BOOK REVIEW

Précis of narrow content

Juhani Yli-Vakkuri[1] · John Hawthorne[1,2]

Chapter One of the book addresses the foundational question of what narrow content is. Answering this question involves first, making explicit a conception of content that is adequate to engaging with the relevant philosophical debates, and second, making precise various internalist slogans to the effect that contents 'are in the head'. On the first issue we opt for a conception of content according to which contents are (or at least determine) functions from indices to truth values. This allows for a wide variety of approaches: It can accommodate a picture in which contents determine truth values at worlds. But it can also accommodate pictures in which truth values are only determined at richer indices (that may, for example, include times and agents). For maximal engagement with the literature we assume that contents are assigned to *thoughts* (i.e., thought tokens), and we define a narrow content assignment as one that supervenes on the qualitative agential profiles (QAPs) of thoughts, which are the properties that encode both how the thoughts are intrinsically as well as how they relate to the inner life of the agent (throughout the history of the agent). So there we have it: a content assignment is an assignment of entities ('contents') that determine functions from indices to truth values to all possible thoughts (and nothing to non-thoughts) in such a way that different contents never assigned to thoughts possible thoughts with different QAPs. (We tried very hard in the book to make all of this—and indeed all of what we say—very precise. We can't be as precise here, since that would take us well over the word limit we were given.)

Is there a narrow content assignment? Obviously the answer is 'Yes'. (Any function that takes every possible thought to the same content is a narrow content assignment.) Matters become more interesting when we ask whether there are

✉ Juhani Yli-Vakkuri
ylivakkuri@gmail.com

1 Australian Catholic University, Melbourne, Australia

2 University of Southern California, Los Angeles, USA

narrow content assignments that satisfy various constraints. In the book we discuss two families of constraint that can give narrow content debates life. One family specifies candidate explanatory roles. E.g.: Is the content assignment we describe using ordinary attitude ascriptions—what we call the 'ur-content assignment'— narrow? Is there a narrow content assignment that can explain *a priority*? A second family imposes structural constraints inspired by formal semantics. Examples include various compositional principles and the *truth-conditionality* condition that a content assignment should assign to a thought a content that determines what are in fact the truth values of thoughts at the indices associated with those thoughts. Narrow content debates get interesting when some combination of constraints is chosen: for many combinations, there are interesting questions whether any narrow content assignment satisfies those constraints. We think that the way forward for narrow content debates is to be very clear about which constraints are in play when we say 'there is a kind of content that is in the head'. Without clarity about the constraints, such claims are apt to be uninteresting.

Chapter Two shows how, with the truth conditionality constraint in play, a narrow content assignments must make use of some unusual indices. Functions from *worlds* to truth values won't do: Within a single world, a thought in one agent may be true but a thought while a thought with the same QAP in another agent is false. Functions from *worlds and agents* to truth values won't do. Loop Lady's inner life has the characteristics of a Nietzchean eternal return world, with thoughts with the same QAP but different truth values occurring at different times. Functions from *worlds, agents, and times* to truth values won't do. Mirror Man is symmetric and has two thoughts with different truth valies but the same QAP along the lines of 'That is human' (demonstrating a human on one side and a wax figure on the other). More non-standard indices will do the trick. Suppose we allow the thought token itself to be an index. Then we will be out of the woods. But this raises further questions. Do we really want to be so liberal about what can count as an index? And even more importantly, what explanatory ends will be served by contents with indices like this? We could satisfy truth-conditionality (e.g.) by assigning a *single content* to all thoughts, explaining variation in truth value not by variation in content but by variation in index. What we learn is that if we are relaxed about what counts as an index, we need more constraints than truth-conditionality in order to get an interesting debate going. It was beyond the scope of our book to explore all combinations of constraints from the two families mentioned above. We devote our attention to a sample: From the first family, we look at various combinations of truth-conditionality and compositionality. From the second we look at the idea that ur-content is narrow and the idea that a narrow content assignment can explain a priority (and more briefly at the question whether narrow content assignments are better suited to explain behavior).

