

PHILOSOPHICAL INDIVIDUALISM

John A. Keller

*If two people agree, one of them isn't a philosopher.*¹

1. INTRODUCTION

It is natural to think that successful philosophical arguments are *ideally publicly recognizable*: recognizable by all rational, informed, and fair-minded inquirers. Peter van Inwagen, for example, has defended a view according to which successful philosophical arguments are those that would be convincing—when ideally presented in the company of an ideal opponent—to an audience of ideal neutral agnostics about the disputed question. Whether an argument meets this criterion is ideally publicly recognizable: insofar as the merits of an argument are not publicly recognized, it is because of ignorance, irrationality, or bias.²

Of course, things are not ideal, and the actual public is bound to misclassify some arguments. Still, we might feel comfortable formulating coercive public policy on the basis of arguments that meet van Inwagen's criterion, since anyone that rejects the conclusion of such an argument is suffering (perhaps non-culpably) from one of the aforementioned maladies. The hope that such arguments provide the means to peaceably settle public disputes has been pervasive and influential. John Rawls, for example, holds that

...citizens are to conduct their public political discussions of constitutional essentials and matters of basic justice within the framework of what each sincerely regards as a reasonable political conception of justice, a conception that expresses political values that others as free and equal also might *reasonably be expected reasonably to endorse*. (Rawls 1993, p.xlvii, emphasis added)

Along related lines, Timothy Williamson notes that:

¹ This quip appears on p.3 of Rescher 1985. I have been unable to determine its provenance.

² The idealization of the audience ensures that they respond *rationaly*; the presence of an ideal opponent ensures that the audience is *informed*; and the audience's neutral agnosticism ensures that they are not *biased* for or against the conclusion. See van Inwagen 2006, ch.4. Of course, such an ideal audience can't be informed in ways that undermine its neutrality.

...we might hope that whether a proposition constitutes evidence is *in principle* uncontentionably decidable, in the sense that a community of inquirers can always in principle achieve common knowledge as to whether any given proposition constitutes evidence...³ (Williamson 2007, p.210)

But alas, according to van Inwagen, this hope is in vain: *there are no* arguments for substantive philosophical conclusions that meet his criterion, or any other publicly recognizable criterion, even in ideal circumstances. If this is correct, we face a dilemma: conclude that *there are no* successful arguments for substantive philosophical theses, or adopt a non-ideally publicly recognizable conception of success. But how could successful arguments fail to be even *ideally* publicly recognizable? Is such a watered down notion of “success” really success at all? My aim in this essay is to argue that it is. There are good *independent* reasons to adopt a (non-publicly recognizable) “private” or “individualistic” conception of success: the fact that we thereby gain a few successful arguments for substantive philosophical conclusions is just icing on the cake.

2. PHILOSOPHICAL SUCCESS

Most philosophers do not think that all philosophical arguments are failures. Most of us find *some* philosophical arguments convincing, our own if no others. But how many of the arguments we deem successful are for *substantive* conclusions? How many are constructive, as opposed to destructive, telling us the way the world is, rather than how it isn't? And if we think we can point to some successful constructive arguments, what proportion of other philosophers agree with us? Since philosophers disagree about almost everything, every argument for a substantive philosophical thesis can be rejected without manifest irrationality.⁴ If successful philosophical arguments must convince all rational comers, it would appear that there are none. According to van Inwagen, however, this sets the bar too high. He holds that an argument that does not convince *all* rational comers can be successful if it would convince all *ideal* rational comers. More carefully, he holds that an argument for conclusion *c* is successful just if it would be convincing—when ideally presented in the company of an ideal

³ Note that common knowledge about what the evidence *is* wouldn't necessarily lead to common knowledge about what the evidence *supports*. In any case, Williamson argues that it is not even commonly knowable what the evidence is. See §3.3.

⁴ The PhilPapers survey and common experience both testify to the lack of agreement amongst philosophers. See Bourget and Chalmers (2014) for a summary of the results of the survey. What it means for a conclusion to be substantive is trickier than it might seem: see §2 of Kelly and McGrath 2015 for a nice discussion.

opponent—to an audience of ideal neutral agnostics about c . Ideal neutral agnostics about c have no opinion or leanings concerning the truth or falsity of c , and are as intellectually virtuous as is humanly possible. Since it is easier to convince an audience of neutral agnostics than it is to convince an audience of die-hard opponents, this account sets the bar for philosophical success lower than an account that requires successful arguments to convince all rational comers.

The primary motivation that van Inwagen gives for his view hinges on this insight. Van Inwagen argues that any plausible account of success other than his own is bound to run afoul of

The Desideratum A criterion of success *should* not entail that there are *no* successful arguments for substantive philosophical conclusions.

The Desideratum expresses a plausible constraint on theories of success. But there is an even stronger—indeed, non-negotiable—constraint, which may be expressed as follows:

The Weak Desideratum A theory of success *must* not entail that paradigmatically successful arguments—arguments for *established* scientific, historical, and non-substantive philosophical conclusions—are not successful.

Any account of success that runs afoul of The Weak Desideratum is clearly inadequate. Another such non-negotiable desideratum on a theory of success is:

The Requirement A criterion of success must entail that those in possession of a successful argument for conclusion c are, *ceteris paribus*, not rationally criticizable for accepting c .

It is plausible that The Requirement is not only necessary, but also *sufficient* for success. When someone believes something *outré*, we think they had better have an argument for it. And not just a clever or interesting argument. (This is especially clear if their belief is both *outré* and *objectionable*.) Whatever else successful arguments must do, then, they must make it rational, or at least not irrational, to accept their conclusions.⁵ For reasons that will become clear in §4.3, however, I do not think that satisfying The Requirement is in fact sufficient for

⁵ See §4 and §5 for further discussion. Note that, for reasons discussed in §2.2.1, §2.3 and §4.4 (and in Keller 2015b), we should not require successful arguments to be such that one could *come to* know their conclusions *on the basis* of their premises.

success. But that is, I think, a surprising result: The Requirement articulates what I take to be the core of our pre-theoretic notion of success.

2.1 Objective Success. Van Inwagen’s conception of success is *dialectical*—it measures success in terms of its ability to convince some audience. But one might wonder why we shouldn’t adopt an *objective* (attitude independent) conception of success? Why not hold that successful arguments are *sound*? Call this the **Soundness Account**. Isn’t the Soundness Account what we teach our students? Since sound arguments for substantive philosophical conclusions obviously *exist*, the Soundness Account does not run afoul of the Desideratum. It also satisfies the Weak Desideratum. The problem is with The Requirement: sound arguments can have premises that we do not know or even believe to be true. We know *of* sound arguments for substantive conclusions about free will, morality, metaphysics, epistemology, etc. The problem is that we know of unsound arguments too, and we don’t know *which* arguments are the sound ones. For all we know, the zombie argument for property dualism is sound.⁶ But for all we know it isn’t. For all we know, the consequence argument for incompatibilism is sound,⁷ and for all we know it isn’t. And so on.

This isn’t just that we don’t recognize that goodness of some sound arguments. Some sound arguments are *bad*. Consider the following arguments:

- | | |
|---|--|
| 1. The Peano axioms | 1. The Peano axioms |
| 2. The Peano axioms \supset Goldbach’s Conjecture | 2. The Peano axioms \supset Goldbach’s Conjecture is false |
| 3. <i>Therefore</i> , Goldbach’s Conjecture | 3. <i>Therefore</i> , Goldbach’s Conjecture is false |

One of these arguments is sound. Call that the **Bad Sound Argument**. The Bad Sound Argument is not just sound: it is a *transparently valid* sound argument. But it isn’t a successful argument, in any interesting sense of ‘successful’. In fact, it seems to be a manifestly *bad* argument—if the Bad Sound Argument is the only argument I have in support of its conclusion, I should not believe its conclusion. The reason for this is simple: I have no reason to accept its second premise.⁸ The Soundness Account, then, runs afoul of The Requirement. Having a sound argument for a conclusion *c* does not make it rationally defensible to accept *c* if one does not believe, much less justifiably believe, the premises of

⁶ Chalmers 1996

⁷ van Inwagen 1983

⁸ Given that Goldbach’s Conjecture has been neither proven nor disproven, what could that reason be?

that argument. Some sound arguments have premises that are not believed, others have premises that are literally unbelievable, and some have premises that no human could be justified in believing.⁹ So the Soundness Account cannot be correct.

