



Contrastive Intentions

ABSTRACT: *This paper introduces and argues for contrastivism about intentions. According to contrastivism, intention is not a binary relation between an agent and an action. Rather, it is a ternary relation between an agent, an action, and an alternative. Contrastivism is introduced via a discussion of cases of known but (apparently) unintended side effects. Such cases are puzzling. They put pressure on us to reject a number of highly compelling theses about intention, intentional action, and practical reason. And they give rise to a puzzle about rather-than constructions such as ‘I intend to ϕ rather than ψ ’: In side effect cases it can seem wrong to claim that the subject intends to ϕ yet acceptable to claim that they intend to ϕ rather than ψ . This cries out for explanation. Contrastivism provides a unified response to all of these problems.*

KEYWORDS: intentions, intentional action, contrastivism, partial intention, choice

1. A Puzzle

Consider the following case:

POISON: A billionaire offers to pay Ben \$3,000,000 to pump poisonous gas into a house. It is clear from the contract that Ben’s payment turns *only* on his pumping the house full of gas (he will be paid if the house fills up with gas, he will not be paid if it does not). Ben agrees and starts pumping. However, he quickly discovers that there is a family inside and realizes that they will die if he keeps pumping. He does not wish to kill anyone. However, if he stops pumping he will not receive payment. Begrudgingly, he continues, and ends up killing the family (cf. Anscombe 1963: 41–42 and Setiya 2022).

Cases like this—so-called side effect cases—are familiar from the literature on intentional action and the doctrine of double effect. They are also deeply puzzling. Consider the following sentences:

1. Ben intentionally killed the family.
2. Ben intended to kill the family.

This paper has been around for quite a while. It has benefited from helpful reviewer comments from a few different journals (including, of course, the Journal of the APA). I have presented this paper (or related work) at the Universities of Leeds and Oslo as well as at the 2018 Joint Session of the Aristotelian Society in Oxford. Countless people have read and commented on the paper (too many to list), and I am grateful for all their contributions. This research received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement no. 818633).



Here (1) seems true. After all, Ben knowingly killed the family. He was able to do otherwise; yet, he chose to bring about the family's demise. But there are strong reasons to reject (2). First, it is widely taken as a datum that agents in such cases do not intend the known but undesirable consequences of their chosen actions. This judgment is central to the framing of such cases both in the literature on intentional action and in the related literature on the doctrine of double effect. Moreover, this intuition is common not only among philosophers but also among lay people: Knobe (2004) has shown that in similar cases where lay participants are happy to say that a foreseen effect ϕ was brought about intentionally, the majority avoid attributing the intention to ϕ . The most natural explanation for the almost universal judgment that (2) is false is that it *is* false.

Furthermore, there are reasons to reject (2) beyond mere intuition:

- A. Intention is a pro-attitude. Intending to ϕ involves being positively disposed toward ϕ ing. Yet Ben is not positively disposed toward killing the family.

Moreover, Ben's attitude toward the family's demise lacks three important features of intention highlighted by Bratman (1987):

- B. Intentions pose problems for practical reason. The intention to ϕ poses the problem of how one will ϕ . Yet Ben is not posed the problem of how to kill the family. There is no pressure on him to identify a means that will guarantee their death.
- C. Intentions place rational constraints on further intentions. If one intends to ϕ , and an option opens up that is incompatible with ϕ ing, one cannot rationally take that option while retaining the intention to ϕ . Yet Ben is not obligated to avoid courses of action simply because they are inconsistent with killing the family.
- D. Intending to ϕ usually leads to one's endeavoring to ϕ . Yet Ben does not endeavor to kill the family.

However, there are also compelling reasons to accept (2). First, consider the simple view of intentional action (hereinafter: the SIMPLE VIEW): if one intentionally ϕ s, one must intend to ϕ . The SIMPLE VIEW is attractive largely due to its simplicity. It initially strikes many people as obvious. And it seems to be entailed by the combination of other intuitive banalities, such as the claim that the actions we perform intentionally are those we choose to perform and the claim that choosing to ϕ involves forming an intention to ϕ . However, according to the SIMPLE VIEW, (1) entails (2). Thus, if (1) is true and (2) is false, then we must reject the SIMPLE VIEW. Second, (2) is entailed by the conjunction of two highly plausible assumptions about the relationship between deliberation, decision, and intention:

DECISION INTENTION: If an agent decides to ϕ and they are rational, then they will form an intention to ϕ .

HOLISTIC DECISION: If an agent knows that their ϕ ing would result in some additional outcome ψ and they decide to ϕ anyway despite their serious consideration of ψ , then they decide to ψ .

DECISION INTENTION is highly intuitive. The act of deciding to ϕ is widely identified with the act of forming an intention to ϕ . This assumption is made by, for example, Kaufman, (1966), Hall (1978), Davidson (1980), McCann (1986), Frankfurt (1988), Mele (1992, , 2003), Kane (1996), Pink (1996), Searle (2001), Hieronymi (2005, 2006), Holton (2009), and Shepherd (2013). Some theorists reject the act view of decision but nonetheless maintain that deciding to ϕ requires intending to ϕ . Examples include Goldman (1970), Meiland (1970), Raz (1975), Aune (1978) and O'Shaughnessy (1980). A principle very close to DECISION INTENTION is also defended by Goetz (1995). Even authors who deny DECISION INTENTION often endorse weaker principles that generate the same problem. For example, Paul (2012) argues that decision and intention only come apart when an agent's motivational state radically conflicts with their decision. This suggests a weaker *ceteris paribus* version of DECISION INTENTION that generates the same problem.

