On Behalf of Controversial View Agnosticism

J. Adam Carter
University of Glasgow

Controversial view agnosticism (CVA) is the thesis that we are rationally obligated to withhold judgment about a large portion of our beliefs in controversial subject areas, such as philosophy, religion, morality and politics. Given that one's social identity is in no small part a function of one's positive commitments in controversial areas, CVA has unsurprisingly been regarded as objectionably 'spineless.' That said, CVA seems like an unavoidable consequence of a prominent view in the epistemology of disagreement—*conformism*—according to which the rational response to discovering that someone you identify as an epistemic peer or expert about *p* disagrees with you *vis-à-vis* *p* is to withhold judgment. This paper proposes a novel way to maintain the core conciliatory insight without devolving into an agnosticism that is objectionably spineless. The approach offered takes as a starting point the observation that–for reasons that will be made clear—the contemporary debate has bypassed the issue of the reasonableness of maintaining, rather than giving up, representational states *weaker* than belief in controversial areas. The new position developed and defended here explores this overlooked space; what results is a kind of controversial view agnosticism that is compatible with the kinds of commitments that are integral to social identity.

INTRODUCTION

Suppose you believe that Marilyn Monroe was born before Queen Elizabeth II, and—as you just now find out—I believe that Queen Elizabeth II was born before Marilyn Monroe; prior to our disagreement, we each reckoned the other to be equally likely to be right on this matter–*viz.*, we each took the other to be *epistemic peers*[1] on matters of 20th century celebrity birthdates[2]. A popular position in recent social epistemology insists that, in situations like the one just described, rationality requires that both parties should 'move to the center'. Call here, following Lackey (2008), positions that embrace this general insight about the epistemic significance of peer disagreement *conformist* views, and the general position embraced *conformism*. In a bit more detail:

---

[1]A standard way to think about epistemic peerhood is in terms of cognitive and evidential parity, *vis-à-vis* the target question. See here Lackey (2008). *Cf.* Conee (2009).

[2]Queen Elizabeth II was born in April, 1926, while Marilyn Monroe was born just a few months later in June, 1926.

> *Conformism*: In a revealed peer disagreement over *p*, each individual should give equal weight to her peer's attitude such that neither is justified in staying as confident as she initially was regarding whether *p*.[3]

Conformist (e.g. Elga (2007), Christensen (2007), Feldman (2006)) approaches to peer disagreement, construed in this inclusive way, are popular in recent social epistemology[4], though there is a well-known trouble waiting in the wings.

As Littlejohn (2013, 171) puts the problem:

> Most of the interesting things we believe (i.e., most of what we believe about epistemology, ethics, metaphysics, politics, and religion) are controversial. Much of this controversy seems to involve peers who disagree with each other fully aware of the fact that there are peers that they disagree with. Because [conformism] is correct, we cannot rationally remain committed to these controversial propositions. Thus, few of the interesting things we believe we believe rationally…The pessimistic conclusion is that we should suspend judgment on most of the interesting things we believe.

But subtract from an individual most of her interesting and controversial beliefs, and what's left is a picture where an important chunk of one's social identity is missing. This is a *prima facie* bad result[5].

Thus, the skeptical worry waiting in the wings for conformism is that it entails precisely this bad result, by entailing something along the lines of *controversial view agnosticism*[6]:

> *Controversial View Agnosticism* (CVA): We are rationally obligated to withhold judgment about a large portion of our beliefs in controversial subject areas, such as philosophy, religion, morality and politics.

As with any case where a *prima facie* intuitive thesis entails an unwelcome result, we meet with a nest of questions: which versions, precisely, of the conformist thesis most clearly generate the CVA result[7]? What moves can be made to block CVA,

---

[3]I am following here Ballantyne and Coffman (2012) in my characterisation. For a typical kind of case used to motivate this position, see Christensen (2007)'s 'Mental Math' example.

[4]For some notable opposition, see Kelly (2005) and Inwagen (1996).

[5]See Inwagen (1996) for a range of reasons to find this kind of result worrying.

[6]Conformism, if true, of course has implications that go beyond the implications its truth would have for our rational requirements in controversial areas. What I'm calling 'CVA' is, to be clear, a thesis just about what rationality requires in controversial areas, and is as such compatible with conformism's having applicability (e.g., especially in individual cases) outside of controversial subject areas. Thanks to a referee at the *European Journal of Phlosophy* for suggesting clarificaiton on this point.

[7]Matheson and Carey (2013) and Grundmann (2013) explore this point in some depth.

by proponents of conformism[8]? Is CVA really so bad[9]? These issues continue to be explored[10], though in what follows I'll be taking a different tack. I'll argue that controversial view agnosticism–or at least, a version of it–is true, but that facts about expert (rather than peer) disagreement ultimately provide the most compelling kind of support for the position. I'll then proceed to show how the version of CVA I think we should accept is not as objectionable as we would initially be let to think. In order to get the argument up and running, it's important to first look more closely at some of the assumptions lurking in the background of the argument from conformism to controversial view agnosticism–which will be our starting point.

THE TRIAD VIEW

The philosopher perhaps most responsible for generating debate about controversial view agnosticism is Richard Feldman[11]. Feldman is aware that the conformist line he's advanced implies we must give up a lot of our controversial beliefs, and Feldman (2006, 236) thinks we're rationally required to do so 'even if suspending judgment in such cases might be extremely difficult to do.'

Feldman's thinking here is undergirded by a commitment to what Turri (2012, 356) calls the Triad View of doxastic attitudes:

> *Triad View*: (i) There are only three doxastic attitudes: belief, disbelief, and withholding and (ii) once you've considered a proposition, there are, intellectually speaking, only three options open to you: you either believe it, disbelieve it, or withhold judgment.

