common ancestry as valid grounds for classification: in practical terms, monophyletic clades are the “real” groupings. Yet cladists use all sorts of evidence, including evidence about shared phenotypic and molecular traits, in order to infer facts about common ancestry. Indeed, for all but the coarsest distinctions among extinct species, shared phenotypic traits from the fossil record are likely to be all we have.

The important distinction between cladistics and phenetics (which classifies based only on similarities in phenotypic traits) or between cladistics and traditional taxonomy (which classifies based on *both* phenotypic traits and descent) is thus in the grounds for taxonomy, not in the evidence which is used to determine those relationships. Of course, the difference in grounds prescribes a different way of using the same evidence. Furthermore, and to prefigure an argument to which we will return, the relationship between evidence and cladistic hypotheses is itself complicated and iterative.

The job ahead

So too, we will claim, with the taxonomy of delusions. Any facts about symptoms and their similarities—including facts about the contents of beliefs—can (in principle) be relevant to sorting out aetiology. But these are ultimately just evidence. Furthermore, there is no separate taxonomy of symptoms themselves: symptoms are valuable as evidence and evidence alone. What we really want is a taxonomy of disorders that give rise to the various features of delusional symptoms (including belief content).

We assume throughout that there is at least one objective classification of delusions, and that it is the sort of thing that can be empirically discovered in the course of gaining clinical and research understanding of the processes that lead to delusions. There are other reasons that one might want to group or distinguish symptoms and other behavioural manifestations of delusions. That is common in medicine. Physiologists might want to distinguish wet and dry coughs because of their potential relation to different pathological causes. General practitioners might also want to distinguish coughs based on the further ill effects they cause (urinary incontinence, sleeplessness, etc.), as this is important for patient care (Chung et al., 2009). Similarly, we may have reason to distinguish delusions that are likely to give rise to (e.g.) violent behaviour, because these have important legal and social consequences. We think that these interest-driven classifications are important but of a different sort than the empirically driven classifications with which we are concerned. Briefly, in the former, the categories, relations, and membership conditions are completely given by our interests; the only scientific question is what fits where. For empirical taxonomies, by contrast, the categories themselves are part of what needs to be uncovered.

The argument to follow is somewhat involved, in part because there are three separate audiences that we address. We introduce what is to come by saying a few words about each.

First, there are readers who already find our claim unremarkable and cannot see what the fuss is about. According to strong versions of the “medical model”, psychiatry ought to follow general medicine in the shift from defining and classifying disorders based on

symptoms (which might incidentally arise through multitudinous pathways), to defining and classifying disorders based on unique causal pathways: each disorder is a distinct cause (some type of dysfunction) that explains the manifest symptoms (Murphy, 2009, pp. 107–108). This fits the picture we present here in regards to delusions. To these readers, the primary interest is likely to be our demonstration of the potential pitfalls of competing approaches, seen in the section “Taxonomic operations”.

Second, there are readers who are suspicious of the very idea of taxonomising mental disorders by underlying (dysfunctional) aetiology. Wakefield (1998, p. 847), for example, suggests that there is no interesting taxonomy of mental dysfunctions, because for any given function, “there are many ways it can fail, and there may be no systematic theory of these diverse breakdowns”. We are suspicious of this claim: breakdowns clearly lump together in interesting ways.³ Indeed, it is an important fact that delusions admit of an interesting taxonomy at all. It might have been the case that all delusions arose by a combination of ordinary cognitive processes combined with idiosyncratic life experiences (simple psychodynamic and behaviourist accounts arguably work like this). Yet this is not the case.

So part of our goal is to argue for the utility of a good taxonomy. We argue for this throughout, but see the sections “Two taxonomic principles” and “Fleshing out the taxonomic process” for some detailed consideration of these points.

The third, and likely most numerous, group, are those who think that content is part of a taxonomic ground—perhaps because they are coming from a tradition in which the syndrome is the functional unit (the so-called neo-Kraepelian approach (Murphy, 2009, p. 107)), or because they think that content and aetiology are on a par and intertwined; or because they are pluralists who think that there are a variety of equally plausible carvings of delusions (Chang, 2017, Rhodes, Jakes, & Robinson, 2005, p. 385). To such readers, we offer a two-pronged argument. First, in the section “Two taxonomic principles”, we argue that aetiology alone seems defensible as grounds for taxonomy. Second, in the section “Evidence and iteration”, we suggest that the attraction of pluralism stems from the undeniable role that content plays in an iterative process of classification. But that process should ultimately, when the full story is told, lead to a taxonomy grounded in aetiology.

Two taxonomic principles

Taxonomy by content

Many theorists talk in ways which presuppose that delusions ought to be delineated by the *content* of delusional beliefs. That is, what makes two patients count as having *somatoparaphrenia* (SP), rather than some other delusion, is just that they believe that their limb belongs to someone else.

Delusional content is typically what brings a patient to clinical attention and is often the most striking aspect of clinical cases. Yet delineation by content is common among cognitive neuropsychiatrists as well. As examples, consider:

- “Cotard’s syndrome is a rare condition of which the central symptom is a nihilistic delusion which, in its complete form, leads the patient to deny his own existence

and that of the external world” (Enoch & Trethowan, 1991, p. 162; taken from Young & Leafhead, 1996).

- “Patients with Capgras syndrome regard people whom they know well such as their parents or siblings as imposters” (Hirstein & Ramachandran, 1997, p. 437).
- “Somatoparaphrenia consists in abnormal or bizarre verbal reports about some parts of the body” (Cogliano, Crisci, Conson, Grossi, & Trojano, 2011, p. 758).
- “Fregoli syndrome is a disorder in which a person holds a delusional belief that different people are in fact a single person who changes his or her appearance or is in disguise” (Atta, Forlenza, Gujski, Hashmi, & Isaac, 2006, p. 58).

In each case, the type of delusion is determined by facts about the content of the unusual belief possessed by the patient. Furthermore, the grounds for debate in taxonomic disputes often involve appeal to similarities of content between different cases. For example, McKay and Ciolotti (2007, p. 356) distinguish Capgras from other delusions by using the subtle distinction between belief that the target is an *imitation* rather than a *duplicate*.

One might think that there are good reasons for this emphasis on content when it comes to taxonomy. After all, when a subject comes to clinical interest, surface features (i.e. content) are all that is available to taxonomise with. As such, we must concede that content clearly plays an indispensable role in the taxonomic process. Further still, when taxonomising disorders, content appears to be the more neutral principle. One-factor theorists (Maher, 1974), two-factor theorists (Coltheart, Langdon, & McKay, 2011), and predictive coding accounts (Corlett, Taylor, Wang, Fletcher, & Krystal, 2010; Gerrans, 2014) can all agree on how to pick out Capgras—by the particular strange belief it involves—while disagreeing on what causes it or even the type of causal explanation that would be appropriate. Even anti-doxastic views, which deny that delusions are beliefs, might find suitable hedges of the content principle in terms of assertions and the like (all Capgras patients assert...). Aetiological taxonomy, by contrast, would abandon the appearance of neutral ground.