Can the Bad Sound Argument can be ruled out as fallacious? If so, it would not run afoul of the **Soundness+ Account**, according to which *non-fallacious* sound arguments are successful. But what fallacy might the Bad Sound Argument illustrate? The only plausible candidate is the fallacy of begging the question (if there *is* such a fallacy), but if there is some subtle sense in which the Bad Sound Argument is question begging, that sense will almost certainly not be objective: it will have to do with how one could *know* the premises and conclusion of the argument, or with the context in which an argument is used. For example, Sinnott-Armstrong 1999 argues that “to avoid begging the question one’s reason to believe the premise must be independent of both (a) one’s belief in the conclusion and also (b) one’s reason to believe the conclusion.” Perhaps some uses of the Bad Sound Argument beg the question in this sense, but this sense of begging the question is not objective: it is relativized to the beliefs and reasons possessed by the individuals assessing the argument.

Objective criteria like soundness may be necessary for success, but they are not sufficient. To satisfy The Requirement, successful arguments must at least have premises that are rationally *believable*. But then it is obvious why we need a dialectical conception of success: different things are rationally believable to different people. A sound argument may fail to convince me—it may be *rational* for it to fail to convince me—if I have no reason to believe its premises. So, since even paradigmatically non-deductive arguments can be easily formalized (by adding premises), a theory of success is primarily a theory about the class of individuals that need to be convinced by an argument in order for it to be a success. In other words, a theory of success is primarily a theory about the proper *audience* for philosophical arguments.¹⁰

⁹ Trivially, any premise that is literally unbelievable is a premise that we could not justifiably believe, but there are believable truths for which we could not have justification: e.g., many truths about events outside of our light cone.

¹⁰ The idea that argumentative success is objective seems to be a legacy of the unfortunate idea that successful arguments are *proofs*: transparently (at least step-wise) valid arguments with communally apodictic premises. But if there are premises that are communally apodictic, we are obviously unable to prove much of scientific or philosophical interest from them.

2.2 Dialectical Success. One could take a number of different views about the nature of this audience.

2.2.1 Everybody. One view is that the proper audience for an argument is *everybody*, or perhaps everybody rational. As noted above, however, no argument for a substantive philosophical conclusion has convinced all comers. The problem isn't just that this criterion fails to satisfy The Desideratum—it fails to satisfy the Weak Desideratum as well. As Williamson notes,

Having good evidence for a belief does not require being able to persuade all comers, however strange their views, that you have such good evidence. No human beliefs pass that test. (Williamson 2007, p.212)

Logical heretics reject standard mathematical proofs, but this does not show that those proofs are not successful arguments. This is especially obvious if the heretics are wrong. But a correct proof is a paradigmatically successful argument, and (even heterodox) logicians are paradigmatically rational. So successful arguments cannot be defined as arguments that everyone (or everyone rational) will accept.¹¹

Indeed, as the sociologists of science delight in pointing out, scientific theories often become dominant only because their opponents *die*. It is much easier to convince (initially agnostic) graduate students of some theory *t* than it is to convince those whose careers have been defined by their opposition to *t*. Einstein famously (or infamously) never fully endorsed quantum mechanics, but that hardly entails that there weren't successful arguments for the standard theory of elementary particles prior to Einstein's death. Ernst Mach went to his grave dissenting from the atomic theory of matter, etc.

This fact—that successful arguments do not always convince everybody, or even everybody rational—is recognized in certain special cases. It is widely conceded, for example, that no one has devised an anti-skeptical argument that is convincing to skeptics. But that does not mean that there are no successful responses to skepticism. As Williamson 2007 notes,

[I]f one uses only premises and forms of inference that a skeptic about perception will allow one...one has little prospect of reaching the conclusion that one has hands. But that does not show that we should not be confident that we have hands. (p.238)

¹¹ See Williamson 2007, ch.4, for a discussion of related issues.

Of course, skeptics will say that...claims about our environment [like “dreams with the sustained coherence of waking life are very rare”] merely beg the question...But the claims were not addressed to skeptics, in a futile attempt to persuade them out of skepticism. Instead, they figure in our appraisal of skeptical arguments, from our current non-skeptical point of view.¹² (p.249)

On this view, non-skeptics can successfully defend their knowledge by demonstrating that the cost of skepticism is unacceptably high, even if skeptics themselves couldn't come to know that we have knowledge *on the basis of* those arguments, since those arguments *presuppose that* we have knowledge (e.g., knowledge that “dreams with the sustained coherence of waking life are very rare”).¹³ Many non-skeptics *know* that they have such knowledge, so this presupposition is (at least according to Williamson) a legitimate one. Still, it is not a presupposition skeptics can coherently accept. Such anti-skeptical arguments are *essentially defensive*, but they do satisfy The Requirement. This, in conjunction with our stipulation that the non-skeptics *know* the premises of their arguments to be true, strongly suggests that these non-skeptical arguments are successful.

Now, one might think that those who hold onto their pet theories in the face of successful arguments to the contrary are either irrational or uninformed, and hence that the view that successful arguments must convince *idealizations* of one's opponents does not run afoul of our desiderata. But however much we might wish it were so, skeptics are not always irrational or uninformed. Similarly, Einstein's certainly wasn't uninformed about quantum mechanics, and it is far from obvious that his resistance to it was *irrational*.

In any case, there is an obvious reason why different people will be convinced by different arguments. An argument for *c* is a presentation of some reasons or evidence for *c*. But whether an argument gives one a compelling reason to accept *c* depends on one's other evidence, one's priors, etc. If, as some people believe, no consistent set of priors is more rational than any other, it follows that no argument is guaranteed to make it irrational for one's interlocutors to reject *c*. For if one's interlocutor's priors related to *c* can be arbitrarily low, it will take an argument of arbitrary strength to make it irrational for her to reject *c*.

¹² See Pryor 2000, Pryor 2004, and Kelly 2005 for a nice discussion of related issues.

¹³ C.f. the Introduction to Lewis 1983 on arguments being tools for “measuring the cost” of denying their conclusions. See also Keller 2015b.

Similarly, since one's interlocutor might have arbitrarily strong evidence (however misleading) against c , it will take an argument of arbitrary strength to make it irrational for her to deny c .¹⁴ If this is correct, The Requirement and The Weak Desideratum jointly entail that successful arguments cannot be those that convince all rational comers.

2.2.2 *Agnostics*. And so we are led to a conception of success like van Inwagen's, where a successful argument for c must only be convincing to ideal neutral *agnostics* about c , even if it is not convincing to *everyone*, including rational opponents of c . But does this criterion do any better than the others with respect to our desiderata? It plausibly satisfies The Requirement and The Weak Desideratum, but are there arguments for substantive philosophical conclusions that are convincing, when ideally presented in the presence of an ideal opponent, to all rational agnostics? It is hard to believe that there are, since, over time, one would think that the existence of such arguments would lead to significantly more convergence on substantive philosophical theses than we in fact observe. As van Inwagen puts it,

If any reasonably well-known philosophical argument for a substantive conclusion had the power to convert an unbiased ideal audience to its conclusion (given that it was presented to the audience under ideal conditions), then, to a high probability, assent to the conclusion of that argument would be more widespread among philosophers than assent to any substantive philosophical thesis actually is.¹⁵ (van Inwagen 2006, p.53)

But if van Inwagen's criterion fails to satisfy The Desideratum, why doesn't he adopt a more liberal one? "Alas," van Inwagen writes, "there is no more liberal criterion. The criterion I have proposed is the most liberal possible criterion." (p.160, n.5). If all criteria of philosophical success are partners in guilt when it comes to The Desideratum, however, one might wonder why we should favor van Inwagen's criterion over any other.

There are two aspects of van Inwagen's criterion that make it preferable to the others we have discussed so far: as noted above, it seems to satisfy The Requirement and The Weak Desideratum, and it also does a better job of assigning the "burden of proof" than these

¹⁴ I take belief and knowledge to be more fundamental than credences and priors, but talk of credences and priors is sometimes illuminating. Having "high" or "low priors" may be paraphrased as being inclined to believe or disbelieve for those who prefer sticking to a single idiom.

