

Trying Cognitivism: A Defence of the Strong Belief Thesis

by

AVERY ARCHER

George Washington University

Abstract: According to the *Strong Belief Thesis* (SBT), intending to **X** entails the belief that one will **X**. John Brunero has attempted to impugn SBT by arguing that there are cases in which an agent intends to **X** but is unsure that she will **X**. Moreover, he claims that the standard reply to such putative counterexamples to SBT—namely, to claim that the unsure agent merely has an intention to try—comes at a high price. Specifically, it prevents SBT from playing the kind of explanatory role the cognitivist requires. This paper meets Brunero’s challenge to SBT by offering an account of trying and intending to try that not only saves SBT from Brunero’s criticism, but does so in a way that preserves the explanatory significance cognitivist typically take SBT to have.

Keywords: cognitivism, intending to try, means-ends coherence, intention, practical reason

1. Introduction

According to the *Strong Belief Thesis* (SBT), intending to **X** entails the belief that one will **X**.¹ SBT is often employed as a premise in the cognitivist argument that the rational norms governing intention may be *derived from* (Bratman 2009a: 30), *is a special case of* (Bratman 2009b: 15), or *explained in terms of* (Ross 2009: 243) the rational norms governing belief.² For example, it is argued that if intending to **X** entails the belief that one will **X**, then we may make sense of why it is irrational to have inconsistent intentions in terms of the rational norm against inconsistent beliefs.³ Hence, the cognitivist recasts what at first glance appears to be two separate questions as a single

¹ Brunero (2014: 18).

² For the sake of simplicity, I will henceforth employ Bratman’s formulation of the cognitivist thesis. However, the arguments in this paper may be employed, *mutatis mutandis*, in defence of the other version of the thesis as well.

³ Cognitivism is typically taken to entail that the consistency requirement and the Means-Ends Coherence requirements governing intention may be derived from the belief consistency and belief closure requirements, respectively. However, there is some variation among cognitivists about which and how many of the rational norms governing belief are derived from the rational norms governing intention, with some taking the thesis to apply only to the Means-Ends Coherence requirement but not the consistency requirement, and others taking it to apply not only to the consistency norm and Means-Ends Coherence norm, but also to any other norm governing intention. The Means-Ends Coherence requirement is also commonly referred to as the “instrumental rationality” requirement. (Bratman 2009: 14) Theorists who buy into some version of cognitivism include: John Broome (2009), Wayne Davis (1984), Gilbert Harman (1976, 1986), Kieran Setiya (2007), David Velleman (1989, 2007), and R. Jay Wallace (2006).

question.⁴ The theoretical advantages of such an account, if it could be made to work, are obvious. Not only would it be more parsimonious than accounts that treat the norms governing intention as independent of those governing belief, but it would also significantly reduce the explanatory work required of normativity theorists by replacing two potential explananda with one.

However, in his paper “Cognitivism about Practical Rationality”, John Brunero argues that SBT is either false or unable to play the explanatory role the cognitivist requires. Here, in Brunero’s own words, is a description of his argumentative strategy:

I’ll argue that those argumentative manoeuvres needed to save the Strong Belief Thesis thus face a dilemma: she must either concede that the Strong Belief Thesis is false or maintain that its true but unable to play the explanatory role the cognitivist needs it to play. Either way, the cognitivist explanations employing the Strong Belief Thesis fail.⁵

Call the preceding argument *Brunero’s dilemma*. The aim of this paper is to provide the cognitivist with the resources necessary to escape Brunero’s dilemma.

2. Brunero’s Objection

One set of putative counterexamples to SBT involves cases in which an agent intends something difficult and is therefore unsure they will succeed. In such cases, it is claimed, an agent intends to **X** but does not believe that they will **X**. Brunero describes one such case as follows:

[L]et’s consider the case of someone who intends to lift a heavy log that has fallen onto his front porch. Plausibly, he intends to lift the log but doesn’t believe that he will. It’s not that he believes he won’t; he’s simply agnostic about whether he will.⁶

Typically, defenders of SBT respond to potential counterexamples like the one just adumbrated by arguing that the agent in question intends to try to **X**, rather than intends to **X**.⁷ By claiming that an agent who is unsure she will **X** has an intention to try to **X** rather than an intention to **X**, the cognitivist is able to consistently maintain that an agent intends to **X** only if they believe they will **X**.

In order to see why this might constitute an appropriate response from the point of view of the cognitivist, it is important to keep in mind that the cognitivist is not attempting to offer a description of our ordinary linguistic practice as it relates to the word ‘intending’. Indeed, it can hardly be disputed that as far as our ordinary linguistic practice is concerned, we often say that people intend to do things they are unsure they will successfully accomplish. Hence, if accepting cognitivism required denying this easily verifiable fact, it would be puzzling that the view had any adherents at all. This

⁴ For a particularly clear discussion of how the strong belief thesis features in the cognitivist argument, see Brunero (2014: 20-25).