That said, these *prima facie* advantages are not as strong as they seem. The idea that content provides a common, non-theory-laden starting point is problematic, for one still needs to choose which parts of the content are important (replacement of wife versus husband, replacement versus duplicate, etc.). When practitioners do this, we think that their choices will be theory-laden in some way (typically in relation to aetiology, where the relevant aetiology might be psychodynamic, life course trajectory, and so on). Indeed, the very idea of focusing on content is closely tied to psychodynamic approaches to delusions (Marková & Berrios, 1994): if one has a theory about the dynamics of love-hate or spousal relationships, then it makes sense that the relevant category in Capgras delusion focuses on the relation “loved one” or “spouse”, rather than on “any nearby living being with a face”, or “any nearby object”. Similarly, consider the proposed divisions of delusions categorised by whether or not a person thinks they deserve to be a target of harm (Trower & Chadwick, 1995). These are based on an assumption about personal-historical trajectory, that is, the division maps onto treatment in childhood (Rhodes, Jakes, & Robinson, 2005).

Furthermore, there are compelling arguments to the effect that a-theoretical taxonomising is not necessarily something to be strived for. Such a taxonomy may serve a

purpose in the immature stages of a science, but when it leads to endless proliferation, as we will suggest content sometimes does for delusions, then it is time to seek deeper causal principles by which to ground and constrain the taxonomy. That, according to some, has been the general process of numerous maturing sciences (Follette & Houts, 1996).

Also consider that taxonomy by content is somewhat problematic even on its own terms. On what logic can we group two instances of purported Capgras together based on their contents? Any two given cases will involve *some* difference in contents. At the very least they will refer to different people. Of course, one might reasonably respond that this is an obvious grouping: if two patients each believe that their respective *spouse* has been replaced with an imposter, they are clearly similar. But what of cases involving not a spouse or loved one, but a cat (Darby & Caplan, 2016), plants (Islam, Piacentini, Soliveri, Scarone, & Gambini, 2015), inanimate objects (Islam et al., 2015), or oneself in a mirror (Diard-Detoeuf, Desmidt, Mondon, & Graux, 2016)? Should these be grouped together as Capgras, as a separate group, or should each be distinguished individually? Perhaps we can group them all by abstracting further: these all involve the belief in replacement-by-imposter, and upon this basis we can group them. But why prefer that grouping to the more specific?

Our point is not that these questions do not have answers. Perhaps they do. Rather, we claim that the content itself gives us no principled reason for preferring any of these. If there was no alternative, perhaps that is where things would stand. But there is.

Taxonomy by aetiology

An aetiological taxonomy treats the underlying *cause* of a delusion as the important factor for its categorisation. The principles in favour of taxonomising by aetiology are very general ones that philosophers of science have appealed to for taxonomy across the sciences. A good taxonomic kind implies a stable cluster of powers which can manifest in a variety of different circumstances (Boyd, 1999; Khalidi, 2015). A good taxonomic kind ought to support projectable generalisations: finding out something about one instance of the kind should, at least sometimes, give you some reason to think it applies to the others (Kim, 1998). Projectable kinds are the meat and potatoes of science, because they allow for generalising from individual cases. It might have been the case that each delusion was *sui generis*, arising from idiosyncratic events in early childhood. If so, there would be little use in studying groupings of delusions as such. But the fact that there *are* generalisable patterns to be discovered about delusional groupings suggests that this is not the case. It is this focus on generalisability across cases that will ultimately guide our argument.

In aid of the task of taxonomy by aetiology, there is a rich literature devoted to fleshing out causal accounts of delusional beliefs (Antrobus & Bortolotti, 2016; Coltheart, 2007; Corlett et al., 2007, 2010; Gadsby, 2017; Kapur, 2003). For example, consider the claim that schizophrenia is caused by the presentation of stimuli as unduly salient, during the prodromal phase of psychosis (Kapur, 2003). As is, this serves as a causal explanation for the symptoms of an existing category of disorder—heightened salience *explains* how schizophrenic beliefs arise. However, under the taxonomic principle we are proposing, this processing of stimuli as hyper-salient helps (re)define the

very category “schizophrenic”. This holds for all types of causal explanations. Antrobus and Bortolotti (2016) recently discuss the hypothesis that depressive delusions arise from a distorted processing of information pertaining to the self. Subjects are biased towards attending to and encoding negative information about the self, whereas positive self-information is unable to be assimilated. Again, under our model this is more than just an *explanation* of how depressive delusions arise. The processing bias is a necessary condition for one to be categorised as suffering from depressive disorder (of the delusional subtype) in the first place.

Furthermore, it is possible to have delusions that involve the same content, but which involve distinct underlying causes. Intuitively, these ought to be kept separate in theorising. By way of example, consider cases of mirrored-self misidentification. This is typically classified as involving the belief that “The person I see when I look in the mirror isn’t me, but some stranger who looks like me” (Coltheart, 2007, p. 1043). However, researchers generally agree that this belief can arise through one of two possible cognitive deficits: either *mirror agnosia* (a loss of knowledge about how mirrors work) or a deficit in face processing (2007, pp. 1044–1045). Classified by belief content, Coltheart’s two patients are suffering from the same disorder. When classified by aetiology, they belong to two distinct categories. Indeed, mirrored-self misidentification due to deficits in face processing might potentially have more in common with Capgras syndrome than with mirrored-self misidentification due to mirror agnosia, despite similar belief contents. We suggest that this same principle should be applied to all disorders leading to delusional beliefs.

Interim conclusion

We have proposed that when it comes to taxonomising delusional disorders, aetiology is the more respectable option. In what follows, we consider the practical outcomes of classifying a delusional disorder in terms of belief content and argue that an aetiological taxonomy is preferable.

Of course, the argument above is not a knockdown. For example, one might subscribe to a pluralist view, on which both aetiology and belief content are grounds for taxonomy. We think that the attraction of pluralism is over-rated, and stems in part from running together grounds with evidential considerations (we return to this theme in the section “Evidence and iteration”). If content is meant to be an independent ground for taxonomy, then theories about what counts as the same delusional content cannot ultimately depend on aetiological considerations. That is, we cannot say that beliefs about spouses and parents being impostors should be counted together because they stem from the same underlying deficit in face processing—for that is ultimately to ground via aetiology, not content. But as we have suggested, without referring to aetiology, there is no basis on which to prefer one content carving over another. Content by itself admits of arbitrary groupings, and cannot independently ground the kinds of inductive generalisations we want from a taxonomy.

