

Individualism Entails Concept Particularism

Cressida Gaukroger

DRAFT. DO NOT CITE WITHOUT PERMISSION.

March 2018

Abstract

Individualism about mental content states that only elements internal to the mind can contribute to the content of thought. While accounts of individualism vary, individualists agree that two individuals who are molecule-for-molecule identical will have the same mental content, regardless of differences in their environments or histories. This paper argues that, if one accepts the individualist premise, one must also accept that concepts must be individuated internally, for to do otherwise is inconsistent with the motivation for adopting the individualist position. However, internal individuation of concepts entails concept particularism – two concept tokens that differ in the slightest cannot be members of the same concept type.

Keywords

Individualism; Externalism; Concepts; Mental Content.

Individualism about mental content is experiencing a quiet resurgence. Among others, Gabriel Segal (2000), Katalin Farkas (2008), Joseph Mendola (2008), and Nicholas Georgalis (2015) have revived the thesis that mental content supervenes on internal, or intrinsic, properties. Individualists' positions vary, but they agree that mental states

depend entirely on what is, roughly speaking, ‘in the head’ – brain states, phenomenal experiences, intrinsic properties – and not on an individual’s history or features of their social or physical environment.¹

The individualist/anti-individualist debate focuses on whether concepts should be individuated on the basis of their intrinsic or relational properties. Competing ontological pictures of concepts are judged and compared depending on which account of concept individuation they commit us to: “Twin Earth” style scenarios are about how we individuate concepts – the question of whether doppelgängers located in different environments share concepts is a question of what it means for two concepts to be the same or different.²

While individualism is primarily built on the argument that two internally identical individuals must have the same concepts, individualists often also talk as if, or outright argue (Mendola, 2008:227) that they can individuate concepts so that two internally *different* individuals could nonetheless possess the same concept. I will argue here that the reasons for adopting mental content individualism necessitate that concepts *cannot* be shared between two internally different individuals. This means that individualism entails concept particularism.

¹The debate is sometimes phrased in terms of mental content rather than concepts specifically. However, the central anti-individualist cases that contemporary internalists discuss – Hilary Putnam’s (1975) Twin Earth and beach/elm hypotheticals, and Tyler Burge’s (1979) arthritis case, among others – are specifically focused on concepts: namely, the concepts ‘water’, ‘beech’, ‘elm’, ‘arthritis’. The point of such examples is not just to stir anti-individualist intuitions. They also focus on the possession conditions for concepts, particularly where concept-possessors’ understanding is incomplete. What is being argued about in the debates over such cases is concepts.

²The debate between competing anti-individualist accounts of concepts is similarly concerned with individuation. Some argue that concepts should be individuated with reference to a concept possessor’s linguistic community (e.g. Putnam (1975); Burge (1979; 1982), others their teleological role (e.g. Millikan (1984)), others their history of interaction with natural kinds (e.g. Kripke (1981); Putnam (1975)).

1. Concept Particularism

Internal lives, dispositions, phenomenal experiences, and physical states are never identical across different individuals. Hypothetical doppelgängers share the identical intrinsic properties of their mental states, but their relationship is not a good model for the rest of us – they are, after all *specified* as being intrinsically *identical*. Concept particularism is the thesis that, as the internal properties of our concepts are never qualitatively identical, two concept tokens will never be members of the same type.

Concept particularism is problematic for many reasons. First, it undermines a basic general intuition about concepts, namely that they are the kinds of things that can be (and are) shared. Indeed, we ascribe shared-concept possession across people with *extremely* different internal experiences and states. Fodor (1998:29) captures this intuition:

It seems pretty clear that all sorts of concepts (for example, DOG, FATHER, TRIANGLE, HOUSE, TREE, AND, RED, and, surely, lots of others) are ones that all sorts of people, under all sorts of circumstances, have had and continue to have... Barring very pressing considerations to the contrary, it should turn out that people who live in very different cultures and/or at very different times (me and Aristotle, for example) both have the concept FOOD; and that people who are possessed of very different amounts of mathematical sophistication (me and Einstein, for example) both have the concept TRIANGLE; and that people who have had very different kinds of learning experiences (me and Helen Keller, for example) both have the concept TREE; and that people with very different amounts of knowledge (me and a four-year-old, for example) both have the concept HOUSE. And so forth.³

Second, concept particularism threatens to undermine linguistic communication as it

³This point is also made by Frege (1948:212): “[O]ne can hardly deny that mankind has a common store of thoughts which is transmitted from one generation to another.” Stephen Stich (1983:66-67) originally made the point about Helen Keller.

makes thought incommunicable. Successful linguistic communication requires of a level of shared understanding, which requires concept types to be common across different individuals.⁴ If I don't share concepts with others then it is both trivially true when I ascribe possession of any concept to myself (as the meaning of any word I use in self-ascription just refers to whatever is in my head) and literally false when I ascribe possession of any concept to anyone else (as my concept-related words are indexed to my own concepts and no-one else can share my concepts).⁵

Third, concept particularism undermines the success of intentional explanations of behaviour. This is true both for folk explanations and explanations in psychology (Rey (1985); Burge (1986; 1989); Fodor (1987)). The explanations that we give of each other's behaviour invokes talk of their mental states either explicitly, eg. 'Peter ran away because he has a fear of snakes', or implicitly eg. 'Lara opened the fridge to get a beer' (which involves assumptions about Lara's beliefs and desires concerning the contents of the fridge). However, concept particularism holds then Peter's SNAKE and Lara's BEER concepts would not be of the same type as any other concept tokens. Peter's concept, for example, will be realised by a different pattern of neurone firing and will be represented, presented or contemplated with a (possibly subtly) different phenomenology, to the concepts of any other person. To refer accurately to Peter's concept, if concepts are idiosyncratic in the way described by concept particularism, we need to introduce a neologism SNAKE_P which picks out a concept that only he possesses (and so on for every one of his and everyone else's concepts). This would prevent the generalisations necessary for identifying or describing psychological laws, and it would prevent our being able to understand or explain any of Peter's behaviour.⁶ As argued by Sarah Sawyer

⁴See, for example, Frege (1948); Putnam (1970); Prinz (2002).

