

Interpretivism and norms

Devin Sanchez Curry

Forthcoming in *Philosophical Studies*

Abstract: This article reconsiders the relationship between interpretivism about belief and normative standards. Interpretivists have traditionally taken beliefs (and thus veridicality conditions for belief attribution) to be fixed in relation to norms of interpretation. However, recent work by philosophers and psychologists reveals that human belief attribution practices are governed by a rich diversity of normative standards. Interpretivists thus face a dilemma: either give up on the idea that belief is constitutively normative or countenance a context-sensitive disjunction of norms that constitute belief. Either way, interpretivists should embrace the intersubjective indeterminacy of belief.

Keywords: Belief; Normativity; Intersubjective Indeterminacy; Relativism; Davidson; Dennett

1. Parable

The scene is the sacristy of a small Roman Catholic church in the southeastern United States. A parishioner has been having some doubts. The young priest and middle-aged deacon of the parish are discussing how best to counsel her. Over the course of their discussion, it becomes clear that they disagree about how to frame the parishioner's state of mind. They do not disagree about what she is thinking at any given moment, or how she is feeling. They both know how she behaves—not only how she moves her muscles, but the intentions with which she acts.¹ Nevertheless, they disagree about what she believes.

The deacon emphasizes that “the faithful receive with docility the teachings and directives that their pastors give them” (*Catechism*: 87). The parishioner has not been receiving the teachings of her pastor with docility. On the contrary, she has been asking a lot of impertinent questions, especially on the subject of the resurrection. During bi-weekly bible study, she has stated that she is “not so sure about the whole Easter thing.” The deacon takes her at her word and interprets her as disbelieving that Jesus rose again on the third day of his death. He thinks the goal of counsel should be to indoctrinate her: to get her to believe that Jesus was resurrected, from a place of disbelief.

The young priest interprets his parishioner differently. Yes, she has been asking a lot of questions. She has difficulties with the whole Easter thing. But “ten thousand difficulties do not make one doubt” (*Catechism* 157), and “faith seeks understanding” (158). More importantly, the parishioner attends mass every Sunday and professes her faith during the Apostle's Creed. (She also never misses a bible study or parish barbecue.) The priest therefore interprets the parishioner—a good practicing Catholic—as believing that Jesus rose from the dead.² After all,

¹ For ease of exposition, I am stipulating that the priest and deacon agree about how the parishioner thinks, feel, and behaves on both thin descriptions—they agree about her tokens of mental speech, raw feels, and muscle movements—and thick descriptions (Ryle 1949, 1979; Geertz 1973)—they agree about how she wholeheartedly genuflects, and skeptically arches an eyebrow while listening to certain gospel passages.

² Boudry & Coyne (2016) defend the view that religious dogmas are beliefs (as opposed to imaginings).

he takes thoughtful participation in the life of the Church to be the surest sign of belief. Even though she has some questions about how, exactly, Jesus could have risen again, the parishioner is “carried by the faith of others” in the parish to believe (166). So, the priest thinks the goal of counsel should be to strengthen the parishioner’s extant belief in the resurrection.

The parishioner believes one thing from the deacon’s point of view and another, contradictory thing from the priest’s point of view. It is (at least) not obvious that one of them is mistaken.

Most philosophers of mind and cognitive scientists type-identify beliefs like the parishioner’s with (computational, pure, or teleo)functional states (or their neural realizers) within the cognitive systems of believers. Functional states are individuated with respect to their functions, and functions are not altered based on how they are interpreted from outside of the functional system.³ Most theorists thus deny that beliefs could vary between belief attributors. For example, they would deny that the parishioner could believe that Jesus rose again for the priest yet disbelieve it for the deacon. Instead, most functionalists would insist that there is a single fact of the matter about whether the parishioner has the belief that Jesus rose again.⁴ The parishioner either believes, disbelieves, withholds belief, or both believes and disbelieves (Mandelbaum 2016) intersubjectively—in relation to every truth-tracking point of view. For functionalists, idiosyncratic models of belief fail to capture what people actually believe precisely to the extent that they fail to describe intersubjectively determinate cognitive states of belief.

Interpretivists about belief reject the type-identification of beliefs like the parishioner’s with cognitive states individuated by their functional roles in cognitive systems. Instead, interpretivists take attitudes of belief—the beliefs that people routinely attribute to each other in ordinary life—to exist in relation to interpretive practices. Nevertheless, interpretivists traditionally agree with functionalists about the intersubjective determinacy of belief. Rather than achieving intersubjectivity by conflating attitudes of belief with functionally individuated cognitive states, interpretivists have sought to achieve intersubjectivity by fixing beliefs in relation to objective norms of interpretation.

In this article, I will argue that interpretivists have been wrong to countenance constitutive norms of belief attribution. Interpretivists should admit that what the parishioner believes might well be intersubjectively indeterminate: it might be that the deacon and priest both attribute belief veridically. Belief attributors’ models of belief fix veridicality standards for belief attribution

³ Systemic capacity teleofunctionalists (Cummins 1996)—who take cognitive states to emerge within analyses of cognitive systems that the analyzer treats as having some overarching purpose or other—might reject this claim. But whatever the priest and deacon are disagreeing about, they are not disagreeing about the overarching purposes of the parishioner’s cognitive system. So we can safely assume that the parishioner’s functional states are held constant in relation to the two belief attributors.

⁴ According to functionalists, there is a single fact of the matter even if that fact is less than black and white: even if, for example, the fact is that the parishioner has a .55 credence that Jesus rose again, or that the parishioner has a pro-resurrection belief fragment and an anti-resurrection belief fragment which play distinct functional roles. Indeed, even if it is *objectively* indeterminate what the parishioner believes—maybe cognitive states of belief are metaphysically vague—functionalists would hold that it is still *intersubjectively* determinate: it is metaphysically vague from every truth-tracking point of view.

without relying on constitutive norms of interpretation. The mere fact that some models better serve belief attributors' purposes than others does not suffice to establish that normatively worse models fail to veridically model beliefs.

In §2, I will explain how the interpretivists Donald Davidson, Daniel Dennett, and Bruno Mölder all appeal to constitutive norms of interpretation in order to set intersubjective standards for veridical belief attribution. In §3, I will marshal empirical evidence in support of the theses that belief-ascribing practices vary across cultures and rely on a rich diversity of norms of interpretation. In §4, I will use these findings to undermine traditional interpretivist attempts to achieve intersubjectively stable veridicality standards, and argue that interpretivists should instead take beliefs to exist relative to the idiosyncratic models wielded by actual, individual belief attributors like the priest and deacon. In §5, I will conclude by offering the distinction between attitudes of belief and cognitive states of belief as an olive branch intended to reconcile philosophers concerned with understanding the intersubjectively indeterminate objects of quotidian social cognition and philosophers concerned with understanding how cognitive systems function.

2. Interpretivisms and norms

Donald Davidson wrote that his anomalous monism led him (biographically speaking) to "the recognition of an irreducibly normative element in all attributions of attitude" (2001a: 241). However, some scholars deny that the details of Davidson's interpretivism sustain his claim to have offered a truly normative theory of belief. I will argue that these scholars have misinterpreted Davidson; serendipitously, however, their misinterpretation can be co-opted by a relativistic interpretivism. Timothy Schroeder's distinction between normative categorization schemes and normative force-makers will help illuminate the path towards a genuinely nonnormative interpretivism that relativizes beliefs to particular models of belief wielded by particular belief attributors.

2.1. Davidson on intelligibility as the norm of belief attribution

Schroeder has written an influential article titled "Donald Davidson's Theory of Mind is Non-Normative." Schroeder's central argument hinges on the assertion that although Davidson took believers to be (necessarily) rational, he did not take it to be constitutive of belief that believers *ought to be* rational. In Schroeder's (2003: 1) terminology, Davidson's theory of mind boasts a normative "categorization scheme" — beliefs are categorized as mostly true, rational, and coherent — but it lacks a normative "force-maker" — there is nothing intrinsic to Davidson's theory of mind that gives truth, rationality, or coherence their normative oomph. In other words, a presumption of a rationality is built into the Davidsonian conception of belief, but normatively forceful reasons to be rational are not. Nothing in Davidson's theory of mind makes rationality inherently good.

I will grant that Schroeder is right that Davidson's theory of mind does not include a force-maker that makes truth, rationality, or coherence inherently good.⁵ At the very least, Davidson

⁵ Verheggen (2016) provides reason to think, contra Schroeder, that Davidson's theory of mind does include a normative force-maker, though she admits that the constitutive normativity of belief is hypothetical

did not adequately explain how norms of *belief* acquire normative force.

However, Davidson did explain how norms of *interpretation* acquire normative force, and took beliefs to exist only in relation to veridical interpretations. In particular, Davidson argued that what somebody believes is fixed by the most charitable interpretation of their behavior: the interpretation that would render the believer most intelligible as a rational agent (2001c: 215). This intelligibility norm boasts both a categorization scheme—belief attributions are ranked in terms of how intelligible they render believers—and a normative force-maker—the irreducibly normative structure of the communicative practice of triangulation. On Davidson’s view, the very act of belief attribution presupposes that interpreters aim to render believers intersubjectively intelligible. This built-in aim imbues the intelligibility norm of interpretation with normative oomph. Together with Davidson’s interpretivism, it ensures that beliefs are normatively constituted.