The Triad View is, as Turri notes, 'a common view in epistemology'[12]; for one thing, the view turns up (at various junctures) in the work of an impressive range of thinkers[13]. With the Triad View in play, it looks like a relatively seamless step from the thought that conformism is a reasonable response to peer disagreement to the unwanted implication that we must withhold judgment *vis-a-vis* most of our interesting views.

---

[8]Bogardus (2013) has recently argued against the received suggestion that the equal weight version actually has this implication.

[9]See Feldman (2006, 236) for an attempt to embrace this kind of conclusion.

[10]For a recent collection of work on this topic, see especially Machuca (2013).

[11]A particularly influential paper in generating debate on the topic is Feldman (2007).

[12](Turri 2012, 356). The Triad View is slightly amended for presentation.

[13]For some representative statements of this view, see Feldman (2006), Bergmann (2005), Conee and Feldman (2004), Chisholm (1976), and Sosa (1991).

The Triad View is surely right if qualified so as to say: *with respect to the particular attitude of belief* the options are: believe, disbelieve or withhold[14]. The problem with the Triad View (qualified in this way) isn't that it's false. It's that it is deployed in a way that is *misleading*; with the Triad View in play, the available reasonable avenues to take in response to disagreement seem artifically restricted. To bring this point into sharp relief, consider first a set of examples.

1. It is raining right now.
2. O.J. Simpson was judged by jury in 1994 to be innocent of murder.
3. It it will continue to rain for the next five minutes.
4. O.J. Simpson was wrongly acquitted.

(1-2) are things I downright believe. While I wouldn't say I *believe* either (3) or (4), I do suspect that (3) and (4) are true.

Now consider (5 and 6).

5. There are an even number of spiders in the room.
6. Within the past five minutes, some golfer, somewhere in the world, has hit a hole-in-one.

I don't believe (5-6), but neither do I suspect them to be true. I don't suspect (5) to be true in part because I have good reason to believe that (5) is the sort of claim that will be just as likely true as false at any given time. Regarding (6), I simply can't (sitting here, with my current evidence) competently assess whether it's more likely or not to be true, and thus I don't suspect it's true[15].

Now with the Triad View in play, it will be natural to lump (1-2) into the 'belief' category and (3-6) into the 'withholding' category, glossing over any interesting sense in which (3-4) are views I *do* hold (even if weakly)[16]. I'll return more carefully to what I mean when I say I hold (3-4) weakly.

---

[14]One alternative to this picture, considered by Turri, is to reject the Triad View for a Quartet View, which is Triad + the option of 'withdrawing' from the proposition. I should add here that, while some philosophers (e.g. Friedman (2013b)) draw more fine-grained distinctions between 'withholding' and 'suspending', I am for the present purposes going to be using these terms interchangeably–as picking out one of the three kinds of options adverted by the Triad View. I am open to the suggestion that there are multiple ways one might count as withholding or suspending. The argument I'll develop here won't at any rate turn on any distinction here.

[15]Maybe a hole-in-one is hit every two hours somewhere in the world, in which case (6) is probably false in most contexts of evaluation. Maybe a hole-in-one is hit every two minutes, in which case (6) is probably true in most contexts of evaluation. I haven't a clue.

[16]*Cf.* Friedman (2013b), especially §4, for some helpful discussion on conditions for suspending judgment.

By way of emphasis, I am not denying that the options for *believing* are just belief, disbelief and withholding[17]. The problem is that when we follow Feldman's lead and approach the available responses to peer disagreement through the lens of the Triad View, the game looks rigged so as to exclude from discussion representational attitudes *weaker* than belief from the range of reasonable options[18]. And it might not be unreasonable to maintain, rather than give up, these weaker attitudes in controversial areas[19].

THE 'TWIN GOALS' MODEL

It's tempting to wonder: can't the threat of controversial view agnosticism sketched in §1 be circumvented entirely by simply *swapping out representational state talk for credence talk*? As this line of thought will go, something like the 'Real Number View' is more fundamental (when characterising our possible epistemic stances *vis-a-vis* propositions) than any view that trades only in representational states.

However, as Friedman (2013a) has recently argued, it's not at all obvious that suspension of judgment should be modeled simply by assigning a standard credence[20]. But even if it could be modeled unproblematically in this way, a moment's reflection

[17]I have been interpreting the Triad View as a view that expresses our options *vis-a-vis* a three-place relation between an agent, a *belief* and a proposition. As such, the view is very plausible. Compare this articulation of the Triad View with the following expression of the idea offered by Bergmann (2005): 'Three attitudes one might take towards a proposition p are believing p, disbelieving p (i.e., believing p is false), and withholding p (i.e., refraining from either believing or disbelieving p'. Bergmann (2005, 420). On one way of reading this claim, Bergmann is just saying that these three attitudes are among the possible attitudes for an agent to take *vis-a-vis* a proposition, and this is true. But Bergmann's expression of the view also lends itself to a stronger reading according to which these three attitudes comprise the only available representatal states on the table. Setting aside what Bergmann actually thinks, I want to make clear that this strong reading of the Triad View would be not *merely* misleading in the context of the peer disagreement debate, but manifestly false.

[18]Note that some philosophers have approached peer disagreement by assessing primarily the reasonableness of higher-order epistemic attitudes, as opposed to first-order attitudes. For an example of this kind of approach, see Hazlett (2012).

[19]To emphasise, the sense in which we should accept this triad view is just *relative to the representational state of belief*. With respect to *belief*, the options are: believe, disbelieve and withhold. However, it is a mistake to infer from the fact that these options are available, with respect to the representational state of belief, that they are the exhaustive representation-state options in response to revealed peer disagreement. This is, in short, because (as I will develop in the next section) there are representational states other than belief which we might take up with respect to a proposition even while not believing that proposition.