⁵Concept particularism appears to entail that all concept-terms would operate as 'introspective demonstratives', which was a difficult enough thesis to defend when applied only to phenomenal concepts (Gertler (2012)) and raises the problems associated with the idea of a private language (Wittgenstein (1953)).

⁶"How could we begin even to describe, for example, the Müller-Lyer illusion unless people share a concept of LONGER THAN; or the gambler's fallacy, without them sharing MORE LIKELY? Concepts seem

(2007):

The primary objection to the view is that the introduction of neologisms to express idiosyncratic concepts threatens to undermine both the ordinary practice of ascribing psychological states by means of standard terms and a scientific psychology that appeals to concepts expressed by standard terms.

A further problem (the last I will discuss here, though this list is not exhaustive) is that individuating concepts such that any two differing (even if only *very slightly* differing) concept tokens cannot fall under the same concept type does not just rule out interpersonal concept sharing. It also rules out the *same individual* retaining any concept over time. In “Content and Self-Knowledge”, Paul Boghossian (1989) argues that in ‘slow-switching’ cases where Oscar unknowingly gets transported to Twin Earth and lives there for some time interacting with XYZ as he once did with H₂O there is, according to the anti-individualist, a change such that Oscar’s WATER concept is replaced with a TWATER concept. There is no way Oscar could know this change occurred, nor could he distinguish his new TWATER concept from his old WATER concept, Boghossian argues, which undermines Oscar’s ability to have knowledge of his own mental states. Many have taken the threat to self-knowledge posed by anti-individualism to be a reason to reject it, though it is worth noting that, if concepts particularism holds, after a slow switch Oscar would also be left with a different concept to the one he possessed prior to the switch. In fact it is likely that Oscar has run through *hundreds* of new ‘watery-stuff-directed’ concepts during the time of the switch. Our experiences affect the internal states relevant to concept possession, so according to concept particularism, neologisms would be required not just relative to the person who possessed the concept in question, but to the time they possessed that concept.⁷

to be natural kinds at least to the extent that they are the entities over which psychology generalises.”
Rey (2010)

⁷The inability to share concepts with others would necessarily make the contents of our thoughts

In Fregean terms, conceptions are token instantiations of (or representations of) concepts which vary from person to person. Whenever you eat lunch, this is likely to affect your lunch conception – a very slight, seemingly insignificant effect. If concepts are defined and individuated internally, then conceptions are just concepts, and, if you are a concept particularist, each time you have a change in conception you have a *new concept*. Which means not just that you do not have a LUNCH concept in common when you eat lunch on Monday, and Tuesday, and Wednesday, but that on every day, *with each new bite of your lunch*, an old concept is destroyed and a new one created.

Many of the arguments relating to the fear of concept particularism given by Fodor and others (such as Rey) are in response to particular accounts of concepts, often those developed in psychology, such as prototype theory. My arguments in this paper differ in their target. The individualists I will discuss aren't defending particular theories found in psychology, and the reasons they give in favour of individualism about mental content, while possibly supported by the empirical success of some of these theories (the role psychological states play in causal explanations of behaviour *is*, after all, one reason to be an internalist),⁸ neither assume nor depend on the success of any theory of concepts in psychology.⁹ Rather I will argue here that it is not in the repertoire of tools available to *any* individualist to individuate concepts non-idiosyncratically and that this problem cannot simply be solved by talk of similarity of concepts. Individualism entails concept particularism. This may, in fact, explain both the difficulties the individualist has in answering many anti-individualist critiques, while also explaining why theories of concepts found in psychology that seemingly individuate them internally also suffer from the problems raised by the likes of Fodor and Rey.

private. The inability to share the same concept with oneself over time, however, would end up ruling out even the possibility of a private language.

⁸See section 4.

⁹Indeed, perhaps, if Burge (1986; 1989) is right and psychology actually assumes broad content, then necessarily so.

2. Individualism as a Minimal Constraint on Concept Individuation

One way of understanding individualism as it relates to concepts is that it sets a minimal constraint on concept individuation: at the very least, if two people are ‘internally identical’ then their concepts will be identical. Views about what it means for people to be internally identical, and which internal features, components or states are relevant to determining mental content, will differ depending on the kind of individualist theory one endorses.

According to Farkas (2003, 2008), mental content is nothing over and above how a thought appears to the bearer of that thought. Farkas’ account of individualism holds that the relevant quality that doppelgängers share in Twin Earth thought experiments is that their thoughts “would appear the same” (Farkas, 2003:7). Farkas believes that this, the relevant ‘internal’ feature, can be captured by imagining the thoughts of Putnam’s (1975) Oscar and Twin Oscar being indistinguishable *to them* – that is, to the individuals having those thoughts. She refers to this criterion for an individualist theory as ‘subjective indistinguishability’.

In contrast, Segal (2000) defines individualism in terms of intrinsic and relational properties. Anti-individualism, he argues, is the theory that cognitive content is a relational property – the content of a mental state is a property that mental state has in virtue of the relationship between a mind and the world. Considered in terms of Putnam’s Twin Earth example, the anti-individualist concludes from such a case that Oscar’s and Twin Oscar’s thoughts are different, not because of any intrinsic bodily differences between them, but because of the relationship they bear to their different environments. Individualism, which Segal defends, is the denial that mental content supervenes on anything but intrinsic properties.

The minimal constraints on concepts according to these positions, therefore, are:

F1) If two individuals have subjectively indistinguishable concepts, then they have

the same concepts.

S1) If two individuals have intrinsically identical concepts, then they have *the same* concepts.

Assuming, for the sake of argument, that an individualist minimal constraint is initially appealing, the challenge for these theories is to provide an account of concept individuation that steps beyond these minimal constraints. What the individualist wants to avoid, if they want to avoid concept particularism, is committing themselves to one of the following:

F2) If two individuals have subjectively different concepts, then they have *different* concepts.