Davidson argued that triangulation—in which two radical interpreters openly reflect on their own beliefs, the other’s beliefs, and the objects in the world that both sets of belief are about—dissolves concerns about indeterminacy. Communicating with other belief attributors “forces us to the idea of an objective, public truth” (2001b: 170), “for the triangulation which is essential to thought requires that those in communication recognize that they occupy positions in a shared world” (2001c: 213). With this general notion of objectivity in hand, interpreters are then able to triangulate on particular objects in their shared world, including each other’s beliefs. As Davidson wrote, “gauging the thoughts of others requires that I live in the same world with them, sharing many reactions to its major features” (220). Triangulators like the priest and deacon are able to argue about what the parishioner believes because they recognize that they occupy a shared world with intersubjectively stable properties. In particular, when the priest and deacon triangulate on the parishioner’s beliefs, they are forced to the idea that there is an objective, determinate, empirically discernable fact about what she believes. The goal of belief attribution is to locate this fact. Davidson would chalk up the priest and deacon’s disagreement to either a lack of evidence or a failure of rationality (on at least one of their parts). In order to qualify as belief attributors, they must agree that somebody is mistaken.

The process of triangulation gives normative force to the interpretation that would render the parishioner maximally intelligible. Indeed, according to Davidson, what triangulators believe cannot be articulated without reference to this ideal interpretation. Insofar as they are belief attributors at all, the priest and deacon both ought to interpret the parishioner in a particular, maximally charitable manner, and what the parishioner believes is determined by how her church leaders (and everybody else) ought to interpret her. Thus, contra Schroeder, Davidson’s theory of mind is normative, not because what somebody believes depends on how they ought to believe, but because what somebody believes depends on how belief attributors ought to interpret them as believing.⁶

Schroeder anticipates this objection.

rather than categorical. Glüer (2000) and Engel (2008) provide discussions that are friendlier to Schroeder’s interpretation.

⁶ Schroeder also overlooks Davidson’s view that it is precisely this norm of intelligibility that makes it such that most beliefs must be true, rational, or consistent (2001c: 215).

First objection: Davidson's theory of mind is genuinely normative, because there is a right way to interpret organisms and a wrong way to do so according to the theory. More subtly, some interpretations are better and some are worse; since the theory of mind is (in part) a theory of interpretation, normativity enters the theory here, at least.

Response: There is a right way and a wrong way to ascribe any property. There are better and worse interpretations of radio signals as being caused by quasars, for example; it does not follow that there is something especially normative about the domain of study of radio astronomy. Likewise, the fact that there are better and worse interpretations of the propositional attitudes of organisms in no way implies that propositional attitudes themselves are normative entities. All it implies is that theorizing (about anything) is a normative enterprise, which is a point in the philosophy of science, perhaps, but not necessarily one of significance to the philosophy of mind. (2003: 10)

Schroeder's response falls flat. Davidson argued against precisely this analogy between belief attribution and physical property attribution. In early work on the principle of charity, he explained that "each interpretation and attribution of attitude is a move within a holistic theory, a theory necessarily governed by concern for consistency and general coherence with the truth, and it is this that sets these theories forever apart from those that describe mindless objects, or describe objects as mindless" (2001b: 154). The attribution of a belief differs from the attribution of a physical property insofar as it is *necessarily* governed by the principle of charity. Upon developing his notion of triangulation, Davidson doubled down.

It is here that the irreducible difference between mental concepts and physical concepts begins to emerge: the former, at least insofar as they are intentional in nature, require the interpreter to consider how best to render the creature being interpreted intelligible, that is, as a creature endowed with reason. As a consequence, an interpreter must separate meaning from opinion [read: must separate veridical belief attribution from nonveridical belief attribution] partly on normative grounds by deciding what, from his point of view, maximizes intelligibility. (2001c: 215)

No such requirement *necessarily* governs the attribution of physical properties. Physical property attributions may be normatively evaluated. Successful science depends on normative principles. Schroeder is right that it does not follow that there is something especially normative about radio signals and quasars. But Davidson considered beliefs to differ from radio signals and quasars partly insofar as belief attributions do not just happen to be guided by norms. Instead, belief attributions are necessarily governed—they are belief attributions only if they are governed—by an intelligibility norm of interpretation.

I suspect that Schroeder fails to recognize this distinction between these two senses in which attribution can be a normative enterprise because he pays too little attention to the crux of interpretivism: the thesis that to believe is to be aptly interpretable as believing. Although Davidson held that interpreters come to *know* physical properties via triangulation (just like

mental properties), he insisted that physical properties exist—subject to the laws of physics—whether or not interpreters discover them. The tree falls in the forest even if no radical interpreters are around to hear it. Belief, on the other hand, emerges in the world alongside interpretation. Davidson steadfastly held that “there is no propositional thought without communication” (2001c: 213), and that events and attitudes “are mental only as described” (2001a: 215). A creature who is not (yet) a believer may be disposed to behave in ways consonant with taking the world to be a certain way, but this pattern of dispositions is not a belief unless triangulators reify it as such. Davidson’s theory of mind is normative because he takes beliefs, unlike radio signals and quasars, to metaphysically emerge—rather than simply be discovered—via the irreducibly normative practice of triangulation.⁷

2.2. Dennett on prediction as the norm of belief attribution

Daniel Dennett agrees that beliefs emerge in relation to a constitutive norm of interpretation. For Dennett, what somebody believes depends on the most usefully and voluminously *predictive* interpretation of their behavior (1987: 15), whether or not that interpretation renders the believer maximally intelligible. Dennett (1987: 343) takes a baseline “Assumption of Rationality” to form a backdrop for successful belief attribution, but only because evolutionary forces designed the intentional stance to predict the behavior of *rational* creatures. Dennett agrees with Davidson that charity is required for predictive success, but maintains that belief attributions are (non)veridical relative to how predictable (rather than intelligible) they render believers: “our power to *interpret* the actions of others depends on our power—seldomly explicitly exercised—to predict them” (1998: 98).⁸

Dennett’s predictability norm of interpretation boasts its own, Darwinian force-maker. The intentional stance has the teleofunction of enabling humans to predict—and only thereby interpret—behavior. People capable of adopting the intentional stance ought to provide the most predictive belief attributions because the intentional stance was designed (by evolution) to be more usefully and voluminously predictive of intentional behaviors than interpretations from the design and physical stances. For Dennett, if the priest is able to predict more of his parishioners’ behaviors than the deacon, over the course of their respective ministries, then the priest’s intentional strategy is to be preferred to the deacon’s, and his predictive belief attributions are therefore more apt. If the priest’s intentional strategy is maximally usefully and voluminously predictive, then the parishioner objectively believes that Jesus rose from the dead.

Dennett and Davidson posit different governing norms of interpretation in part because they disagree about whether it can be fully indeterminate what somebody believes. Dennett allows that “there could be two interpretation schemes that were reliable and compact predictors over the long run, but which nevertheless disagreed on crucial cases” (1998: 117). In these cases,

⁷ Insofar as Schroeder’s equivocation of physical states and mental states holds water, it does so by highlighting Davidson’s (2001b: 183–198) flirtation with the idea that all truth is relative to interpretation. But if this were the case, then Davidson would allow that physical states are constitutively normative, not deny that mental states are constitutively normative.

⁸ Davidson, for his part, criticized Dennett’s emphasis on “the issue of prediction [a]s something of a red herring” (2001c: 81).

belief is fully intersubjectively indeterminate. If the priest and deacon are equally good at predicting the parishioner's behavior over the long haul, then there is no determinate fact of the matter about whether she believes in the resurrection: "no deeper fact of the matter could establish that one was a description of the individual's *real* beliefs and the other not" (118).⁹ Nevertheless, Dennettian intersubjective indeterminacy arises only in rare—if not merely conceivable and practically irrelevant¹⁰—cases of competing, substantively different yet equally predictively ideal intentional strategies. For Dennett, the predictability norm of interpretation provides a clear veridicality standard for belief attribution, despite allowing for the possibility of intersubjective indeterminacy.

On Davidson's view, competing ideal interpretations that result in full indeterminacy are inconceivable. Following Quine (1992: 50), Davidson contrasts "full indeterminacy" with the sort of "weak indeterminacy" between whether the freezing point of water is 32°F or 0°C. Just as there is a single correct answer to the question whether or not the temperature is below freezing, no matter what temperature scale the measurer employs, "there is a correct answer to the question whether or not someone has a certain attitude" (2001c: 82), no matter what interpretive strategy the belief attributor employs. In triangulating, the priest and deacon must agree that there is a single publicly accessible truth—fixed by the maximally charitable interpretation—about what the parishioner believes. They just disagree, temporarily, about what that truth is. Davidson wrote that "Dennett has urged that the answer to the [question whether there are objective grounds for choosing among conflicting belief attributions] is that there are no such grounds; but I do not think he has given any reason to accept this answer" (2001c: 82). Davidson rejected Dennett's hypothesis that two competing belief attributions could (even conceivably) be equally veridical. Even if two competing intentional stances could be equally predictive, Davidson denied that they could render believers equally intelligible. Triangulation relies on the shared, factive conviction that if an interpretive disagreement is substantive, then one interpreter must be wrong. Every substantive difference between Davidsonian belief attributions corresponds to a difference in the intelligibility of the believer. For Davidson, then, the intelligibility norm of interpretation provides a monolithic veridicality standard for belief attribution, thereby dissolving the bogey of full indeterminacy.