¹⁵ See Kelly and McGrath 2015 for a critique of this line of reasoning, and Ballantyne 2015 for an argument that there are "knockdown arguments" for substantive philosophical conclusions. (But see Keller 2015b for a critique of Ballantyne's argument.)

other criteria. According to van Inwagen, arguments for c may properly rely on premises that *opponents* of c reject, as long as those premises would be acceptable to *agnostics* about c . As van Inwagen says,

Norma the nominalist need not worry about whether Ronald the realist will accept her premises. She is perfectly free to employ premises she knows Ronald will reject; her only concern is whether the audience of agnostics will accept these premises. Suppose, for example, that she uses the premise, “We can have knowledge only of things that have the power to affect us.” It may well be that no realist...would accept that premise...[But if] Ronald thinks that there is any danger of the agnostics accepting this premise, it will do him no good to tell the audience that of course no realist would accept this principle and that it therefore begs the question against realism. He’ll have to get down to the business of convincing the agnostics that they should reject, or at least not accept, this premise. (van Inwagen 2006, p.46)

The idea that we may properly use premises that our opponents reject is an important insight that an adequate account of success should preserve. But this reason for preferring van Inwagen’s account is beside the point, since van Inwagen’s is not the most liberal possible criterion. **Philosophical Individualism**—roughly, the view that an argument \mathcal{A} for conclusion c is successful for individual i if and only if \mathcal{A} is convincing to i (regardless of i ’s previous attitude towards c)—is more liberal, in that it allows some arguments for substantive philosophical conclusions to be successes. And so, it would seem, The Desideratum supports philosophical individualism over both van Inwagen’s view and the other more demanding views he rejects.¹⁶

2.3 Philosophical Individualism. What makes philosophical individualism “individualistic” is that it relativizes success to individuals, holding that ‘ x is a successful argument’, like ‘ x is in motion’, has implicitly relativistic truth-conditions, such that an argument might be a

¹⁶ Fischer and Tognazzini 2007 suggests a liberalization of van Inwagen’s view that would only require successful arguments to *sway*, rather than *convince*, an idealized neutral audience—to get them to change their *credences*, even if they don’t change their *minds*. A similar account is proposed in Chalmers 2015, and van Inwagen has indicated to me (pc) that he now endorses this modification of his view. There are problems, however: if the account merely requires successful arguments to raise the credences of ideal neutral agnostics *at all*—say, from .50 to .51—it runs afoul of The Requirement, since not all such arguments articulate satisfactory reasons to accept their conclusions. If the account requires successful arguments to raise the credences of neutral agnostics to the threshold of rational belief, however, it satisfies The Requirement but collapses back into van Inwagen’s original account. See §3 for further reasons to be skeptical of this approach.

success for you without being a success for me. Philosophical individualism relativizes the success of an argument to the person evaluating it. There are several different flavors to choose from:

The Belief Account: an argument is successful for an individual just if she *believes* it is sound (and non-fallacious)

The True Belief Account: an argument is successful for an individual just if she *correctly* believes it is sound (and non-fallacious)

The Justification Account: an argument is successful for an individual just if she *justifiably* believes it is sound (and non-fallacious)

The Knowledge Account: an argument is successful for an individual just if she *knows* it is sound (and non-fallacious), or

The KK Account: an argument is successful for an individual just if she *knows that she knows* it is sound (and non-fallacious).¹⁷

One way to think about these criteria is as specifying the type of *epistemic gain* typically produced by successful arguments: new beliefs, new *true* beliefs, new *justified* beliefs, new *knowledge*, etc. An argument that one, e.g., *knows* to be sound (and non-fallacious) typically could be used to gain knowledge of its conclusion. But I want to remain neutral about whether successful arguments must *always* be able to provide one with some epistemic gain, since that would rule out essentially *defensive* arguments—arguments that can only be used to defend conclusions one already accepts (see §2.2.1 and Keller 2015b)—and arguments for things we know *essentially*, such as, perhaps, that we exist. While it isn't clear that there *is* anything we know essentially, there *could* be an essentially omniscient being, or at least a maximally epistemically good being—an epistemic *god*—for whom no argument could provide anything of epistemic value. Still, an epistemic god could have successful arguments—if anyone has successful arguments, epistemic gods do! Those arguments would not provide the god with any epistemic gain, but they would still satisfy The Requirement—they would still demonstrate the rationality of the god's beliefs.

¹⁷ Recall that we are focusing on deductive arguments, and note that when I say that *i* knows/believes/etc. an argument *A* to be sound, this should be read as shorthand for “*i* knows/believes/etc. *A*'s premises and *i* either knows/believes/etc. that *A* is valid, or *A* is transparently valid”. This is important since someone could know, by testimony, *that A* was sound without having a clue what the argument was. Since such a person may not even *believe A*'s premises, they do not know *A* to be sound in the relevant sense.

In any case, satisfying the Knowledge Account seems clearly *sufficient* for success. This account is more liberal than van Inwagen's—it satisfies The Desideratum—since there are some arguments for substantive philosophical conclusions that are known by some people to be sound. For example, since I know that there are anatomical properties shared by spiders and insects, and I know that *there are anatomical properties shared by spiders and insects* entails *there are properties*, I have a successful argument for the existence of properties. Successful for me, that is. It isn't *generally* successful: most nominalists don't believe, and so don't know, that that entailment holds. And since I know this argument to be sound, it is clearly rational for me to accept its conclusion—thus satisfying The Requirement. The criterion also seems to satisfy the Weak Desideratum: can an argument be paradigmatically successful if it is not known to be sound?

In Section 4 I argue that the Knowledge Account is correct: that knowing an argument to be sound is also *necessary* for it to be successful. But determining *which* properties of an individual success is relativized to is less important than realizing *that* success is relativized to individuals in the first place. The idea that correctly evaluating an argument does not depend on anyone's attitude towards the argument but one's own is a significant departure from—and improvement upon—both our natural way of thinking about success and the dominant trend in the literature. An account of success that relativizes it to something else about individuals—say, their justified beliefs—would still be better than one that relativizes success to the properties of some third party.

The picture of philosophy as the *first-person pursuit of understanding*, found in the work of philosophers like Aristotle, Descartes, and Marcus Aurelius, aligns well with such individualistic conceptions of success. We are often assisted in this pursuit by interaction with and input from others, but it is our own epistemic position that is the arbiter of success. Even if we typically hope that our arguments will be of use to other people struggling with the same or similar problems, the success of our arguments is to be measured in terms of their ability to solve *our* problems. As Robert Nozick says,

Some of the things the skeptic says or points out...I accept; these are or become part of *my own* belief system. *My* problem is that I don't see (or no longer see, after the skeptic has spoken) how these things go along with yet other things in my belief system...My task here is to remove the conflict, to put *my own* beliefs in

alignment, to show how those of the things the skeptic says which I accept can be fit in with other things I accept.¹⁸ (Nozick 1981, p.16)

All that is required to be a philosophical individualist is to think that the dialectical standards Nozick claims are appropriate for arguing with the skeptic apply in general (recall §2.2.1). Our goal is not—or should not be—to convince our opponents that they are wrong. It is to see for *ourselves* why they are wrong—to uncover the flaws in their arguments. Arguments can be successful for some people and not others, and the aim of philosophical inquiry is producing arguments that are successful for *us*.

3. IN SUPPORT OF PHILOSOPHICAL INDIVIDUALISM

We saw above that The Desideratum, The Weak Desideratum, and The Requirement seem to support (some species of) philosophical individualism. But there are a number of additional reasons for thinking that success is individualistic.

3.1 Essentially First-Person Arguments. Sometimes I give arguments of the form “I remember that p , so p ”. In many such cases I couldn’t have a better argument for p —vividly remembering that p trumps, *rationality* trumps, almost any degree of external objective evidence that not- p . Sometimes my “memory arguments” are convincing to others, but there are cases where others ought not be convinced, even though it is rational for me to be. For example, let p be the claim that I was not at the scene of Jones’s murder, and suppose I have been exquisitely framed for Jones’s murder. It might then be rational for me to maintain my innocence in the face of the objective evidence to the contrary, since I vividly remember doing something else at the time of Jones’s murder. But it might *not* be rational for the jury to agree with me. Memory arguments are much *more* compelling when it is one’s own memory that is being appealed to—in the first person case they are almost always sufficient to produce knowledge.¹⁹ Other such essentially first-person arguments seem relevant to debates about phenomenology, motivational internalism/externalism, the significance of religious experience, and one’s own existence. Even if Descartes didn’t mean to be giving an *argument* for his own existence (since he was suspending belief about the reliability of his reasoning), *we* can formulate *Cogito*-esque arguments—essentially first-person arguments for our own

¹⁸ Here Nozick seems to be assuming that *believing* an argument to be sound is sufficient for success. This assumption is widespread: see, e.g., Sinnott-Armstrong 1999.