S2) If two individuals have intrinsically different concepts, then they have *different* concepts.

The problem with individuating concepts in the way set out in F2) and S2) is clear: every person has different psychological states in virtue of their dispositions and experiences so no two people have concepts that share qualitatively identical internal properties. If two concepts have to be internally identical to count as the same concept, then, aside from the examples where internal indistinguishability is specified, such as it is in the Twin Earth scenario, it would never be the case that two people could possess the same concept, or even that one individual could have the same concept twice. The question is: Does individualism necessitate that concepts must be individuated purely internally? To reject F2) and S2) the individualist must present a theory of how to individuate concepts such that they are the same when the relevant internal features are the same,

but they can also be the same when these internal features are different (even if only slightly).

An individualist attempting to avoid concept particularism, can approach concept individuation three ways. The first is to use the minimal mechanism provided by their particular version of individualism to individuate concepts across persons. The second is to turn to something external, such as reference, to individuate concepts. The third is to identify a feature of concepts that is defined in terms of their narrow properties, but is more general than these properties. I will show below that none of these methods are both workable and consistent with the principles that motivate individualism.

3. Defining an ‘Individualist’ Maximal Constraint on Concept Individuation

One way to determine the conditions under which two concepts are members of different concept types would be to use the same criteria one used for determining when two concepts are members of the same concept type, i.e. the minimal criteria outlined in the previous section. There are reasons to want to do this – the reasons that motivated being an individualist in the first place. To individuate concepts using the minimal criteria is to individuate them *internally*.

Consider Farkas’ individualism. Farkas argues that two concepts that are subjectively indistinguishable are the same concept. However, while the principle of subjective indistinguishability can be used for concept individuation, such a principle would rule out cases of two people with subjectively *distinguishable* thoughts sharing the same concept. That is to say – if subjective indistinguishability is a necessary condition for concept sharing, then this would entail concept particularism. If one wants to individuate concepts on the basis of subjective distinguishability alone then they will get F1), but they will also have to accept F2).

Alternatively, one might argue that subjective indistinguishability is merely a guide to subjective similarity. In the cases where two concepts are subjectively indistinguishable they are clearly similar enough to be the same concept, thereby keeping a commitment to F1). Furthermore, on this account, if Oscar and Twin Oscar had very slightly different ‘watery stuff’ experiences, their concepts might be very slightly subjectively distinguishable. However, such concepts could be understood as being subjectively similar enough for them to count as the same concept. This would mean that F2) could be rejected.

The problem with this approach is that the fact of two concepts being subjectively similar cannot account for the identity of any two concepts so long as they are in any way subjectively dissimilar. Consider the following example: you and I are aware that Polly and Penny are identical twins. When I think about Polly it is very subjectively similar to when you think about Polly. There are some differences – we had different experiences of, and interactions with Polly. For example, we have seen Polly in different circumstances, from different angles, wearing different clothing etc. However, our Polly-related thoughts have enough subjective similarities that we feel warranted in saying that we share a ‘Polly’ concept. Taking Po_M to be my Polly concept and Po_Y to be your Polly concept, it follows that:

$$\text{a) } Po_M = Po_Y$$

However, my thoughts about Polly are also very subjectively similar to your thoughts about Penny. Once again, there are some differences – we have seen Polly and Penny under different circumstances, from different angles, wearing different clothing – but there is as much subjective similarity between my Polly concept and your Penny concept as there is between my Polly concept and your Polly concept.¹⁰ If subjective similarity

¹⁰I have used the names ‘Polly’ and ‘Penny’ here to differentiate the concepts to make things clearer for the reader, but we could assume that the people we are considering do not know (or perhaps, remember) the names of the different women. This shouldn’t be thought of as a puzzle about names. This is also

(but not identity) were enough to say that my ‘Polly’ concept and your ‘Polly’ concept are the same, then it would also be enough to say that my ‘Polly’ concept and your ‘Penny’ concept are the same. With Pe_M as my Penny concept and Pe_Y as your Penny concept, it follows that:

$$b) Po_M = Pe_Y$$

But, presumably, my Penny concept is also going to be very similar to your Penny concept, meaning:

$$c) Pe_M = Pe_Y$$

Which entails:

$$d) Po_M = Pe_M$$

That my Polly and Penny concepts are the same concept is a troubling consequence. Similarity might not be transitive, but identity is, so if you define identity in terms of similarity then you will be faced with a slew of similarly problematic consequences.¹¹

Segal’s account of individualism fares no better at providing an account of concept individuation that retains S1) while rejecting S2). By being committed to S1) Segal’s theory entails that two concepts are always the same if their corresponding intrinsic properties are the same. To reject S2) Segal must show that two concepts can be the same even in cases where their corresponding intrinsic properties differ.

Segal’s account faces the same problems as Farkas’ – if the minimal conditions are the only conditions relevant to concept individuation, then one will end up with con-

not a problem that only applies to concepts that pick out individuals – the same argument could be just as easily run by swapping out ‘Polly’ and ‘Penny’ for ‘beech’ and ‘elm’.

¹¹Furthermore, for reasons repeatedly argued by Fodor and Lepore (1992), similarity-based accounts of concepts have a difficult time not importing circular identity-based mechanisms to make sense of what a robust account of concept similarity could be.

cepts being idiosyncratic because the intrinsic properties of mental states vary across individuals and within the same individual over time. For the individualist who accepts S1), something like a neural correlate, for example, may serve as both intrinsic and yet realisable across multiple individuals. However, it would be highly surprising if there were the same neural correlates for most or indeed any concepts across people. And, even if all concepts turned out to have neural correlates, to explain why they were neural correlates for the *same* concept would presumably require a non-internal or non-intrinsic explanation.

One might want to argue that, while concept tokens vary between individuals, intrinsic properties can be understood as types that can be instantiated in many individuals. The problem with this approach is that to group together intrinsic properties of mental states into types must be to group them together according to some characteristic that goes beyond mere intrinsic properties. One could do this by using some external mechanism for grouping together concept types (such as causal history, functional role, or extension/reference), but if one wanted to do it *internally* the only option would be to talk in terms of similarity of intrinsic states.