In the next two sections, we work through a range of delusional examples to demonstrate our point, focusing on instances where theorists have either split or lumped together delusional examples based on various principles. We show that a focus on content can often lead us astray, and aetiology might provide a superior organising principle.

Taxonomic operations

Lumping

In the section “Taxonomy by aetiology”, we considered the case of mirrored–self misidentification, whereby aetiology distinguishes two different conditions that can lead to beliefs with the same content. Even if one is inclined towards taxonomising by aetiology, one might consider this a merely semantic quibble. Why care whether there is one condition or two?

In fact, however, this sort of question makes a considerable difference when one turns to the *inferences* that one can make from delusional disorders. A key reason to taxonomise is to promote projectability: that is, we would like empirical discoveries about some patients with a disorder to be a guide to what we can expect about other patients with the same disorder. If we decide that mirrored–self misidentification arises from a disorder of face processing in addition to a general deficit in belief evaluation (on a two-factor view), then we could make novel predictions about other circumstances that might lead such patients to delusions (Breen, Caine, & Coltheart, 2001). We might predict that such patients would also form delusions when looking at themselves in a family photo album, for example.⁴ In actual practice, the relation between such delusional symptoms and the underlying disorder is complex. Yet the point of making taxonomic claims is to find stable groups that (at least sometimes) support projectable inferences.

Furthermore, cognitive neuropsychiatrists are interested in delusions not just because of what they predict, but because of what these disorders can show us about the structure of ordinary cognition (Coltheart, 2007). Many of these inferences depend upon both association and dissociation of symptoms across distinct patients. So, for example, if we find a group of patients with delusion *X* that have feature *A* and some other patients with delusion *X* that have feature *B*, we are (defeasibly) entitled to infer that patients with *X* have *A* and *B*. But, crucially, that inference is only valid if our set of patients belong to the same taxonomic category.

Consider, by way of example, a recent debate over the perception of pain by patients with somatoparaphrenia (SP). In paradigm cases of SP, patients assert that one of their limbs belongs to someone else. SP typically presents as a monothematic delusion, in that patients express strange beliefs only about this one topic (their limbs) and appear otherwise to have ordinary cognitive capacities.

Briefly, Klein (2015a, 2015b) argues that motivation is an intrinsic feature of pain, and that motivation and pain come apart (always and only) in relatively rare instances where there is lack of care for the body. de Vignemont (2015) offers SP as a counterexample: SP patients both feel pain and are motivated by it, despite their lack of care for the affected limb. Klein (2017) responds that this is not in fact a counterexample if SP is a form of restricted depersonalisation, characterised by a lack of care. Complicating matters, both have empirical cases that appear to fit their mould. Crucially, both sides accept that SP is delineated by the content of belief, and that *therefore* there must be some other shared set of features (motivation, depersonalisation, and so on) in response to pain for each patient with that belief content.

Yet there is actually considerable evidence that the aetiology of SP is heterogenous: that is, different conditions can converge on the same belief. Many implicate paralysis as the first factor in SP (Coltheart, Langdon, & McKay, 2007, p. 644; Coltheart, 2007, p. 1047;

Rahmanovic, Barnier, Cox, Langdon, & Coltheart, 2012), and paralysis is near-universal in Vallar and Ronchi's (2009) review of cases. Yet temporary remission of delusional symptoms (e.g. by caloric vestibular stimulation) can be achieved without affecting paralysis (Bisiach, Rusconi, & Vallar, 1991; Fotopoulou et al., 2011; Grabherr, Gianluca, & Bigna, 2015; Rode et al., 1992; Salvato et al., 2016; Schiff & Pulver, 1999; Spitoni et al., 2016; Van Stralen, Van Zandvoort, & Dijkerman, 2011). Others relate SP to a breakdown in a *sense of bodily ownership* over the affected limb (de Vignemont, 2011; Van Stralen, van Zandvoort, Kappelle, & Dijkerman, 2013). Yet "sense of bodily ownership" is likely itself a heterogeneous phenomenon, subject to distinct breakdowns (Kilteni, Maselli, Kording, & Slater, 2015). For example, breakdown might occur due to a failure at the stage of multi-sensory integration (Vallar & Ronchi, 2009, p. 548) or through a deficit in the signals themselves (Paqueron et al., 2003). A final set of case reports link SP to depersonalisation (Critchley, 1953; Dieguez, Staub, & Bogousslavsky, 2013; Schilder, 1935/1950).

Thus there are a large array of aetiologies that appear to plausibly give rise to the belief that one's limb is not one's own. The problem may simply be insufficient cleverness: there is a common factor, but it has either not been found or else not described in a sufficiently perspicacious way. An alternative—and one which we find more plausible—is that the class *itself* is the problem. That is, delusions which involve the same coarse-grained propositional content can arise from distinct aetiologies. According to the aetiological view, then, there is not one delusion, SP, which arises from different causes. Rather, there are different delusions (the SPs, if you will), which happen to coincide in their associated propositional content.

This has real scientific upshot. We noted above that, in the debate between Klein and de Vignemont, both appeared to accept a content-based taxonomy. Yet the argument may actually be moot. If the SP patients who feel pain are also depersonalised, and the non-depersonalised patients also do not feel pain, then there is no *single* group of patients who could constitute a counterexample. The distinct aetiologies we have identified for SP might well be expected to result in different relationships to pain.

Splitting

The section "Two taxonomic principles" considered a case where the content approach might lump together a class which ought to be split. The opposite is possible as well. Taxonomising according to content rather than aetiology can lead to instances where one might proliferate the number of delusions unnecessarily, because of their superficially different contents.

Proliferation is as bad as mistaken lumping. Overly fine-grained taxonomies which do not include superordinate categories overlook useful generalisations. Treating as wholly distinct the delusion that one's *husband* has been replaced from the delusion that one's *wife* has been replaced would obscure the important similarities that make delusions worth studying in the first place.

Furthermore, needlessly fine-grained taxonomies risk conflating properties that define a delusion with properties that are idiosyncratic to a patient. Most delusional beliefs have a structure where a common core is elaborated by individual circumstances. Focusing on

content can result in needless proliferation, driven by these theoretically superfluous idiosyncrasies (Marková & Berrios, 1995).

Consider De Clérembault's syndrome, sometimes called erotomania. It was originally described by De Clérembault as involving a delusion in which the patient believes that someone is secretly in love them. In the original descriptions of the malady, the sufferer is usually a woman, and the man in question is older, usually of some higher social standing, often either famous or in a prominent public position. This belief often leads to unwanted attention—such as stalking—directed towards the non-reciprocating object of the delusion. This description of the delusion is still common in the literature today (Brüne, 2001; Olojugba, de Silva, Kartsounis, Royan, & Carter, 2007).