¹⁹ This is not to say that knowledge based on memory is typically formed *via* memory arguments: normally such knowledge is non-inferential.

existence—and those arguments are clearly successful. ‘I think, *therefore* I am’ is a *superb* argument for my existence: indeed, it is rationally compelling. Rationally compelling for *me* that is. It is much less compelling to others, including neutral agnostics about my existence.

3.2 Against Neutrality. The primary appeal of criteria like van Inwagen’s over philosophical individualism is that, by removing the requirement that successful arguments be convincing to a *neutral* audience, philosophical individualism allows *bias* to (partially) determine argumentative success or failure.²⁰ While an agnostic about *p* neither accepts nor rejects *p*, a *neutral* agnostic about *p* assigns *p* and not-*p* equal probability. There are, however, a number of problems with thinking that successful arguments must convince neutral agnostics. First, neutral agnosticism is infectious. If neutral agnostics about *p* only had to be agnostic about *p* itself, some such agnostics might firmly believe that almost all arguments for or against *p* are sophistical. Such agnostics would be unlikely to be convinced by *any* argument for or against *p*, successful or not.²¹

But once we see that neutral agnostics about *p* must be agnostic about more than *p* itself, it is hard to know where to stop. Would neutral agnostics about the existence of free will have to be compatibilists? Moral realists? Consequentialists? Physicalists? Theists? Scientific realists? Even if our agnostics were strictly neutral about the existence of free will, consider how differently “primed” towards free will someone would be who was an incompatibilist, moral nihilist, atheist, physicalist, and scientific realist compared with someone who was a compatibilist, moral realist, theist, dualist, and scientific anti-realist. Of course, one *can* believe in free will as an incompatibilist, moral nihilist, atheist, physicalist, and scientific realist without manifest irrationality: Mark Balaguer accepts all of those except moral anti-realism,²² and there is no reason to think that adding moral anti-realism to the mix would make his position inconsistent. But convincing Balaguer to reject free will seems much easier than convincing someone who is agnostic about free will, but who is a compatibilist, moral realist, theist, dualist, and scientific anti-realist. Explicit opposition to a claim doesn’t guarantee that the totality of one’s views doesn’t “mesh” better with that claim than does another totality that is neutral about the claim.

²⁰ Where by ‘bias’ I mean simply non-neutrality. Such bias needn’t be irrational: it might be rational, or it might be non-rational without being irrational. (E.g., according to subjective Bayesians, having consistent but loaded priors makes one non-rationally but not irrationally biased.)

²¹ Thanks to John Hawthorne for this point.

²² See, e.g., Balaguer 2010. I should note that Balaguer’s acceptance of moral realism, free will, and incompatibilism is hedged.

Of course, to be truly neutral with regard to free will one would also have to be agnostic about determinism, consequentialism, reductionism, and much else. If the members of a jury of “ideal neutral agnostics” were not agnostic about these theses, they might come to different verdicts on the basis of their differing opinions. But an audience that was agnostic about all of these topics would believe so little that it is hard to imagine convincing them of anything at all.

An additional problem is that agnostics are not neutral in debates about agnosticism itself, religious or otherwise. Pyrrhonian skeptics aim only to get us to *suspend belief* about whether there is an external world, so someone who is agnostic about whether there is an external world already accepts the skeptic’s conclusion.²³ When it comes to arguments that are skeptical in this sense, there is no neutral vantage point. But if anti-skeptical arguments do not need to appeal only to premises acceptable to the skeptic, and they do not need to appeal only to premises acceptable to someone agnostic about skepticism, what is left? The only plausible answer is that they can appeal to premises that non-skeptics themselves accept. But that supports philosophical individualism over views that claim that successful arguments must utilize premises acceptable to everyone, a neutral audience, etc.²⁴

3.3 Against Idealization. There are also problems with the idea that the success of an argument depends on what some *idealized* group of people think or feel. Most obviously, it is unclear how the effect of an argument on such an audience would be relevant to the argument’s *actual* satisfaction of The Requirement. Whether an argument produces some epistemic gain (justification, knowledge, etc.) depends crucially on the beliefs of the *actual* person to whom it is addressed. An idealized version of myself might not have the problem that the argument I am considering is attempting to remedy. Much philosophy aims at solving problems that arise only because of our non-ideality, after all. Philosophers aim to discover the entailments of claims like ‘Aquinas admired Aristotle’, ‘Lois believes that Clark Kent cannot fly’, and ‘Joe freely chose to lie to Mary’. Does the first entail that there are objects that are not present? Does the second entail that Lois believes that Superman cannot

²³ Can one coherently be an agnostic about whether we can know the truth-value of x without being an agnostic about x ? If so, one could coherently claim to know that, say, God exists while being agnostic about whether we can know if God exists. And that does not seem coherent.

²⁴ Note that it is not clear to what extent bias accounts for disagreement in philosophy, since there is persistent apparently rational disagreement about philosophical questions about which no one (or hardly anyone) has a preconceived opinion: Newcomb’s Paradox, the nature of mental representation, quantifier variance, whether properties are parts of their instances, and so on.

fly? Does the third entail that Joe was not determined? These are paradigmatically philosophical questions that a logically ideal audience would have no reason to ask.²⁵

A logically ideal audience would not only fail to face such philosophical problems, such an audience could not be *neutral* about the solutions to such problems, since they would *know* the solutions. Combining ideality and neutrality generates problems that ideality and neutrality do not face individually.²⁶ Finally, note that many people are persuaded by arguments aiming to show that there are no ideal persons. The inability of those arguments to convince an audience of ideal persons is clearly beside the point.²⁷

3.4 Non-Neutrality. Further support for philosophical individualism may be derived from “evidence non-neutrality”. Williamson 2007 argues that we should not expect all (rational) parties to be able to agree on what the evidence is, or what it supports. On Williamson’s view, one’s evidence is one’s knowledge, and so, since different people know different things, an argument might be compelling evidence for someone who knows its premises but not for someone who doesn’t. The non-neutrality of evidence is not tied to Williamson’s account of evidence, however: taking one’s evidence to be one’s justified beliefs makes evidence no more neutral than taking it to be knowledge.²⁸ And on *any* account where evidence is non-neutral, there are striking implications for how we evaluate arguments—implications that lend support to individualistic conceptions of success:

How much do failures of Evidence Neutrality threaten the conduct of philosophy? From an internal perspective, they make consensus harder. Each of many conflicting theories may be the one best supported by the evidence by its own lights. The role of evidence as a neutral arbiter is undermined. From an external perspective, both the good fortune of being right and the misfortune of being wrong are magnified. If your theory is true, so are its consequences for which propositions constitute evidence...If your theory is false, it may have false

²⁵ Above, I glossed van Inwagen’s ideal audience as being as rational as is *humanly possible*, rather than *logically omniscient*. This may help with the objection here, but it opens space for the reaction of such an audience to be the result of logical mistakes or shortcomings.

²⁶ Kelly and McGrath 2015 argue along similar lines that ideal agnostics are not the relevant audience for determining success.

²⁷ Thanks to John Hawthorne for this point.

²⁸ If, as Williamson 2000 argues, no or virtually no conditions are *luminous*, there is nothing capable of playing the evidence role that we can expect all parties to be in a position to recognize.

consequence for which propositions constitute evidence...(Williamson 2007, p.213)

A related view is defended by Michael Bergmann, who argues that *believed defeaters are defeaters*: that one cannot know that p if one believes, even incorrectly, that one has a defeater for p .²⁹ If we think of defeaters for p as a form of evidence against p , this view also entails that evidence is non-neutral, since different people—even people who have the same “objective” defeaters for p —may *believe* they have, and so have, different defeaters for p .³⁰

Robert Nozick has also defended a position along these lines regarding the resolution of *antinomies*: apparently inconsistent sets of apparently true sentences. We resolve an antinomy when we come to see that one of those appearances is illusory—that the apparently inconsistent sentences are not inconsistent after all, or that the apparent truth of one or more of the sentences is an illusion. As Nozick puts it,

Given the (apparent) incompatibility between the apparent [truths] and p , there are two ways to continue to maintain... p . First, one of the apparent [truths] can be denied, or there can be a denial of their conjunction all together...Second, each of the apparent [truths] can continue to be maintained, while their apparent incompatibility with p is removed, either by close scrutiny showing the reasoning from them to not- p to be defective, or by embedding them in a wider context or theory that specified how p holds in the face of these apparent [truths].³¹ (Nozick 1981, p.10)

Since different apparent truths appear true to different people, and different claims appear incompatible to different people, this way of thinking about philosophical inquiry supports an individualistic conception of success. It is the apparent (to *me*) truths that appear (to *me*) to be inconsistent with my other beliefs that I have to worry about—you may have a different set of worries, or no worries in the neighborhood at all. An argument that alleviates my worries might not alleviate yours, and *vice versa*.