Now, Segal (2007:15) disagrees that “conceptual variation” is necessarily problematic as, he argues, we only need to introduce neologisms “when small conceptual variations matter”, and, presumably, the rest of the time we can continue to talk as if concepts are shared (even though they are not). According to Segal (2000:81; 2007:26) whether it is appropriate to classify two different concepts as being the same depends on their being *similar enough* where “similar enough” is based on “pragmatic factors” relating to context.¹²

However, this reply misses the point. The problem of neologisms wasn’t a problem about the lack of sufficient terms in public language to cover a wide variety of concepts,

¹²Segal’s point appears to be that linguistic practice needn’t have a bearing on the nature of mental content. But, if this is right then Segal cannot look to our use of neologisms when discussing Frege cases as support for his view, which he does in Segal (2007).

rather it is meant to demonstrate our inability to talk about the contents of other minds *at all*. Segal appears to be relying on the fact that we already group concepts together, but, as I will argue below, we *must* be doing so in virtue of some relational properties, so he cannot just say that we have our current practice to fall back on without endorsing an anti-individualist system of concept individuation.¹³ At one point Segal (2000:78) describes concept change in the following terms: “It doesn’t matter whether we describe this as one concept changing its semantics over time or as an ancestor concept ceasing to be and being replaced by a descendent concept. The point is that there is a natural process that involves some kind of mental entity persisting through a semantic change.” But the very point at issue is whether we can say that an individual or type of mental entity *did* survive the change. The problem is that the only way to give an account that captures the same entity existing through that change is to individuate it relationally. Put another way, the fact that there are no internal (or, I will argue below, no individualist-consistent) grounds for grouping internally similar concepts together is a reason for the *individualist* not to use general language to apply to concepts.

The easiest way for the individualist to provide an account of concept identity would be in purely individualistic terms, using the minimal mechanism provided for them by their respective definitions of individualism. However, to do so would be to end up with an idiosyncratic account of concepts, as their minimal accounts do not provide the tools for defining concepts as types that extend beyond tokens. For an account of concept individuation that avoids concept particularism, therefore, the individualist must look to the external.

¹³Loar (1988:161) uses the example of reading a diary which was written by either Oscar or his twin, to illustrate that we can not only understand what is written, but use it for psychological explanation, without needing to know which twin owns the diary: “We understand the diarists’ explanations because we know how they conceive things.” This, however, serves to highlight my point that arguments for individualism often rely on the belief that we share mental content with the people being described. See also: “...narrow contents are not ineffable... We ascribe narrow content by producing words that have the same narrow content for us.” (Loar, 1988:163)

4. Reference as a Maximal Constraint on Concept Individuation?

If the individualist wants to resist concept particularism, one option would be to reject an exclusively internal model for individuating concepts. This can be done by combining the internal minimal constraint with an external (or relational) maximal constraint on concept individuation. A maximal constraint would be one that defines what it takes for two non-internally-identical concepts to nonetheless count as members of the same concept type. To be compatible with individualism, any maximal constraint on concept individuation must adhere to three basic conditions:

The Minimal Preservation Condition (MPC): Any account of concept individuation must preserve identity of concepts between physical replicas.

The Sharing Restriction Condition (SRC): Any account of concept individuation must not entail that two people have the same concept where this would undermine *MPC*.

The Consistency Condition (CC): Any account of concept individuation must be consistent with the broader goals of individualism.

MPC is necessary given the primary specification that defines the individualist position.

SRC requires that accounts of concept individuation do not entail that those people who have *very* internally-different concepts share those concepts. It must be the case, for example, that non-experts do not always share concepts with experts. Why? Because in all the hypothetical doppelgänger scenarios there are experts who do not share concepts: e.g. future experts on Earth and Twin Earth; current experts in Burge's (1979) arthritis and 'tharthritis' worlds. If the doppelgängers share concepts with the experts

in their respective worlds, and experts across worlds do not share concepts, then the doppelgängers cannot share concepts with each other.

CC is needed to respect the arguments that motivate the individualist beyond intuitions about doppelgänger scenarios. The individualist position is motivated by a range of arguments including the ability to account for concepts of non-existent or impossible entities; the ability to account for Frege Puzzles; the possibility of self-knowledge; and the causal inefficacy of relational properties. If a maximal constraint on concept individuation is inconsistent with the individualist position on these issues, then accepting it would undermine the reasons for being an individualist in the first place.

Providing a relational account of concept individuation that respects the above conditions, however, is difficult. Consider reference as a way of individuating concepts. To argue that concepts are ontologically narrow does not entail commitment to the position that concepts do not refer. Could reference as a constraint on concept individuation be made compatible with individualism? Anyone attempting to do so would face several serious problems.

One problem with reference serving the maximal constraint role is finding a theory of reference that preserves *MPC* and *SRC*. It is open to the individualist to argue for individualism about *mental content*, while accepting that we may need a different account of reference. Reference itself is a relation that holds between two things – mind and world – and it is not inconsistent to understand the mind part of that relation purely in terms of its internal (phenomenal/intrinsic) properties. However, if the individualist is to individuate concepts in accordance with reference then they cannot understand reference itself as something outside the remit of an individualist account of mental content. To be consistent with *MPC*, the individualist must reject the idea that it is possible to have the same concept with different referents, if they want to individuate concepts by their referents. Consider the Oscar/Twin Oscar case. The individualist is committed to the position that Oscar and Twin Oscar have the same watery-stuff-related

concepts in virtue of their mental states being subjectively indistinguishable or intrinsically identical. If the individualist concludes that, while Oscar and Twin Oscar have indistinguishable thoughts, their concepts have different referents, then reference cannot be used to individuate concepts if *MPC* is to be retained. What the individualist needs is a robust account of reference where, if two concepts are subjectively indistinguishable or intrinsically identical, they *always* refer to the same thing. As such, *MPC* precludes the individualist using any of the traditional or currently dominant theories of reference in individuating concepts.