⁵ Brunero (2014: 22).

⁶ Brunero (2014: 22).

⁷ The *locus classicus* of this kind of response is Gilbert Harman (1986: 90-94).

should be one of our first hints that the cognitivist is attempting to do something other than describe how we ordinarily use the word ‘intending’.

I believe the cognitivist is most charitably seen as offering a formal or technical definition of ‘intending’; one that relates to the ordinary usage of the word in a manner analogous to how the formal definition of the word ‘insect’ relates to the informal usage of the word. According to the term’s informal usage, an insect is “any of numerous small invertebrate animals (such as spiders or centipedes) that are more or less obviously segmented.”⁸ However, spiders and centipedes would fail to qualify as insects in a formal or technical sense, owing to the fact that they have more than three pairs of legs. Thus, the formal use of the word ‘insect’ turns out to be narrower than the informal usage. A similar point may be made about the cognitivist use of the word ‘intending’. Insofar as the cognitivist is committed to SBT, she is committed to saying that an attitude towards a certain action counts as an intention only if its possessor believes they will perform the action. This would naturally exclude those cases in which an agent is unsure she will succeed in performing an action, thereby making the cognitivist usage of ‘intending’ narrower than the quotidian usage of the term. However, this no more delegitimizes the cognitivist’s usage than the fact that we sometimes call spiders ‘insects’ delegitimizes the term’s formal usage.

Furthermore, the formal usage of ‘insect’ may be seen as more precise (and hence, more theoretically useful) than the informal usage. For example, while the informal usage of ‘insect’ leaves it unclear whether one is referring to an invertebrate animal with three pairs of legs or four pairs of legs, the formal usage of the term has no such ambiguity. Analogously, from a cognitivist point of view, the quotidian usage of the word ‘intending’ leaves it unclear whether the speaker is referring to an action an agent believes or does not believe she will perform. The cognitivist usage of ‘intending’, by contrast, has no such ambiguity. Given that the demands of theorizing often require greater precision than everyday discourse, it should be unsurprising that the cognitivist views her more restrictive notion of intending as better suited for theorizing than the quotidian one.

The takeaway from the immediately preceding observations is that if we are going to effectively impugn the cognitivist conception of ‘intending’, it would not be enough to point out that we often describe agents who are unsure they will perform a certain action as having an intention to perform that action. Our criticism of cognitivism will need to proceed on theoretical rather than linguistic grounds. This is precisely what Brunero sets out to do in his paper. To this end, he claims that even if the cognitivist is able to preserve SBT by claiming that the log-lifter merely has an intention to try to lift the log, she does so at the expense of the intuition that the log-lifter is rationally criticisable for not intending to bend his knees. By Brunero’s lights, the log-lifter is rationally criticisable for not intending to bend his knees because he thereby violates the following rational norm:

⁸ “Insect.” *Merriam-Webster.com*, 2018. <https://www.merriam-webster.com/dictionary/insect>. (Accessed 21 March 2018).

MEANS-ENDS COHERENCE:

Rationality requires that [if one intends to **E**, and believes that one will **E** only if one intends to **M**, then one intends to **M**].⁹

The log-lifter intends to lift the log and believes that he will lift the log only if he intends to bend his knees. If he does not intend to bend his knees, he would be in violation of Means-Ends Coherence. However, the log-lifter may believe that intending to bend his knees is necessary for lifting the log without believing that it is necessary for *trying* to lift the log. Hence, if the cognitivist holds that the log-lifter is merely intending to try to lift the log, she is unable to maintain that the log-lifter violates Means-Ends Coherence because he does not intend to bend his knees. Brunero puts the point thus:

Intuitively, if this man were to believe that he'll lift the log only if he bends his knees when he lifts, and were to fail to intend to bend his knees when he lifts, he would be criticizable as means-ends incoherent. But he might think that bending one's knees, while necessary for lifting the log, isn't necessary for trying to lift the log. After all, we could suppose that the last time he didn't bend his knees, he tried and failed to lift the log, but didn't fail to try to lift the log. So, if his intention is merely to try to lift the log, he is no longer criticizable as means-ends incoherent in failing to intend to bend his knees.¹⁰

In sum, the challenge facing the cognitivist is to explain why the log-lifter is criticisable as means-end incoherent for failing to bend his knees, given that (by the cognitivist's lights) he lacks the intention to lift the log. Moreover, even if the cognitivist could invent some explanation of why the log-lifter violates Means-Ends Coherence, it is difficult to see how SBT could play a role in such an explanation. This is significant because the cognitivist's primary motivation for advocating SBT is that it allows us to explain why it is rationally problematic to violate Means-Ends Coherence. If a defence of SBT were unable to preserve its role in explaining why it is problematic to violate Means-Ends Coherence, then it would amount to little more than a Pyrrhic victory. Hence, the challenge facing the cognitivist is two-fold: (i) to explain why the log-lifter is criticisable as means-end incoherent, and (ii) preserve the cognitivist tenet that SBT plays a role in the aforementioned explanation.