Moreover, there are good independent reasons to endorse DECISION INTENTION: in deciding to ϕ we settle the practical question of whether or not to ϕ in favor of ϕ ing. Intention is the state that results from deliberation about what to do, much as belief is the state that results from deliberation about the facts. Just as believing that p consists in settling on p and being committed to its truth, intending to ϕ involves settling on ϕ and being committed to ϕ ing (Hieronymi 2005, 2006; Shah 2008, 2013; McHugh 2013). Moreover, even in cases where the question of whether to ϕ and whether to intend to ϕ have been thought to come apart (for example, in Kavka's toxin puzzle [Kavka 1983]), reasons in favor of deciding and intending to ϕ are indistinguishable. Indeed, it is hard to make coherent sense of a request to decide to ϕ without intending to ϕ .

DECISION INTENTION has been challenged by Bratman (1987), Rozeboom (2015) and McGuire (2016). Each of these authors points to side effect cases as counterexamples. None addresses the considerations adduced here in favor of DECISION INTENTION. Bratman offers an alternative conception of the relationship between decision and intention: decision is holistic—we decide upon whole scenarios. Likewise, intentions usually encompass the whole scenario as well. However, sometimes intending certain aspects of a scenario will conflict with prior self-governing policies we have (for example, Ben's killing the family may conflict with a policy he has against killing innocents). In such cases the conflicting aspects of a scenario are not intended. There are a number of problems with Bratman's proposal. First, as Bratman himself maintains, policies admit of exceptions. One might have a policy to wear a seat belt and nonetheless rationally retain that policy while violating it in special circumstances. Likewise, policy violations in side effect cases may constitute exceptions. Second, if one genuinely has a policy against ϕ ing and envisages that in ψ ing one will thereby ϕ and chooses to ψ (and thereby ϕ) anyway, it is not clear that one can avoid the charge of irrationality or inconsistency simply by filtering ϕ ing from one's intentions.

Rozeboom (2015) also takes decision to be holistic, but denies that intentions must likewise be holistic because he denies that intentions constitute the conclusions of practical reason. He also holds that considerations of cognitive economy speak against holistic intentions. However, the claim that intentions do not constitute conclusions of practical reason is a bitter pill to swallow. If we wish to hold on to the view that intentions conclude practical reason (as I certainly do), we should be skeptical of Rozeboom's position. Moreover, it is not clear that accepting apparent side effects when choosing a course of action (as Rozeboom rightfully insists we must) is any less cognitively demanding than intending them (when these intentions do not place many new and distinctive demands on our practical reason).

HOLISTIC DECISION holds that our decisions encompass more than just our ends abstracted from their contexts and consequences. We are not able to and should not screen off the known consequences of our decisions. Rather, we choose between the situations we have the power to bring about. Like DECISION INTENTION, HOLISTIC DECISION has a great deal of intuitive support. As Bratman says of his own version of the principle:

What the principle requires is only a certain clear-headedness and intellectual honesty—an absence of 'bad faith', if you will. Once I seriously consider *A*'s anticipated effect, *E*, in my deliberation about whether to *A*, I should see that the issue for my deliberation concerns a complex scenario, one that includes, inter alia, *A together with E*. If I am clear-headed and intellectually honest about this, my conclusion should concern this complex scenario, and not merely my *A*-ing simpliciter. (Bratman 1987: 151)

Moreover, as Bratman points out, if we do not compare and decide upon whole scenarios but rather weigh our actions against their consequences in a piecemeal way failing to track the total utility of the scenarios our actions will likely bring about, then we risk sanctioning irrational decision-making. Beyond this, Rozeboom (2015) emphasizes that decisions resolve practical unsettledness. If we know that ϕ ing will also involve ψ ing and we resolve to ϕ while remaining unsettled on ψ , then we are irrationally treating ϕ as settled and ψ as unsettled while recognizing that they come together.

Unfortunately, HOLISTIC DECISION and DECISION INTENTION together with the facts of POISON entail (2). Ben knew that by pumping the house full of gas he would kill the family. This is something he seriously considers; yet, he chooses to keep pumping anyway. That is, by HOLISTIC DECISION he decides to kill the family. However, together with DECISION INTENTION, this entails that Ben intends to kill the family.

This creates a puzzle. On the face of it, it seems clear that (2) is false. However, HOLISTIC DECISION and DECISION INTENTION, not to mention the SIMPLE VIEW itself, are highly plausible. Something has to give. The usual response is to reject the SIMPLE VIEW together with either HOLISTIC DECISION or DECISION INTENTION. However, I propose that we can keep the SIMPLE VIEW, HOLISTIC DECISION, and DECISION INTENTION while capturing the almost universal intuition that (2) is false. In order to do so,

we must embrace the existence of contrastive intentions. Contrastivism allows us to deny the entailment from ‘A intentionally ϕ ed’ to ‘A intended to ϕ ’ without denying the most plausible versions of HOLISTIC DECISION, DECISION INTENTION, and the SIMPLE VIEW.

This paper is structured as follows: section 2 sketches the basic idea behind contrastive intentions, provides a basic argument for contrastivism, and presents a suggestive argument for the hypothesis that all intentions are contrastive. Section 3 applies contrastivism to the side effect puzzle. Section 4 gives the beginnings of a unified account of intentions within the contrastivist framework. Finally, section 5 further illustrates the structure and explanatory power of contrastivism by showing how it solves a number of other problems for the SIMPLE VIEW.

2. Contrastive Intentions

2.1 Contrastive Intention Ascriptions

Not all intention ascriptions have the surface form ‘I intend to ϕ ’ (or, ‘I will ϕ ’). We can also make explicitly contrastive intention claims like ‘I intend to eat salad rather than chips’, ‘I intend to go to the library rather than the party’, or ‘I intend to quit rather than be held accountable’.

We could treat such sentences as meaning ‘I intend to ϕ and not ψ ’. Thus, when I say ‘I intend to go to the library rather than the party’, I mean ‘I intend to go to the library and not go to the party’ (call this the ‘and-not’ reading). On the other hand, we could take them to be genuinely contrastive. That is, we could take intention to constitute at least sometimes a three-place relation between an agent, an action, and an alternative. And we could take such sentences to make explicit both the intended action and the alternative.