However, the more troubling problem is that individuating concepts in terms of reference risks violating *CC*. Any attempt to individuate concepts relationally is going to face problems with concepts that do not refer to anything that exists. It is possible to have concepts associated with unicorns, phlogiston, and Thor, and yet no concepts could refer to such animals, substances or individuals, as they do not exist. As these concepts all share a referent - nothing - if they are individuated on the basis of reference, this would entail that they are all the same concept. Furthermore, individuating concepts on the basis of reference will not allow two individuals to have different concepts that have the same referent, thereby failing to individuate concepts in a way that can handle Frege Puzzles. Importantly, while such problems have been raised with anti-individualist theories of concept individuation, the internal individuation of concepts was a way of avoiding them. One of the great virtues of individualism is that, if concepts are internal, then it is obvious how ‘unicorn’ concepts and ‘phlogiston’ concepts are different, and how ‘hesperus’ concepts and ‘phosphorus’ concepts are different: they differ internally. To individuate concepts so as to conclude otherwise would, therefore, be to violate *CC*.

5. Concept Individuation on the Basis of Extension Conditions

There is another alternative for the individualist: Segal (2000) proposes that internally defined concepts should be individuated in terms of their extension conditions. Extension conditions will have to be defined in a non-standard way if they are to be compatible with individualism. However, this approach appears initially appealing – extension conditions need not be as coarse-grained or context-specific as reference, nor as fine-grained or idiosyncratic as the narrow properties of concepts. Given the right account, this could save individualism from entailing concept particularism.

Whatever account of extension conditions the individualist chooses to use in individuating concepts, the first condition it must uphold is *MPC*. In other words it must be the case that, in every instance where two concepts are internally identical, they have identical extension conditions. As Segal (2000:28) argues, this means that extension conditions will be determined depending on the internal properties of any concept: “[E]xtension conditions are themselves narrow... It means that the extension conditions of a thinker’s concepts are determined by intrinsic features of the thinker.”¹⁴

5.1 *Extension Conditions Determined by Reference*

On Segal’s (2000, 2004) account the extension conditions for a concept pick out a set that extends across possible worlds (such as the set that contains all possible watery-stuff).

¹⁴I am going to treat Segal’s position as stating that extension conditions, in determining conceptual content, are the basis for concept individuation: when extension conditions are the same, concepts are the same, and when they are different, concepts are different. At points, however, Segal appears not to hold this exact position: “...we need to think of concepts as organic entities that can persist through changes of extension. Alf takes it that after correction he still deploys the same concept he had earlier. In a sense he is perfectly correct. It is the same concept in the sense that it is the same organic unity that has survived the conversation with the doctor. However, it has undergone a change of cognitive content and even of extension conditions.” (Segal, 2000:77) If this is to be understood as demonstrating that concepts are not, in fact, to be individuated on the basis of extension conditions, but rather individuated on the basis of their intrinsic (in this case physical or organic) properties, then the problem of individuating concepts purely internally returns, as discussed in section 2. It is for this reason that in this section I will consider only the argument that concepts should be individuated in accordance with extension conditions, using Segal as a guide to what such a position would look like. Additionally, I will take as my target a position that is exclusively individualist, avoiding discussion of the fact Segal (2004) actually argues that the extension conditions for concepts which refer to individuals are anti-individualist.

On such an account (which is the account Segal endorses), Oscar and Twin Oscar would (at the very least) have *both* H₂O and XYZ in their extensions.¹⁵ In this regard, Segal identifies an account of extension conditions that upholds *MPC*. This account also has the capacity to maintain *SRC*. In the Twin Earth scenario we want to say that future experts on Earth and Twin Earth do not share concepts, but Oscar and Twin Oscar do share concepts. All that the extension conditions method of concept individuation needs to accommodate this is to make sure that neither Oscar nor Twin Oscar share concepts with the future experts on their respective worlds. The future-Earth expert's concept's extension conditions only pick out H₂O in all possible worlds because being H₂O is part of her water concept. Oscar's concept's extension conditions, in contrast, pick out much more than H₂O. As different extension conditions mean different concepts, *SRC* is maintained.

Similarly, in Burge's (1979) arthritis case, Alf's concept picks out ailments of the joints and ailments of the thigh, so he does not share his concept with the experts in his world whose concepts extend only to ailments of the joints. In contrast, as the experts in Twin Alf's world do have a concept which extends to ailments of the thigh as well as ailments of the joints in just the way that Twin Alf's and, indeed, Alf's concepts do, Alf, Twin Alf and the Twin-experts all share a 'tharthritis' concept. Once again, this account is able to avoid violating *SRC*.

When it comes to *CC*, however, Segal's account of extension conditions faces several difficulties. The problems come from the fact that this account ties extension conditions to actual extension. If extension conditions supervene on actual extension (even cross-world extension) then the individualist loses the ease they had in explaining concepts with empty referents (or empty extensions). One argument in favour of adopting individualism was that it is good at handling concepts of non-existent entities. If part of my

¹⁵Segal (2000:19): “[M]y view is that both Zowies’ diamond concepts apply to both diamonds and twin diamonds, so contrary to what some might initially think, if Zowie pointed to a twin diamond and said “That’s a diamond,” she would be saying something true in her idiolect.”

‘Father Christmas’ concept is that he doesn’t exist, and part of my ‘unicorn’ concept is that unicorns don’t exist, then, if extension conditions are determined internally, the extension conditions will presumably be the same for both concepts, as they must both pick out nothing in all possible worlds. But if the extension conditions for concepts with empty referents are the same (at the very least in cases where the concept includes non-existence as one of its internal or intrinsic features) then, on the extension condition account, the concepts are the same, violating *CC*.