2.3. Mölder on community standards as norms of belief attribution

More recently, Bruno Mölder (2010) has argued that community norms set the veridicality standards for belief attribution. There is an objective (if abstract) fact of the matter about how most people in a community construe any given belief; most belief attributors wield roughly similar stereotypes of how people possessing particular beliefs live, such that sociologists could

⁹ Or, to be more precise, there is one fact of the matter about what the parishioner believes in relation to the priest's intentional stance, and another, incommensurable fact of the matter in relation to the deacon's intentional stance. Both of these facts are perfectly objective, as well as (in principle) publicly accessible via empirical study of the priest and deacon's respective intentional stances, as well as the parishioner's thoughts, feelings, and behaviors.

¹⁰ Dennett stresses severally that "the bogey of radically different interpretations with equal warrant is ... metaphysically important ... but practically negligible" (1987: 29).

construct abstract community-level stereotypes of those beliefs. Mölder considers these community standards to constitute both normatively innocuous *statistical* norms and normatively forceful *prescriptive* norms (Antony 2013: 532): they capture both how most members of a community *do* construe any given belief and, *therefore*, how the members of that community *ought* to construe that belief.¹¹

Mölder argues that a believer possesses a belief if and only if said belief is “canonically ascribable” to said believer. Mölder rejects Davidson and Dennett’s “thought that perhaps we can specify the canonicity condition through the notion of an ideal ascriber,” because “it makes all of our ordinary ascriptions noncanonical [and thus nonveridical], for ordinary folk cannot instantiate the ideal ascriber” (2010: 172). However, Mölder also preemptively rejects an interpretivism that relativizes beliefs to arbitrary attributors, on the grounds that such a theory would fail to provide an intersubjectively discernable veridicality standard for belief attribution. Mölder insists that in order to provide a veridicality standard that actual belief attributors can meet, “we need to locate the canonicity in the space between an arbitrary and an ideal ascription” (173).

Enter community norms. According to Mölder, “what is required for the application of mental terms [like belief] lies out in the open, in our common-sense psychology and it can be mastered by anyone who masters folk psychology” (2010: 2). Attributions of belief are canonical—and thus veridical—insofar as they would have been made by ordinary folk psychologists with access to public evidence about believers. In particular, Mölder identifies canonical ascriptions with attributions that ought not be revised in light of evidence about the believers in question and “facts about how ordinary ascribers interpret” believers (178). Actual belief attributors can ensure that an attribution meets the latter condition by ascertaining that competent members of their folk psychological community lack “warranted objections” (174).

Mölder thus claims that community norms partly determine the possession of belief. To decide whether the priest or deacon interprets the parishioner correctly (given that they both possess as much evidence as practicable), we have to figure out which man of faith interprets her in line with the community standard stereotype for belief in the resurrection. Mölder does not specify how to determine the community in question—is it the congregation? the diocese? Western culture? the set of all folk psychologists?—though he abstractly notes that canonicity is always indexed to a particular situation, and that “the features of that situation determine what sort of information is relevant and which revisions may be needed” (174).

On Mölder’s view, “the range of ascriptions that are coherent and do not require revision can be quite large” in any given situation (175), but not so large as to result in intersubjective indeterminacy. Instead, he avers that “the fact that the canonical ascriber is faced with two different or incompatible ascriptions is a good reason for seeking further revisions” (184). That the priest and deacon are arguing is reason enough to assume that one of them falls short of the community norm for attributing belief in the resurrection. For Mölder, the parishioner (objectively and intersubjectively) believes whatever she would be interpreted as believing by a

¹¹ Mölder (2010: 151-158) acknowledges an intellectual debt to Lynne Rudder Baker’s (1995: Ch. 8) notion of the “common-sense conception”. In some moods, Eric Schwitzgebel (2013: 75) goes in for the same view.

normal interpreter with all of the accessible evidence about her thoughts, feelings, and overt behaviors.

3. Is there a normatively forceful constitutive norm of belief attribution?

Mölder, Dennett, and Davidson all introduce normativity into their interpretivist theories of mind in the guise of veridicality standards for belief attribution. Their shared insinuation is that the norms of interpretation they respectively introduce are theoretically crucial because they reveal how interpreters can be wrong about what people believe. As Mölder puts it, they reveal “the is/seems difference” that makes attributions of belief veridical or nonveridical (2010: 170).

3.1. Schwitzgebelian interpretivism and the notion of an appropriate match

However, interpretivists need not countenance normative force-makers as constitutive of belief in order to count some attributions as veridical and others as nonveridical. Consider Eric Schwitzgebel’s (2002) dispositionalist account of belief, which coheres nicely with both interpretivism and recently influential model-theoretic accounts of the cognitive mechanisms of belief attribution (Maibom 2003; Godfrey-Smith 2005; Spaulding 2018b). According to Schwitzgebel, somebody believes that p just in case they live in a manner that appropriately matches a belief attributor’s stereotype—the pattern of dispositions to act, react, think, and feel that an interpreter models as identical with taking p to be true. Schwitzgebel’s conceptual machinery provides interpretivists with a normatively forceless parameter for judging the veridicality of belief attributions, and thus the building blocks for a genuinely non-normative interpretivism. On a Schwitzgebelian interpretivism, any arbitrary belief attribution is veridical just in case the believer really does possess a pattern of dispositions that appropriately matches the model of belief applied by the belief attributor.¹²

Davidson, Dennett, or Mölder might suspect that the notion of an ‘appropriate match’ at play in Schwitzgebel’s definition is normatively loaded. The requirement that believers appropriately match particular models of belief seems to amount to constitutive interpretive normativity.

However, as Schroeder would point out, the apparently normative term ‘appropriate match’ provides a categorization scheme with no force-maker. Some matches of patterns of dispositions to models of belief are categorized as appropriate, making belief attributions veridical; others are categorized as inappropriate, making belief attributions nonveridical. However, there is nothing inherently good about attributions that appropriately match believers to models of belief. That a believer lives in a pattern that appropriately matches a model of belief entails nothing beyond the fact that the believer has the modeled belief. In particular, it entails

¹² I will argue that the Schwitzgebelian interpretivist need not embrace a constitutively normative theory of mind, without suggesting that Schwitzgebel himself does not. In addition to the fact that he sometimes flirts with constitutive community norms (2013: 75), it is unclear whether Schwitzgebel takes the notion of an appropriate match to be normatively forceful. Relatedly, it is unclear what Schwitzgebel ultimately thinks of relativism. He relegates mention of the fact that his dispositionalism might engender relativism to a footnote—otherwise proceeding “as though there were one stereotype for every belief, though strictly speaking this is not true” (2002: 251)—and immediately downplays it therein (2002: 271, fn. 6).

nothing about how the belief attribution fares in light of any norms with prescriptive oomph: a veridical Schwitzgebelian interpretation does not *necessarily* render the believer the slightest bit more intelligible, predictable, or normal from an intersubjective point of view.

On the traditional interpretivist views canvassed above, beliefs exist relative to inherently normatively forceful norms of interpretation. For Davidson, beliefs exist relative to the intelligibility norm, which all belief attributors ought to meet insofar as they are belief attributors at all, because it is the heart of the communicative practice of triangulation from which beliefs emerge. For Dennett, beliefs exist relative to the predictability norm, which all belief attributors ought to meet insofar as they are belief attributors, because prediction is the proper teleofunction of the intentional stance. For Mölder, beliefs exist relative to community norms, which all belief attributors ought to meet insofar as they are belief attributors, because they undergird the intersubjective nature of folk psychology. Unlike Mölder, Dennett, and Davidson, a Schwitzgebelian interpretivist can deny that beliefs exist relative to normatively forceful norms of interpretation. Interpretivists need not consider belief attribution to be constitutively normative, in any sense that carries more normative oomph than the sense in which attributions can be sorted according to the categorization scheme of (in)appropriateness of match to a model of belief.

3.2. Diversity in models of belief

The priest and deacon wield substantively different models of the belief that Jesus rose again: they associate different dispositions with that belief, and perhaps differ on which degrees and respects of match to their models they take to be appropriate. As such, they interpret the parishioner differently, even though they agree entirely about how she is disposed to think, feel, act, and react.

I suspect there is more diversity in models of belief than other interpretivists allow. The cognitive capacity to attribute beliefs, while (nearly) universal among humans, develops at variable rates across cultures (Slaughter & Perez-Zapata 2014). More to the point, Cecilia Heyes and Chris Frith (2014) compellingly argue that ethnographic and experimental research on explicit mindreading tutelage (Pyers et al 2009; Taumoepeau & Ruffman 2006, 2008; O'Brien et al 2011) reveals that culture has a strong and pervasive effect on how humans learn to construe beliefs.