[20]Friedman (2013a, 58–59) aims to establish, more generally, the failure of what she calls *straightforward reduction*, the 'reduction of the traditionalist's doxastic attitudes to formalist's degrees of belief that says that believing p, disbelieving p, and suspending about p are just matters of having (some specified) *standard credences* for p'.

suggests that credence talk hardly disarms the worry; rather, to talk about credences will just be to turn our attention away from the problem[21].

In this section I want to briefly sketch, with reference to a background commitment to the twin epistemic goals of maximizing truth and avoiding error, a simple picture that shows how we can think of belief as occupying a particular kind of space in a heirarchy of stronger and weaker representational attitudes. Once this picture is established, I'll argue that a version of controversial view agnosticism should be embraced, but that the version we should embrace does not carry with it some of the kinds of agnostic *implications* that other commentators have found objectionable.

That said, the simple model I want to sketch here will (without presupposing any particular account of belief) locate 'belief' between two other representational attitudes: suspecting-that-p and being certain-that-p. Because I'll be featuring just three attitudes, it's not a comprehensive picture; but it'll be enough to make the point I'm after.

It's intuitive enough to locate 'believing-that-p' in between 'suspecting-that-p' and 'being certain-that-p.[22]' Though, as is indicated by the examples (1-4) in §2, there is a sense in which my believing that p is correct when p is true[23], but *so is* the weaker attitude of merely suspecting-that-p[24], as well as the stronger attitude of being certain that p. But what explains belief's situatedness *between* these two other representational states that share the same aim?

While all three of these attitudes share the same aim, we pursue the truth aim— the aim of maximizing truth and avoiding error—*differently* depending on which attitude we take up[25], and as a function of how we weight, in taking up each of these attitudes, respectively, the twin goals of possessing truths and avoiding error. Consider that truly representing and not misrepresenting are clearly *competing aims*, de-

---

[21] After all, we might agree that the real numbers 0-1 are adequate to represent our possible epistemic positions for any proposition (*Cf.* Easwaran (2014)). But it remains that credences *aren't* representational states, and unless we should think that one who takes herself to be rationally required to give up most of her interesting beliefs will be comforted when it is pointed out that she continues to have credences, we shoudn't think the threat of controversial view agnosticism can be so easily dispelled.

[22] The claim that 'suspecting-that-p' is a weaker representational state than 'believing-that-p' is compatible with granting that we can (and often do) hold suspicions maximally *vigorously*. The sense in which suspecting-that-p is more weakly affirmative than believing-that-p is (as is articulated further in this section) this: that an agent, in taking up the attitude of suspecting that p, affords more weight to possessing truth than avoiding error, than does one in, by comparison, taking up the attitude of believing-that-p. This point is sharpened further in this and later sections.

[23] See Engel (2005) for an overview. For a defence of truth as a constitutive norm of belief, see Williams (1976), Shah and Velleman (2005). *Cf.* Gibbons (2013).

[24] For more on this point, see Owens (2009), who notes that it is because one's suspecting that p is incorrect when false that suspecting is like belief but unlike attitudes such as imagining.

[25] This is just a sketch of some of the core elements of a more detailed view that has been defended in Carter, Jarvis, and Rubin (2016).

spite jointly comprising the aim maximizing truth and avoiding error. This is because one can only fulfill the first aim (truly representing) by putting oneself at risk of not fulfilling the second (not-misrepresenting). For example, weighting the second aim more than the first would naturally lead one to be more cautious in order to avoid possible misrepresentation[26]. Correspondingly, weighting the first aim more would lead one to be riskier in order to possess more truths[27].

Being certain-that-p (in a sense that will soon be clearer) corresponds with affording much more weight to *avoiding error* than to possessing truth, and *vice versa* for suspecting-that-p. To see this idea in action, suppose there are 100 questions on a multiple-choice test, with *A*, *B*, *C*, and *D* as the available options, for each question. Suppose for each of the 100 questions, you have ruled out at least one or two of *A*, *B*, *C*, and *D* and are, in each case, leaning just slightly toward one of the remaining options.

The idea that *being-certain-that-p* is an attitude that affords much more weight to avoiding error than possessing truth can be appreciated when considering that being-certain-that is not an attitude you take up toward *any* of the 100 questions. You thus forego the chance to possess truth, with *that* attitude, in each case. You do however rightly *suspect* in each of the 100 cases that the answer toward which you're (modestly) leaning is right. It is in this sense that, in taking up the attitude of suspecting-that, you're affording more weight to possessing truths than avoiding error.

Belief, on the simple picture sketched, falls somewhere between suspecting-that and being certain-that. Where *exactly* belief falls in this space is beyond what I can argue for here. Brief reflection, though, indicates that, in taking up the attitude of belief, we are giving more weight to avoiding error than to attaining truth–that this much is so would explain why, for instance, 'belief' is not an attitude I'd take up *vis-a-vis* (3-4), in §2, even though I *suspect* (3-4) are true[28].

[26]*Ibid.*, 2336-7.

[27]A related expression of this point is embraced, albeit within a two-tiered account of knowledge, by Sosa (2015). In particular, on Sosa's model, two types of belief—what he calls *guessing* as well as what *judging*—involve affirming a proposition with the alethic aim of believing truly. However, the two states differ with respect to the level of risk they tolerate. Taking up the attitude of guessing (as when we are asked to read the smallest row of letters in an eye exam) is an affirmation that endeavours to attain truth while tolerating significant levels of risk. Judgment, according to Sosa, is not the kind of attitude we take up when such risk is tolerated. Rather, judgment involves affirming with the alethic aim of believing truly, but also with the epistemic aim of believing aptly—viz., believing truly because of competence (Sosa (2015 Ch. 3)). While I am sympathetic to this proposal, I am not inclined to regard 'guessing' as a variety of belief. Given the level of risk it tolerates, it would be located (on the taxonomy I'm proposing) somewhere closer to suspicion, perhaps riskier. See Carter (forthcoming) for a recent discussion of varieties of guessing, from educated guessing to blind guessing.