Indeed, if the extension conditions are determined by what a thought actually extends to or ‘picks out’ then why talk in terms of extension conditions at all, as opposed to mere extension? If extension conditions reduce to (or supervene on) extension, then no two concepts could have the same extension conditions and yet also have different extension. Segal’s account, therefore, faces the same problems as any account that tries to individuate concepts on the basis of reference, as discussed in section 4.¹⁶

5.2 A Modal Account of Extension Conditions

Given the limitations of Segal’s account, it is worth considering another individualist account of extension conditions, which may be used to individuate concepts. A modal account could be presented where concepts include or contain an (internally-determined) formula that picks out their referents, dependent on the context of the concept-possessor. On such an account, the extension conditions for concepts work much like the reference-fixing conditions for indexicals – what concepts pick out may vary across worlds, while their extension conditions remain the same. To give an example, Oscar’s concept would pick out the ‘watery stuff’ in the world he is in, meaning that it actually only picks out H₂O. However, if Oscar were transported to Twin Earth, his concept would pick out XYZ, because in that world the ‘watery stuff’ is XYZ. This account satisfies *MPC* –

¹⁶Furthermore, insofar as Segal argues that the extension conditions for ‘single’ concepts (those that refer to individuals as opposed to sets, kinds, or categories) are anti-individualist this account will not have any individualist advantage in explaining Frege Puzzles.

Oscar and Twin Oscar have the same ‘watery stuff’ concepts in spite of the fact that their concepts have different substances in their actual extensions, because they would refer to all and only the same substances in any world inhabited by both doppelgängers (i.e. even where extension varies, extension conditions remain the same). The account differs from that of Segal because extension conditions are not determined by what a concept actually refers to (meaning, for example, that Oscar’s concept needn’t refer to all ‘watery stuff’ in all possible worlds), but rather that reference is determined by extension conditions, which take into account modal context.

The modal account has no trouble dealing with *SRC* as, while Oscar and the future experts in his world have ‘water’ concepts that actually share an extension, were they to be transported to Twin Earth, the expert’s ‘water’ concept would continue only to have H₂O in its extension, while Oscar’s corresponding concept would now have XYZ in its extension. This means that extension conditions of the two concepts, and therefore the two concepts themselves, are different.

Furthermore, the modal account has some advantages over Segal’s account of extension conditions in dealing with *CC*. In response to Frege Puzzles, it can say that the concepts of those people who did not equate the morning star with the evening star were different because they had different extension conditions. Their ‘morning star’ concept had in its extension only the star that appeared in the sky in the morning. In the world that they were in, this happened to be the star that also appeared in the evening. But in a counterfactual world where there were two stars which appeared in an identical way to our one star, their concepts would only pick out the one that appeared in the morning. In the world where the star only showed itself in the evening, their ‘morning star’ concept would pick out nothing. As our present-day concepts pick out the morning star and evening star as the same star, we do not share a ‘morning star’ concept with those people in the past because our concepts have different extension conditions. A modal account of internally-determined extension conditions, therefore, looks like it

could satisfy *CC* when it comes to being able to explain Frege Puzzles.

This account will still struggle to explain concepts which the concept possessors themselves do not believe refer – such as the case of ‘unicorn’ given above. The individualist who does not rely on extension conditions for concept individuation has a good way of explaining how my ‘unicorn’ and ‘Father Christmas’ concepts are different, even if I believe both of unicorns and of Father Christmas that they (necessarily) don’t exist. But, if I believe that neither of these beings exist then nothing can fall in the extension of either of my concepts, even in other possible worlds, which, if one individuates concepts in accordance with extension conditions, would mean that these are actually the same concept. In this way, even the modal account – the most promising individualist account of concept individuation provided so far – still cannot help but violate *CC* insofar as it comes to (at least some) non-referring concepts. It is possible, however, that this is a price the individualist is willing to pay – being weaker than traditional individualist theories of concepts on just this one aspect may not be reason enough to reject a theory that otherwise saves individualism from concept particularism. In the following section, however, I will show that there are deeper problems with theories that attempt to fix the extension conditions of concepts in a way that is compatible with individualism.

6. Fixing Extension Conditions: Behaviour and the Trouble with Multiple Realisability

Leaving aside the question of consistency with *CC*, it is worth noting that consistency with *MPC* and *SRC* is, on its own, no reason to adopt an extension condition model of individualist concept individuation. After all, concept particularism satisfies *MPC*, *SRC*, and indeed *CC*. If the extension condition model of concept individuation is meant to do more than idiosyncratic individualist individuation, it must also be *robustly* non-idiosyncratic. Segal believes that his model is robustly non-idiosyncratic due to the fact

that extension conditions are multiply realisable: (some) concepts that differ internally will nonetheless have the same extension conditions. If this is correct, then, while internal states determine extension conditions, the extension conditions themselves, individuated more broadly than idiosyncratic internal states, could avoid the problem of concept individuation that is too fine-grained. But to say only this much is to say too little.

It is not enough to argue that two concepts are the same in virtue of their sharing the same extension conditions. It must be shown how we can work out when extension conditions are the same. Individualism itself is not committed to any position on when two internally different individuals share the same concepts. It is ad hoc to argue that those cases where we intuitively take two internally different people to be sharing concepts are cases where those people's concepts share extension conditions, if it is not explained *why* this is the case, or what unifies such instances.

The modal account provides a basic framework for fixing extension conditions by arguing that extension conditions are a formula that picks out referents based on the context of the concept possessor. However, without knowing the formula itself (and indeed, how contexts are defined) we can say neither whether the use of extension conditions to individuate concepts is consistent with the fundamental conditions of individualism, nor whether, even if it were individualism-compatible, using such extension conditions would individuate concepts any more broadly than concept particularism. After all, there need only be one world in the infinite number of possible worlds where the concepts of two people who are identical in every way except for one tiny internal difference would pick out different things (even *slightly* different things), for this to mean that these concepts are different. Just as it was impossible to think of a way to individuate concepts internally such that they would not end up being idiosyncratic, it may be impossible to individuate extension conditions internally such that they could be shared where there were any internal differences between concept-possessors. Something more still has to be brought in to determine when two internally different concepts nonetheless share

extension conditions.