Heyes and Frith's case is further bolstered by evidence of cross-cultural variation in mental state vocabulary (Lillard 1998; Lomas 2016), the relationship between attributions of belief and the cultural identities of believers (Perez-Zapata et al 2016), and divergent cultural predilections for perspective-taking (Wu & Keysar 2014). Members of collectivist cultures are less likely to attribute abnormal beliefs to their compatriots than members of individualist cultures, and members of cultures with few mental state terms appear to wield coarser models of beliefs than members of cultures with many mental state terms. There is also robust evidence that, from infancy, people are much more likely to imitate (and otherwise attend to the social cognitive teachings of) in-group members than out-group members (Kinzler et al 2007; Buttelmann et al 2013; Gruber et al 2017). In short, it is plausible that the models—and judgments of appropriateness of match—that interpretivists take to determine what people believe vary from

cultural group to cultural group.

There is likely variance within cultures too. Shannon Spaulding argues that “deep mindreading disagreements [are] common among neurotypical adults” (2018a: 4014); referencing the Eric Garner case, Spaulding points out that “despite having access to the same evidence, intelligent, rational, well-meaning people profoundly disagree about how to interpret” believers (2018a: 4011). Spaulding puts these disagreements down to a panoply of factors, several of which have to do with how different interpreters model beliefs (and other attitudes). Models of the beliefs of in-group members are more nuanced than models of the beliefs of out-group members (Ames 2004; Haslam 2006), and in-group membership is highly context-specific, turning on things like shared hobbies as well as ethnic identities (Tajfel 1974). In quotidian, practical contexts, lay belief attributors from different subcultures are prone to model the beliefs of a particular believer differently. As Spaulding writes, “this is not simply a matter of theoretical discussion not capturing the messy, empirical details. Individuals with different social backgrounds generate different mindreading judgments in predictable ways” (2018a: 4022).¹³

This evidence coheres well with the emerging consensus that a rich diversity of norms governs belief attribution. Davidson presupposes that people attribute beliefs in order to render believers intelligible, Dennett presupposes that people attribute beliefs in order to predict behaviors, and Mölder presupposes that people attribute beliefs in order to explain behaviors in ways that are decipherable by the folk psychological community. They are each right, sometimes. Belief attributions function not only to render believers intelligible and predict and explain behavior, but also to contextualize (Ryle 1949; Tanney 2013; Curry 2018a), judge (Morton 2003; Monroe & Malle 2017), regulate (McGeer 2007; Hrdy 2009; Andrews 2015), and manage impressions of (Malle, Knobe & Nelson 2007; Zawidzki 2013; Bohl 2015) how believers lead their lives. It is far from obvious that these diverse functions are always best fulfilled by first rendering believers maximally intelligible, or maximally predictable, or maximally explanatory to normal folks.¹⁴

All interpretivists acknowledge that what we count as a good—and not merely veridical—belief attribution depends on what norms are at play in the context of attribution. I want to suggest that there is no reason to expect that belief attributors always identify the exact same patterns of dispositions with beliefs when they aim to predict behavior as when they aim to look

¹³ Even individuals with similar backgrounds may operate with substantively different models of belief. It might be, for example, that philosophical theorizing has had a top-down influence on how philosophers like Schwitzgebel and Tamar Gendler construe beliefs. For example, Schwitzgebel (2010) construes Juliet the implicit racist as in-between believing that black students are intellectually inferior to white students, whereas Gendler (2008) construes Juliet as believing that black students and white students are intellectual equals, but alieving that they are not.

¹⁴ Perhaps there are epistemological or ethical norms that happen to seriously constrain how belief attributors ought to model certain beliefs across contexts, given the belief attributors’ general epistemological and ethical ends. But unlike traditional interpretivists, I see no reason to take these norms to be constitutive of belief attribution. (Perhaps they are constitutive of rational agency, or decent personhood.) In denying that belief attributors always ought to wield particular models across normative contexts, I mean to deny only that such a norm is built into the metaphysics of belief attribution (and thus built into the interpretivist metaphysics of belief).

cool in front of their friends. Moreover, there is no reason to expect that belief attributors always *ought to* identify the exact same patterns of dispositions with beliefs in different normative contexts. As Spaulding writes, when “individuals have different goals—as is often the case in real world social interpretation—not only will they take different information as input, they will process that information very differently as well” (2018a: 4025). In some normative contexts—like when we are trying to express a nuanced, personal understanding of a loved one—slower processing that exploits models of particular believers’ particular beliefs (and perhaps even constructs them on the fly) is preferable; in other normative contexts—like when we are trying to put a suspicious stranger in their place—faster processing that exploits stereotypes is preferable.

If models of belief that fix the veridicality standards for belief attribution can vary across normative contexts, then Davidson, Dennett, and Mölder are wrong to hold that what somebody believes is always determined by any single, privileged norm of interpretation. Neither functionalisms nor traditional interpretivisms have the flexibility to countenance widespread, practically significant variation between how competing belief attributors veridically construe beliefs. If such variation is rampant—because of (inter- and intra-cultural) differences in the models of belief deployed by idiosyncratic attributors working towards idiosyncratic ends—then interpretivists should embrace this variation rather than attempting to explain it away.

3.3. Variation in belief between belief attributors

My relativistic interpretivism countenances the possibility of widespread and practically significant variation, between belief attributors, in what somebody believes. Each belief exists in relation to a particular model wielded by a particular belief attributor.¹⁵

If two different color perceivers—myself and a black bear, say—stably perceive an object as colored in two different ways, then there is no intersubjective standard to appeal to in order to prove that one of us has misperceived the surface of the object (Hatfield 2009). If we have both seen the object in many lights and from many vantage points, we have respectively perceived two different relational properties of the object: the color the object truly has in relation to my perceptual system, and the color the object truly has in relation to the bear’s perceptual system. These two perceiver-relative colors are nevertheless objectively discernable (setting aside the problem of other minds). I can come to appreciate that the bear veridically perceives a genuinely different color than the color that exists in relation to my own perceptual capacities.

Analogously, there is no intersubjective standard that could probatively settle the debate between the deacon and the priest about what the parishioner believes. By hypothesis, they are both deeply familiar with the parishioner’s dispositions. They differ only on what patterns of dispositions they identify with the belief that Jesus rose again. The two men make sense of the parishioner’s thoughts and feelings, her actions and reactions, in different ways. She has the belief that Jesus rose again for the priest, because that is how the priest understands mental profiles like

¹⁵ This is relativism about belief according to Paul Boghossian’s definition, which states that “the relativist about a given domain, D, purports to have discovered that the truths of D involve an unexpected relation to a parameter,” with the rider that “no one of [the values of this parameter] is more correct for the purposes of determining the facts about [the domain] than any of the others” (2007: 13). It is not relativism *qua* assessment-sensitivity (Macfarlane 2014: 24).

her's. She lacks that belief for the deacon for the same reason. There is no constitutive norm of interpretation—no way in which the two men both ought to understand the parishioner's mental profile *insofar as they count as belief attributors*—to appeal to in order to decide between the two interpretations.

The parishioner's two attributor-relative beliefs are nevertheless objectively discernable. Without revising their models of belief, the priest and deacon will (rightly) refrain from agreeing about what the parishioner believes: she really does believe one thing for the priest and another for the deacon. Nevertheless, the two men may well come to agree about the fact *that* the parishioner believes one thing for the priest and another for the deacon. Through long conversation, the priest may come to appreciate—if not adopt—the deacon's old-fashioned, by-the-books stance on faith, and thereby come to appreciate that the parishioner lacks belief in the resurrection in relation to the deacon's model. By achieving this understanding of the deacon as a belief attributor, the priest can even come to see the deacon's belief attribution as veridical. It *is* veridical, given that the deacon has an accurate grasp of the parishioner's dispositions and has ascertained that she fails to appropriately match his model. Even in this harmonious scenario, it is fully intersubjectively indeterminate what the parishioner believes, so long as the priest stands by his new-fangled practice-first model of belief in the resurrection.¹⁶

This claim is less radical than it first appears. In most scenarios, belief attributors differ only subtly on which patterns of dispositions they associate with any given belief, and thus avoid talking past one another. Strong evolutionary and societal pressures ensure that most humans model beliefs in a normalized manner (McGeer 2007; Sterelny 2012; Heyes 2018). As Tad Zawidzki (2013) has persuasively argued, "mindshaping" practices—including imitation, pedagogy, conformity to norms, and narrative self-constitution—function to make both beliefs themselves and models of belief homogenous across human populations. Being members of the same culture, the priest and deacon will not only agree about the vast majority of what this particular parishioner believes; they will also agree about whether most other parishioners in the parish believe in the resurrection. Like the subtle differences in our perceptual systems, the subtle differences between the priest and deacon's respective models of belief result in conflicting belief attributions only in idiosyncratic cases.

Moreover, there are variations in style of belief between believers. Whether the parishioner believes in the resurrection varies from stereotype to stereotype. However, belief attributors like the priest and deacon often go beyond stereotypes—general-purpose models—and construct believer-specific models that capture the particular styles in which believers take the world to be certain ways (Godfrey-Smith 2005: 4–6; Maibom 2009: 374–375). Thus, variation in belief between belief attributors is sometimes merely variation in the general-purpose models through which belief attributors get an initial grasp on complex, individualized beliefs.