3.5 Bayesianism. Bayesians hold that what is rational to conclude on the basis of new arguments or evidence depends on one’s priors. But people have different priors, and

²⁹ See, e.g., Bergmann 2005.

³⁰ What I say here applies most straightforwardly to rebutting defeaters, but a similar (albeit more complicated) story could be told about undermining defeaters.

³¹ Nozick’s embedding strategy seems like an instance of the close scrutiny strategy, but since the two kinds of argument can have different “feels” to them I have left the passage intact.

according to (subjective) Bayesian orthodoxy no set of coherent priors is more rational than another. Hence, two people exposed to the same evidence and arguments may reach different conclusions. Indeed, if they are rational, they *must* reach different conclusions, each in accordance with her priors. But then there can be no objective fact about whether an argument makes acceptance of its conclusion rational, and hence whether the argument is successful. For this will depend on one's priors, and priors are individualistic. Hence, so is argumentative success.³²

3.6 Reflective Equilibrium. It is widely thought that philosophy aims at bringing our judgments about principles and cases into reflective equilibrium. Since our initial judgments about cases and principles often differ, we will often come to different points of reflective equilibrium. As David Lewis says:

The reader in search of knock-down arguments in favor of my theories will go away disappointed...when all is said and done, and all the tricky arguments and distinctions and counterexamples have been discovered, presumably we will still face the question which prices are worth paying...On this question we may still differ. And if all is indeed said and done, there will be no hope of discovering still further arguments to settle our differences...Our "intuitions" are simply opinions; our philosophical theories are the same. Some are commonsensical, some are sophisticated; some are particular, some general; some are more firmly held, some less. But they are all opinions, and a reasonable goal for a philosopher is to bring them into equilibrium. Our common task is to find out what equilibria there are that can withstand examination, but it remains for each of us to come to rest at one or another of them...

If you say flatly that there is no god, and I say that there are countless gods but none of them are our worldmates, then it may be that neither of us is making any mistake of method. We may each be bringing our opinions to equilibrium in the most careful possible way, taking account of all the arguments, distinctions, and counterexamples. But one of us, at least, is making a mistake of fact. (Lewis 1983, p.x)

³² Yes, priors can often be "swamped" by the evidence. (See, e.g., Doob 1971 and Gaifman and Snir 1982.) But they are not *always* swamped, and so as long as we hold that any consistent set of priors is rational—or even a wide range of consistent priors—exposure to the same arguments will not necessarily lead to agreement, even between people who respond to evidence perfectly.

If the goal of philosophical inquiry is to reach reflective equilibrium, we succeed in philosophy if we reach reflective equilibrium. Different people can come to different points of reflective equilibrium, however, even if they are familiar with the same arguments. But if it is possible for you and me to know the same arguments and distinctions, and think they support different conclusions, and be making no mistake of method, then the success of an argument must be relativized to individuals.³³

Some object to the method of reflective equilibrium because it is insufficient to produce *justification*. Van Inwagen, for example, agrees with Lewis that equally competent and informed philosophers often reach different states of equilibrium. But he takes this to be *prima facie* incompatible with their being *justified* in holding those positions, and suggests that the fact that many of one's peers do not share one's point of reflective equilibrium should knock one out of it, at least if one is aware of this disagreement.³⁴ But it is not clear why we should think that knowing that one's peers do not share one's point of reflective equilibrium should knock one out of it. Reaching different states of reflective equilibrium is a species of peer disagreement, and it is far from clear that rational peer disagreement is impossible.³⁵

Of course, that doesn't entail that being in reflective equilibrium is sufficient for justification. But this is only important if successful arguments are required to provide justification for believing their conclusions—if satisfying The Requirement is actually necessary (and perhaps sufficient) for success. It is to this question that we now turn.

4 ALTERNATIVE FORMS OF INDIVIDUALISM

I claimed above that *i*'s knowing that *A* is sound is *sufficient* for *A*'s being a success for *i*. It is tempting to think that this is also a necessary condition. Call the resulting view the *Knowledge Account*. Knowledge is paradigmatically epistemically valuable, so it is easy to see why arguments that can (typically) be used to produce knowledge are successes. But a number of things other than knowledge are also sometimes held to have epistemic value: belief, true belief, justified belief, etc. So let us consider accounts of success framed in terms of an argument's potential to produce these.

³³ Cappelen 2012 has a nice discussion of van Inwagen and Lewis's view that what we call "intuitions" are just things we believe. This view seems to support the idea that an argument being a success for *i* simply requires *i* to *believe* it to be sound. See §4.1.

³⁴ In van Inwagen 2004. See also Kelly and McGrath 2010.

³⁵ Partly for reasons outlined by van Inwagen himself: see, e.g., van Inwagen 1996 and Kelly 2010.

4.1 The Belief Account. The weakest plausible account of success holds that an argument A for conclusion c is successful for someone just if they *believe* it to be sound. Call this the *Belief Account*. On this account, successful arguments will often have the power to produce belief in their conclusions. There are two reasons why we might be drawn to such an account. First, the Belief Account delivers more successful arguments than the Knowledge Account. Second, the Belief Account seems *phenomenologically* correct: when I evaluate an argument, I ask myself whether I *believe* it is sound. If I do, then, all else being equal,³⁶ I count it a success and accept its conclusion.

There are four main problems with the Belief Account. First, it runs afoul of The Requirement: simply *believing* an argument to be sound does *not* justify one in accepting its conclusion. Second, merely acquiring new beliefs is not, except perhaps incidentally, a form of epistemic gain—gaining new false and unjustified beliefs is not epistemically valuable. Third, the Belief Account makes successful arguments implausibly easy to come by. And fourth, it allows there to be successful arguments for false or contradictory conclusions.

4.2 The True Belief Account. So let us turn to the *True Belief Account*, according to which an argument A is successful for i just if i *correctly* believes A to be sound. According to the True Belief Account, successful arguments will (typically) be able to produce true beliefs, and true beliefs are (plausibly) epistemically valuable. The True Belief Account also fixes various other problems with the Belief Account: the True Belief Account does not make successful arguments implausibly easy to come by, nor would it allow successful arguments for false or contradictory conclusions. True beliefs seem to be at least *prima facie* valuable, and since there is no special phenomenology attached to *true* beliefs, phenomenological considerations do not give us a reason to prefer the Belief Account to the True Belief Account. Finally, the True Belief Account, like the Belief Account, makes successful arguments for substantive philosophical conclusions more common than the Knowledge Account. Despite all of these advantages, however, the True Belief Account runs afoul of The Requirement. Since The Requirement is non-negotiable, the True Belief Account cannot be correct.

4.3 The Justified Belief Account. What we might call the *Justified Belief Account* seems tailor-made to satisfy The Requirement, and will satisfy The Desideratum and The Weak Desideratum at least as well as the Knowledge Account. Is the Justified Belief Account

³⁶ If the conclusion is less plausible than the conjunction of the premises, all else is not equal.

therefore preferable to the Knowledge Account? There are three main problems. First, if the Justified Belief Account is not to collapse into the Knowledge Account, we must assume that justification is fallible. But then the Justified Belief Account entails that there could be successful arguments for contradictory or false conclusions. As noted above, this seems like a serious problem. Indeed, it suggests an additional desideratum on a theory of success:

The Constraint A theory of success should *not* entail that there are successful arguments for contradictory/false conclusions.