There is a boring version of contrastivism according to which whenever we intend to ϕ , we always intend to ϕ rather than $\neg\phi$. This approach gains intuitive support from the thought that when one forms an intention, one is committed to not failing. This contrastivism is very much in the spirit of standard approaches to intention as involving what I will call all-out commitments. An all-out commitment to ϕ , as I understand it here, can be thought of as a commitment to ϕ *simpliciter*. If one has an all-out commitment to ϕ , one is committed to ϕ ing, and one cannot do anything inconsistent with ϕ ing without thereby violating one’s commitment. This is not to say that subjects who intend to ϕ *simpliciter* and thereby undertake an all-out commitment to ϕ never rationally perform actions inconsistent with ϕ ing. It does not commit the subject to ϕ ing come what may. However, if such a subject were knowingly to perform an action inconsistent with their ϕ ing, this would involve abandoning their intention (and thus their commitment) to ϕ . We can rationally revise our intentions. But we cannot rationally perform actions inconsistent with our intentions while retaining those very intentions.

A more ambitious contrastivism would hold that we can intend to ϕ without undertaking an all-out commitment to ϕ . It would hold that one can intend to ϕ rather than ψ while not intending to ϕ rather than χ for some known alternative χ that is incompatible with ϕ . This would allow that one could intend to ϕ rather

than ψ while not intending to ϕ rather than $\neg\phi$. This robust contrastivism can be thought of as follows: When forming intentions, we settle practical questions. That is, we settle the question of what we will do and thereby undertake a commitment to a particular outcome. However, we often fall short of completely resolving these practical questions. An analogy with theoretical questions will be helpful: Jonathan Schaffer (2007: 387) points out that one may be able to answer the question ‘Is that George Bush or Janet Jackson?’ without being able to answer the question ‘Is that George Bush or Will Ferrell?’ (famous at the time for his George Bush impersonation). In this case, one can only partially resolve the question ‘Is that person George Bush?’. One is able to rule out some possibilities incompatible with the relevant individual’s being George Bush, but one is not able to rule out all possibilities incompatible with the individual’s being George Bush. Likewise, Ben is only partially able to answer the practical question of whether to kill the family. He has resolved to kill them rather than stop pumping. But this is as far as his commitment goes. Such contrastive intentions are less like desires and more like preferences. That is, they involve a preference-like choice between alternatives viewed as possible outcomes of an action. The intention to ϕ rather than ψ involves a ranking of alternatives that places ϕ ahead of ψ . This is consistent with one’s ranking other alternatives ahead of or on a par with ϕ . For example, one may also rank χ as on a par with ϕ . This ranking is committal. It would commit the agent to ϕ over ψ . But it would not commit the agent to ϕ over χ .

To be clear, the analogy with preferences here is purely structural. Traditional conceptions of intention can be seen as treating intention as like desire in the following sense: it is possible to desire some outcome *simpliciter*. One need not rank that outcome relative to alternatives in order to desire it. Similarly, it may be thought that one always intends an outcome *simpliciter*. The ranking of the outcome above alternatives is not part of the attitude of intention itself (although it may be part of the process that gives rise to an intention). Preferences are not like this. It makes little sense to talk of a preference for ϕ *simpliciter*. Rather, we always prefer ϕ over some alternative ψ . Likewise, the contrastivist about intention says that at least some (potentially all) intentions are contrastive in structure: rather than intend to ϕ *simpliciter*, we intend to ϕ rather than some alternative ψ . This is as far as the analogy goes. Traditional approaches to intention do not treat desire for ϕ as necessary for an intention to ϕ . Likewise, the contrastivist about intention does not treat a preference for ϕ over ψ as necessary for an intention to ϕ rather than ψ . Instead, it merely involves a commitment to ϕ ing rather than ψ ing. Moreover, preferences and desires, unlike intentions, do not require a commitment to the desired or preferred outcomes.

The ‘and-not’ and ‘boring contrastivism’ approaches encounter problems when we consider cases like the POISON case. Consider the following:

3. Ben intends to kill the family *rather than stop pumping*.

Unlike (2), (3) does not seem obviously false. Indeed, this is a natural way to describe Ben’s intention and provides some inconclusive support for the truth of (3), and it will suffice as a starting point. I will assume that (3) is true and outline

the consequences of this assumption. Following this, I will show how the resultant view is able to do significant work: it is able to reconcile plausible versions of HOLISTIC DECISION, DECISION INTENTION, and the SIMPLE VIEW with the standard intuitions about side effect cases, and it is able to resolve a number of other challenges to the SIMPLE VIEW. These considerations provide weighty independent support for contrastivism.

2.2 Taking Contrastive Intention Ascriptions Seriously

If (3) is true, then we must reject the and-not reading of rather-than constructions. On the and-not reading, (3) says that Ben intends to kill the family and not stop pumping. This entails (2). But (2) is false. Similar considerations carry over to boring contrastivism. For the boring contrastivist, the intention to ϕ rather than ψ involves an all-out commitment to ϕ ing. But it is clear that Ben does not undertake an all-out commitment to the family's demise. If the option arose to keep pumping while allowing the family to live, he would take it. This would not require revision of his prior commitments.

Thus, if (3) is true, then we must allow that Ben intends to kill the family relative to some alternatives but not others. For example, we should allow that (3) is consistent with the falsity of (4):

4. Ben intends to kill the family rather than let them live.

Thus, if (3) is true, then we must embrace robust contrastivism. If contrastivism is to work, it should be seen as a thesis about the structure of intentions rather than about their contents. It is unclear what contrastive contents could be. Contents of attitudes, be they propositions or relatives such as plans or maps, are typically individuated in terms of their satisfaction conditions—the demands they place on the world (perhaps in combination with modes of presentation and structure). It is unclear what the satisfaction conditions could be for the content [ϕ rather than ψ]. I can see only two obvious options. The first is for reality itself to be contrastive and to be such that [ϕ rather than ψ]. The alternative is for [ϕ rather than ψ] to demand that ϕ obtain and ψ not obtain. The first option is hard to swallow. Yet, the second would render [ϕ rather than ψ] equivalent to the proposition that ϕ and not ψ . Thus 'I intend to ϕ rather than ψ ' would simply be equivalent to 'I intend to ϕ and not ψ '. But we have already ruled out this reading.