This description incorporates some rather specific content: the target is an older person, a man, and most intriguingly, someone of higher social standing. Why should *these* specific features pick out a particular kind of delusion? Indeed, case reports routinely flaunt these criteria. De Clérembault's is reported in men, for starters. Sometimes the definition of famous (prominent/of higher standing) is stretched, and sometimes completely discarded: a teacher at a local conservatory, a friend met on holiday, or even “a man living in California” have all been reported (Doust & Christie, 1978, p. 103).

This is an instance where a focus on content might have led to an overly specific classification of a delusion (Ellis & Mellsop, 1985). In fact, there are good reasons why we might suspect that the highly specific definition of this delusion is the result of factors entirely *unrelated* to the underlying nature of the disorder itself. For a start, we should consider the fact that De Clérembault's is often brought to clinical attention via legal processes that arise from the unwanted behaviour of the sufferer. One could see why—especially during the time-period of the original formulation of the disorder—it might be common for harassment of older men (and people of high social status more generally) to be reported, and for those reports to be actioned. This selection bias of the patient group could have led to the clinical observation that the condition mostly involves male (and high social status) targets, even though versions of the disorder in the non-clinical population, with the same causal features, may not have this gender- (and status-) based skew.⁵

Similarly, part of the way we recognise delusions seems to be the way in which delusions violate folk-epistemological norms (Murphy, 2006, 2012). A friend who is obsessed with a favourite singer and reads personal declarations of love into her Top 40 hits is troubled: things do not work that way. By contrast, if your friend makes similar comments about the woman at the coffee cart, her secret love expressed by a heart-frothed milk and a smile as she calls your friend's name, you might shrug this off as a harmless crush. The former more strongly violates our sense of the mechanisms by which beliefs can be formed, even if both beliefs actually stem from the same underlying cause.

Another area worth examining involves delusions that are contingently linked to a particular time, place, or culture. As Gold and Gold (2014) note, The Truman Show Delusion appears to depend on a popular movie with the same name (Gold & Gold, 2014). Delusions about surveillance similarly keep track with modern technology. The contingency of the contents of these delusions is clear. As Marková and Berrios (1995) point out, content sits at such remove from the underlying aetiology that it is often of dubious use as a classifying principle. Of course, many theorists would not advocate for taxonomising by content in cases like these: at the very least, most would offer some kind of generalising

categories (persecutory, misidentification, etc.). But on the position we propose, the important point is that we should exercise care when allowing proliferation of delusional classification based on contents like this, and look to the causal processes to provide the organising principles to impose order on the various contents.

Final thoughts on Capgras

The aetiological approach we have proposed has played a role in some contemporary discussions of delusions. Capgras is an interesting case in this regard. Many delusions have content which appears similar to typical Capgras (replacement by imposter), but where the question arises as to whether they really should be classed as Capgras. For instance, there have been cases of what has been humorously termed “Cat-gras” or misidentification concerning pets (Darby & Caplan, 2016). Other examples have been reported involving other animals, plants, and inanimate objects (Islam et al., 2015). Finally, one curious example involved what might sound initially like mirrored-self misidentification, but where the subject thought his reflection was not a stranger but a double of himself, and which the authors classified as a type of Capgras (Diard-Detoeuf et al., 2016).

These cases are so readily brought together that it is worth pointing out that there is nothing in the content which demands it. Beliefs about spouses and pets are importantly different in many contexts (e.g. pets are typically replaced more frequently than spouses). Psychodynamic theories of delusions would find ample reason to distinguish them. What grounds are there for abstracting away to a category like “beloved person” or even “beloved object”?

We have proposed that it is aetiology that fills this gap. In the case of so-called Cat-gras, for example, Darby and Caplan (2016) argue that this delusion arises from a failure to link up a current percept with an autobiographical memory of whatever is being perceived (in this case, the cat), thus resulting in misidentification. They go on to suggest that this is true of all the misidentification syndromes, including Capgras, Fregoli, etc. If this were true, it would point to an even more coarse-grained taxonomy than has usually been considered in this area, placing together the misidentification syndromes generally, or at least the set of them that arise from such causes, with little taxonomic use in distinguishing misidentification of loved ones, friends, animals, plants, and objects. The authors indeed argue along these lines, proposing that one mechanism should explain the misidentification syndromes (Darby & Caplan, 2016, p. 255). While one may take issue with the particular proposal, Darby and Caplan’s reasoning—moving from what we know of content, down to the neural/cognitive bases, and then back up to re-group and taxonomise based on causal knowledge—is exactly the kind of process we advocate in this paper.

Fleshing out the taxonomic process

We have argued in favour of an aetiological approach to the taxonomy of delusions. Such an approach taxonomises delusional disorders in terms of their underlying causes. As with any causal account, complex issues arise when we turn to the actual practice of disentangling causes. Causation is itself a philosophically complex topic. Insofar as there are difficulties for the aetiological account, they fall into three broad categories. None bar an aetiological account, but deserve special attention lest they detract from our story.

General causal issues

Any aetiological account must disentangle taxonomically relevant causes from mere correlates and background conditions. These are very general problems which arise when trying to disentangle any sort of complex causal relationships. We believe that the philosophical literature on causation in psychology has hit upon a rough consensus on how to solve these problems.

We adopt an *interventionist* account of causation. (Craver, 2007; Woodward, 2003). On this view, causation is expressed as an invariant relationship between variables under intervention on one of them. Causes are thus *difference-makers*: (some) ways of manipulating the value of a cause result in a change to the effects. Note that “intervention” is used in a broad sense. Interventions are not limited simply to medically available interventions: as far as disentangling causation goes, all that matters is the *possibility* of some manipulation that is consistent with the laws of nature.

In the case of delusions, a factor—a pattern of brain activity, a cognitive mechanism of belief formation, a neurotransmitter level, a gene—is aetiologically relevant to a delusion just in case there is some way to change it while keeping everything else fixed, which would thereby change whether or how the patient has the delusion. This is a relatively permissive account of causation. Most of the time, we are interested in the *actual difference-makers* for a delusion: that is, the factors that actually vary in the population, which give rise to the delusion in question (Waters 2007).

Interventionist accounts are powerful in part because difference-making is a fundamentally contrastive notion (Woodward 2003). So for example, we might care why a patient has Capgras’ delusion *rather than* some other delusion, or why he is delusional rather than merely possessed of an odd perceptual experience, or even why he has delusional beliefs rather than no beliefs at all. Each contrast gives a distinct question with a distinct answer.