Of course, the Justified Belief Account wouldn't allow arguments for contradictory conclusions to be successful for the same people (at the same time), at least if we assume that one can't be simultaneously justified in believing both p and $\text{not-}p$.³⁷ But this only partially ameliorates the problem: it is at least somewhat implausible that I could have a successful argument for p while you have a successful argument for $\text{not-}p$. And of course, the Justified Belief Account still allows arguments for false conclusions to be successes.

The second problem with the Justified Belief Account is that, if there is a class of beliefs in no need of justification—one way of understanding the notion of a “foundational” belief—then the Justified Belief Account will misclassify arguments based on such foundational beliefs. If my beliefs about the reliability of sense perception, the validity of *modus ponens*, and the reflexivity of identity need no justification to be knowledge, I can successfully argue to other conclusions on the basis of such premises, despite their lack of justification.³⁸

Finally, for paradox of the preface type reasons, the Justified Belief Account doesn't actually ensure that one is justified in believing the conclusions of successful arguments: I may be justified in believing each claim on a long list, and I may be justified in believing that conjunction introduction is valid, but not be justified in believing the conjunction of those claims. Indeed, I might be justified in believing the *negation* of that conjunction!³⁹ Perhaps this

³⁷ This will be false if propositions are individuated coarsely—e.g., if propositions are sets of worlds.

³⁸ Of course, we might say that foundational beliefs are immediately or non-inferentially justified, rather than unjustified, so this point is not decisive. However, it is natural to describe foundational beliefs as those that need no justification. See, e.g., Sinnott-Armstrong 1999, “sometimes one has *no* reason at all to believe [p], because [p] is self-evident and *needs* no justification.”

³⁹ Note that the True Belief Account is not afflicted by this kind of worry. If I *correctly* believe an argument to be sound, then, assuming I believe the conclusion, I correctly believe it. It is unclear that the paradox threatens the Knowledge Account. If I *know* the premises of an argument, and I know it to be valid and non-fallacious, then it is very plausible that, assuming I believe the conclusion, I know it. See Hawthorne 2004, p.49 for discussion.

falls under the *ceteris paribus* clause of The Requirement, but it is an unlovely result nonetheless.

4.4 Summing Up. In light of the above, I see no reason not to accept the Knowledge Account. I should stress, however, that the primary aim of this paper is to argue that philosophical individualism is true, not to defend any particular species of the view.

5 SUCCESS AND PROGRESS

It is natural to think that successful arguments should lead to *agreement* about the truth of their conclusions, and that such agreement is required for there to be real *progress* in a discipline.⁴⁰ So progress, agreement, and success appear to be closely related, if not the same topic under different guises. Philosophical individualism, however, holds that *agreement* is fundamentally distinct from *success*: that an argument can be successful for someone who knows it to be sound even if no one else agrees.

But while the separation of agreement and success can be jarring, it is hard to see a reason why they *must* be linked. What about success and progress? Once we realize that it is possible for someone to know premises that others reject, it is difficult to see how the production of a sound argument for some interesting philosophical conclusion using those premises could fail to be a kind of progress for that person, at least if producing such an argument was her goal. Surely nothing more is required for progress than getting closer to one's goal. And so, if success is individualistic, so is progress. Agreement and convergence, on the other hand, are essentially social.

5.1 Essentially Social Success. Some might insist, however, that *success* is essentially social—that philosophical success must be recognized by the philosophical community. In certain contexts, such as high school debate, success *is* essentially social. One's goal in a high school debate is to convince the judges of one's conclusion—or at least to convince them that one's argument is better than one's opponents. A debater's opinion of her own arguments is irrelevant. Likewise, one might think, with philosophical arguments and the philosophical community: successful arguments are those that the *community*, or some privileged subset of the community, deems successful. Call this the Communitarian Account of success.

⁴⁰ As noted by van Inwagen in his discussion of philosophical failure. See also van Inwagen 2004, Wilson 2013, and Chalmers 2015.

Note, however, that if we are willing to attribute propositional attitudes to collections of individuals, then philosophical individualists can give a nice account of social success: an argument \mathcal{A} is successful for collection of individuals c just in case c knows that \mathcal{A} is sound.⁴¹ Call collections of individuals that can bear propositional attitudes ‘social individuals’. There is no reason that a philosophical individualist cannot count such social-individuals as individuals. But she still can and must relativize success to *proper* individuals as well. If, for example, I were shipwrecked on a deserted island, I might spend my time thinking about the Sorites paradox. While on the island, I might solve the puzzle, coming to know that an argument with, say, epistemicism as its conclusion was sound. I would then no longer be bothered by the paradox, and would have a better understanding of vagueness than I began with. It is clear that by any reasonable standard I have made philosophical progress, and that the argument for the solution I have reached is a success.

Now suppose that after I am rescued from the island I cannot find anyone willing to publish my results. Indeed, suppose that no one else finds my argument convincing. Would that mean that the argument that *was* successful is *now* a failure? How could that be? A sound argument for epistemicism, which I *know* to be sound, is successful for me even if it convinces no others. While philosophy is social in many important ways, it is not *essentially* social. It can be practiced on a deserted island, and it can be practiced there successfully.⁴²

It is worth noting that while it is possible to rob a person p of her knowledge that an argument \mathcal{A} is sound—if, say, p ’s interlocutors refuse to take the idea that that \mathcal{A} is sound seriously—is it much *easier* to rob someone of their due influence in the philosophical community. While this latter form of epistemic injustice is an important one, I take it to be a virtue of the Knowledge Account that it allows marginalized philosophers—philosophers to whom few listen—to produce successful arguments. If we thought (perhaps implausibly) that successful philosophical research involved the production of successful arguments, a social conception of success would entail that by refusing to take the arguments of marginalized

⁴¹ Everyone, I suppose, would be willing to attribute knowledge of p to a collection c (at least in a loose sense) if everyone in c knows that p . The more interesting question is whether a collection c can know that p even if some members of c don’t know that p . It is often claimed, for example, that the Greeks knew that the Earth was round, but obviously not *every* Greek knew that the Earth was round. Perhaps this example can be accommodated by treating ‘the Greeks’ as a generic (see, e.g., Leslie 2007), but other examples attributing knowledge to Congress, one’s students, etc. may not be so easily handled.

⁴² This is of course consistent with the claim that truth is *best* pursued in an intellectual community.

philosophers seriously we thereby *make it the case* that the research of those philosophers is not successful. That is surely not true.⁴³

In any case, there are paradigmatically successful arguments that are, in principle, successful only for the person giving them. Consider:

1. Only John knows that he is in pain.
2. John is tired.
3. *Therefore*, there are things that no non-tired person knows.

This is a *great* argument: assuming that there are no omniscient beings, I know that it is sound. But it is not an argument that could be used to (rationally) convince anyone else of its conclusion, including social individuals, since no one other than I can know its first premise.⁴⁴

6 SUCCESS PLURALISM

The arguments of the previous section aimed to show that there must be an individualistic notion of success. But some who grant that might still insist that *in practice* philosophy is almost always social, and that when we argue in a social context, our arguments have social success conditions. After all, we often want our arguments to convince others, or at least to be well received. So some might hold that we should be *success pluralists*, holding that there are (at least) two kinds of success, corresponding to two different goals we might have: *personal* success, which is individualistic, and *public* success, which is not.

We have already seen the problems with several versions of non-individualism, and the Communitarian Account, introduced in the previous section, fares no better. One problem is that it runs afoul of The Desideratum: no arguments for substantive philosophical conclusions have convinced the philosophical community. A second problem is that, to make the account plausible, we are almost certainly going to need to put some constraints on who belongs to the philosophical community, and weigh the opinions of different members differently—convincing “the community” had better not require convincing everyone who

⁴³ See, e.g., Fricker 2007 for a nice discussion of epistemic injustice. While this essay is pushing back *a bit* against the “social turn” in epistemology—a turn for which Fricker is the standard bearer—it is worth noting that much or all of what Fricker says makes sense within a classical epistemological framework, and that much of what I say here is compatible with her view.

⁴⁴ Premise 1 is what we might call a “Fitch-style” proposition, but rather than being true and unknowable, it is something that only I can know. Thanks to John Hawthorne for suggesting this argument. As noted in §3.1, there are a variety of other arguments that are (plausibly) essentially first-personal, including the *Cogito*.

happens to belong to the APA or teach philosophy, else it will run afoul of The Weak Desideratum. A natural thought would be to give extra weight to the opinions of *experts*, but that requires a criterion for expertise. And in any case, there is often *more* disagreement amongst the experts than there is in the general philosophical population: the opinions of logicians are more divided about the validity of *modus ponens* than are the opinions of philosophers as a whole (c.f. §2.2.1).