3. Responding to Brunero's Objection

For the remainder of this paper, I will be attempting to meet the above challenge head-on. Here is a brief roadmap of the arguments to come. I begin (in section 3.1) by offering an account of *trying* that enables the cognitivist to preserve the intuition that the log-lifter violates Means-Ends Coherence. I then (in section 3.2) attempt to show that my explanation of why the log-lifter violates Means-Ends Coherence ultimately depends on SBT. To this end, I offer a necessary condition for *intending to try* that is grounded in

⁹ Brunero (2014: 20). See and cf. Broome (1999).

¹⁰ Brunero (2014: 23).

SBT. My paper concludes (in section 4) by responding to three potential objections to my proposed solution to Brunero's dilemma.

3.1. A Necessary Condition for Trying

I take as my point of departure the account of trying defended by Jennifer Hornsby, who defines trying to do something as roughly "doing what one can to do the thing."¹¹ I will call this the *Hornsby account* of trying.¹² The central insight of the Hornsby account is that trying to do something requires that one make a good faith effort to do that thing.¹³ Building on this insight, I propose the following necessary condition for trying to **X**:

GOOD FAITH REQUIREMENT (GFR):

S is trying to **X** only if S does everything she believes to be necessary and correctly believes to be entirely up to her to do **X**.

Complications aside, GFR asserts that an agent counts as trying to **X** only if that agent does everything she believes to be necessary to **X**. This is the "good faith" aspect of the *Good Faith Requirement*. However, GFR includes an important qualification; namely, that the agent correctly believes that doing **X** is entirely up to her. This qualification ensures that trying to do something is always up to us, even when successfully doing that thing is not.

I am committed to the following necessary condition for someone believing that something is entirely up to them: one believes that doing **X** is entirely up to one only if there is no obstacle to doing **X** that one is not confident one will overcome. For example, suppose that I wish to purchase a book from the bookstore, but that I am aware that there is a gang of ruffians who will attempt to prevent me from entering the bookstore. Insofar as I believe that the ruffians constitute an obstacle to my entering the bookstore that I am not confident I will overcome, I do not believe that entering the bookstore is entirely up to me. In the above example, the gang of ruffians constitutes an external obstacle to my entering the bookstore. However, there may also be internal obstacles to my doing **X**. For example, in place of a gang of ruffians nearby, let us substitute my being incredibly forgetful. Here, there is not an external obstacle to my stopping by the bookstore. But there is an obstacle nonetheless, albeit an internal one. Moreover, an internal obstacle—like a limitation of memory, concentration, skill, or willpower—may be just as effective as an external obstacle at preventing someone from successfully completing an action. Since my awareness of a certain internal obstacle may be sufficient to undermine my confidence that I will perform a certain action, said awareness may prevent me from believing that doing **X** is not entirely up to me. The

¹¹ (Hornsby (2010: 19).

¹² While suggestive, Hornsby's account of trying remains largely undeveloped. In what follows, I attempt to fill in some of the details of such an account in a way that is friendly to cognitivist aims. For a more fully fleshed out account of intending to try that builds on the Hornsby account, see §6 of Archer (2017).

¹³ For a particularly illuminating discussion of the notion of effort at play in such discussions, see: Brent (2017).

upshot is that since GFR only requires that one do those things one believes to be entirely up to one, it would not require that one do something one is unsure one will do due to external or internal obstacles.¹⁴

With GFR in hand, we may now explain how the log-lifter violates Means-Ends Coherence, given the assumption that he has an intention to try to lift the log. According to GFR, trying to lift the log requires doing all one can to lift the log. This means that trying to lift the log entails doing everything one believes to be necessary and correctly believes to be entirely up to one's to lift the log. Since intending to bend his knees is something the log-lifter believes to be necessary for lifting the log and correctly believes to be entirely up to him, it follows that the log-lifter can try to lift the log only if he intends to bend his knees. The upshot is that insofar as the log-lifter intends to try to lift the log, he violates Means-Ends Coherence by failing to intend to bend his knees. The preceding point becomes clear when we substitute "do all that one can to lift the log" for **X** and "bend one's knees" for **Y** in Brunero's formulation of Means-Ends Coherence, yielding the following application of the norm:

MEANS-ENDS COHERENCE (LOG-LIFTER)

Rationality requires that [if one intends to **do all that one can to lift the log**, and believes that one will **do all that one can to lift the log** only if one intends to **bend one's knees**, then one intend to **bend one's knees**].

Notice that Means-Ends Coherence applies to the log-lifter in virtue of the fact that he has a certain intention. It just so happens that the content of the log-lifter's intention is extensionally equivalent to trying, as described by GFR. Hence, the application of Means-Ends Coherence to the log-lifter is structurally no different from the standard application of the norm. This should be sufficient to show that GFR equips the cognitivist with the resources necessary to affirm that the log-lifter violates Means-Ends Coherence by failing to intend to bend his knees.