Stereotypes matter. People frequently attribute beliefs using general-purpose models in order to classify—and thereby make practically significant judgments about—folks qua believers

¹⁶ This full intersubjective indeterminacy differs from the indeterminacy between whether the temperature is 32°F or 0°C because the priest and deacon's models of belief in the resurrection track substantially different patterns of dispositions, rather than merely representing the exact same pattern of dispositions on different scales.

that *p*. Nevertheless, clashing interpretations of what somebody believes do not necessarily involve differing reads on the details of how that person is disposed to act, think, or feel. They sometimes just involve models that detect distinct patterns amongst the believer's dispositions. If the priest and deacon were to construct believer-specific models of the parishioner's idiosyncratic attitude towards the resurrection, they might well construct functionally equivalent models, and agree that the parishioner has a specific skeptical-but-reverent attitude towards the resurrection. Even so, they would continue to wield incompatible stereotypes of belief in the resurrection, and continue to assess what the parishioner believes differently (though equally veridically) in light of these general-purpose models.

As told, the parable about the priest and deacon concerns third-person attributions of belief. What if the parishioner herself were to step forward and declare her belief (or disbelief) in the resurrection? Many philosophers—Davidson (2001c: 3–38, 205–220) among them—have argued that believers have privileged access to their beliefs, and that there is therefore a strong presumption that believers do not misattribute beliefs to themselves. It might be inferred that, when attributing belief to the parishioner, the priest and deacon should defer to the parishioner's own model of belief in the resurrection.

I am generally skeptical of claims to first-person authority about beliefs. There is, undeniably, an asymmetry between first- and third-person belief attribution. However, I follow Peter Carruthers (2011) in taking this asymmetry to come down to differing evidence rather than any strong form of privileged access to beliefs. People do have access to some exclusively first-person indirect evidence about what they themselves believe—introspective evidence about the occurrent thoughts and feelings associated with their beliefs. Nevertheless, as Carruthers argues, there is no good evidence that people have direct introspective access to beliefs themselves. Moreover, the asymmetry between first- and third-person ascriptions cuts both ways: there is very good evidence that people have systematic blindspots concerning their own attitudes (Nisbett & Wilson 1977; Zajonc 2001; Wilson 2002). So there is no guarantee that a believer's proprietary models of belief track her own beliefs extraordinarily well. To be precise: for any given belief, it is guaranteed neither that the believer actually possesses the pattern of dispositions which she attributes to herself by applying her own model of belief, nor that her believer-specific model of her own belief will be normatively superior to other belief attributors' models of her belief.

My view is that the parishioner has one belief in relation to the priest's point of view, a second belief in relation to the deacon's point of view, and a third belief in relation to her own point of view. The parishioner's own idiosyncratic model of belief in the resurrection might be more finely attuned to her unique conception of faith. And she might have access to some introspective evidence, which the priest and deacon lack, that her dispositional profile does in fact neatly fit her own model. These asymmetries in model construction and evidence make it plausible, though far from guaranteed, that the parishioner's self-attribution is not only veridical, but also renders her more intelligible (or predictable or normal) than do other veridical attributions. But these asymmetries do not, in and of themselves, make the priest or deacon's interpretations nonveridical.¹⁷

¹⁷ My skepticism notwithstanding, I should note that some forms of first-person authority are compatible

4. A dilemma for interpretivists

If this is right, then interpretivists are left with two options. First option: deny that norms of interpretation are constitutive of beliefs. Second option: insist that norms of interpretation play constitutive (as well as normative) roles in belief attribution, but allow that different norms reign over belief attributions made for different purposes. Either way, the possibility of intersubjective indeterminacy inevitably follows on interpretivism's heels.

4.1. *Horn one: rejecting constitutive norms of belief attribution*

The first option, which I favor, is to deny that interpretivism must be a normative theory of belief. Davidson, Dennett, and Mölder argue that norms must be constitutive of belief in order to provide veridicality standards for belief attribution. But a Schwitzgebelian notion of models of belief, with their (non-normatively-forceful) veridicality standards, rids interpretivists of the need to mix the normative with the metaphysical in this fashion. Appeal to a normatively ideal or canonical attribution might dictate which models belief attributors ought to adopt, but it does not dictate which models belief attributors actually wield—it does not dictate which models actually metaphysically determine beliefs.

Recall Schroeder's claim that although "there are better and worse interpretations of radio signals as being caused by quasars ... it does not follow that there is something especially normative about the domain of study of radio astronomy" (2003: 10). Schroeder gets Davidson wrong, but Schroeder's pseudoDavidson is right about the normativity of belief attribution. There are better and worse ways to model beliefs, given attributors' goals; it does not follow that there is something especially normative about belief attribution. On the contrary, my relativistic interpretivism takes belief attribution to involve a categorization scheme with no constitutive normative force-maker.

Consider Mölderian community standards. In any given interpretive context, there may be an objective fact of the matter about the patterns of dispositions that *most* people (in the parish, diocese, city, civilization, or world) consider to be constitutive of any given belief. Or there may be a fact about the general shape of the models wielded by *normal* people, or *experts*. For example, perhaps the parishioner disbelieves in the resurrection, according the models of most—or the most normal, or most expert—members of the parish. Such community and expert standards, presumably developed partly on the basis of long-term interpretive success, play crucially

with relativistic interpretivism. For example, agentialists about self-knowledge (Moran 2001) take first-person authority to be realized by the active commitment to certain patterns of thoughts, feelings, and behaviors, rather than by privileged access to beliefs. Agentialists could vindicate first-person authority by showing how the parishioner's model necessarily captures her own belief—as fixed in relation to her own point of view—without inferring that there are no other viable models. Indeed, were he willing to renounce intelligibility as the univocal constitutive norm of interpretation, Davidson's (2001c) account of first-person authority could be understood as a variety of agentialism that fits nicely with a relativistic interpretivism. The only forms of first-person authority ruled out by intersubjective indeterminacy are those advocated by direct acquaintance and inner sense accounts of privileged access (Gertler 2011), according to which I infallibly (or nigh-infallibly) know what I believe, and if you attribute a different belief to me, then you must be wrong.

important roles in setting norms, the enforcement of which causally influences the models of belief that members of the community (both actually and ought to) develop and adopt. But the fact that community (and expert) folk psychological standards play these crucially important *normative* and *causal* roles in interpretive practices does not entail that they directly fix the *metaphysical* constitution of beliefs.

Of course, it is perfectly coherent to speak abstractly of the models of beliefs wielded by heterogeneous groups like Americans, or Catholics, or whomever. Social scientists and lay people often talk fruitfully about the common features of the general-purpose models wielded by most (or the most salient) members of a community. However, these group-level models are merely abstractions from (and heuristic stand-ins for) the idiosyncratic models wielded by the particular belief attributors collected by the group under discussion.¹⁸ These abstractions are metaphysically derivative of the particular models from which they are abstracted. Individual belief attributors do often have strong pragmatic reasons to alter their models of belief to more closely approximate the models of belief that are normal in their communities. Moreover, folks often exert themselves to cause others to adopt models that more closely approximate their own, by way of social activities like teaching, praising, coercing, and teasing. However, these strong normative and causal pulls towards homogeneity in belief attribution do not make recalcitrant outliers' abnormal belief attributions any less veridical (though they may make them significantly less practical when outliers attempt to communicate with normative belief attributors).

Much the same response applies to Dennett and Davidson's respective appeals to the norms of intelligibility and predictability. There may be an objective fact of the matter about which attribution of belief is most conducive to rendering believers maximally intelligible or predictable. Again, appeal to an ideal (strategy of) belief attribution might play an important role in dictating which models of belief people ought to adopt. But it does not dictate which models actually fix the constitution of beliefs in our nonideal world.¹⁹

The analogy with color relationalism remains informative. Maybe it would be ideal, practically speaking, for me to perceive a ripe wild raspberry as a more vibrant, brighter shade of red—even easier to spot against the green leaves. But I do not; the raspberry actually sports a dull hue in relation to my perceptual capacities. If the black bear perceives the raspberry as more vibrant, her color-perceptual point of view may be normatively superior to mine. She is better at

¹⁸ If group-level models are emergent from particular attributors' models, they are transformationally (as opposed to synchronically) emergent (Humphreys 2016).

¹⁹ If interpretivists must pick an overarching norm of interpretation, Davidson's choice is perspicuous. Some intelligibility stems from any useful interpretation, no matter what other norms are in play. But this does not vindicate Davidson's intelligibility norm of interpretation. First, interpretivists need not pick an overarching norm of interpretation; not all interpretations have to be useful in order to be veridical. Second, the notion that all belief attributions must provide *some* intelligibility does not entail that only the belief attribution that renders believers *maximally* intelligible is veridical. Third, Davidson makes an error in assuming that the intelligibility of believers as believers is something that can be fixed independently of the psychologies of particular belief attributors and the norms at play in the particular social contexts in which belief attribution occurs. One main thrust of my argument is that just *how* people are intelligible as believers depends, in large part, on both belief attributors' models and contextually variable norms of interpretation.

spotting raspberries amongst leaves, because of her way of perceiving color. But this normative superiority does not make a metaphysical difference. The raspberry remains a dull shade of red relative to me, even if I would be better off were it more brightly colored (as it is relative to the bear). Neither the bear nor I descriptively misrepresent the color of the raspberry. Likewise, neither the priest nor deacon descriptively misrepresent what the parishioner believes, even if one of their models of belief in the resurrection allows for greater predictability or intelligibility.