[28]A more comprehensive account of this proposal is offered in Carter, Jarvis, and Rubin (2016) and Carter, Jarvis, and Rubin (2015).

EXPERT DISAGREEMENT AND EPISTEMIC DANGER

Against the background of the 'twin goals' model, it should be evident that (for different reasons) some *subject areas* are ones where risk of error is (as it was in the case of the 100 questions featuring in the examination example) simply too great for *belief* to be the appropriate default attitude to take up. Call such subject matters, for ease of reference, *epistemically dangerous*. To offer an easy (albeit, near-trivial) case right out of the blocks, take *astrology*–where the subject matter in question is ways that cosmic alignment affects fortune. To say that believing astrology propositions (i.e. specifying different celestial-fortune relations) is epistemically 'risky' is an understatement. The risk of error is almost maximally high (astrology beliefs, so construed, will be true only in far-off worlds, where the laws of nature are very different.)

Astrology would thus count as epistemically dangerous, for a given thinker, because of properties *internal* to astrology–its own principles. Let's hereafter set astrology (and similar such subjects) aside. The interesting cases I'll turn to will be ones that are, like astrology, epistemically dangerous (e.g. too risky to ordinarily traffic in beliefs), but (unlike astrology) epistemically dangerous because of *relational* or extrinsic properties of the subject matters in question, rather than properties intrinsic to or constitutive of the subject matter[29].

Take here, as a pet case, the subject matter of attention defect hyperactivity disorder (ADHD) management. What combination of counseling, medication and lifestyle choice are most effective in managing ADHD? Here it turns out that the recognised experts are sharply divided, and fundamentally so–those best suited to getting it right on the topic, *if anyone is*, disagree strikingly[30].

This fact looks to be epistemically significant. There are a number ways one might attempt to motivate the point. Matheson and Carey (2013, 131), for instance, have suggested that:

> …significant disagreement amongst the experts indicates that either our evidence on the topic is not very good or that we aren't very good at evaluating it. In either case, agnosticism toward the disputed proposition seems called for.

[29]There is of course also a 'subjective' sense in which a subject matter ϕ could be epistemically dangerous; for instance, ϕ might be dangerous for S_1, but not for S_2, and because of a difference in S_1 and S_2's ϕ abilities. The subject matter of tree- (rather than bird-) identification might in this way be epistemically dangerous *for* the ornithologist, but not the for the arborist. The cases of epistemically dangerous subject matters I'm interested in will not be agent- or ability- relative in this subjective sense. Rather, I'll be exploring cases that are *objectively* epistemically dangerous, given facts independent of what abilities we bring to the table.

[30]See Jadad et al. (1999). Note that there is also controversy over the status of ADHD, apart from issues to do with its management.

I'll outline why I think the kind of expert disagreement we find in cases like ADHD management is epistemically significant in §4.2. I want to leave it open, to be clear, whether and to what extent discovered *peer* disagreement is epistemically significant[31]. What I want to suggest is that the ADHD management case turns out to be closely akin to a range of other controversial subjects[^cs], including subjects that fall within philosophy, politics, religion, and other controversial (or, as Littlejohn puts it, 'interesting', areas[32]). Roughly, the idea I'll pursue now is that the risk of error in these areas is (with reference to our 'twin goals' model) higher than will be tolerated by the attitude of belief.

Centrality and symmetry conditions

It's important to make clearer what I take to be going on in cases where I say 'the experts disagree' in a way that will make the subject area epistemically dangerous (and thus, make trafficking in beliefs problematic). For one thing, the mere *fact* of expert disagreement *vis-à-vis* a given subject, ϕ, isn't *always* epistemically significant. The kind of disagreement matters. Take, for instance, the subject matter of early 20th-century Antarctic exploration. Experts in the area might disagree about whether Shackleton's desperate crew ate more seals or more penguins to stay alive. However, experts unanimously agree *that* Shackleton attempted to reach the South Pole in the early 1900s. We aren't inclined to think that 20th century Antarctic exploration is epistemically dangerous in virtue of the kind of expert disagreement that we find here.

This speaks to a *centrality constraint* on the epistemic significance of expert disagreement; consider that in the ADHD management case, experts are divided on issues *central* to the subject matter in question, whereas, in our case of 20th Cen-

---

[31] One way to see the comparative significance of expert versus peer disagreement will be to contrast the subject of ADHD management with the hot-wire issue of anthropogenic climate change. Many individuals whom I regard as epistemic peers are divided (passionately) here, though experts are not. According to a 2013 metastudy conducted by Cook (2013) in which 11,000 abstracts published in peer-reviewed scientific journals were evaluated, it was found that over 97% of those articles embraced the 'consensus position that humans are causing global warming'. A similar metadata study conducted in 2010 and published in the Proceedings of the National Academy of Science in the US reached a similar conclusion (97%-98%) and also noted that 'the relative climate expertise and scientific prominence of the researchers unconvinced of [Anthropogenic Climate Change] are *substantially below* that of the convinced researchers' So anthropogenic climate change is an issue where my peers are divided, but the experts are not. ADHD management is also a subject where my peers are divided, but so are the experts. Perhaps the fact that many of my peers reject anthropogenic climate change is epistemically significant *vis-à-vis* my beliefs on the topic, even though I am aware that expert opinion is not divided. I won't weigh in here.