In sum, GFR allows the cognitivist to respond to the first part of Brunero's challenge: to explain why the log-lifter is criticisable as means-ends coherent. The log-lifter is criticisable as means-ends coherent because he intends to try to lift the log, but fails to intend to do something that he believes to be necessary for trying to lift the log. This is because trying to lift the log requires doing all that one can to lift the log, and there is something that is necessary for doing all that he can to lift the log that the log-lifter fails to intend to do; namely, bend his knees. The log-lifter therefore fails to intend a necessary means to his intended end.

3.2. The Cognitivist Analysis of Means-Ends Coherence

Let us now turn our attention to the second part of Brunero's challenge: demonstrating that SBT plays the essential role that the cognitivist typically takes it to play in my proposed explanation of why the log-lifter violates Means-Ends Coherence.

¹⁴ I expand on this point in §4.3.

Accomplishing this task will require that we briefly examine the standard cognitivist picture of how SBT is supposed to explain Means-Ends Coherence. According to the cognitivist, Means-Ends Coherence is either an instance of or explainable in terms of something along the lines of the following closure principle:

BELIEF CLOSURE:

Rationality requires that [if one believes that **P**, and believes that **P** only if **Q**, then one believes that **Q**].¹⁵

Suppose I intend to travel to New York, believe that I will travel to New York only if I intend to buy an airplane ticket, but fail to intend to buy an airplane ticket.¹⁶ Since, according to SBT, intending to **X** entails the belief that I will **X**, I may be characterised as believing I will travel to New York, believing that I will travel to New York only if I intend to buy an airplane ticket, but failing to believe that I will buy an airplane ticket. Hence, given SBT, my violation of Means-Ends Coherence may be re-characterised as a violation of Belief Closure.

It is worth noting that there is a slight mismatch between Means-Ends Coherence, as described by Brunero, and Belief Closure. Recall, Brunero characterises Means-Ends Coherence as follows:

MEANS-ENDS COHERENCE:

Rationality requires that [if one intends to **E**, and believes that one will **E** only if one intends to **M**, then one intends to **M**].

Were we to replace every occurrence of ‘intends to’ in Means-Ends Coherence with ‘believes that’, what we would get is not Belief Closure, but rather the following norm:

BELIEF CLOSURE*:

Rationality requires that [if one believes that **E**, and believes that **E** only if one believes **M**, then one believes that **M**].

However, Belief Closure* is a rather odd norm, for it only applies to cases in which an agent thinks that believing **M** is a necessary condition for the truth of **E**. Since believing **M** does not entail that **M** is true (i.e., one may have a false belief) and since **M** being true does not entail believing **M**, Belief Closure* does not entail Belief Closure and Belief Closure does not entail Belief Closure*. Hence, the two norms turn out to be quite distinct from each other. Moreover, there are relatively few cases in which an agent

¹⁵ There is a great deal of disagreement about how norms relating to belief closure are to be formulated, and also about how Means-Ends Coherence is supposed to be reduced to or explained in terms of closure. While these are all interesting questions in their own right, these are not issues I can address here. Instead, I will simply proceed under the assumption that these issues can be resolved. Since Brunero does the same (if only for the sake of the present argument), this seems like a dialectically appropriate approach to take.

¹⁶ I borrow this example from Brunero (2014: 41).

would satisfy the antecedent of Belief Closure*. For example, let us substitute 'I am in Paris' for **E** and 'I am France' for **M**. We get the following application of Belief Closure*:

BELIEF CLOSURE* (PARIS):

Rationality requires that [if I believe that **I am in Paris**, and believe that **I am in Paris** only if I believe **I am France**, then I believe **I am in France**].

However, it is false that I am in Paris only if I believe I am in France. Suppose that I was kidnapped, blindfolded, and taken to Paris, and that I had no beliefs about my current whereabouts. It would not follow from the fact that I did not *believe* I was in France that I was not in Paris. Hence, unless I had the bizarre belief that I could be at a certain location only if I believed I was, I would fail to satisfy the antecedent of Belief Closure*. The upshot is that the potential application of Belief Closure* turns out to be significantly limited.

In light of considerations like the one just limned, many cognitivists would take issue with the formulation of the means-ends coherence norm that features in Brunero's paper. For example, Jacob Ross offers the following alternative formulation:

MEANS-ENDS COHERENCE (ROSS)

One ought rationally to be such that [if one intends to **X** and one believes that **Y-ing** is a necessary means to **X-ing**, then one intends to **Y**].¹⁷

Notice that by omitting the intention that falls within the scope of the second belief operator in Brunero's version of Means-Ends Coherence, Ross's version avoids the mismatch between Means-Ends Coherence and Belief Closure just adumbrated. Hence, in at least this respect, Ross's version of Means-Ends Coherence would be deemed preferable by the cognitivist.