The final question asks about general background conditions—like being alive—that are necessary for delusions because they are necessary for cognition generally. Mere background conditions are indicated by strongly asymmetric and nonspecific patterns of dependence (Craver, 2007; Woodward, 2010). The heart plays a causal role in all delusions by keeping patients alive. Yet there is no way to intervene on hearts that affects delusions *alone*, and the only delusion-relevant interventions on hearts are crude and nonspecific ones. That is why we know that the presence of a heart is a mere background condition, though it is part of the aetiology of delusions.

Focus on the other two contrasts shows how we might arrive at a hierarchical taxonomy of aetiological factors. Factors common to all delusions distinguish delusional beliefs from others, whereas specific factors differentiate particular delusions from one another. To see how this might work in an idealised example, suppose we have two delusional patients expressing the belief that their spouse has been replaced by an imposter. Suppose that through an iterative process (discussed further in the section “Evidence and iteration”), we have formulated a two-factor theory about this delusion. Our proposed first factor is a breakdown in face-processing function, with a general belief formation deficit as a second factor. The assumption at this stage is that both patients share the same type of delusion.

Now suppose we could make a targeted intervention on the relevant cognitive function, perhaps through the use of direct electrical stimulation aimed at the fusiform face area.⁶

We would expect that such an intervention would modulate the delusional symptom: modulating a cause should modulate the effect. Depending on the exact nature of the intervention, the resulting modulation might be a temporary remittance of the delusion (since now the general background deficit in belief formation is not being fed any aberrant sensory data), or perhaps, for more *recherché* interventions, the modulation of various dimensions of the delusion (conviction, content, intensity, etc).

However, if this intervention modulated one patient's delusion but not the other, it would suggest a taxonomic split into two aetiologically distinct delusions, with face processing being a difference-maker in only one of these cases. On the view we have put forward, these are the kinds of data that would ground a taxonomy of delusions. Conversely, if the same intervention led to the remittance of symptoms in a wide variety of delusional patients, we might conclude that it targets only factors common to all delusions.

In sum, interventions on aetiological factors can be used to determine the *particular* role that a factor plays. These causal distinctions can be further used to uncover hierarchical structure amongst aetiological factors, and thereby refine the categories that we already possess.

Levels of causation

The second set of issues are those which arise when we think about causation at different levels of description. Causes may be distinguished at many different levels of description; "aetiology" serves as a placeholder for whatever level proves most fruitful.

Broadly speaking, we think that the natural place to look for aetiological causes is at the cognitive level. Delusions should be classed together if they arise from breakdowns in similar cognitive functions, howsoever those functions are realised in the brain.⁷ This is arguably the approach to delusions embodied in current versions of the two-factor theory. It has its roots in neuropsychology, particularly of the ultra-cognitive sort (Langdon & Coltheart, 2000). On such a view, delusions are interesting because they show the consequences of breakdowns of specific cognitive modules.

We emphasise the cognitive level because of the empirical successes of cognitive neuropsychiatry. Several alternatives exist. The two most familiar would be delineation by *neural* aetiology and by *molecular* aetiology. On the former, delusions should be counted as the same if they arise from (e.g.) damage to the same brain region. On the latter, delusions should be counted as the same if they involve (e.g.) the same breakdowns in genetic expression or neurotransmitter levels. Call these options collectively the *sub-cognitive* level.

Sub-cognitive aetiologies have some intuitive attraction. Everyone in the debate thinks that delusions must ultimately arise from *some* organic cause. Whatever organic dysfunction gives rise to the delusion will be a difference-maker for the delusion, in the sense that some suitable manipulation of that organic state will let one manipulate the cognitive state itself. So sub-cognitive aetiologies have an advantage of always being present and playing the relevant role.

Yet sub-cognitive aetiologies have obvious downsides. For one, it is not clear what would count as a "similar" neural aetiology. The brain admits of multiple carvings along a wide variety of properties. What counts as the same region is something of an

arbitrary choice: one might choose anatomical landmarks, or cytoarchitectural features, or common blood supplies, or any other way of parcelling up the brain. With any of these, the boundaries will be vague. Nor is one restricted to regions: one might choose spatially connected networks, or functionally connected networks, or matrices of similar cell types, or any other way of connecting up diverse parts that has been found scientifically useful. The point is not (merely) that it is hard to draw a non-arbitrary carving, though that is a problem. The real problem is that in most cases, our brain taxonomies are themselves driven by functional considerations across a variety of scales (Bassett & Gazzaniga, 2011).

So suppose two patients have lesions a centimetre apart. Should we count them as having the same delusion? The answer would seem to depend, at least at first pass, on what the lesioned parts are *doing*: they ought to be grouped together just in case they have the same function.⁸ But that is just to return to something like the cognitive sense of aetiology. Conversely, causes at the genetic and molecular levels have not, in practice, been specific enough: few if any disorders map neatly onto such factors.

That said, we emphasise that this is ultimately an empirical question. Carving up cognitive faculties is itself a complex problem (as we emphasise in the next section). Furthermore, it may be that there are other levels of description that best capture some delusions, such as the phenomenological (Clutton, *in press*).

Evidence and iteration

A third and final problem concerns the actual practice of classification. Recall that our goal has been to give a story about what grounds an appropriate taxonomy. If we had a complete story about the aetiologies of delusions, sorting out their taxonomy would be straightforward. We do not. Worse, taxonomies are interesting in part *because* they help to build theories about the causes of delusions themselves: so if we had a complete story about aetiologies, taxonomy would arguably be of less interest.

Again, we emphasise that there is an important distinction between the *grounds* of a taxonomy on the one hand and the *evidence* about appropriate taxonomies on the other. We have suggested that a taxonomy of delusional beliefs ought to ultimately be based on aetiological grounds, but all manner of evidence might be appropriate to use in discovering those grounds. Furthermore, the actual practice of taxonomising is likely to have a complex character: taxonomic claims are often intermediate hypotheses, which are used as the basis of experiments and refined as more evidence comes in (Wimsatt, 1976).

This iterative process is common in other sciences. Even basic physical properties like temperature were the product of iterative refinement (Chang, 2004). As we noted in the section “Grounds and evidence”, cladistic methods involve continual refinement of hypotheses about common ancestry. This has a straightforward looping structure. Cladistic methods in biology require initial hypotheses (in the form of outgroup choices for rooting trees) which can be refined by further exploration (O’Brien & Lyman, 2003). Furthermore, the evidence base itself is complex and iterative. Homologous phenotypic traits, as opposed to non-homologous similarities, are the proper evidence about systematic groupings of species. But whether we consider a trait homologous across species itself depends on how we think those species are related. This is not a vicious circle: the whole practice of systematics developed precisely to find principled ways to sort out such relationships.