Of course, nothing I have said rules out the possibility of analyzing contrastive intentions as, for example, a complex of conditional and/or disjunctive intentions. The simplest possibility along these lines would be to hold that an intention to ϕ rather than ψ is a conditional intention to ϕ if one's only options are ϕ and ψ . But this cannot be right. If you mistakenly believe that ϕ and ψ are your only options and you ϕ in light of this, then your conditional intention will not be satisfied. Yet, as long as you do not also ψ , your intention to ϕ rather than ψ will be satisfied. There are more complex variants of the conditional intention view that avoid this problem. For example, it might be suggested that one intends to ϕ rather than ψ if one intends to ϕ so long as one *believes* one's only

options are ϕ and ψ . Such views face problems of their own. But rather than list every conceivable problem for every conceivable modification of the conditional view, I will simply point to a general problem faced by the majority of such approaches: Conditional intentions produce actions by issuing all-out intentions when the antecedent is satisfied. But Ben acts on his intention to kill the family rather than stop pumping without forming an all-out intention to kill the family. This tells against the possibility of straightforwardly reducing contrastive intentions to conditional intentions.

It would be impossible to rule out the reduction of contrastive to conditional intentions a priori. Indeed, it would not be disastrous if contrastive intentions could be reduced to conditional intentions. Certainly this would render contrastive intentions easier to understand. That said, it would perhaps weaken the force of my suggestion that admitting both contrastive and non-contrastive intentions results in an uncomfortably disunified picture of intention. After all, it looks like we already have to draw a distinction between conditional and nonconditional intentions. Adding in contrastive intentions would not introduce any further structural disunity to the class of mental states we call intentions.

Nonetheless, no conditional analysis seems obviously forthcoming, and a simpler alternative is available. Attitudes can be structurally contrastive. Preferences constitute an uncontroversial example of such an attitude. Thus, we should take contrastive intentions to be contrastive in structure, rather than a structurally noncontrastive attitude that can take contrastive contents.

2.3 Are All Intentions Contrastive?

The argument presented so far suggests that at least some intentions are structurally contrastive. As we will see, even this limited contrastivism is able to do substantial work: it is able to save the SIMPLE VIEW along with plausible versions of HOLISTIC DECISION and DECISION INTENTION. However, it is worth briefly considering a stronger view: that all intentions are contrastive. The reason for this is simple: the view that only some intentions are structurally contrastive is disunified. It holds that there are two kinds of intention and that they differ not only in content but also in structure. The view that all intentions are contrastive is unified. It sees all intentions as structurally alike. In the remainder of this section I will present a linguistic argument that picks up on this theme.

There is a well-known strategy for generalizing contrastivist analyses beyond explicitly contrastive cases: If 'intend' is not always at least implicitly contrastive, then 'intend' is ambiguous. Yet, there is little independent reason to believe that 'intend' is ambiguous, and 'intend' fails standard tests for ambiguity, such as the coordination across conjunction test (employed by Schaffer [2007] and Snedegar [2013, 2017] in defense of contrastivism about knowledge and reasons respectively). To apply the coordination across conjunction test we conjoin two sentences containing potentially different uses of a term while explicitly using the term only once. If the term is ambiguous, then the sentence should sound infelicitous or at least require some interpretative effort. Consider the following examples (cf. Sennett 2016):

- 2.3.1 The colors are light.
- 2.3.2 The feathers are light.
- 2.3.3 The colors and the feathers are light.

The conjunction here is semantically anomalous. It requires some interpretative effort. We cannot take the same sense of ‘light’ to apply to both the colors and the feathers. By contrast, consider the following:

- 5. I intend to pay my debts and to kill the family rather than stop pumping.

This sentence is neither infelicitous nor anomalous. It does not require any interpretative effort. Yet, the single occurrence of ‘intend’ attaches to both conjuncts. And the second conjunct is explicitly contrastive. That is, ‘intend’ is contrastive in this case. However, the first conjunct is not explicitly contrastive. Thus, ‘I intend to pay my debts’ must be implicitly contrastive. If this is right, then we have a reason to believe that all intention claims are contrastive. However, this does not entail that all intention claims are interestingly contrastive. When somebody asserts ‘I intend to ϕ ’ without qualification and without any contextual cues to suggest otherwise, we usually take them to be undertaking an all-out commitment to ϕ ing. This suggests that the default contrast for ‘I intend to ϕ ’ is ‘ $\neg\phi$ ’. That is, ‘I intend to ϕ ’ will be interpreted as meaning ‘I intend to ϕ rather than $\neg\phi$ ’. I will, from here on, call such intentions all-out intentions. Whereas a merely contrastive intention to ϕ involves a partial ranking where ϕ is ranked above some alternative, an all-out intention involves a total ranking where ϕ is ranked above all alternatives. All-out intentions, unlike merely contrastive intentions, involve all-out commitment.