Now, it may be that the deacon ought (again, normatively speaking) to revise his model of belief in the resurrection to make his belief attributions more useful (for fostering a sense of community in the parish, for instance).²⁰ If carried out, this revision would make it such that the parishioner believes that Jesus rose again for the deacon. Nevertheless, it would not retroactively make it such that the parishioner believes according to the deacon's previous model. In other words, it would not make the deacon's previous belief attribution nonveridical. The deacon—having revised his model to match the priest's—is correct that the parishioner believes, and was previously—before revising his model—correct that the parishioner did not believe. This diachronic inconsistency is possible not because the parishioner changed—we can assume that all of her dispositions remained constant—but because the deacon changed. The parishioner really did not believe in light of the deacon's previous point of view, and really does believe in light of the deacon's revised point of view. The deacon's revised point of view may be normatively superior (given his practical ends), but it is no more veridical: the previous point of view made no descriptive errors.

Alternately, the deacon might come to appreciate the fact that the priest's model of belief is preferable to his own for some purposes, yet still refuse to alter his conservative model. (Perhaps it is a better model for the purpose of solidifying his identity as a God-fearing man, even while being a worse model for the purpose of shepherding the parishioner.) In that case, the deacon would still be right to describe the parishioner as disbelieving *for him*, even though he knows that it would have been better (for his immediate purpose of counseling the doubtful) if he had modeled belief differently.²¹

²⁰ If agentialists about privileged access are right, then part of why the deacon ought to revise his model may be that it conflicts with the parishioner's own model, which is likely to be normatively superior just in virtue of being the parishioner's. In this sense, the parishioner's own model may, intrinsically, hold more normative weight than anybody else's (if agentialism is true). But this built-in normative weight still would not entail that the parishioner's model is the only one that metaphysically fixes belief. Instead, it may be that people necessarily understand themselves better—in richer, more complex and fine-grained detail, for example—than anybody else can, even though third-person attributions of belief are frequently veridical.

²¹ Reflection on an analogous debate in the metaphysics of race literature may be helpful. Several philosophers have converged on the view that disagreements about the reality of race hinge on a shared understanding of the empirical facts but differing metaphysical definitions of race (see also Riddle 2013 on disability). If one defines race in terms of biological essences, then Kwame Anthony Appiah (1985) is right that race is illusory. If one defines race in terms of biological populations, then Quayshawn Spencer (2014) is right that race is real but largely irrelevant to questions of social justice. If one defines race in terms of a history of discrimination, then W.E.B. Du Bois (1897) is right that race is real and relevant to questions of social justice. Paul Taylor (2000), Ron Mallon (2006), and Sally Haslanger (2012) have

Thus far, I have argued that normatively forceful constitutive norms of interpretation are not required to—and as a matter of fact do not—set veridicality standards for belief attribution. I have not yet addressed the other major reason that traditional interpretivists have adopted their preferred constitutive norms. Among the desiderata for an account of belief is that it sets beliefs apart from other attitudes, and other aspects of believers' personalities, by making sense of belief's distinctive properties (such as its representational content and mind-to-world direction of fit). It might be objected that by dispensing with constitutive norms and instead fixing beliefs in relation to the particular models wielded by particular belief attributors, a nonnormative relativistic interpretivism will fail to fulfill this desideratum. If models of belief are causally shaped by the social enforcement of a wide variety of norms, yet constitutively constrained by none, then in virtue of what is it ensured that they remain models of *belief*?

In developing their respective interpretivisms, Davidson, Dennett, and Mölder all emphasize that humans' very conception of others as minded depends on the attribution of belief. In Davidson's (2001b: 135) language, radical interpretation can only get off the ground if radical interpreters can reliably attribute "the attitude of holding true" to radical interpretees (including themselves). According to Davidson, the principle of charity provides the only reliable means of tracking attitudes of holding true; Dennett thinks only the combination of a presumption of rationality and focus on maximizing predictability will do; Mölder thinks rigorously aligning our belief attributions with those of others in our community does the trick. All three interpretivists agree that these norms serve not only to provide veridicality standards for belief attribution, but also to ensure that interpreters (necessarily guided by normatively forceful constitutive norms) actually latch onto *beliefs*—attitudes of holding true—rather than other, less folk psychologically fundamental patterns of dispositions. Interpretees are rendered maximally intelligible (or predictable or normal) only insofar as interpreters can get a grasp on how they represent the world.

By giving up on normatively forceful constitutive norms of interpretation, and allowing an indefinite range of norms to causally shape models of belief, it might be thought that the relativistic interpretivist gives up on what makes belief *belief*. Could a model largely shaped by norms that have very little to do with how the believer represents the world—*aesthetic norms*, for example, or norms of *etiquette*—really be a model of belief? If the priest's model were to dictate that somebody believes in the resurrection just in case they are disposed to root for the Alabama Crimson Tide, refrain for laughing when the priest's toupee is askew, and so on, then in virtue of what would that count as a model of belief (as opposed to a model of comradeship or

suggested that we move beyond purely theoretical discussions of what race is, and argue instead about how we *ought to* define race, given anti-racist ethical ends (as opposed to the purely epistemic principles supposedly guiding descriptive metaphysical inquiry). Unlike metaphysicians of race, interpretivists have failed to cleanly distinguish metaphysical questions and pragmatic questions. Metaphysically speaking, the parishioner believes for the priest, but disbelieves for the deacon (similarly: race is real for Du Bois, but not real for Appiah, since they model race differently). A distinct question is pragmatic. Ought the deacon revise his stereotype of belief in the resurrection, just as (I think) Appiah ought to revise his understanding of what would make race real? One can give a principled answer to that normative question without assuming that belief attribution is irreducibly normative.

some other attitude or personality trait)?

As it happens, people do not pervert their models of belief to such egregious degrees. Beliefs do not emerge relative to any old cognitive models; they emerge relative to the models of taking the world to be some way—models of the Davidsonian attitude of holding true—which social psychologists have demonstrated to be almost universally wielded by modern humans. The capacity to construct models of belief has co-evolved (and occurrently co-develops) with mindshaping practices which serve to make believers' patterns of dispositions more cleanly modellable. As Victoria McGeer puts it, "folk-psychology involves a primary capacity for forming and regulating our mental states in accordance with a rich array of socially shared and socially maintained sense-making norms" (2015: 260). My rejection of constitutive norms of belief attribution is accompanied by the recognition of pervasively normative social practices which shape individuals' models of belief as much as they shape individuals' beliefs themselves. Both mindshaping and belief-ascribing practices are robustly (albeit contingently) normatively guided; students and teachers of folk psychology frequently aim to maximize their own and others' intelligibility, predictability, and normalcy (among myriad other ends). When their models of belief fail to serve these ends, folks revise their models, often under relentless—sometimes even coercive—social pressure from their peers.

Models of belief that have been overly influenced by aesthetic norms or norms of etiquette simply fail to serve belief attributors' ends. Indeed, models of belief serve most folk psychological practices best when they faithfully subsume dispositions to act, react, think, and feel in manners that fulfill believers' goals if and only if the world is actually the way that believers are being modeled as taking it to be. Interpretivists have no need to posit constitutive norms of belief attribution because the norms that contingently but reliably guide belief-ascribing practices have (phylogenetically and ontogenetically) ensured that belief attributors are responsive to how believers represent the world.

All realist accounts of belief—interpretivist or otherwise—construe beliefs to be special phenomena in that they realize distinctive perspectives on the world (as opposed to literal representations—reproductions—of the actual properties of the world). My relativistic version of interpretivism recognizes an additional special aspect of attitudes of belief: they realize distinctive perspectives on the world, *as appreciable from another distinctive perspective*. In other words, beliefs are particular believers' ways of representing the world, as modeled by particular belief attributors. By my reckoning, this theoretical convolution is more than made up for by relativistic interpretivism's unique capacity to pin down the slippery objects of everyday belief-ascribing practices.

For example, it is illuminating to grasp, in one synoptic picture, how the deacon and priest respectively model the parishioner's perspective on the resurrection. By understanding both the deacon's perspective and the priest's perspective, we can come to understand the parishioner's own distinctive perspective better. She lives both as if Jesus rose again (in relation to the priest's practice-oriented model of belief) and as if Jesus did not rise again (in relation to the deacon's creed-oriented model of belief). This conjunction of attributor-relative beliefs sets her perspective apart from those of most people, who either believe or disbelieve in relation to both models. Without appreciating that the priest and deacon wield incompatible yet equally viable models of

belief, we would not be positioned to fully appreciate the details of the parishioner's distinctive doxastic relationship to the resurrection. Indeed, the priest might construct a singularly subtle and profoundly informative believer-specific model of the parishioner's unique style of belief (and the deacon might construct a similarly commendable believer-specific model of the parishioner's unique style of disbelief), precisely by taking into account the pattern of dispositions identified by the other church leader.