One attempt to reconcile the demands of individuating extension conditions in a way that is determined by the internal features of concepts, while itself not being idiosyncratic, is to fix extension conditions according to the actual or potential behaviour of the concept possessor.¹⁷ One of the arguments for individualism focuses on the causal inertness of relational properties – if it is the narrow properties of mental states that have causal power over behaviour, then we would expect conceptually-relevant behaviour to vary with the internal properties of concepts.¹⁸ Furthermore, behaviour is often used to work out concept extension: behaviours that use concepts (such as categorisation behaviour, and linguistic communication) can indicate how those concepts actually extend. For these reasons, behaviour is the most plausible contender for the ‘something more’ that needs to be brought in to determine when internally-determined extension conditions remain the same, even in cases where the internal properties of concepts vary. Mendola (2008:227) similarly uses individuation on the basis of behaviour to reject the argument that internal properties are as private and idiosyncratic as they are usually portrayed. According to Mendola’s account, individualism individuates the contents of thoughts as being the same if they lead to the same behaviour (defined a-contextually such that behaviour is determined exclusively by bodily movements).¹⁹

However, behaviour will not work as a way of individuating concepts or extension conditions in an internally-consistent way. Consider, for example, the position that states that extension conditions are fixed by a disposition to identify. When asked to identify something as water, so long as they were in the same immediate context Oscar and Twin Oscar would always pick out the same things. Extension conditions understood in terms of disposition to identify would explain how the doppelgängers’ concepts were different from the experts in their respective worlds. In some contexts (i.e.

¹⁷Segal (2004) himself draws on behaviour to argue that the extension conditions of natural kind concepts are not as Putnam describes them.

¹⁸Fodor (1987); Gaukroger (2017).

¹⁹Mendola (2008:229) for some problems with this account see Ebbs (2013); Richard (2013).

on Earth) Oscar and future Earth experts would identify the same substance as water, but in other contexts (i.e. on Twin Earth) Oscar and the experts would diverge in their identification practices. This approach works nicely insofar as it is both consistent with, and potentially motivated by the principles that underpin the adoption of individualism – the argument that it is exclusively the intrinsic properties of mental states that have causal powers and that this explains the fact that people who have the same concepts will behave in the same ways. However, to know whether or not someone has classified a substance under their ‘water’ concept we already need to know what concept is in play in their classificatory practice. To know this we must already have settled the question of how to individuate concepts – so what we are left with is a regress. Disposition to identify just cannot do what we need it to – provide an independent means to individuate concepts.

A similar problem will befall an alternative account that states that extension conditions are determined by public language behaviour. On this approach it is the disposition to say “this is water” which determines the extension conditions of your water concept. However, to say something you need more than just to utter a particular sound – to determine whether the word Oscar is using is actually “water” and not merely a homonym we must establish what concept is associated with his speech utterances. Even if you don’t accept the pragmatic picture that speakers’ intentions determine speakers’ meaning, it is difficult to make sense of an account of how disposition to say particular things could be used as a way of individuating concepts if the speakers did not need to have the concepts we associate with the meanings of the words they are disposed to use.

Another concern for the individualist is that they must avoid the conclusion that any likely or future behaviour of concept-possessors is going to indicate the extension of their concepts. Taking future behaviour as a guide to current concept extension conditions (and, therefore, the nature of our current concepts) risks violating *MPC*. If Oscar found out that the watery stuff on his planet was H_2O and Twin Oscar found

out that the watery stuff on his planet was XYZ, they would both be inclined to classify H₂O and XYZ as different substances, on the basis of which (future) behaviour we could conclude that they had had different concepts all along. In fact, looking to future behaviour is exactly what the *anti*-individualist argues we should do. Burge (1979:94-5) believes this intuition to be so fundamental that it is written into the language we use to discuss concepts: “The patient does not say (or think) that he had thought he had some-category-of-disease-like-arthritis-and-including-arthritis-but-also-capable-of-occurring-outside-of-joints in the thigh instead of the error commonly attributed. This sort of response would be disingenuous. Whatever other beliefs he had, the subject thought that he had arthritis in the thigh.”

A further challenge, if individualist concept individuation is to be based on behaviour, is to present an account of behaviour individuation that is not itself idiosyncratic. Concept-driven behaviour is often varied, vague, and inconsistent, and therefore hard to individuate in a clear-cut way. Consider, for example, the fact that for pretty much any category people will ‘rank’ items as being more or less typical members of that category (Rosch (1973, 1999)). Penguins are considered less typical birds than sparrows, for example. Having an anvil dropped on your head is considered a less typical way of being killed by the mafia than being shot in an Italian restaurant (Barsalou (1987)). This is referred to as concepts having ‘graded structure’. Graded structure is rejected by some anti-individualists as not being about the metaphysics of concepts themselves, but rather the epistemology or behaviour of individual concept possessors (Rey (1999)). However, the individualist does not have the same motivation to reject the idea that concepts themselves are graded. In fact, it can be seen as a strength of individualism that it can account for this feature of how concepts are actually used by those who possess them (Prinz (2002)). There is good reason to believe that our typicality judgments relating to category membership are represented in the internal structure of our concepts, not least because categorising behaviour is understood to be a typical example

of conceptually-driven behaviour.

Where a difference in graded structure is reflected in behaviour, should the behaviour be understood to be different? If so then, if extension conditions are determined by behaviour (such that different behaviour means different extension conditions), it becomes very unlikely that two instances of conceptually-relevant behaviour will ever be the same, entailing that extension conditions, and therefore concepts, are idiosyncratic. If behaviours reflecting different graded structures of concepts are not understood as being different, then it must be for reasons that are not reflected in the internal properties of concepts. As shown by Barsalou (1987), people's graded categories change with context; there is a lot of variation between individuals (something not revealed in averages across groups); and the way people structure and apply their concepts does not remain the same in the same individual over time.