In Brunero's defence, he explicitly acknowledges that his formulation of Means-Ends Coherence is "no doubt in need of further refinement".¹⁸ Moreover, there may be drawbacks to Ross' version of Means-Ends Coherence as well, so I do not wish to claim that Ross's formulation is the final word on the matter. However, the question of what version of Means-Ends Coherence the cognitivist would need to accept in order for their attempted assimilation with Belief Closure to get off the ground is not one Brunero takes up in his paper. Nor is it a question I wish to take up here. Instead, I will simply proceed under the assumption that Ross's version of Means-Ends Coherence is a viable way to map Means-Ends Coherence onto Belief Closure in the way the cognitivist envisions.

¹⁷ Ross (2009: 250).

¹⁸ Brunero (2014: 20).

3.3. A Necessary Condition for Intending to Try

With the standard cognitivist explanation of how SBT allows us to assimilate Means-End Coherence with Belief Closure now in view, I propose the following necessary condition for *intending to try*:

REQUIREMENT FOR INTENDING TO TRY (RIT)

S intends to try to **X** only if S believes that she will do everything she believes to be necessary and correctly believes to be entirely up to her to do **X**.

The necessary condition for intending to try specified by RIT is the result of combining GFR with SBT. Recall, according to GFR, one tries to **X** only if one does everything one believes to be necessary and correctly believes to be entirely up to one, to do **X**. This means that one intends to try to **X** only if one intends to do everything one believes to be necessary and correctly believes to be entirely up to one, to do **X**. Furthermore, according to SBT, one intends to try to **X** only if one believes one will try to **X**. When this requirement is combined with GFR, we arrive at RIT; the claim that one intends to try to **X** only if one believes one will do everything one believes to be necessary and correctly believes to be entirely up to one, to do **X**.

Significantly, the necessary condition for intending to try specified by RIT is a constraint on an agent's beliefs. This comports with the general cognitivist strategy of explaining the norms governing intention in terms of the norms governing belief. We may reconstruct the cognitivist analysis of the log-lifter as follows: Suppose that the log-lifter intends to do everything he believes to be necessary and correctly believes to be within his power to lift the log, and that he believes that bending his knees is necessary to lift the log and entirely within his power, but that he does not intend to bend his knees. According to SBT, this example of Means-Ends Coherence violation is an instance of the following pattern of Belief Closure violation: S believes that all the members of some set **X** have some property **Q** and believes that **Y** is a member of set **X**, but does not believe that **Y** has property **Q**. For example, suppose that I believe that everyone in my class will get a passing grade, and that I believe John Snow is a member of my class, but that I do not believe that John Snow will get a passing grade. In such a case, I may be characterized as believing that all members of the set <everyone in my class> have the property <will get a passing grade>, and believing that <John Snow> is a member of the set <everyone in my class>, but failing to believe that <John Snow> has the property <will get a passing grade>.

The log-lifter example may also be assimilated to the aforementioned pattern of Belief Closure violation. The log-lifter believes that all members of the set <actions that I believe to be necessary for lifting the log and correctly believe to be within my power> have the property <will be done by me>, and he also believes that <bending his knees> is a member of the set <actions that I believe to be necessary for lifting the log and correctly believe to be within my power>, but fails to believe that <bending his knees> has the property of <will be done by me>. Hence, given SBT, it is possible to assimilate

the log-lifter to a standard case of Belief Closure violation. I take this to show that the account of intending to try offered in this paper preserves SBT's role (as standardly conceived by cognitivists) in explaining why the log-lifter violates Means-Ends Coherence.

4. Objections and Replies

I will conclude my discussion of Brunero's dilemma by responding to three potential objections to my proposal.

4.1. Objection I

It may be claimed that my response to Brunero's dilemma alters a crucial part of his example; namely, that the log-lifter tried to lift the log on previous occasions without bending his knees. Brunero puts the point as follows:

But he might think that bending one's knees, while necessary for lifting the log, isn't necessary for trying to lift the log. After all, we could suppose that the last time he didn't bend his knees, he tried and failed to lift the log, but didn't fail to try to lift the log.

Since GFR makes doing all that one can a requirement for trying, it appears to preclude the possibility that the log-lifter previously tried to lift the log.

In response to this objection, I maintain that whether or not the log-lifter qualifies as having previously tried to lift the log will depend on how we fill in the details of the case. There are two salient possibilities as far as his previous log-lifting related activities are concerned: either the log-lifter did or did not believe that bending his knees was necessary for lifting the log. (Note: the log-lifter being agnostic about the matter falls under the umbrella of him *not believing*.) I will consider both possibilities in turn.