4.2. Horn two: embracing a pluralistic normative interpretivism

I will not press the point that normative standards play no constitutive role in the metaphysics of belief any further. Nonnormativism is not, strictly speaking, crucial to my case for a relativistic interpretivism about belief. The interpretivist's second option is to double down on the normativity of belief attribution, and argue that different norms constitutively and forcefully govern belief attributions made for different purposes.

How belief attributions are properly normatively evaluated depends on which goal(s) the attributions subserve. Thus, even if norms of interpretation are constitutive of belief, different belief-ascribing practices (governed by different norms) will constitute beliefs in different ways. The predictability norm sets veridicality standards when the goal of attribution is to predict behavior; however, the attribution that renders believers most predictable is not veridical in every context. For example, it may not be the veridical attribution when the belief attributor's goal is to regulate behavior. The deacon may wield the best model of belief in the resurrection for the purpose of prediction, whereas the priest wields the best model for the purpose of regulation. If the deacon and priest are at cross-purposes—perhaps the deacon wants to predict which pointed questions the parishioner will ask, whereas the priest wants to unite his parish in faith—then both of their attributions can be veridical, even stipulating that belief is determined with reference to normative ideals.

If normative standards play a constitutive as well as normative role, then this may be evidence of more, not less, variation in belief. A normative relativistic interpretivism would dictate that individual believers have different beliefs depending on what the attributions of belief are for, as well as who is doing the attributing. I myself find this pluralism about constitutive norms implausible, if only because I see no reason to countenance constitutive norms at all. Many models of belief are general purpose tools that people employ across a wide variety of normative practices. These tools are causally shaped by a thousand norms—including community standards, intelligibility, and predictability—but, by my lights, metaphysically constituted by none. Nevertheless, if interpretivists persist in taking normatively forceful norms of interpretation to be constitutive of beliefs, then they must grapple with the normative profligacy inherent in actual practices of belief attribution.

5. Conclusion: belief as attitude, belief as cog

Insofar as the normative and the metaphysical collide, interpretivists should take beliefs to be determined relative to the diverse goals of belief-ascribing practices. Insofar as the normative and the metaphysical should be kept distinct, interpretivists should take beliefs to be determined relative to the diverse conceptions of belief inherent in belief-ascribing practices.

Either way, interpretivists should take beliefs to be constituted in relation to each of the multifarious standards by which actual people attribute beliefs.

Many philosophers will respond to this dilemma by smirking “so much the worse for interpretivism.” Their dismissal will be too quick. Relativistic interpretivism is the only extant theory of belief with the built-in flexibility to countenance the multitudinous messy ways in which idiosyncratic belief attributors, across cultures and normative contexts, get a grip on idiosyncratic believers. Like Dennett but unlike Davidson and Mölder, the relativistic interpretivist must embrace the full intersubjective indeterminacy of belief. Unlike even Dennett, the relativistic interpretivist must also embrace the possibility that full intersubjective indeterminacy is sometimes practically significant. For example, the indeterminacy about what the parishioner believes make a substantive difference to the counsel provided by the priest and deacon. Biting this bullet allows the relativistic interpretivist to pin down the slippery objects of actual lay belief attribution practices. This unique capacity strikes me as a worthy reward, though I know some readers’ teeth are aching at the mere thought.

As a salve for aching teeth, I offer a distinction between two related but independent mental phenomena that are usually labeled ‘belief’; I term them ‘attitudes of belief’ and ‘cognitive states of belief’, respectively. Attitudes of belief are the attitudes of holding true that people routinely attribute to each other (and other animals) in everyday life. Cognitive states of belief are the cogs in cognitive systems that (some) cognitive scientists theoretically posit in order to explain how minds function to represent their environments and produce appropriate behaviors. According to my relativistic interpretivism, to have an attitude of belief is to live—to be disposed to act, react, think, and feel—in a pattern that an actual belief attributor models as sufficient for taking the world to be some way. I am agnostic, meanwhile, about the nature—and, indeed, existence—of cognitive states of belief. But one major upshot of the distinction between attitudes and cognitive states is that committed functionalists (and other realists about cognitive states of belief) need not see interpretivism about attitudes of belief as conflicting with their preferred cognitive architectures.²²

Elsewhere, I have offered several reasons for refusing to conflate attitudes of belief with cognitive states pitched at any level of psychological explanation (Curry 2018a, 2018b). Here, I will restrict myself to pointing out that intersubjective indeterminacy provides one especially powerful reason to suspect that philosophers primarily concerned with uncovering cogs in cognitive systems are pursuing different objects of inquiry than philosophers primarily concerned with uncovering the targets of lay belief-ascribing practices.²³ There is an

²² For example, Jake Quilty-Dunn and Eric Mandelbaum’s (2018) recent piece in this journal offers a critique of dispositionalism (and, by extension, interpretivism) that crucially relies on the conflation of attitudes of belief with cognitive states of belief.

²³ While discussing the relationship between attitudes and cognitive states, Peter Godfrey-Smith wisely echoes Wilfrid Sellars (1962), stressing that “one of the roles for philosophy ... is to describe the coordination between the facts about interpretations and the facts about wirings-and-connections” (2004: 149). In other words, philosophers should not only give a metaphysics of attitudes of belief and a distinct metaphysics of cognitive states of belief. They should also explain the connections between the two. I wholeheartedly agree with Godfrey-Smith, but maintain that the proper first step towards a clearheaded

intersubjectively determinate fact about whether or not the parishioner has a cognitive state of belief in the resurrection. Cognitive states of belief are intersubjectively determinate objects of descriptive scientific investigation par excellence. Attitudes of belief are not.

This difference has convinced many philosophers—including but not limited to interpretivists like Davidson and Mölder—that attitudes of belief must be constitutively normative, and thus unamenable to descriptive cognitive scientific inquiry. I have argued that interpretivists should not construe beliefs as constitutively normative. I will conclude by adding that, despite being intersubjectively indeterminate, attitudes of belief are perfectly amenable to descriptive cognitive scientific inquiry: researchers simply have to study the models of belief wielded by belief attributors as well as believers' patterns of dispositions.

Acknowledgements

I have received invaluable feedback on the ideas presented in this article from Kristin Andrews, Ben Baker, Marie Barnett, Liz Camp, Amelia Curry, David Curry, Louise Daoust, Zoltan Domotor, Nabeel Hamid, Gary Hatfield, Brian Huss, Raffi Krut-Landau, Errol Lord, Lisa Miracchi, Bob Myers, Bjørn Ramberg, Quayshawn Spencer, Jordan Taylor, Claudine Verheggen, Michael Weisberg, Rob Willison, Ben Winokur, and two anonymous referees for this journal. The opening parable was inspired by a conversation with the Rev. Dr. Geoffrey Lentz, Pastor of the First United Methodist Church of Port St. Joe, Florida. Thanks also to participants in the 2017 conference on The Legacy of Donald Davidson at York University, as well as audiences at SUNY Potsdam, Kenyon College, and The College of Wooster. Any remaining flaws in my argument should be blamed on the bomb cyclone that prevented philosophers from attending my talk at the 2018 Eastern APA on an island in Savannah, Georgia.

This research was partially supported by a Mellon/ACLS Dissertation Completion Fellowship from the American Council of Learned Societies.

References

- Ames, D.R. (2004). "Inside the mind reader's tool kit: projection and stereotyping in mental state inference", *J Pers. Soc. Psychol.*, 87(3): 340-353.
- Andrews, K. (2015). "The Folk Psychological Spiral: Explanation, Regulation, and Language", *Southern Journal of Philosophy*, 53: S1, 50–67.
- Antony, L. (2013). "'Human Nature' and its Role in Feminist Theory" in *Arguing about Human Nature*, (Eds. Downes & Machery), Routledge, 521-542.
- Appiah, K. A. (1985). "The Uncompleted Argument: Du Bois and the Illusion of Race" *Critical Inquiry*, Vol. 12, No. 1, 21-37.
- Baker, L. R. (1995). *Explaining Attitudes*. MIT Press.
- Boghossian, P. (2007). *Fear of Knowledge: Against Relativism and Constructivism*. OUP.
- Bohl, V. (2015). "We read minds to shape relationships", *Philosophical Psychology*, 28:5 674-694.
- Boudry, M. & Coyne, J. (2016). "Disbelief in belief: On the cognitive status of supernatural

account of how attitudes and cognitive states are coordinated consists in recognizing that the former ought not be conflated with the latter.