[32] As Christensen (2009, 756) observes, 'On many factual questions–examples from politics or economics are easy to think of – widespread disagreement, even among experts, is the norm.'

tury Antarctic exploration, expert disagreement lies merely at the periphery. Indeed, *most* subject matters feature expert disagreement at the periphery–it is thus expert disagreement on matters at the very heart of a subject matter that will be of interest here[33].

Along with a centrality constraint, there will also need to be a kind of *symmetry constraint*; the interesting cases will be cases where there is thought to be *roughly symmetrical* expert disagreement on issues central to a subject matter[34]. Let's assume for a moment that the individuals taking the recent Phil Papers survey are 'experts'. (This is an idealisation, but it will work for illustrative purposes). In that survey, 27 / 931 (2.9%) endorse a truth-relativist semantics for knowledge attributions[35]. Here, the relevant experts are more convinced–(and even more so than they are that anthropogenic climate change is true)–that a truth-relativist semantics for knowledge attributions is false. Thus the experts here 'disagree' in the weak sense that they are not unanimous, but they do not approach anything near roughly symmetric disagreement.

Things are much different though, with respect to the issue of Platonism versus nominalism with respect to abstract objects. Here the Phil Papers survey participants are divided about 50%-50%—a statistic that intuitively tracks a famous division amongs the experts on an enduring metaphysical issue[36]. In the Platonism/nominalism case, it looks like both the symmetrical and the centrality constraints are easily met: the roughly symmetric disagreement concerns not peripheral beliefs, but issues at the very heart of the subject matter–what *kind* of thing is an abstract object; do such things exist? Nothing here is remotely settled.

---

[33]What distinguishes beliefs peripheral to a subject matter from beliefs central to a subject matter? The best account I'm aware of here is found in the literature on objecutal understanding–the state one attains when one counts as understanding a subject matter (e.g. 'Tim understands algebraic geometry'.). As Kvanvig (2003) notes, the distinction between central and peripheral beliefs (with respect to a subject matter) is key to a plausible conception of the how understanding a subject matter should be held to a *factivity* constraint. For instance, two historians can both understand why Caesar was assassinated, even if one (falsely) believes that the co-assassin Brutus was born in late (rather than, correctly, early) June, 85 BC. This false belief is at the periphery of the subject matter in question. This same historian surely fails, however, to understand why Caesar was assassinated if he falsely believes that Brutus killed Cicero rather than Caesar—a proposition central to the subject matter. Accordingly, we can think of the central/peripheral distinction as mapping closely with those beliefs that are essential for understanding a subject matter, and those which are not. This characterisation of the distinction of course identifies the central and the peripheral by their *epistemic* properties. For some further discussion on the factivity condition on understanding, as it stands to illuminate the distinction between central and peripheral beliefs, see also Riggs (2007), Carter and Gordon (2016), Gordon (2017) and Gordon (2012). *Cf.* Elgin (2009).

[34]For discussion on this point, see Grundmann (2013, 73).

[35]See here http://philpapers.org/surveys/results.pl.

[36]See here http://philpapers.org/surveys/results.pl.

Expert Disagreement and Epistemic Risk

Recall now the examination example in §3. In those cases where, for instance, you lean just slightly towards one of the options on a multiple choice test (say, answer *B*), you don't believe *B* is the correct answer, though you do suspect that *B* is right. The attitude of suspecting that, weaker than belief, was argued to differ from belief in that suspecting-that-*p* is an attitude that (on the simple model defended in §3) weights attaining truth more so than avoiding error. One way to think of this point is as follows: the attitude of suspecting-that-*p* *tolerates epistemic risk* to a greater extent than does the attitude of belief. (And, furthermore, the attitude of being certain-that-*p* tolerates little to no epistemic risk[37]).

Plenty of epistemic risk featured in our exam case—enough to preclude the appropriateness of forming beliefs on any of the 100 questions where one merely leans slightly in one direction. But, as was suggested, the level of risk at play did *not* preclude the taking up of the attitude of *suspecting-that*, for any of the 100 questions. Moreoever, given the extent to which suspecting-that-*p* is an attitude that places not just a low premium on avoiding error, but a rather high premium on getting truth, it is fitting (given this fact about suspecting-that) to positively suspect that *p* in cases like the examination example. (Compare: one who leans toward one of the four answers in each case but never purports to suspect any such answers to be right[38].)

Now: does the attitude of *belief* tolerate the kind of epistemic risk at play *vis-a-vis* subject matters where expert disagreement satisfies the symmetry and centrality conditions? I'll suggest not. To see why, consider some parallels between the exam case (featuring 100 questions) and subject matters–like ADHD management, and Platonism/nominalism–that seem to feature symmetrical expert disagreement on issues *central* to the subject matter in question. In the exam case–where it was appropriate to suspect, but not to believe–notice that, firstly, the agent *does* have favouring epistemic support such that: on the matter of whether ϕ, the favouring epistemic support recommends ϕ over its denial[39]. (Compare: this was *not* the case in examples

---

[37]While decision theorists tend to use 'risk' to distinguish decisions under risk (where probabilities of outcomes are known) from decisions under uncertainty (where probabilities of outcomes are not known), this is a technical meaning of the term which I am not relying on. Rather, I'm using 'risk' in the non-technical sense, where 'risk' is taken to be present in situations where it is possible but not certain that an unwanted outcome will matererialise. See here Hansson (2014), §1 and §4 for an overview.

[38]The same point can be made with respect to the appropriateness of certainty. Not only is it inappropriate to take up the attitude of being-certain-that p in the face of epistemic risk, but moreover, we would be inclined to find an agent in error should she purport to be able to rule out all error possibilities, eliminating epistemic risk altogether, but refraining from claiming certainty.