A third problem is that people can belong to more than one philosophical community. If success were a matter convincing one's philosophical community, contradictions would arise whenever one belonged to philosophical communities that evaluated arguments differently. Since this happens, the Communitarian Account flirts with incoherence. If we avoid incoherence by explicitly relativizing success to philosophical communities, two other problems arise.

First, if the evaluator of an argument A belongs to communities that disagree about the merits of A , she will not be able to say whether A is successful. She will be able to say that it is successful relative to one of her communities and a failure relative to another. But she will not be able to univocally say that the argument is, or is not, a success. This lack of a unique result is for most practical purposes equivalent to arriving at no result at all. Second, in making success explicitly relational, we have taken a big step towards philosophical individualism. Once we have reached this point, we should ask whether it is preferable to relativize success to communities or to individuals. Reasons for relativizing to individuals were given in §3, but one that I have not discussed is that it is possible for one to know that the evaluation of one's community is *wrong*. This can happen in ordinary cases, I think, but even if not it is clearly possible for there to be *exceptionally* good philosophers—people who are significantly better at philosophy than everyone else in their community. Such a philosopher might be able to formulate arguments that her coworkers cannot fully understand or appreciate. But it wouldn't follow that her arguments were not successful.⁴⁵

Finally, note that philosophical individualism is perfectly compatible with the view that an argument A can be a success relative to one community and not to another—i.e., if the members of the first community know that A is sound while the second do not. The philosophical individualist merely insists that one's arguments can be successes—for

⁴⁵ Does this undermine our reason for thinking that van Inwagen's criterion runs afoul of the Desideratum? To some extent, yes. But there are other reasons for preferring philosophical individualism to van Inwagen's account: see §3.

oneself—even if they are not success relative to one’s philosophical community, or indeed any (interesting) philosophical community at all.

6.1 Pluralism Revisited. Even if the Communitarian Account is unsatisfactory, one still might think that there are different legitimate notions of success, only one of which is individualistic. While we may have reason to think that *some* success is individualistic, why think that this is the *only* kind of success? Indeed, given our proclivity to speak as if being successful is a *property* of successful arguments, as opposed to a relation they stand in, and given our proclivity to speak as if all and only sound (non-fallacious) arguments are successful, one might think that we have a powerful reason to think that some objective or at least non-relativistic notion of success is the primary one.⁴⁶

To answer this objection we must distinguish between descriptive theses about how ‘success’ is used from normative theses about how it should be used. Philosophical individualism is fundamentally a normative thesis: it is about the *proper* standards for evaluating arguments, the standards to which arguments *should* be subjected. So it cannot be refuted simply by noting that it contradicts certain entrenched patterns of use. But if the account of success provided by philosophical individualism does not fit with our use of ‘success’, one might wonder how it can claim to be an account of *success* at all? The short answer is: just as the account of addition provided by Peano arithmetic is the correct account of adding, despite its failing to perfectly fit our (imperfect) use of ‘add’. The long answer would require a lengthy discussion of semantic externalism and meaning magnetism that would be inappropriate here, but let me gesture at the shape that this answer would take. Semantic externalism plausibly entails that *if* the normative claim that ‘success’ should be used individualistically is correct, it follows that philosophical individualism is also an account of the standard of success that we have (largely unwittingly) been appealing to all along. Since the appeal will have been largely unwitting, this account will not fit perfectly with use: it will be a case where use is being trumped by other meaning-determining considerations.⁴⁷ But actual use of ‘success’ is a mess, with the most common use of ‘success’ in philosophical contexts assuming that successful arguments are proofs, or at least valid arguments with premises that are commonly known. These standards are untenable.

⁴⁶ I thank David Chalmers, Lorraine Juliano Keller, and Steve Petersen for pressing me on this issue.

⁴⁷ See Sider 2012 for a general discussion of such cases, and Keller 2015, §4.2 for a discussion of how use can be trumped.

This, by itself, doesn't refute success pluralism. Even if there is *a* sense in which success is individualistic, there could be others. Why not say that (non-fallacious) sound arguments are *objectively* successful, that arguments that convince our opponents are *dialectically* successful, that arguments that would convince an idealized audience are *ideally* successful, that arguments with apodictic premises are *demonstrably* successful, etc.? The short answer is that this multiplication of senses of 'success' is more trouble than it is worth. Consider the notion of "objective success", according to which sound arguments are successful. Is this any kind of success at all? Is there *any* sense in which the Bad Sound Argument is successful? It is difficult to see what that sense could be. If someone is asked why she believes an unproven arithmetical theorem, giving the Bad Sound Argument is a wholly unsatisfactory response, and satisfying this Requirement seems like a non-negotiable condition on being a successful argument.

What about dialectical success: convincing some audience, whether one's opponents or a neutral or idealized audience? Note first that philosophical individualism can accommodate the idea that successful arguments convince one's audience, since it entails that an argument is successful *for one's audience* just if one's audience knows that the argument is sound. The advocate of dialectical success says that this is not enough: that we need two species of success, "audience success" and "arguer success".⁴⁸ But doesn't it make more sense to analyze audience success and arguer success in terms of success simpliciter? A philosophical individualist can say that an argument is an audience success just if it is a success for the audience, and an argument is an arguer success just if it is a success for the arguer. This seems more parsimonious than, and hence preferable to, thinking that there are two distinct notions of success in play.

Of course, there *are* a plurality of kinds of 'success' in the following sense: convincing one's opponents, or the philosophical community more generally, are all goals someone might have, even if it is not the goal by which we measure the success of an argument *qua* argument. Convincing people other than oneself might be a professional goal, or a personal goal, or some other kind of goal. That's all true, but irrelevant: whether an argument achieves such non-epistemic goals is distinct from whether it is successful as an argument. It is not impossible, or even uncommon, to use bad arguments because they will be convincing to

⁴⁸ C.f. Sinnott-Armstrong 1999 on "audience justification" and "arguer justification". But aren't arguers part of their own audience, at least if they are arguing in earnest? Yes: this is another reason to accept philosophical individualism, with a univocal notion of "evaluator success".

others: that would be one way to define ‘propaganda’. But we shouldn’t dignify propaganda by requiring successful arguments to be useful propaganda tools. Successful arguments don’t always make for successful propaganda, and *vice versa*. An argument’s lack of ability to produce converts may point to a flaw in the argument as a political tool, but it doesn’t necessarily point to any flaw in the argument itself. The idea that convincing others is a kind of *epistemic* goodness for arguments presupposes that the people one is trying to convince will respond rationally to the argument, and are in a position to recognize the truth of its premises. But alas, that is not always the case.

What, finally, about van Inwagenian *ideal* success? This kind of success would avoid the problem afflicting non-ideal species of dialectical success: an ideal audience will presumably respond rationally to arguments, and may well know any knowable premises. But there are other problems with this view, as we saw in §3. Once we have idealized away the limitations that undermined the other dialectal conceptions of success, we may well have idealized away the limitations that give rise to philosophical problems in the first place. An audience that could just *see* that Goldbach’s Conjecture followed from the Peano axioms might find the Bad Sound Argument convincing. That doesn’t mean that the Bad Sound Argument is good, it just means that such idealized audiences aren’t a barometer of success.

6.2 Arguments and Proofs. Perhaps the best argument for pluralism about success comes from the fact that successful mathematical arguments are *proofs*, while successful arguments in philosophy and other disciplines may fall short of this standard. This would seem to indicate that there are at least two notions of argumentative success.