So too, we expect, in the taxonomy of delusions. Though aetiology ought to be the ultimate grounds of our taxonomy, a variety of evidence—including delusional content—may be relevant along the way. Consider that content can often give us clues about underlying aetiology. When delusions represent the world as more desirable, for example, this gives us reason to think motivational factors may be involved.⁹ Insofar as content may be relevant, we think that some elements are likely to be less so than others. Hyper-specific content, like that seen in the definition of De Clerebault's, is likely to hinder more than it helps, obscuring potential aetiological clues with the clutter of contingent content. Where we have reason to think some proposed content might be completely unrelated to aetiology, we should be especially willing to revise the category by bringing to bear whatever relevant aetiological knowledge we may have.

Appreciation of this point brings insight to the issue of pluralism. Content-based groupings are useful, but only in so far as content is driven by aetiology. Content by itself admits of entirely arbitrary groupings, and cannot ground the kinds of inductive generalisations one desires of a scientific taxonomy.

Future directions

Several outstanding questions remain. For example, we have presented taxonomy as primarily a matter of delineating a large number of distinct delusions. But delusions likely form *hierarchies* as well, and it is unclear how to sort out hierarchical questions. For example, ought we consider Capgras and certain kinds of mirrored-self misidentification the same delusion, as suggested above? Or might the relationships be more complicated, for example, different species of an overarching genus?

We think that the aetiological account has the resources to answer these questions. Indeed, causal factors can stand in a similar hierarchical relationship: one might consider a broad class of sensory-limbic disconnection syndromes as falling under the same genus, for example, with particular disconnections giving the species. That said, the possibility of hierarchy complicates some of the arguments given above. For example, it remains unclear whether one ought to take the evidence about the diversity of causes of SP as indications that one should absolutely distinguish multiple classes, or instead find some superordinate causes under which more specific aetiologies might be organised.

That said, we conclude by noting that these are still ultimately questions about how to think *about* aetiologies, rather than problems with the framework as such. If anything, such problems re-emphasise the need to taxonomise according to aetiology rather than content. The content of a delusional belief is a bit like a nonspecific symptom in other parts of medicine. A hacking cough may give some clue to what has gone wrong—but many different diseases can give rise to hacking coughs, and the same disease can cause hacking cough in one person but not another. We have to start somewhere, and content is the most easily accessible place to start. But having begun the taxonomic project, we should expect to kick away the ladder and focus exclusively on causes.

Notes

1. For some examples of previous work, see Marková and Berrios (1994) and Rodrigues and Banzato (2006). We endorse this kind of overt discussion of the taxonomy of individual

delusions and aim to advance the discussion with some theoretical and practical principles regarding the use of the aetiological approach we advocate.

2. When discussing biological species, “taxonomy” is often used to refer to methods other than the cladistic. We follow that convention when discussing cladistics.
3. Consider the analogy that Wakefield draws to the process of automobile repair. Wakefield claims that auto repair manuals are “strikingly unparsimonious” because “problems do not hang together the way the car’s parts do” (1998, p. 847). Yet standard repair manuals *do* organise symptoms by the failures of the systems which cause them (Choate & Haynes, 1992, 20ff).
4. As Breen et al. (2001, p. 250) point out, their patient did not in fact have difficulty recognising faces in photographs, which led them to suggest that there might be something functionally unique about recognising faces in the mirror.
5. Our point here can be more succinctly stated: the content definition of De Clérembault’s may suffer from Berkson’s (1946) fallacy, in which a contingent factor that influences which patients are brought to clinical attention is incorrectly taken as an important part of the disorder.
6. Note that we do not currently have sufficient understanding or sufficient technology to do so. The point is illustrative: it shows the *kind* of evidence that would be relevant.
7. Note that when we refer to cognitive functions and their breakdown, we have in mind a causal role sense of “function” rather than the role that something was selected to do. So for example, pure alexia is undoubtedly a breakdown of some of the functional underpinnings of visual word recognition, even if the underlying machinery was not evolutionarily selected for reading. For aficionados, we here use “function” in the sense of Cummins (1975) rather than of Wright (1973), though we are ultimately pluralists about this issue (Godfrey-Smith, 1993).
8. Of course, determining whether two brain regions have the same function is a nontrivial task in its own right; further, the best taxonomies of brain regions for many neuroscientific disciplines may not be in terms of function but (e.g.) connectivity or cytoarchitecture. The point is merely that a relatively “pure” neural classification for delusions seems somewhat rare. In practice, lesions are grouped together based on their functional effects, rather than the other way around. Thanks to an anonymous reviewer for pressing us to clarify this point.
9. Thanks to an anonymous reviewer for this point.

Acknowledgements

Thanks to Tim Bayne and Max Coltheart as well as a reading group in philosophy and the belief formation Group in the CCD for helpful feedback on previous drafts. A portion of this content developed out of material presented at “The Value of Suffering” workshop, jointly funded by the Value of Suffering Project and the Macquarie University Research Centre for Agency, Values, and Ethics (CAVE).

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

This work was partially supported by the Australian Research Council [grant number FT140100422 (to C.K.)], an Australian Government Research Training Program (RTP) Scholarship (to P.C and S.G), and membership of the ARC Centre of Excellence in Cognition and its Disorders (P.C).