[39]The distinction between favouring and discriminatory epistemic support has been drawn out most clearly in the literature on perceptual knowledge (e.g. Pritchard (2010).). Consider that, even if I can't distinguish a goldfinch from a chaffinch, and I see a bird that looks like it could be either, I might

(5-6) in §2–there I had no favouring epistemic support, and hence, taking a stance would be too risky to be tolerated *even* by the attitude of suspecting-that.)

Secondly, in conjunction with the possession of this (albeit, modest) favouring support in the exam case, the agent finds herself in a situation where, even armed with her favouring support, she appreciates that she can easily be wrong across a wide range of perfectly *normal* circumstances. For instance, when the test-taker is barely leaning toward *B*, suspects that *B*, but finds out *C* is correct, she will not normally be *surprised* to be wrong. And this is just what we should expect. As Scheffler (1991, 12) puts it, 'surprise' is a kind of cognitive emotion 'resting on the epistemologically relevant supposition that what has happened conflicts with prior *expectation*'[40]. When one *believes* that *p*, and comes to find out that not-*p*, it is natural to be surprised–as what has happened does indeed conflict what a kind of *expectation* that is not shared equally when one merely suspects-that something is so.

Accordingly, then, this latter point speaks to the suggestion that what is present in light of the kind of risk we find here is not *belief*; moreover, the former point about favouring epistemic support speaks to the idea that the exam case is *not* one where withholding *entirely* is called for (e.g. refraining from *any* representational attitude) as in the case of (5-6) in §2.

Consider now cases of subject matters that feature expert disagreement that satisfy the centrality and symmetry conditions. Does Platonism give a correct account of abstract objects? Suppose a bright doctoral student in metaphysics–Mette–has read a lot of classic and recent work in the area and thinks that, on balance, the Platonist offers a more compelling case to make than the nominalist. What is the appropriate representational attitude for Mette to take toward the proposition:

> 7. Platonism about abstract objects is true.

In such cases, it's clear that Mette *does* have some epistemic support that favours Platonism over its denial, and so the case at hand recommends (at least) an attitude of suspecting-that Platonism is true, over the alternative of forming no representational attitude whatsoever despite the possession of favouring support.

---

nonetheless possess some *favouring* support, which is some evidence that goldfinches are somewhat more likely to be found than chaffinches in this area. This is support that *favours* one alternative to another, despite one's not being in a position to otherwise discriminate the two. Extrapolating from the perceptual case, we can helpfully think of favouring support (in our exam case) as support that favours one alternative to another despite our lacking the relevant abilities and evidence to warrant a positive expectation that alternative be correct. (Note that this extrapolated sense of favouring support is wider than Pritchard's use of the term specifically with respect to perceptual alternatives.)

[40]See also Lorini and Castelfranchi (2007) for a detailed defence of this kind of cognitive structure of the emotion of 'surprise'. See here also Engel (2014) for a helpful overview.

But, given that experts are in roughly symmetrical disagreement about central issues *vis-a-vis* Platonism and nominalism, Mette has good reason to believe that *experts* aren't more than likely to be correct than incorrect on this issue. Mette should not normally be *surprised* if it turns out she is wrong (just as around half the experts will be)[41]; put another way, Mette finds herself in a situation where, even armed with her favouring support, she appreciates that she can easily be wrong across a wide range of perfectly *normal* circumstances. Mette after all can expect herself to err, despite her favouring evidence, in contexts where she can expect at least half of those most likely to be correct if anyone is, are in error.

With reference to the 'twin-goals' model sketched in §3, then, we have cause to think that Mette should neither (i) form *no* representational attitude whatsoever; nor (ii) form a belief. Moreoever, the kind of epistemic risk present in issues featuring roughly symmetrical expert disagreement is just the sort that we saw to be tolerable by the attitude of suspecting-that. Extrapolating from the case of Platonism and nominalism, the reasons Mette should suspect *p* (when her favouring support points that way) in the Platonism/nominalism case generalise *mutatis mutandis* to many other interesting and entrenched issues where experts are in roughly symmetrical disagreement.

MAKING PEACE WITH CONTROVERSIAL VIEW AGNOSTICISM

Controversial view agnosticism, as was indicated in §1, looked hard to swallow. As Christensen (2009, 758) rightly notes, 'many are quite adverse to thinking that they should be agnostic about all such [epistemically dangerous] matters. And the aversion may be even stronger when we focus on our opinions about politics, economics, or religion'. But we don't merely *dislike* the prospects of controversial view agnosticism; moreoever, it can look like a downright bad outcome for independent reasons. As van Inwagen (1996, 139) puts it, in the arena of philosophy:

> What people have believed about the philosophical theses advanced by– for example–Plato, Locke, and Marx has had profound effects on history. I don't know what the world would be like if everyone who ever encountered philosophy immediately became, and thereafter remained, a philosophical skeptic, but I'm willing to bet it would be a vastly different world[42].

[41] See here Matheson and Carey (2013).

[42] Van Inwagen's remarks—as well as the 'spinelessness' critique of CVA more generally—could be interpreted as *practical* objections. However, there is also an epistemic dimension to these critiques. It is a widespread datum in contemporary epistemology that if one's theory has as a consequence that have much less of a given epistemic state than is ordinarily attributed, that this is a *prima facie* problematic

Van Inwagen here, no less than Feldman, is tacitly exluding the possibility of weaker representational states from the discussion entirely when moving seamlessly from the not-belief result to the kind of skeptical result that is intuitively objectionable to accept for the kinds of reasons indicated. As I've alluded to, this move bypasses the matter of whether the kind of 'agnosticism' featuring a positive weaker representational attitude–in the absence of belief–is really so objectionable.