Chapter Three addresses what we call *sectarian internalism*—the thesis that the ur-content is assignment narrow. Not all internalists are sectarians. (There are also *ecumenical internalists*, who concede that ur-content is not narrow but maintain some narrow content assigment plays some other explanatory role.) It may seem completely obvious that sectarian internalism is doomed. Suppose Juhani is hungry. Juhani believes Juhani is hungry. (If this précis had just one author, 'I am hungry'

would have made the same point even more clearly.) But if Juhani has a duplicate in a far-off land, the duplicate does not believe that Juhani is hungry. So on the face of it sectarian internalism is a non-starter (and we don't need to fuss about natural kinds to see this): Isn't it clear that when we claim correctly that Juhani believes that Juhani is hungry, there is a content—that Juhani is hungry—that is the ur-content of one of Juhani's thoughts but is not the ur-content of any of his duplicate's thoughts? The *prima facie* case is pretty compelling. But do the best we can to find ways for the sectarian internalist to keep their project alive by extending the resources introduced by David Kaplan in 'Quantifying In' (where, roughly, Kaplan was trying to reconcile the felicity of such constructions as 'There is an x such that Juhani believes x is hungry' with the idea that all contents are, roughly speaking, conceptual). To cut a long story short, the basic idea is to analyze 'Juhani believes that Juhani is hungry' along the lines: 'There is a content c such that Juhani believes c and c represents for Juhani that Juhani is hungry (note that we use a three place relation: c represents p to x). Now there is wiggle-room for the sectarian internalist. Maybe Juhani and Twin-Juhani have thoughts with the same content but that content represents the proposition that Juhani is hungry to Juhani but not to Twin-Juhani. We critically examine the prospects for this (admittedly rather desperate) sectarian reconstruction and arrive at a largely negative verdict.

Chapter Four explores the prospects for using a narrow content assignment to explain both *a priority* and a priori connections between thoughts (and, more generally, to explain rationality-theoretic relations). The prospects do not look. This is especially evident when it comes to connections between thoughts. For example, the thought on Mirror Man's left has intimate a priori connections to itself, but not to the thought on the right. But this difference cannot be explained by a difference in the contents assigned by some narrow content assignment which, ex hypothesi, have to be the same. Matters are a little more delicate when it comes to simple *a priority* (*a priority* as a property of single thoughts). But once any of various natural compositionality constraints are imposed, the project of explaining simple *a priority* by narrow content assignments is demonstrably doomed.

In Chapter Five we look at the prospects of retreating to a different conception of narrowness and thus a different conception of what it takes for a content assignment to be narrow. Narrow content, as standardly conceived, requires different contents to thoughts with the same QAP. But what kept the idea that narrow content is determined by "what is inside the head" but construed the "insides" as including not only qualitative inner goings-on but also *haecceitistic* inner goings-on? To make this idea precise, we introduce then notion of *quasi-narrowness*. Quasi-narrowness is a very new game. A quasi-narrow content assignment need not, for example, assign the same content to Putnam and Twin Putnam's thoughts, since they relate differently to Putnam's and Twin Putnam's inner *haecceities*. This raises the worry that the new game does not impose any interesting constraints at all and hence is not of much theoretical interest. We argue that this suspicion has considerable merit.

Chapter Six contrasts two models in the internalist spirit for explaining the felt security of thoughts associated with speeches like 'Dave = Dave'. In the *narrow content assignment* model one explains the felt security by invoking matching a narrow assignment of content constituents to the thought constituents associated

with each occurrence of 'Dave'. In the *relational narrowness* model one explains it by invoking a special relation between the thought constituents that is relationally narrow in that it also relates any corresponding pairs of thought constituents in any other that are qualitatively the same on the inside. A key insight is that even if the obtaining of a relation between a pair of tokens supervenes on how the pair relates to the qualitative inside of the agent, it may not supervene on the QAPs of the tokens. To get a picture of this, imagine a system of four spheres, all duplicates, with A connected to B and only B and C connected to D and only D. Suppose there is a plane of symmetry with A, B on one side and C, D on the other, and that the spheres are at the vertices of a square. Consider the *A–B–C–D-profile* of a sphere, the strongest qualitative property that encodes how it relates to the system, so that (e.g.) the *A–B–C–D-profiles* of A entails that A is connected to exactly one other sphere, etc. All four spheres have the same A–B–C–D-profile, but some interesting facts about how spheres in the system relate to each other are not captured by their A–B–C–D-profiles. For analogous reasons, some interesting facts about how tokens in a symmetric agent relate to each other are not captured by their QAPs. (Perhaps perfectly symmetric agents are impossible, but the lesson carries over to possible cases where any asymmetries are irrelevant to psychology.) We argue that the relational narrowness model has more going for it than the narrow content assignment model. But even here we have some reservations: this relational narrowness model requires either a highly restrictive conception of the *security* of '$a = a$' thoughts (allowing that the security that is *felt* is often not there) or else abandoning the view that that kind of security guarantees truth.