It might be considered unrealistic to suggest that we ever form all-out intentions in this sense—we do not rank our selected alternative over *all* alternatives. We do not even consider most alternatives. There is too much uncertainty involved in decision-making for us to ever fully commit to a single course of action. Similar considerations apply to the notion of intention *simpliciter* and have led some (e.g., Ferrero 2009) to suggest that all intentions are fundamentally conditional. There are two responses available here. The first is to treat the above account as an idealization: when we form an all-out intention to ϕ , we do not really rank our ϕ over *all* alternatives. We merely rank it above all the alternatives in some privileged set (e.g., realistic or relevant alternatives). A merely contrastive intention to ϕ would rank ϕ above some but not all of these alternatives. This strategy is similar to Ferrero’s, and the reasoning in the following sections would be unaffected by a shift to this de-idealized approach. The second approach, to which I am sympathetic, points out that our ranking of alternatives depends on how fine-grained our partition of the space of possibilities is. Under constraints of time and cognitive load we might conceive of our space of options in a coarse-grained way—perhaps partitioning the space of possibilities into the ϕ possibilities and the $\neg\phi$ possibilities. Relative to such a partition, we might rank ϕ above $\neg\phi$. Since all the $\neg\phi$ possibilities are contained within the $\neg\phi$ cell of our

partition, this will involve ranking ϕ above all $\neg\phi$ possibilities. It would involve undertaking an all-out commitment to ϕ . However, circumstances may change. Previously unconsidered possibilities may arise that induce a more fine-grained partition of the space of possibilities. And relative to this new partition, the subject may no longer rank ϕ above all $\neg\phi$ alternatives. This would trigger a revision in the subject's intentions and thus a revision in their commitments. This seems to be an ordinary feature of decision-making: we often make plans and then revise or relinquish them as previously unconsidered possibilities become salient.

The considerations adduced in this section suggest that not only some, but all intentions may be structurally contrastive. This is, of course, a strong hypothesis. As we will see in section 3 and section 5, much of the work that can be done by contrastive intentions does not require us to endorse anything so strong. Nonetheless, I believe that for the sake of maintaining a unified theory of intention, we should take seriously the possibility that all intentions are contrastive. We will return to this issue in section 4 where I provide the beginnings of a unified view of intentions in the contrastivist framework. However, before doing so we must return to the puzzle with which we started because the solution informs the positive view I present in section 4.

3. Contrastivism, Decisions, and Side Effects

We are now in a position to re-examine side effect cases in light of contrastivism. It is easy for the unified contrastivist (i.e., those who see all intentions as fundamentally contrastive) to explain the intuition that (2) is false. Unqualified statements of the form 'S intends to ϕ ' are, by default, interpreted as meaning 'S intends to ϕ rather than $\neg\phi$ '. Thus, (2) is most naturally understood as attributing an all-out intention to kill the family. Yet, although Ben has an all-out intention to keep pumping, he does not have an all-out intention to kill the family. After all, were the family to leave the house, Ben would make no attempt to stop them. And this would require no revision of his prior intentions. Thus, he is committed to killing the family only insofar as this is a consequence of his continued pumping. His commitment goes no further. Thus, (2), on its most natural reading, is false.

Contrastive intentions also help us with HOLISTIC DECISION and DECISION INTENTION: Consider the type of decision Ben makes when he resolves to keep on pumping. He realizes that if he keeps pumping, he will kill the family. He considers this fact, and he decides to continue pumping anyway. There is, thus, a clear sense in which he decides to kill the family. If asked 'what have you decided to do, kill the family or stop pumping?', he would answer 'I have decided to kill the family'.

However, his decision to kill the family is essentially tied to his decision to continue pumping. He is only settled upon killing the family insofar as this is a necessary consequence of his continued pumping and only insofar as he maintains a commitment to his continued pumping. That is, he chooses the death of the family over the cessation of pumping. But this is the full extent of his commitment. His decision is, thus, qualified. It is more naturally described as a decision to kill the family *rather than stop pumping*.

This distinction between all-out and qualified decisions allows us to capture the sense in which Ben decides to kill the family without bringing HOLISTIC DECISION into conflict with the truth of DECISION INTENTION and the falsity of (2). In order to see this, we must consider two versions of HOLISTIC DECISION (where left unqualified I take ‘decision’ to mean ‘all-out decision’ in these principles):

Holistic All-Out Decision (HAOD): If an agent knows that their ϕ ing would likely result in some additional outcome ψ and they decide to ϕ anyway despite their serious consideration of ψ , then they make an all-out decision to ψ .

Holistic Qualified Decision (HQD): If an agent knows that their ϕ ing would likely result in some additional outcome ψ and they decide to ϕ anyway despite their serious consideration of ψ , then they at least make a qualified decision to ψ .

We can see that HAOD is false. Ben seriously considers the fact that his continued pumping will result in the family’s demise. But his decision to kill is not all-out. HQD is all we need to capture the sense in which Ben decides to kill the family.

Now consider DECISION INTENTION: We are told that if an agent decides to ϕ and the agent is rational, then the agent will form an intention to ϕ . However, this principle is ambiguous between all-out and qualified decisions as well as between all-out and merely contrastive intentions. Since Ben’s decision is qualified, we need only consider the analogues of DECISION INTENTION for qualified decisions:

Qualified Decision, All-Out Intention (QDAI): If an agent makes a qualified decision to ϕ and they are rational, then they will form an all-out intention to ϕ .

Qualified Decision, Qualified Intention (QDQI): If an agent makes a qualified decision to ϕ and they are rational, then they will form a qualified intention to ϕ .

QDQI is the most plausible rendering of DECISION INTENTION for qualified decisions. QDAI is too strong. If one settles upon ϕ ing only insofar as one maintains a commitment to ψ ing, then it makes little sense to undertake an all-out commitment to ϕ ing as this would hold even in circumstances where one’s ψ ing was precluded.

Ben’s decision to kill the family is a qualified one. Thus, by QDQI, we should infer only that Ben intends to kill the family rather than stop pumping. His decision does not bring with it an all-out commitment to the family’s demise. Thus, the contrastivist is able to maintain HOLISTIC DECISION and DECISION INTENTION without endorsing (2) (on its most natural reading).