- beliefs", *Philosophical Psychology* 29, 601-615.
- Buttelmann, D., Daum, M., Zmyj, N. & Carpenter, M. (2013). "Selective Imitation of In-Group over Out-Group Members in 14-Month-Old Infants", *Child Development*, Vol 84: No 2, 422-428.
- Carruthers, P. (2011). *The Opacity of Mind: An Integrative Theory of Self-Knowledge*. OUP.
- Catechism of the Catholic Church. (1994). http://www.vatican.va/archive/eng0015/_index.htm
- Cummins, R. (1996). *Representations, targets, and attitudes*. MIT Press.
- Curry, D.S. (2018a). "Beliefs as inner causes: the (lack of) evidence", *Philosophical Psychology* 31:6, 850–877.
- Curry, D.S. (2018b). *How Beliefs are like Colors*. Publicly Accessible Penn Dissertations. 2991. <https://repository.upenn.edu/edissertations/2991>
- Davidson, D. (2001a). *Essays on Actions and Events*. OUP
- Davidson, D. (2001b). *Inquiries into Truth and Interpretation*. OUP.
- Davidson, D. (2001c). *Subjective, Intersubjective, Objective*. OUP.
- Dennett, D. (1987). *The Intentional Stance* MIT Press.
- Dennett, D. (1998). *Brainchildren*. MIT Press.
- Du Bois, W.E.B. (1897). "The conservation of races". American Negro Academy.
- Engel, P. (2008). "Davidson on Epistemic Norms" in *Knowledge, Language and Interpretation: On the Philosophy of Donald Davidson*, Amoretti and Vassallo (eds.), Ontos Verlag: 123-146
- Geertz, C. (1973). *The Interpretation of Cultures*. Basic Books.
- Gendler, T. (2008). "Alief and Belief", *Journal of Philosophy*, 105 (10):634-663.
- Gertler, B. (2011). *Self-Knowledge*. Routledge.
- Glüer, K. (2000). "Dreams and Nightmares" in *Interpreting Davidson*, ed. P. Kotatko/P. Pagin/G. Segal, CSLI Publications: 53-74.
- Godfrey-Smith, P. (2004). "On Folk Psychology and Mental Representation" in *Representation in Mind: New Approaches to Mental Representation*, H. Clapin, P. Staines, and P. Slezak (eds.), Elsevier.
- Godfrey-Smith, P. (2005). "Folk Psychology as a Model", *Philosophers' Imprint* 5:6, 1-16.
- Gruber, T., Deschenaux, A., Frick, A., & Clement, F. (2017). "Group Membership Influences More Social Identification than Social Learning or Overimitation in Children", *Child Development*, Early View.
- Haslanger, S. (2012). *Resisting Reality: Social Construction and Social Critique*. OUP.
- Hatfield, G. (2009). *Perception & Cognition*. OUP.
- Haslam, N. (2006). "Dehumanization: an integrative review", *Personality and Social Psychology Review* 10:3, 252-264
- Heyes, C. & Frith, C. (2014). "The cultural evolution of mind reading", *Nature*, Vol 344, 6190.
- Heyes, C. (2018). *Cognitive Gadgets: The Cultural Evolution of Thinking*. HUP.
- Hrady, S. (2009). *Mothers and Others*. HUP.
- Humphreys, P. (2016). *Emergence*. OUP.
- Kinzler, K. D., Dupoux, E., & Spelke, E. S. (2007). "The native language of social cognition". *Proceedings of the National Academy of Sciences, USA*, 104, 12577–12580.
- Lillard, A. (1998). "Ethnopsychologies: Cultural variations in theories of mind". *Psychol. Bull.* 123,

3–32.

- Lomas, T. (2016). "Towards a positive cross-cultural lexicography", *The Journal of Positive Psychology*, 11:5, 546-558
- Macfarlane, J. (2014). *Assessment Sensitivity: Relative Truth and its Applications*. OUP.
- Maibom, H. (2003). "The Mindreader and the Scientist", *Mind & Language* 18:3, 296–315.
- Maibom, H. (2009). "In Defence of (Model) Theory Theory", *Journal of Consciousness Studies*, 16, 360.
- Malle, B., Knobe, J. & Nelson, S. (2007). "Actor-observer asymmetries in behavior explanations: New answers to an old question" *Journal of Personality and Social Psychology*, 93, 491–514.
- Mallon, R. (2006). "'Race': Normative, Not Metaphysical or Semantic." *Ethics* 116(3): 525-551.
- Mandelbaum, E. (2016). "Attitude, inference, association: On the propositional structure of implicit bias", *Noûs* 50(3), 629–658.
- McGeer, V. (2007). "The regulative dimension of folk psychology." In *Folk Psychology Re Assessed*, D. Hutto, M. M. Ratcliffe, Eds. (Springer, Dordrecht, Netherlands, 2007), pp. 137–156.
- McGeer, V. (2015). "Mind-making practices: the social infrastructure of self-knowing agency and responsibility", *Philosophical Explorations* 18:2, 259–281.
- Monroe, A. & Malle, B. (2017). "Two paths to blame: Intentionality directs moral information processing along two distinct tracks." *Journal of Experimental Psychology: General*, Vol 146(1), 123-133.
- Moran, R. (2001). *Authority and Estrangement: An Essay on Self-Knowledge*. PUP.
- Morton, A. (2003). *The importance of being understood: folk psychology as ethics*. Routledge.
- Mölder, B. (2010). *Mind Ascribed: An Elaboration and Defence of Interpretivism*. John Benjamins Publishing.
- Nisbett, R. & Wilson, T. (1977). "Telling more than we can know: Verbal reports on mental processes", *Psychological Review* 84(3): 231–259.
- O'Brien, K., Slaughter, V., & Peterson, C. C. (2011). "Sibling influences on theory of mind development for children with ASD". *J. Child Psychol. Psychiatry* 52, 713–719
- Perez-Zapata, D., Slaughter, V., & Henry, J. (2016). "Cultural Effects on Mindreading" *Cognition* 146, 410-414
- Pyers, J., Gollan, T. & Emmorey, K. (2009). "Bimodal bilinguals reveal the source of tip-of-the-tongue-states" *Cognition*, 112, 323-329,
- Quilty-Dunn, J., & Mandelbaum, E. (2018). "Against dispositionalism: belief in cognitive science", *Philosophical Studies* 175, 2353–2372.
- Quine, W.V.O. (1992). *Pursuit of Truth*. HUP.
- Riddle, C. (2013). "Defining Disability: Metaphysical Not Political." *Medicine, Health Care, & Philosophy* 16:3, 377-384.
- Ryle, G. (1949). *The Concept of Mind*. CUP.
- Ryle, G. (1979). *On Thinking*. Rowman & Littlefield.
- Schroeder, T. (2003). "Donald Davidson's Theory of Mind is Non-Normative", *Philosophers' Imprint*, Vol 3, No. 1, 1-14.
- Schwitzgebel, E. (2002). "A Phenomenal, Dispositional Account of Belief", *Nous*, 36, 249–275.
- Schwitzgebel, E. (2010). "Acting Contrary to Our Professed Beliefs, or the Gulf Between Occurrent

- Judgment and Dispositional Belief", *Pacific Philosophical Quarterly*, 91, 531– 553.
- Schwitzgebel, E. (2013). "A Dispositional Approach to Attitudes: Thinking Outside of the Belief Box", In Nikolaj Nottelmann (ed.), *Belief* (75–99). Palgrave.
- Sellars, W. (1962). "Philosophy and the Scientific Image of Man", *Frontiers of Science and Philosophy*, R. Colodny (ed.), University of Pittsburgh Press.
- Slaughter, V. & Perez-Zapata, D. (2014). "Cultural Variations in the Development of Mindreading", Vol 8., no. 4., 237-241
- Spaulding, S. (2018a). "Do You See What I See? How Social Differences Influence Mindreading". *Synthese* 195 (9), 4009–4030.
- Spaulding, S. (2018b). *How We Understand Others: Philosophy and Social Cognition*. Routledge.
- Spencer, Q. (2014). "A Radical Solution to the Race Problem", *Philosophy of Science* 81:5, 1025-1038.
- Sterelny, K. (2012). *The Evolved Apprentice*. MIT Press.
- Tajfel, H. (1974). "Social Identity and intergroup behavior", *Soc. Sci, Info.*, 13:2, 65-93.
- Tanney, J. (2013). *Rules, Reason, and Self-Knowledge*. HUP.
- Taumoepeau, M. & Ruffman, T. (2006). "Mother and infant talk about mental states relates to desire language and emotion understanding". *Child Dev.* 77, 465–481.
- Taumoepeau, M. & Ruffman, T. (2008). "Stepping stones to others' minds". *Child Dev.* 79, 284–302.
- Taylor, P.C. (2000). "Appiah's Uncompleted Argument: W.E.B. Du Bois and the Reality of Race", *Social Theory and Practice* 26(1): 103-128.
- Verheggen, C. (2007). "Triangulating with Davidson", *Philosophical Quarterly* 57 (226):96-103.
- Wilson, T. (2002). *Strangers to Ourselves: Discovering the Adaptive Unconscious*. HUP.
- Wu, S. & Keysar, B. (2014). "The Effect of Culture on Perspective Taking" *Psychological Science* vol. 18, no. 7, 600-606.
- Zajonc, R. (2001). "Mere Exposure: A Gateway to the Subliminal", *Current Directions in Psychological Science* 10:6, 224–228.
- Zawidzki, T. (2013). *Mindshaping*. MIT Press.