First, let us assume that during his past attempts, the log-lifter did not believe that bending his knees was necessary for lifting the log. It would follow from GFR that bending his knees would not be necessary for trying to lift the log. Thus, GFR would allow that the log-lifter may have previously tried to lift the log without bending his knees. However, during his present attempt, the log-lifter does believe that bending his knees is necessary for lifting the log. Hence, there has been an important change in the log-lifter's doxastic economy between his previous and present attempts. Whereas he previously did not believe that bending his knees was necessary for lifting the log, and therefore (according to GFR) did not need to bend his knees to qualify as having tried to lift the log, he now does believe that bending his knees is necessary for lifting the log, and consequently (according to GFR) qualifies as having tried to lift the log only if he bends his knees. Hence, if we assume that the log-lifter did not believe that bending his knees was necessary to lift the log during his previous attempts, GFR can easily accommodate the intuition that the log-lifter could have previously tried and failed to lift the log without bending his knees.

Second, let us suppose that the log-lifter believed that bending his knees was necessary for lifting the log all along. Let us also assume that the log-lifter correctly believed that whether or not he bends his knees is entirely up to him.¹⁹ According to GFR, the log-lifter did not genuinely try to lift the log if he did not bend his knees. I believe that this is an intuitively plausible position to take. Suppose that I told you that I was trying to turn on my car and that you knew that I believed that inserting the key in the ignition was necessary to turn on the car and also within my power. But suppose I refrained from inserting the key in the ignition. It would be natural for you to insist that I have not really tried to turn on the car. This is why Hornsby insists that trying requires a good faith effort. If we were to discover that someone failed to do something they believed to be necessary and entirely up to them to accomplish some goal, we would rightly conclude that they failed to make a good faith effort and therefore failed to genuinely try. Hence, if we assume that the log-lifter believed that bending his knees was necessary for lifting the log on those previous occasions, the cognitivist has principled grounds to deny that he genuinely tried to lift the log.

4.2. Objection II

The second objection I wish to consider attacks one of the primary assumptions of my account; namely, that the log-lifter believes he will do all that he can to lift the log. It may be claimed that this assumption is shown to be false by Brunero's description of the example. Recall that, according to Brunero, the log-lifter believes that bending his knees is necessary for lifting the log, but he does not intend to bend his knees. Given that he lacks the intention to bend his knees, it is unlikely that he believes he will bend his knees. But if he does not believe he will bend his knees, then it is false that he believes he will do all that he can to lift the log. If this is right, then according to RIT, the log-lifter lacks the intention to try to lift the log.

The first step to responding to the above objection is to register that it proves too much. The objection assumes that the log-lifter could believe that he will do all that he can to lift the log only if he believes he will bend his knees. More generally, it assumes that one can believe that all members of some set, **X**, display some property **Y** only if one believes of each and every member of **X** that it has property **Y**. However, if this were true, it would render most of the commonly observed violations of Belief Closure impossible. Consider my example of the agent who believes everyone in her class will get a passing grade, believes John Snow is in her class, but does not believe John Snow will get a passing grade. According to the present objection, since John Snow is a member of S's class, then it can be true that S believes everyone in her class will get a passing grade only if S believes that John Snow will get a passing grade. It would

¹⁹ Otherwise, GFR would not require that the log-lifter bend his knees in order to qualify as having tried to lift the log. Recall, GFR only requires that an agent perform those actions that they *correctly* believe to be entirely up to them in order to qualify as having tried. This means that if it turned out that X-ing was not entirely up to an agent or that the agent did not believe that X was entirely up to them, then X-ing would not be necessary to qualify as having tried.

therefore be impossible for S to violate Belief Closure by believing that everyone in her class will get a passing grade and failing to believe that John Snow will get a passing grade. However, we all know (from our observations of ourselves and others) that it is indeed possible for someone to violate Belief Closure in the manner just described. Something has clearly gone wrong.

My diagnosis of what has gone wrong is that the objection overlooks the fact that our mental representations of universal generalizations do not necessarily consist in a representation of each and every individual instance that satisfies that generalization. For example, if I believe that all humans are mortal, it is unlikely that my mental representation of <all human beings> would involve a representation of each and every human being. Instead, my mental representation will typically involve words, symbols, and/or images that stand proxy for <all human beings>. This means that mentally representing a universal generalization does not require representing each and every instance that makes the generalization true.

I believe that it is this psychological fact about our mental representations that allows for the possibility of someone believing <everyone in my class will get a passing grade> without believing <John Snow will get a passing grade>, despite the fact that John Snow is a member of the class. Moreover, it is a curious fact about us that we are capable of psychological compartmentalisation. Hence, I may believe that John Snow is a member of my class, without that belief coming into contact with my other beliefs in the manner necessary for completing the inference to the conclusion: <John Snow will get a passing grade>. When the possibility of psychological compartmentalisation is combined with the aforementioned fact about our mental representations, it ceases to be mysterious how someone could violate Belief Closure in the manner just described.