I want to briefly close with two reasons for thinking it is not–the first to do with the 'social identity' objection sketched in §1, and the second to do with rational action. Recall that controversial view agnosticism is considered a problematic implication for a popular thesis in social epistemology, but this was because of unpalatable consequences it was taken to *entail*. One such alleged entailment was that we must give up a hearty chunk of our social identity and/or interesting views. But the variety of controversial view agnosticism I've opted for here avoids this unwanted implication. Consider that, for one whose thinking is artificially restricted by the Triad Model, of course, the entailment will seem to go through because *withholding* (simpliciter) looks very much like it (as Bogardus (2013, 5) puts it) 'requires an unacceptable degree of spinelessness'. The proposal here, however, insists that just as belief is the wrong attitude, *so too is withholding simpliciter*. One's favouring support, in the face of the kind of epistemic risk one finds when trafficking in a subject matter with symmetrical expert disagreement, rationalises that one positively *commit* to the truth of the proposition *via* a weaker doxastic attitude. There is no in principle problem with one's social identity being understood as in part a function of what one suspects about a range of issues–'belief' is not obviously essential in this respect, even if *some* positive representational attitude is.

Secondly, a reason to think the kind of controversial view agnosticism I'm recommending is not spineless is that the kind of representational commitments prescribed by the version of agnosticism being advanced are commitments that play a role in rational action not played by withholding, *simpliciter*. To take a simple illustrative case, suppose my hand is forced such that I *must* act in some circumstance C where I have (albeit marginal) favouring epistemic support for p over its denial (and I must either act on p or its denial); in C it would be rational for me to act on p and irrational of me not to (e.g. to, rather than act on what I'm warranted to *suspect* is true, to say, flip a coin instead[43]).

---

epistemological consequence of the theory. Spinelessness-style criticisms, even if they lend themselves to plausible practically grounded interpretations, also can be understood as grounded in worries that are epistemological in character. Further to that, it should be pointed out that those partial to Thomas Kelly's 'right reason' account will think, on the basis of entirely epistemological considerations, that widespread withholding of judgment by experts will be irrational.

[43]The same cannot be said, however, for the attitude of withholding *simpliciter*, *vis-a-vis p*. Even if we fix, by stipulation, the practical stakes at 'nothing' (such that I will not be better or worse off by, say,

Thus, the weaker doxastic attitudes that the kind of controversial view agnosticism I'm recommending here save are attitudes that (unlike withholding, *simpliciter*) can plausibly do the work belief does in featuring in one's social identity, and further, will plausibly play important roles in rational action that are not played by attitudes where one lacks favouring support entirely.

CONCLUSION

I've argued here that, once we see how a tacit commitment to the Triad View has the effect of artificially restricting the range of reasonable attitudes we might take up in controversial areas, it becomes relevant and interesting whether–rather than to withhold entirely in such cases–it might be reasonable to take up in such areas attitudes weaker than belief. In §3 I offered a simple framework for thinking about how we can locate the attitude of belief in between two other representational attitudes–suspecting-that-*p* and being certain-that-*p*. In particular, I argued that, although each of these attitudes aims at truth, we pursue the truth goal *differently* in each case, as a function of how we weight the twin goals of attaining truth and avoiding error. I developed this point further in §4 by showing how suspecting-that-*p* tolerates *epistemic risk* to a greater extent than does belief. I proceed to show that the kind of epistemic risk at play when we form beliefs in subject areas where expert disagreement is roughly symmetric is not tolerated by belief, but that it *is* by the weaker attitude of suspecting-that. Moreover, I argue why (in such controversial areas) suspecting that (say) Platonism is true, rather than nominalism, is more reasonable than withholding, *simpliciter*. In §5, I offered some considerations for thinking that this variety of controversial view agnosticism does not entail some of the objectionable consequences associated with 'spineless' controversial view agnosticism, as a thesis that recommends withholding, altogether[44].

The principal application of the model is not to the conformism thesis, *per se*, but rather to the presumption about representational-state options that leads individuals to infer what they take to be an objectionable consequence (i.e., spinelessness) from conformism. However, the model does offer resources for directly modifying the conformist thesis, by offering resources for different and more nuanced picture of what we should infer about our representational state requirements (of which requirements

turning right or left), it won't be rational to turn right *rather* than left unless I at least some favouring grounds for doing so (even if it is rational for me to go right *or* left.)

vis-a-vis beliefs are a species) in the case of a revealed peer disagreement. A development of this further application of the model is a task that would be a fruitful one for another day.

REFERENCES

Ballantyne, Nathan, and EJ Coffman. 2012. "Conciliationism and Uniqueness." *Australasian Journal of Philosophy* 90 (4):657–70.

Bergmann, Michael. 2005. "Defeaters and Higher-Level Requirements." *The Philosophical Quarterly* 55 (220). Oxford University Press:419–36.

Bogardus, Tomas. 2013. "Disagreeing with the (Religious) Skeptic." *International Journal for Philosophy of Religion* 74 (1). Springer:5–17.

Carter, J. Adam. forthcoming. "Sosa on Knowledge, Judgment and Guessing." *Synthese*, 1–20.

Carter, J. Adam, and Emma C. Gordon. 2016. "Objectual Understanding, Factivity and Belief." In *Epistemic Reasons, Norms and Goals*, edited by Pedro Schmechtig and Martin Grajner, 423–42. De Gruyter.

Carter, J. Adam, Benjamin W. Jarvis, and Katherine Rubin. 2015. "Varieties of Cognitive Achievement." *Philosophical Studies* 172 (6):1603–23.

———. 2016. "Belief Without Credence." *Synthese* 193 (8):2323–51.

Chisholm, Roderick M. 1976. *Person and Object: A Metaphysical Study*. G. Allen & Unwin (London).