ORCID

Peter Clutton  <http://orcid.org/0000-0002-2040-5107>

Stephen Gadsby  <http://orcid.org/0000-0002-2302-844X>

Colin Klein  <http://orcid.org/0000-0002-7406-4010>

References

- Antrobus, M., & Bortolotti, L. (2016). Depressive delusions. *Filosofia Unisinos*, 17(2), 192–201.
- Atta, K., Forlenza, N., Gujski, M., Hashmi, S., & Isaac, G. (2006). Delusional misidentification syndromes: Separate disorders or unusual presentations of existing DSM-IV categories? *Psychiatry (Edgmont)*, 3(9), 56–61.
- Bassett, D., & Gazzaniga, M. (2011). Understanding complexity in the human brain. *Trends in Cognitive Sciences*, 15(5), 200–209.
- Bayne, T. (2016). Delusions and the norms of rationality. In T.-W. Hung, & T. Lane (Eds.), *Rationality: Constraints and contexts* (pp. 77–94). Amsterdam: Elsevier.
- Berkson, J. (1946). Limitations of the application of fourfold table analysis to hospital data. *Biometrics Bulletin*, 2, 47–53.
- Bisiach, E., Rusconi, M. L., & Vallar, G. (1991). Remission of somatoparaphrenic delusion through vestibular stimulation. *Neuropsychologia*, 29, 1029–1031.
- Bortolotti, L. (2010). *Delusions and other irrational beliefs*. Oxford: Oxford University Press.
- Boyd, R. N. (1999). Kinds, complexity, and multiple realization. *Philosophical Studies*, 95, 67–98.
- Breen, N., Caine, D., & Coltheart, M. (2001). Mirrored-self misidentification: Two cases of focal onset dementia. *Neurocase*, 7, 239–254.
- Brüne, M. (2001). De clerambault's syndrome (erotomania) in an evolutionary perspective. *Evolution and Human Behavior*, 22(6), 409–415.
- Chang, H. (2004). *Inventing temperature: Measurement and scientific progress*. New York, NY: Oxford University Press.
- Chang, H. (2017). Epistemic iteration and natural kinds: Realism and pluralism in taxonomy. In K. S. Kendler, & J. Parnas (Eds.), *Philosophical issues in psychiatry IV: Classification of psychiatric illness* (pp. 229–245). Oxford: Oxford University Press.
- Choate, C., & Haynes, J. H. (1992). *Ford full-size vans 1969-1991*. Haynes Automotive Repair Manual Series, &hash;36090. Newberry Park, CA: Haynes North America.
- Chung, K. F., Bolser, D., Davenport, P., Fontana, G., Morice, A., & Widdicombe, J. (2009). Semantics and types of cough. *Pulmonary Pharmacology & Therapeutics*, 22(2), 139–142.
- Clutton, P. (in press). A new defence of doxasticism about delusions: The cognitive phenomenological defence. *Mind & Language*.
- Clutton, P., & Gadsby, S. (2017). Delusions, harmful dysfunctions, and treatable conditions. *Neuroethics*, 1–15.
- Cogliano, R., Crisci, C., Conson, M., Grossi, D., & Trojano, L. (2011). Chronic somatoparaphrenia: A follow-up study on two clinical cases. *Cortex*, 48, 758–767.
- Coltheart, M. (2007). The 33rd Sir Frederick Bartlett lecture cognitive neuropsychiatry and delusional belief. *The Quarterly Journal of Experimental Psychology*, 60(8), 1041–1062.
- Coltheart, M., Langdon, R., & McKay, R. (2007). Schizophrenia and monothematic delusions. *Schizophrenia Bulletin*, 33(3), 642–647.
- Coltheart, M., Langdon, R., & McKay, R. (2011). Delusional belief. *Annual Review of Psychology*, 62, 271–298.
- Corlett, P. R., Murray, G. K., Honey, G. D., Aitken, M. R., Shanks, D. R., Robbins, T. W., ... Fletcher, P. C. (2007). Disrupted prediction-error signal in psychosis: Evidence for an associative account of delusions. *Brain*, 130(9), 2387–2400.
- Corlett, P. R., Taylor, J. R., Wang, X. J., Fletcher, P. C., & Krystal, J. H. (2010). Toward a neurobiology of delusions. *Progress in Neurobiology*, 92(3), 345–369.
- Craver, C. F. (2007). *Explaining the brain: Mechanisms and the mosaic unity of neuroscience*. Oxford: Oxford University Press.

- Critchley, M. (1953). *The parietal lobes*. New York, NY: Hafner Publishing Company.
- Cummins, R. (1975). Functional analysis. *Journal of Philosophy*, 72(20), 741–765.
- Darby, R. R., & Caplan, D. (2016). “Cat-gras” delusion: A unique misidentification syndrome and a novel explanation. *Neurocase*, 22(2), 251–256.
- de Vignemont, F. (2011). Embodiment, ownership and disownership. *Consciousness and Cognition*, 20, 82–93.
- de Vignemont, F. (2015). Pain and bodily care: Whose body matters? *Australasian Journal of Philosophy*, 93(3), 542–560.
- Diard-Detoeuf, C., Desmidt, T., Mondon, K., & Graux, J. (2016). A case of Capgras syndrome with one’s own reflected image in a mirror. *Neurocase*, 22(2), 168–169.
- Dieguez, S., Staub, F., & Bogousslavsky, J. (2013). Asomatognosia: Disorders of the bodily self. In O. Godefroy, & J. Bogousslavsky (Eds.), *The behavioral and cognitive neurology of stroke* (Vol. 2, pp. 215–253). Cambridge: Cambridge University Press.
- Doust, J. L., & Christie, H. (1978). The pathology of love: Some clinical variants of de Clerambault’s syndrome. *Social Science & Medicine. Part A: Medical Psychology & Medical Sociology*, 12, 99–106.
- Ellis, P., & Mellsop, G. (1985). De Clarembault’s syndrome—a nosological entity? *The British Journal of Psychiatry*, 146(1), 90–93.
- Enoch, M. D., & Trethowan, W. H. (1991). *Uncommon psychiatric symptoms* (3rd ed.). Oxford: Butterworth-Heinemann.
- Follette, W. C., & Houts, A. C. (1996). Models of scientific progress and the role of theory in taxonomy development: A case study of the DSM. *Journal of Consulting and Clinical Psychology*, 64(6), 1120–1132.
- Fotopoulou, A., Jenkinson, P., Tsakiris, M., Haggard, P., Rudd, A., & Kopelman, M. (2011). Mirror-view reverses somatoparaphrenia: Dissociation between first- and third-person perspectives on body ownership. *Neuropsychologia*, 49(14), 3946–3955.
- Gadsby, S. (2017). Anorexia nervosa and oversized experiences. *Philosophical Psychology*, 30(5), 594–615.
- Gerrans, P. (2014). *The measure of madness: Philosophy of mind, cognitive neuroscience, and delusional thought*. Cambridge, MA: The MIT Press.
- Godfrey-Smith, P. (1993). Functions: Consensus without unity. *Pacific Philosophical Quarterly*, 74, 196–208.
- Gold, J., & Gold, I. (2014). *Suspicious minds: How culture shapes madness*. New York, NY: Free Press.
- Grabherr, L., Gianluca, M., & Bigna, L. (2015). The moving history of vestibular stimulation as a therapeutic intervention. *Multisensory Research*, 28(5-6), 653–687.
- Hirstein, W. S., & Ramachandran, V. S. (1997). Capgras syndrome: A novel probe for understanding the neural representation of the identity and familiarity of persons. *Proceedings of the Royal Society B: Biological Sciences*, 264, 437–444.
- Islam, L., Piacentini, S., Soliveri, P., Scarone, S., & Gambini, O. (2015). Capgras delusion for animals and inanimate objects in Parkinson’s disease: A case report. *BMC Psychiatry*, 15(1), 73.
- Kapur, S. (2003). Psychosis as a state of aberrant salience: A framework linking biology, phenomenology, and pharmacology in schizophrenia. *American Journal of Psychiatry*, 160(1), 13–23.
- Khalidi, M. A. (2015). Natural kinds as nodes in causal networks. *Synthese*, 1–18. doi:10.1007/s11229-015-0841-y
- Kilteni, K., Maselli, A., Kording, K. P., & Slater, M. (2015). Over my fake body: Body ownership illusions for studying the multisensory basis of own-body perception. *Frontiers in Human Neuroscience*, 9, 141. doi:10.3389/fnhum.2015.00141
- Kim, J. (1998). *Mind in a physical world*. Cambridge, MA: The MIT Press.
- Klein, C. (2015a). What pain asymbolia really shows. *Mind; A Quarterly Review of Psychology and Philosophy*, 124(494), 493–516.
- Klein, C. (2015b). *What the body commands: The imperative theory of pain*. Cambridge, MA: MIT Press.