The point generalizes to the kinds of beliefs that would be implicated by an intention, if SBT were true. The mental representation implicated by the belief that one will do all that one can to X need not include a representation of each individual action that falls under the scope of doing all that one can. This means that the belief that one will do all that one can need not involve the belief that one will perform some given action that falls within the scope of <doing all that one can>. The upshot is that it is psychologically possible to believe you will do all that you can to X while also failing to believe one will perform one of the actions that fall within the scope of the universal generalization.²⁰

²⁰ It is plausible that intentions are also governed by a version of the good faith requirement. For example, intending to X appears to entail intending to make a good faith effort to X. If someone did not intend to make a good faith effort to X, then it would be natural to conclude that they did not genuinely intend to X. Moreover, if intending to X did not involve intending to do what was necessary and within one's power to X, then it becomes mysterious why someone with an intention to X is rationally obligated to either intend the means necessary for X-ing or give up their intention to X. Hence, the problem raised by the present objection does not only confront the cognitivist. It confronts anyone who thinks that one has a genuine intention to X only if one intends to make a good faith effort to X and who also holds that violations of Means-Ends Coherence are possible. Anyone so committed must explain how it is possible to intend to do all that one can to X without intending something one believes to be necessary for X-ing.

While the above explanation of how Belief Closure violations of a certain kind are psychologically possible is fairly crude, I believe it is at least on the right track. However, one need not buy into my particular explanation of the phenomenon to appreciate that the phenomenon is real. What the preceding discussion illustrates is that the task of explaining how the log-lifter could believe he will do all that he can to lift the log without believing that he will bend his knees is one that also confronts anyone who holds that a certain general pattern of Belief Closure violation is possible. This fact alone should be enough to vindicate the cognitivist. After all, the cognitivist never claimed that every puzzle relating to the norms governing belief has already been resolved. What they claim is that solving the puzzles surrounding Belief Closure violations would be sufficient to solve the problems relating to cases of Means-Ends Coherence violations. What I have attempted to show in my reply to the present objection is that a satisfactory solution to the Belief Closure case would generalize to the Means-Ends Closure case as well. Thus, my reply preserves what I have characterised as cognitivism's primary selling-point—namely, that it reduces the explanatory workload of theorists by replacing two explananda with one.

4.3. Objection III

A third objection to my proposed necessary condition for intending to try is that it is inconsistent with the cognitivist treatment of another widely discussed case—namely, that of the forgetful cyclist.²¹ On the standard retelling, the forgetful cyclist has the aim of stopping by the bookstore on his way home from work. However, the cyclist is aware that whenever he gets on his bicycle, he tends to go into autopilot, which would result in him cycling all the way home, forgetting to stop by the bookstore on the way. Given that he is aware of his tendency to forget in situations like the one at hand, the cyclist is unsure he will stop by the bookstore. The standard cognitivist analysis of the cyclist would be to say that, due to his uncertainty, he has an intention *to try* to stop by the bookstore. Hence, the forgetful cyclist is analysed along similar lines to that of the log-lifter. However, it may be claimed that RIT prevents the cognitivist from employing this strategy. Recall, according to RIT the cyclist has an intention to try to stop by the bookstore only if he believes that he will do everything he believes to be necessary and correctly believes to be entirely up to him to stop by the bookstore. However, since the cyclist is aware that he tends to go into autopilot when he gets on his bicycle, there is no action fitting this description that he thinks he will perform. It therefore follows from RIT that the cyclist does not have the intention to try to stop by the bookstore, or at least so the objection goes.

In response, it should be noted that the belief that one will do all that one can to **X** should be seen as corresponding to a universal generalization rather than to an existentially quantified statement. (This is a point I discuss at some length in §4.2.) On the existentially quantified reading, the belief that one will do all that one can to **X**

²¹ For a paradigm retelling of this example, see Holton (2008).

involves holding that there exists some Y such that Y is an action one believes to be necessary and correctly believes to be entirely up to one, that one will perform. If there is no action that fits the description of Y , then the existentially quantified statement is false. By contrast, on the universal generalization reading, the belief that one will do all that one can to X involves holding that for all Y , if one believes Y is necessary for doing X and correctly believes Y is entirely up to one, then one will do Y . In the case in which there is no action corresponding to Y , it does not follow that the universal generalization is false. On the contrary, if the antecedent of the conditional that falls within the scope of the universal quantifier is false, then the universal generalization is trivially satisfied.

Let us apply the above observations to the forgetful cyclist. If the cyclist does not believe that there is anything that is both necessary and entirely up to him to stop by the bookstore, then the antecedent of the universal generalization is false and the generalization is trivially satisfied. This means that (as far as RIT is concerned) there is no action that the cyclist needs to believe he will perform in order to qualify as having the intention to try to stop by the bookstore. Hence, if we stipulate that it is the cyclist's awareness that he might forget to perform all of the actions necessary for stopping by the bookstore that explains his uncertainty about whether or not he will stop by the bookstore, then (as far as RIT is concerned) there is no action that he believes to be entirely up to him, and therefore no action he needs to believe he will perform in order to qualify as having the intention to try to stop by the bookstore. Rather than making it impossible for the forgetful cyclist to have the intention to try, RIT actually makes it easier for him to do so. The upshot is that RIT is perfectly consistent with the standard cognitivist analysis of the forgetful cyclist.