It is non-trivial to give a satisfactory account of mathematical proof, but we can bypass this complication by noting that there is a glaring difference between mathematical proofs and successful philosophical arguments: the axioms (and theorems) that mathematicians use in their proofs are almost universally agreed upon (within the relevant mathematical community) and perfectly precise. Virtually every mathematician knows that the proofs (in her sub-discipline, at least) are sound, and knows that with something like certainty. Does this mean that there is a different standard for success when it comes to mathematics? No it does not. If mathematical axioms and theorems are commonly known, there is a *de facto* “public” and apodictic conception of success in mathematics (since validity is publicly recognizable). But there is no need to treat this as a different kind of success than the success we have in in other disciplines. Rather, we can attribute the public nature of successful

mathematical arguments, and the certainty of their conclusions, to the public and certain knowledge of their premises. It is perfectly consistent with philosophical individualism that there are some, indeed many, publicly recognizably successful arguments, and even publicly recognizable demonstrations. And if we *can* account for the differences between math and philosophy in this way, considerations of parsimony suggest that we *should* account for them this way. There is no need to postulate two kinds of success, when we already know that premises can be publicly or merely privately known, with or without certainty.⁴⁹

6.4 Against Success Pluralism. The fundamental argument against success pluralism, then, is that there just aren't a plurality of viable notions of success. Every account of success we have considered, other than van Inwagen's Account, the Knowledge Account, and Justification Account, runs afoul of The Weak Desideratum, or The Requirement, or both. Van Inwagen's account is incompatible with the existence of essentially first-person arguments, evidence non-neutrality, and subjective Bayesianism. Furthermore, complications arising from the ideality and neutrality of his audience threaten its satisfaction of The Weak Desideratum and The Requirement. This still leaves room for a weak form of pluralism that countenances just two kinds of success, as outlined by the Knowledge and Justification Accounts. But both of those accounts are both forms of individualism about success, and, perhaps more importantly, The Constraint seems to rule out the Justification Account. There is no reason, then, to be a success pluralist.

7 CONCLUSION

According to philosophical individualism, the success of an argument is relativized to individuals. Seeing that this is the case is an important improvement upon other ways of thinking about success, regardless of the species of philosophical individualism we accept. I have defended the version of individualism that requires someone to *know* an argument to be sound in order for it to be a success for her, but those with different philosophical temperaments might be inclined to accept one of the other accounts discussed here, or perhaps another account altogether.

⁴⁹ Similarly, the difference between mathematical and scientific arguments is that scientific arguments contain premises that are less publicly known, and which are not known by anyone with the kind of certainty with which the mathematical community knows its premises.

It goes without saying that philosophical individualism has radical implications for how we should evaluate the success of philosophical arguments. Van Inwagen suggests that successful arguments would convince a substantive portion of the philosophical community, and we often teach our students that arguments should have uncontroversial premises, and in particular premises that opponents of the argument's conclusion can accept. If philosophical individualism is true, none of this is correct. If we accept the Knowledge Account, we should rather ask whether we know the argument to be sound and, perhaps, whether the person giving it does. Nothing else matters.

Of course, it follows from philosophical individualism that success is not ideally publicly recognizable. But this is not as bad as it seems: it does not, for example, entail that success is *subjective* in any pejorative sense. For the most plausible species of philosophical individualism hold that successful arguments must be sound. And the soundness of our arguments is something that is (typically) independent of what we think or feel. Still, philosophical individualism is not as good as we might have hoped. For if success is individualistic, we cannot always use reason to peaceably settle public disputes.⁵⁰ This is disappointing, but it is unavoidable: it is the human condition. Of course, we should continue to use reason to settle as much as possible. Uncertainty, too, is part of the human condition, but we should try to be as certain as possible about important public policies. Likewise, even if we cannot hope to rationally convince *everyone* to accept some policy, we should try to rationally everyone we *can*. We just need to remember that this is a *political goal*, rather than an *epistemic requirement*.⁵¹

REFERENCES

Balaguer, M. (2010), *Free Will as an Open Scientific Problem* (MIT Press).

Ballantyne, N. (2014), 'Knockdown Arguments', in *Erkenntnis* 79.

Bergmann, M. (2005), 'Defeaters and Higher-Level Requirements', in *The Philosophical Quarterly*, 55(220).

⁵⁰ At least in the most straightforward sense. Reason may dictate that, if success is individualistic, we should accept a democratic conception of success for arguments in the public square. This would not, however, be using reason to settle public disputes. Rather, it would be using reason to find the best non-rational method for settling public disputes.

⁵¹ Thanks to Nathan Ballantyne, David Chalmers, John Hawthorne, Lorraine Juliano Keller, Stephen Kershner, Steve Petersen, Joshua Smart, and Peter van Inwagen for helpful comments and discussion.

- Bourget, D., and Chalmers, D. (2014), 'What Do Philosophers Believe?', in *Philosophical Studies* 170(3):465-500.
- Cappelen, H. (2012), *Philosophy without Intuitions* (Oxford University Press).
- Chalmers, D. (1996), *The Conscious Mind* (Oxford University Press).
- Chalmers, D. (2015), 'On the Limits of Progress in Philosophy', in J.A. Keller (ed.), *Being, Freedom, and Method: Themes from the Philosophy of Peter van Inwagen* (Oxford University Press).
- Doob, J.L. (1971), 'What is a Martingale?', in *American Mathematical Monthly* 78:451–462.
- Fischer, J.M., and Tognazzini, N.A. (2007), 'Exploring Evil and Philosophical Failure: A Critical Notice of Peter van Inwagen's *The Problem of Evil*', in *Faith and Philosophy* 24(4):458–474.
- Fricker, M. (2007), *Epistemic Injustice: Power and the Ethics of Knowing* (Oxford University Press).
- Gaifman, H., and Snir, M. (1982), 'Probabilities Over Rich Languages', in *Journal of Symbolic Logic*, 47:495–548.
- Goldman, A. (1999), *Knowledge in a Social World* (Oxford University Press).
- Hawthorne, J. (2004), *Knowledge and Lotteries* (Oxford University Press).
- Keller, J.A. (2015a), 'Paraphrase, Semantics, and Ontology', in K. Bennett and D. Zimmerman (eds.), *Oxford Studies in Metaphysics*, volume 9 (Oxford University Press).
- Keller, J.A. (2015b), 'On Knockdown Arguments', in *Erkenntnis*.
- Kelly, T. (2005), 'Moorean Facts and Belief Revision, or Can the Skeptic Win?', in *Philosophical Perspectives: Epistemology* 19.
- Kelly, T. (2010), 'Peer Disagreement and Higher Order Evidence', in R. Feldman and T. Warfield (eds.), *Disagreement* (Oxford University Press).
- Kelly, T., and McGrath, S. (2010), 'Is Reflective Equilibrium Enough?', in *Philosophical Perspectives* 24(1):325–359.
- Kelly, T., and McGrath, S. (2015), 'Are There any Successful Philosophical Arguments?', in J.A. Keller (ed.), *Being, Freedom, and Method: Themes from the Philosophy of Peter van Inwagen* (Oxford University Press).
- Leslie, S.-J. (2007). 'Generics and the structure of the mind', in *Philosophical Perspectives: Philosophy of Mind*, 21: 375–403.
- Lewis, D. (1983), *Philosophical Papers*, Volume 1 (Oxford University Press).

- Nozick, R. (1981), *Philosophical Explanations* (Belknap Press).
- Pryor, J. (2000), 'The Skeptic and the Dogmatist', in *Noûs* 34:517–549.
- Pryor, J. (2004), 'What's Wrong with Moore's Argument?', in *Philosophical Issues* 14(1):349–378.
- Rawls, J. (1993), *Political Liberalism* (Columbia University Press).
- Rescher, N. (1985), *The Strife of Systems: An Essay on the Grounds and Implications of Philosophical Diversity* (University of Pittsburgh Press).
- Sider, T. (2012), *Writing the Book of the World* (Oxford University Press).
- Sinnott-Armstrong, W. (1999), 'Begging the Question', in *Australasian Journal of Philosophy* 77(2):174–191.
- van Inwagen, P. (1983), *An Essay on Free Will* (Oxford University Press).
- van Inwagen, P. (1996), 'It is Wrong, Everywhere, Always, and for Anyone, to Believe Anything upon Insufficient Evidence', in E. Stump and M.J. Murray (eds.), *Philosophy of Religion: The Big Questions* (Blackwell).
- van Inwagen, P. (2004), 'Freedom to Break the Laws', in *Midwest Studies in Philosophy*, XXVIII.
- van Inwagen, P. (2006), *The Problem of Evil* (Oxford University Press).
- Williamson, T. (2000), *Knowledge and Its Limits* (Oxford University Press).
- Williamson, T. (2007), *The Philosophy of Philosophy* (Blackwell Publishers).
- Wilson, J. (2013), 'Three Dogmas of Metaphysical Methodology', in M. Haug (ed.), *Philosophical Methodology: The Armchair or the Laboratory?* (Routledge).