- Klein, C. (2017). Pain, care, and the body: A response to de Vignemont. *Australasian Journal of Philosophy*, 95(3), 588–593.
- Langdon, R., & Coltheart, M. (2000). The cognitive neuropsychology of delusions. *Mind & Language*, 15(1), 184–218.
- Maher, B. A. (1974). Delusional thinking and perceptual disorder. *Journal of Individual Psychology*, 30(1), 98–113.
- Marková, I. S., & Berrios, G. E. (1994). Delusional misidentifications: Facts and fancies. *Psychopathology*, 27(3-5), 136–143.
- Marková, I. S., & Berrios, G. E. (1995). Mental symptoms: Are they similar phenomena? *Psychopathology*, 28(3), 147–157.
- Mckay, R., & Cipolotti, L. (2007). Attributional style in a case of Cotard delusion. *Consciousness and Cognition*, 16(2), 349–359.
- Miyazono, K. (2015). Delusions as harmful malfunctioning beliefs. *Consciousness and Cognition*, 33, 561–573.
- Murphy, D. (2006). *Psychiatry in the scientific image*. Cambridge, MA: The MIT Press.
- Murphy, D. (2009). Psychiatry and the concept of disease as pathology. In M. Broome, & L. Bortolotti (Eds.), *Psychiatry as cognitive neuroscience* (pp. 103–117). Oxford: Oxford University Press.
- Murphy, D. (2012). The folk epistemology of delusions. *Neuroethics*, 5(1), 19–22.
- O'Brien, M. J., & Lyman, R. L. (2003). *Cladistics and archaeology*. Salt Lake City: University of Utah Press.
- Olojugba, C., de Silva, R., Kartsounis, L. D., Royan, L., & Carter, J. (2007). De clerambault's syndrome (erotomania) as a presenting feature of fronto-temporal dementia and motor neuron disease (FTD-MND). *Behavioural Neurology*, 18(3), 193–195.
- Paqueron, X., Leguen, M., Rosenthal, D., Coriat, P., Willer, J. C., & Danziger, N. (2003). The phenomenology of body image distortions induced by regional anaesthesia. *Brain*, 126, 702–712.
- Rahmanovic, A., Barnier, A., Cox, R., Langdon, R. and Coltheart, M. (2012). That's Not My Arm: A hypnotic analogue of somatoparaphrenia. *Cognitive Neuropsychiatry*, 17(1), 36–63.
- Rhodes, J., Jakes, S., & Robinson, J. (2005). A qualitative analysis of delusional content. *Journal of Mental Health*, 14(4), 383–398.
- Rode, G., Charles, N., Perenin, M. T., Vighetto, A., Trillet, M., & Aimard, G. (1992). Partial remission of hemiplegia and somatoparaphrenia through vestibular stimulation in a case of unilateral neglect. *Cortex*, 28, 203–208.
- Rodrigues, A. C., & Banzato, C. E. (2006). Delusional misidentification syndrome: Why such nosologic challenge remains intractable. *Psychopathology*, 39(6), 296–302.
- Salvato, G., Gandola, M., Veronelli, L., Agostoni, E. C., Sberna, M., Corbo, M., & Bottini, G. (2016). The spatial side of somatoparaphrenia: A case study. *Neurocase*, 22(2), 154–160.
- Schiff, N. D., & Pulver, M. (1999). Does vestibular stimulation activate thalamocortical mechanisms that reintegrate impaired cortical regions? *Proceedings of the Royal Society B: Biological Sciences*, 266, 421–423.
- Schilder, P. (1935/1950). *The image and appearance of the human body*. New York, NY: International Universities Press.
- Spitoni, G., Pireddu, G., Galati, G., Sulpizio, V., Paolucci, S., & Pizzamiglio, L. (2016). Caloric vestibular stimulation reduces pain and somatoparaphrenia in a severe chronic central post-stroke pain patient: A case study. *PLoS One*, 11(3), e0151213.
- Sterelny, K., & Griffiths, P. E. (1999). *Sex and death: An introduction to philosophy of biology*. Chicago, IL: University of Chicago Press.
- Trower, P., & Chadwick, P. (1995). Pathways to defense of the self: A theory of two types of paranoia. *Clinical Psychology: Science and Practice*, 2(3), 263–278.
- Vallar, G., & Ronchi, R. (2009). Somatoparaphrenia: A body delusion. A review of the neuropsychological literature. *Experimental Brain Research*, 192(3), 533–551.
- Van Stralen, H., Van Zandvoort, M., & Dijkerman, H. (2011). The role of self-touch in somatosensory and body representation disorders after stroke. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 366(1581), 3142–3152.

- Van Stralen, H. E., van Zandvoort, M. J., Kappelle, L. J., & Dijkerman, H. C. (2013). The rubber hand illusion in a patient with hand disownership. *Perception*, 42(9), 991–993.
- Wakefield, J. C. (1998). The DSM's theory-neutral nosology is scientifically progressive: Response to Follette and Houts (1996). *Journal of Consulting and Clinical Psychology*, 66(5), 846–852.
- Waters, C.K. Causes that make a difference. *The Journal of Philosophy*, 104(11), 551–579.
- Wimsatt, W. C. (1976). Reductionism, levels of organization, and the mind-body problem. In G. Globus, I. Savodnik, & G. Maxwell (Eds.), *Consciousness and the brain: A philosophical investigation* (pp. 205–267). New York, NY: Plenum Press.
- Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford: Oxford university press.
- Woodward, J. (2010). Causation in biology: Stability, specificity, and the choice of levels of explanation. *Biology & Philosophy*, 25(3), 287–318.
- Wright, L. (1973). Functions. *The Philosophical Review*, 82(2), 139–168.
- Young, A. (2000). Wondrous strange: The neuropsychology of abnormal beliefs. *Mind & Language*, 15(1), 47–73.
- Young, A. W., & Leafhead, K. M. (1996). Betwixt life and death: Case studies of the Cotard delusion. In P. W. Halligan, & J. C. Marshall (Eds.), *Method in madness: Case studies in cognitive neuropsychiatry* (pp. 147–171). Hove: Psychology Press.
- Young, A. W., Leafhead, K. M., & Szulecka, T. K. (1994). The Capgras and Cotard delusions. *Psychopathology*, 27(3-5), 226–231.