In light of the above observations, there may be a worry that RIT goes too far in the opposite direction by making it too easy for the cyclist to count as having the intention to try. However, it is important to keep in mind that RIT is not offered as a sufficient condition for intending to try. It only represents a *necessary condition*. Hence, the fact that RIT is trivially satisfied under certain conditions doesn't mean that an agent has the intention to try under those conditions. It merely means that RIT poses no obstacle to the agent having the intention to try under those conditions. We may ultimately wish to supplement RIT by requiring that an agent also display certain counterfactual or dispositional properties in order for that agent to count as having the intention to try. For example, we may stipulate that the cyclist only qualifies as having the intention to try to stop by the bookstore if something like the following counterfactual were true of him: if there were some Y that he believed to be necessary for stopping by the bookstore and correctly believed to be entirely up to him, then the cyclist would do Y . Whether the cyclist must satisfy this or some other requirement in order to count as having the intention to try to purchase the book is a question we need not settle here. What I wish to stress is that it is no objection to RIT that it fails to constitute a sufficient condition for intending to try.

5. Conclusion

The aim of this paper has been fairly modest: to provide the cognitivist with the resources necessary to escape Brunero's dilemma. According to Brunero, the standard cognitivist attempt to defend SBT by claiming that the agent who is unsure she will do **X** merely has the intention to try to do **X** suffers from a fatal flaw: it is unable to preserve the intuition that the unsure agent is still subject to the means-ends coherence norms related to doing **X**. This paper has attempted to meet Brunero's objection by offering a necessary condition for intending to try that allows the cognitivist to consistently hold that (i) the agent who is unsure she will do **X** merely has the intention to try to do **X**, and (ii) the agent who has the intention to try to do **X** is still subject to the means-ends coherence norms related to doing **X**. By so doing, the present investigation attempts to go some distance towards establishing that cognitivism is worth giving a try.²²

References

- Archer, Avery. (2017) Do We Need Partial Intentions? *Philosophia* **45** (3): 995-1005.
- Bratman, Michael. (1984) Two Faces of Intention. *The Philosophical Review*, **93** (3): 375-405.
- Bratman, Michael. (2009a) Intention, Belief, Practical, Theoretical. In Simon Robertson. (ed.) *Spheres of Reason: New Essays in the Philosophy of Normativity*. (Oxford, UK: Oxford University Press).
- Bratman, Michael. (2009b) Intention, Belief, and Instrumental Rationality. In David Sobel and Steven Wall (eds.), *Reasons for Action*. (Cambridge, UK.: Cambridge University Press): 13-36.
- Brent, Michael. (2017) Agent Causation as a Solution to the Problem of Action. *Canadian Journal of Philosophy*, **47** (5): 656-673.
- Broome, John. (1999) Normative Requirements. *Ratio*. **12** (3): 398-419.
- Broome, John. (2009) The Unity of Reasoning? In Simon Robertson (ed.), *Spheres of Reason*. (Oxford, UK: Oxford University Press): 62-91.
- Brunero, John. (2014) Cognitivism about Practical Rationality. In *Oxford Studies in Metaethics*, **9**: 18-44.
- Davis, Wayne. (1984) A Causal Theory of Intending. *American Philosophical Quarterly*. **21** (1): 43-54.

²² I will like to thank John Brunero, Albert Hu, Laura Papish, Usha Nathan, and Kieran Setiya for helpful discussion, along with the generous feedback on earlier drafts of this paper from the anonymous referees for this journal.

- Harman, Gilbert. (1976) Practical Reasoning. *Review of Metaphysics*, **29** (3): 431–63.
- Harman, Gilbert. (1986) *Change in View: Principles of Reasoning*. (Cambridge, MA: MIT Press).
- Holton, Richard. (2008) Partial Belief, Partial Intention. *Mind*. **117** (465): 27-58.
- Hornsby, Jennifer. (2010) Trying to Act. In Sandis, T, and C. O'Connor (eds.) *A Companion to the Philosophy of Action*. (New York, NY: Blackwell): 18-25.
- Ross, Jacob. (2009) How to be a Cognitivist about Practical Reason. *Oxford Studies in Metaethics*, **4**: 243-281.
- Setiya, Kieran. (2007) Cognitivism about Instrumental Reason. *Ethics*, **117** (4): 649–73.
- Velleman, J. David. (1989) *Practical Reflection*. (Princeton, NJ: Princeton University Press).
- Velleman, J. David. (2007) What Good is a Will? In A. Leist and H. Baumann (eds.), *Action in Context*. (Berlin: de Gruyter): 193–215.
- Wallace, R. Jay. (2006) *Normativity and the Will*. (Oxford, UK: Clarendon Press).