Christensen, David. 2007. "Epistemology of Disagreement: The Good News." *The Philosophical Review*. JSTOR, 187–217.

———. 2009. "Disagreement as Evidence: The Epistemology of Controversy." *Philosophy Compass* 4 (5). Wiley Online Library:756–67.

Conee, Earl. 2009. "Peerage." *Episteme* 6 (03). Cambridge Univ Press:313–23.

Conee, Earl, and Richard Feldman. 2004. *Evidentialism: Essays in Epistemology: Essays in Epistemology*. Oxford University Press.

Easwaran, Kenny. 2014. "Regularity and Hyperreal Credences." *Philosophical Review* 123 (1). Duke Univ Press:1–41.

Elga, Adam. 2007. "Reflection and Disagreement." *Noûs* 41 (3). Wiley Online

Library:478–502.

Elgin, C. 2009. "Is Understanding Factive?" In *Epistemic Value*, edited by A. Haddock, A. Millar, and D. Pritchard. Oxford: Oxford University Press.

Engel, Pascal. 2005. "Truth and the Aim of Belief." In *Laws and Models in Science*. Citeseer.

———. 2014. "Epistemic Emotions, Nature and Roles." http://www.unige.ch/lettres/philo/episteme/EpistemicEmotions_project.pdf.

Feldman, Richard. 2006. "Epistemological Puzzles About Disagreement." In *Epistemology Futures*, edited by S. Hetherington, 216–36. Oxford University Press.

———. 2007. "Reasonable Religious Disagreements." In *Philosophers Without Gods: Meditations on Atheism and the Secular*, edited by Louise Anthony, 194–214. OUP.

Friedman, Jane. 2013a. "Rational Agnosticism and Degrees of Belief." *Oxford Studies in Epistemology* 4:57–82.

———. 2013b. "Suspended Judgment." *Philosophical Studies* 162 (2). Springer:165–81.

Gibbons, John. 2013. *The Norm of Belief*. Oxford University Press.

Gordon, Emma C. 2012. "Is There Propositional Understanding?" *Logos & Episteme* 3 (2).

Gordon, Emma C. 2017. "Understanding in Epistemology." *Internet Encyclopedia of Philosophy*.

Grundmann, Thomas. 2013. "Doubts About Philosophy? The Alleged Challenge from Disagreement." In *Knowledge, Virtue, and Action*, edited by T. Henning and D. Schweidard, 72–98. Routledge.

Hansson, Sven Ove. 2014. "Risk." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Spring 2014. http://plato.stanford.edu/archives/spr2014/entries/risk/.

Hazlett, Allan. 2012. "Higher-Order Epistemic Attitudes and Intellectual Humility." *Episteme* 9 (3):205–23.

Inwagen, Peter van. 1996. "It Is Wrong, Everywhere, Always, for Anyone, to Believe Anything Upon Insufficient Evidence." In *Faith, Freedom and Rationality*, 137–54. Savage, Maryland: Rowman; Littlefield.

Jadad, Alejandro R, Lynda Booker, Mary Gauld, Ritsuko Kakuma, Michael Boyle, Charles E Cunningham, Marie Kim, and Russell Schachar. 1999. "The Treatment of Attention-Deficit Hyperactivity Disorder: An Annotated Bibliography and Critical Appraisal of Published Systematic Reviews and Metaanalyses." *Canadian Journal of Psychiatry* 44 (10):1025–35.

Kelly, Thomas. 2005. "The Epistemic Significance of Disagreement." *Oxford Studies in Epistemology* 1. Oxford: Clarendon Press:167–96.

Kvanvig, J. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.

Lackey, Jennifer. 2008. "A Justificationist View of Disagreement's Epistemic Significance." In *Proceedings of the Xxii World Congress of Philosophy*, 53:145–54.

Littlejohn, Clayton. 2013. "Disagreement and Defeat." In *Disagreement and Skepticism*, edited by D. Machuca, 169–93. Routledge.

Lorini, Emiliano, and Cristiano Castelfranchi. 2007. "The Cognitive Structure of Surprise: Looking for Basic Principles." *Topoi* 26 (1):133–49.

Machuca, Diego E. 2013. *Disagreement and Skepticism*. Routledge.

Matheson, John, and Brandon Carey. 2013. "How Skeptical Is the Equal Weight View?" In *Disagreement and Skepticism*, edited by D. Machuca, 131–50. Routledge.

Owens, David. 2009. "Freedom and Practical Judgement." In *Mental Actions*, edited by L. O'Brien and M. Soteriou. Oxford University Press.

Pritchard, Duncan. 2010. "Relevant Alternatives, Perceptual Knowledge and Discrimination." *Noûs* 44 (2). Wiley Online Library:245–68.

Riggs, Wayne D. 2007. "Understanding 'Virtue' and the Virtue of Understanding." In *Intellectual Virtue: Perspectives from Ethics and Epistemology*, edited by M. Depaul and L. Zagzebski. Clarendon Press.

Scheffler, Israel. 1991. *In Praise of the Cognitive Emotions*. Routledge.

Shah, N., and J. D. Velleman. 2005. "Doxastic Deliberation." *Philosophical Review* 114 (4):497–534.

Sosa, Ernest. 1991. *Knowledge in Perspective: Selected Essays in Epistemology*. Cambridge University Press.

———. 2015. *Judgment and Agency*. Oxford University Press, USA.

Turri, John. 2012. "A Puzzle About Withholding." *The Philosophical Quarterly* 62 (247). Wiley Online Library:355–64.

Williams, Bernard. 1976. "Deciding to Believe." In *Problems of the Self: Philosophical Papers 1956-1972*, 136–51. Cambridge University Press.