Indeed, even when people seemingly agree on the definition of a concept, they will vary in the way they actually apply it. Barbara Malt (1994), showed that subjects were likely to classify pool water as being water but tea as not being water, even though pool water contains less H_2O than tea. Even when subjects said they believed water was H_2O , this was not borne out by their behaviour. While two internally-identical people may behave in exactly the same way under all the same conditions, those who differ internally, even if only slightly, are likely to differ in their behaviour, if behaviour is individuated internally. Which leaves us with a further conclusion the individualist has to try to get around, very much like F2) and S2), which is that while two people who behave the same have the same concepts, two people that have (even slightly) different behaviour will have different concepts.²⁰

If we cannot individuate behaviour in a way that is both non-idiosyncratic and compatible with individualism, this rules out the last plausible way of trying to determine extension conditions such that they could form the basis for an individualist account of

²⁰It must be the case that, when two people exhibit the same behaviour they have the same concepts if concepts are individuated on the basis of behaviour and *MPC* is to be retained.

concept individuation that did not entail concept particularism.

Conclusion

Individualists are not keen to embrace the idiosyncratic nature of concepts.²¹ However, as we have seen, there is no robust means of individuating concepts internally, or in accordance with individualist principles, that does not entail that concepts are idiosyncratic. No matter how you divide concepts up, if you hold something like F1) or S1), you are going to be committed to something like F2) or S2). Having run through all the possible individualist-compatible accounts of concept individuation, we can see that the individualist has no option but to accept that their position entails concept particularism.

References

Barsalou LW (1987) The instability of graded structure. In: Neisser U (ed) *Concepts and Conceptual Development: Ecological and Intellectual Factors in Categorization*, Cambridge University Press, Cambridge, pp 101–140.

Boghossian P (1989) Content and self-knowledge. *Philosophical Topics* 17(1):5–26.

Burge T (1979) Individualism and the mental. *Midwest Studies in Philosophy* 4(1):73–121.

Burge T (1982) Other bodies. In: Woodfield A (ed) *Thought and Object*, Oxford University Press

Burge T (1986) Individualism and psychology. *The Philosophical Review* 95(1):3–45.

²¹Segal, for example, often brings up examples to support his position that assume both that concepts can retain their identity even after undergoing changes over time (2000:77-8), and that cognitive content is the same in cases where psychology would usefully generalise it as being so (“So, for example, the “water” concepts of twinned Bedouins on Dry Earth, XYZ Twin Earth and Earth share a cognitive content because a good psychology would subsume all of them under the same generalizations.” (Segal, 2000:85)).

- Burge T (1989) Individuation and causation in psychology. *Pacific Philosophical Quarterly* 70:303–322.
- Ebbs G (2013) Mendola’s internalism. *Analytic Philosophy* 54(2):248–257.
- Farkas K (2003) What is externalism? *Philosophical Studies* 112(3):187–208.
- Farkas K (2008) *The Subject’s Point of View*. Oxford University Press, Oxford
- Fodor J, Lepore E (1992) *Holism: A Shopper’s Guide*. Blackwell, Cambridge, Mass.
- Fodor JA (1987) *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. MIT Press, Cambridge, Mass
- Fodor JA (1998) *Concepts: Where Cognitive Science Went Wrong*. Oxford University Press, Oxford
- Frege G (1948) Sense and reference. *The Philosophical Review* 57(3):209–230.
- Gaukroger C (2017) Why broad content can’t influence behaviour. *Synthese* 194(8):3005–3020
- Georgalis N (2015) *Mind, Language and Subjectivity: Minimal Content and the Theory of Thought*. Routledge, Cambridge, Mass.
- Gertler B (2012) Renewed acquaintance. In: Smithies D, Stoljar D (eds) *Introspection and Consciousness*, Oxford University Press, pp 93–128.
- Kripke S (1981) *Naming and Necessity*. Blackwell, Maldon, Mass.
- Loar B (1988) Social content and psychological content. In: Merrill RHGDD (ed) *Contents of Thought*, University of Arizona Press, Tucson, pp 99–139.
- Malt BC (1994) Water is not h₂o. *Cognitive Psychology* 27:41–70.
- Mendola J (2008) *Anti-Externalism*. Oxford University Press, Oxford

- Millikan RG (1984) *Language, Thought and Other Biological Categories*. MIT Press, Cambridge, Mass.
- Prinz J (2002) *Furnishing the Mind: Concepts and their Perceptual Basis*. MIT Press, Cambridge
- Putnam H (1970) Is semantics possible? *Metaphilosophy* 1(4):187–201.
- Putnam H (1975) The meaning of meaning. *Minnesota Studies in the Philosophy of Science* 7:131–193.
- Rey G (1985) Concepts and conceptions. *Cognition* 19:297–303.
- Rey G (1999) Concepts and stereotypes. In: Margolis E, Laurence S (eds) *Concepts: Core Readings*, MIT Press, Cambridge, Mass., pp 279–300.
- Rey G (2010) Concepts versus conceptions (again). *Behavioural and Brain Sciences* 33(2-3):221–222.
- Richard M (2013) Content inside out. *Analytic Philosophy* 54(2):258–267.
- Rosch E (1973) Natural categories. *Cognitive Psychology* 4:328–350.
- Rosch E (1999) Principles of categorization. pp 189–206.
- Sawyer S (2007) There is no viable notion of narrow content. In: McLaughlin JD Brian P; Cohen (ed) *Contemporary Debates in Philosophy of Mind*, Blackwell, Malden, MA, pp 20–34.
- Segal GMA (2000) *A Slim Book About Narrow Content*. MIT Press, Cambridge, Mass.
- Segal GMA (2004) Reference, causal powers, externalist intuitions and unicorns. In: Schantz R (ed) *The Externalist Challenge: New Studies on Cognition and Intentionality*, De Gruyter

Segal GMA (2007) Cognitive content and propositional attitude attributions. In: McLaughlin JD Brian P; Cohen (ed) *Contemporary Debates in Philosophy of Mind*, Blackwell, Malden, MA, pp 5–19.

Stich SP (1983) *From Folk Psychology to Cognitive Science: The Case Against Belief*. MIT Press, Cambridge, Mass.

Wittgenstein L (1953) *Philosophical Investigations*. Blackwell, Oxford