The possibility of merely contrastive intentions also allows us to maintain a version of the SIMPLE VIEW, whereby an agent intentionally ϕ s only if relative to some alternative ψ they intend to ϕ rather than ψ . This version of the SIMPLE VIEW

does not require an all-out intention to ϕ . Thus, it is consistent with the apparent falsity of (2). Importantly, the vindication of HOLISTIC DECISION, DECISION INTENTION, and the SIMPLE VIEW does not require that all intentions are contrastive. Those who favor a disunified view of intention can capture the falsity of (2) by holding that it expresses a noncontrastive intention (in which case it is false). Moreover, those who favor a disunified account can (and should) still draw the relevant distinctions discussed above between HAOD and HQD together with QDAI and QDQI. This reasoning can be replicated with all-out intentions (i.e., intentions to ϕ rather than $\neg\phi$) replaced by unqualified noncontrastive intentions.

4. The Nature of Contrastive Intentions

In the preceding two sections I introduced and argued for contrastivism about intentions. However, the reader may still at this point feel adrift. I have argued that intentions have a particular structure, but I have not said much about what intentions fundamentally are. And without any commitments here, it is easy to worry that the solutions I have provided are hollow. Moreover, I have suggested that all intentions may be fundamentally contrastive. Although this strong hypothesis was not required to solve the puzzle with which we began, it is worth taking seriously. Finally, I have suggested that contrastive intentions constitute a kind of partial commitment. Yet, it may be unclear in what sense subjects like Ben undertake such a partial commitment. Indeed, it might not be obvious that Ben is really committed in any sense to killing the family in our example.

With this in mind I will, in this section, provide the beginnings of a unified approach to intentions within the contrastivist framework. Doing so will also allow me to explain why the reasons we observed in section 1 for rejecting (2) (i.e., the failure of Ben's apparent intention to satisfy the functional roles laid out by Bratman) do not apply to (3).

We can start with some remarks by Bratman (1987), who distinguishes intentions from beliefs and desires by reference to their functional role. As he tells us, they are 'conduct-controlling pro-attitudes, ones which we are disposed to retain without reconsideration, and which play a significant role as inputs to [means-end] reasoning' (Bratman 1987: 20; cf. Setiya 2022).¹

This is not the only role of intentions, of course. Like beliefs, they can also serve as stable points for nonpractical reasoning about the future (and other things, such as assertions about what one will do in the future). The difference is that the intention to ϕ because it involves a commitment to ϕ poses a problem for practical reason: the problem of how one will ϕ .

I depart from Bratman in a number of ways. First, Bratman treats intentions as analogous to desires. They are like desire like in structure, except that they involve commitment. It is the fact that intentions are committing that allows them to play

¹ This is a slight modification of what Bratman actually says, found in Setiya (2022). Bratman (1987) does not, strictly speaking, state that intentions are inputs to means-end reasoning, but rather inputs to reasoning 'to yet further intentions' (20). Nonetheless, the context and later discussion make it clear that he has means-end reasoning in mind.

their volitional role and their role as inputs to practical reasoning. I take intentions to be analogous to preferences rather than desires. That is, when one intends to ϕ rather than ψ , one is committed to ϕ ing rather than ψ ing. But one's commitment need not go any further than this. We can, of course, intend to ϕ rather than $\neg\phi$. This will involve an all-out commitment.

Second, Bratman takes all intentions to be conduct controlling (in the sense that if one intends to ϕ , then one is disposed to ϕ when given the opportunity), and he takes intentions to filter out incompatible alternatives (in the sense that if one intends to ϕ , one is disposed not to intend things incompatible with ϕ). I take these to be features of all-out intentions. But they are not features of merely contrastive intentions.

Given that these features are commonly taken to be central to intention, it might be worried that because merely contrastive intentions do not have these features, they are not *really* intentions. This is at best a terminological worry regarding whether or not anything less than all-out intention is deserving of the label 'intention'. I think it is fitting to treat merely contrastive intentions as intentions because they are a species of mental state that (as I will shortly explain) plays important roles in planning, they are commonly referred to with intention ascriptions (as discussed previously), they are (unlike all-out intentions) necessary for intentional action, and they can be seen as a species of the same genus as all-out intentions. If one prefers to save the label 'intention' for all-out intention, then one can take 'merely contrastive intention' to be a newly identified type of mental state that plays important roles in planning that we commonly refer to with intention ascriptions, that is (unlike all-out intention) necessary for intentional action, and that can be seen as a species of the same genus as all-out intention. Very little of substance turns on this. If we go this way, we do have to say that the SIMPLE VIEW is, strictly speaking, false. The same is applies to HOLISTIC DECISION and DECISION INTENTION. But all three principles turn out to be so close to the truth that we are delivered a straightforward explanation of their intuitive appeal.

The final respect in which I depart from Bratman regards the sense in which intentions frame our practical reasoning. Bratman's thought is roughly this: when I intend to ϕ , I am thereby posed the task of identifying and committing to a way in which I will ϕ . I believe that all-out intentions, due to the fact that they incorporate all-out commitments, do pose such problems for means-end reasoning. All-out intentions can be thought of as a subspecies of a more general kind. They are characterized as involving all-out commitments. It is this feature of all-out intentions that allows them to play the functional roles Bratman assigns them.

Moreover, I am sympathetic to Bratman's general claim that intentions frame practical reasoning. And it is clear that some merely contrastive intentions do pose problems for practical reason. However, it is not clear that all contrastive intentions pose problems for means-end reasoning. Therefore, I deny that it is a core aspect of the functional role of intentions that they serve as inputs to means-end reasoning. Rather, I hold that intentions function as inputs to instrumental practical reasoning and that not all instrumental practical reasoning is means-end reasoning. This can be illustrated by returning to our core example.

In section 1 we observed that Ben was not posed the problem of how to kill the family. That is, he was not forced to identify a means to this end. He was already

fully committed to the means, and he was only committed to the killing insofar as it was necessitated by his commitment to the means. Thus, his intention to kill the family *rather than stop pumping* did not pose a distinct problem for his means-end reasoning.

However, there is another sense in which it does pose a problem for his instrumental reasoning: For most able-bodied agents the act of pumping a house full of gas will be fairly trivial. However, the act of killing innocents is far from trivial. Most people will have a hard time doing this. It requires significant strength of will. Thus, there is a sense in which the question of ‘how will I kill the family?’ is a pressing one for Ben. He is not faced with the problem of identifying a strategy that will guarantee the family’s death. He is, however, faced with the problem of overcoming his moral scruples and following through on a course of action that will result in the deaths of several innocent people. That is, he is faced with the problem of resisting predictable weakness of the will (in the sense of Holton 1999). This is no trivial matter. Without a strategy for overcoming his moral impulses, he will likely fail. Thus, if he is to follow through on his plan, Ben is required to identify such a strategy. His strategy could be something as simple as gritting his teeth and putting all thoughts out of his mind. He may look away, try to block out the family’s screams, or visualize his life after he has received his payment. He may even try to reason his way out of his moral scruples. None of these strategies constitute means to his end, at least not in the standard sense. Looking away or visualizing a happy problem-free life ahead, is not a way of pumping gas or of killing a family. However, the application of such strategies and the identification of the need for them is essential if Ben is to complete his plan. And the need for such strategies is prompted by the fact that his plan involves the killing of the family even though this is not one of his ultimate ends. This constitutes a second way in which intentions frame our practical reasoning.

A worry suggests itself at this point: Are contrastive intentions really needed to play this functional role? Could this functional role not be played just as well by beliefs about side effects? After all, beliefs also frame our practical reasoning, and the belief that continued pumping will bring about the family’s demise places similar demands on Ben’s planning. However, we can modify the case to eliminate the possibility that such beliefs are playing this role. Suppose that Ben expects the family to smell the gas and escape, but that he also considers it a live (but unlikely) possibility that they will remain and perish. Furthermore, suppose that he resolves to keep pumping in this eventuality. The acknowledgement of this possibility will pose a similar problem for his instrumental practical reasoning. However, this version of the case does not involve a belief that a particular side effect *will* occur.

Moreover, unlike beliefs about side effects, contrastive intentions do not merely pose the problem of how to resist weakness of the will—they may also constitute part of the solution. Following Holton (2004), we might hold that we can overcome or plan for weakness of the will by resolving to ϕ . Thus, in addition to triggering further practical reasoning, our intentions may serve as triggers for the will. It is our faculty of the will that explains our ability to stick to our plans (Holton 2003). As already noted, it would not make sense for Ben to resolve or

fully commit to killing the family. But forming a contrastive intention constitutes a sort of partial resolution. He resolves to kill the family rather than stop pumping. And this would be enough to trigger his will and thus allow him overcome weakness of the will. In this sense contrastive intentions might often solve the very problem they pose.

More generally, this discussion illustrates why it makes sense to form at least a qualified intention to ϕ upon deciding to ϕ . Known effects of our intended actions can raise obstacles that must be planned around, especially effects that, in Harman's terms, 'raise alarm bells', meaning that they are factored into our practical reasoning (Harman 1986). The resultant planning may not be means-end reasoning. It may involve identifying means for bringing about background conditions or facilitating conditions. These background conditions may, for example, include being of a mental constitution whereby one can overcome one's moral scruples.

This also illustrates the sense in which Ben really is committed to killing the family. This is not simply some anticipated outcome of his action that he can passively accept. It would be irrational for him simply to accept this consequence, in much the same way that it would be irrational simply to accept that one will ϕ after forming an intention to ϕ without ever planning the means via which one will ϕ .

The general idea here can be summarized as follows: In deciding to ϕ , we thereby accept the considered consequences of ϕ ing. We accept (however grudgingly) the whole situation encompassed by our ϕ ing. But our acceptance of these consequences can place fresh demands on us. And these demands need to be approached via practical reasoning. With this in mind, I propose that we can think of intentions as action guiding preference-like commitments that pose problems for instrumental practical reasoning. All-out intentions are a subspecies of intention involving all-out commitment (akin to a preference over all alternatives), and this feature allows them to play functional roles not played by other types of intention.

We are now in a position to return to the considerations (A)–(D) discussed in section 1 for rejecting (2) and to show that they do not undermine (3).

- A. Intentions as pro-attitudes: Intentions are pro-attitudes. However, all an intention to ϕ requires is, at best, a preference for ϕ over some salient alternative. Ben has no desire to kill the family. However, he does prefer their death to the cessation of his pumping.
- B. Problems for practical reason: Ben is not posed a means-end problem. However, he is posed a problem for practical reason in the broader sense identified above.
- C. Rational constraints on intentions: Ben's intention to kill the family rather than stop pumping requires that he avoid courses of action that conflict with his killing of the family insofar as they also conflict with his continued pumping. Ben's intentions are constrained in this way. His intentions are not constrained beyond this because his commitment is only partial.

- D. Endeavoring: All-out intentions, because they involve all-out commitments, give rise to endeavors. However, merely contrastive intentions do not involve all-out commitments. Thus, they will not inevitably give rise to endeavors.

Accordingly, the reasons provided in section 1 for rejecting (2) do not apply to (3). In the remainder of the paper I will further illustrate the structure and explanatory power of contrastivism by showing that it allows us to resolve two more objections to the SIMPLE VIEW.

5. Contrastive Intentions and the Simple View

5.1 Contrastivism and Consistency

We have seen that contrastivism allows for a version of the SIMPLE VIEW that is resistant to the objection from known side effects. However, the SIMPLE VIEW faces other challenges. By exploring the contrastivist response to these challenges, we are able to grasp the structure and explanatory power of contrastive intentions better.

The most influential objection to the SIMPLE VIEW turns on cases in which an agent intentionally and rationally ϕ s even though it seems irrational for them to intend to ϕ . Consider the following case:

GAMES: Michael is presented with two identical video games. In these games one uses a joystick to guide a missile to a target. Each game is difficult, and there is a prize for hitting one of the targets. However, one cannot hit both of the targets. If one hits the target on game A, then game B shuts down and vice versa. Moreover, both games shut down if one is about to hit both targets at the same time. Michael knows all this. Yet, being ambidextrous, he decides to play both games. He reasons that his chances of hitting either target are low and that by playing both games he maximizes his chance of winning a prize. Michael's strategy ultimately pays off. He hits target A and thereby wins a prize (cf. Bratman 1984: 381–83).

Michael hit target A intentionally. It was no accident that he hit it. His hitting target A was the result of a great deal of concentration and effort. However, as Bratman points out, if we say that Michael intended to hit target A, then we should say that he also intended to hit target B. His attitudes toward the two targets were, after all, symmetrical. However, he knew that it was impossible to hit both targets. Thus, if he intended to hit target A and he intended to hit target B, then he had inconsistent intentions. But this would render Michael's intentions irrational. This is the wrong result. Bratman concludes that the simple view is false. Intentionally ϕ ing does not require an intention to ϕ .

This reasoning assumes that intention is always a two-place relation between an agent and an action. Once we introduce contrastive intentions, it becomes clear that the SIMPLE VIEW is not threatened. Michael would be irrational if he had an all-out

intention to hit target A and an all-out intention to hit target B. Because hitting target B is known to be incompatible with hitting target A, an all-out intention to hit target A would commit Michael to intending to hit target A rather than target B. Michael has symmetrical attitudes toward A and B, thus he would also intend to hit target B rather than target A. This is inconsistent. However, to capture the fact that Michael intentionally hit target A, the proponent of the SIMPLE VIEW need only point out that Michael had an intention to hit target A rather than neither target. This is consistent with his also intending to hit target B rather than neither target. And it is consistent with his having no preference for A over B and vice versa.

It is natural to worry at this point that if one intends to ϕ rather than ψ and to χ rather than ψ , then one must intend to $(\phi \ \& \ \chi)$ rather than ψ . This inference is sometimes acceptable. If I say ‘I intend to go to the gym rather than the party’ and ‘I intend to go to the library rather than the party’, it seems reasonable to infer that I intend to go to the gym and the library rather than the party. If this inference is valid, then we have not resolved Bratman’s challenge: Michael intends to hit A rather than neither target and intends to hit B rather than neither target, and thus he must intend to hit A and B rather than neither target. This would render his intentions inconsistent.

However, this inference is not valid. Contrastive intentions involve preference-like rankings of alternatives viewed as possible outcomes of action. Forming an intention involves making a choice between these alternatives. The fact that ϕ is a possible outcome and that χ is a possible outcome does not entail that $(\phi \ \& \ \chi)$ is a possible outcome. There are cases in which ϕ and χ are both individually possible outcomes, but $(\phi \ \& \ \chi)$ is not a possible outcome due to the incompatibility of ϕ and χ . This is the case for GAMES.

Moreover, the inference breaks down even when the outcomes are compatible. Consider the following:

SUIT: I am going out in my fancy suit. I know there is a chance of wind and rain and that the rain would ruin my suit. I have an umbrella and a coat. The coat does not suit me. Still, I would rather wear the coat than get rained on. My best option is to use my umbrella. However, the wind may prevent my umbrella from functioning. Therefore, I set out with both my umbrella and my coat. I intend to wear my coat rather than get rained on. And I intend to use my umbrella rather than get rained on. However, if I am able to use my umbrella, I have no reason to don my ugly coat. I will wear my coat only if my umbrella is unable to function. I do not intend to wear my coat and use my umbrella both rather than get rained on.

Here we see that one can consistently intend ϕ rather than ψ and χ rather than ψ , without intending $(\phi \ \& \ \chi)$ rather than ψ even when $(\phi \ \& \ \chi)$ is an open possibility. Why, then, does the inference sometimes sound reasonable? My diagnosis is as follows: In SUIT and GAMES the relevant actions were ways of achieving the same end. Moreover, in each case performing one of the actions precludes obtaining the contrasting possibility. If one hits target A, then one has

not hit neither target. And if one utilizes one's umbrella, then one has avoided getting wet. Thus, hitting target A or utilizing one's umbrella effectually renders one's intention to hit target B rather than neither or to wear one's coat rather than get rained on, inert (note that this is not the same as the intention being satisfied; rather, it is akin to a conditional intention such as the intention to ϕ if ψ being rendered inert by the preclusion of ψ).

The situation is different when we consider tasks such as going to the gym or library. Gym and library attendance are standardly associated with very different ultimate ends. One goes to the library in order to become more learned. One goes to the gym in order to get fit. One will not become more learned by going to the gym. And one will not get fit by going to the library. Thus, if one succeeds in going to the library rather than the party, one does not thereby render inert one's intention to go to the gym rather than the party as this latter intention was formed (we will generally assume) as part of a different plan: one with the improvement of one's physical well-being as its end. It is natural to infer that when somebody expresses an intention to go to the library rather than the party and an intention to go to the gym rather than the party, they intend to go to both the library and to the gym, rather than to the party.

Once again, the reasoning in this section can be replicated by the disunified theorist by replacing talk of all-out intentions with talk of unqualified noncontrastive intentions.

5.2 Disjunctive Intentions

Bronner and Goldstein (2018) have presented a new challenge to the SIMPLE VIEW. They offer the following case:

MAIL BOMB: Millie mails a bomb to the Smith residence, intending to kill either Mr. Smith or Ms. Smith. As it happens, Mr. Smith opens the package and is killed (Bronner and Goldstein 2018: 794).

Millie intentionally killed Mr. Smith. However, it seems wrong to claim that she intended to kill Mr. Smith. After all, her intention would seem to be satisfied equally well by killing Ms. Smith. She is, thus, not committed to bringing about Mr. Smith's demise. Her commitment is weaker than this. She is only committed to killing either Mr. Smith or Ms. Smith.

MAIL BOMB is clearly problematic for versions of the SIMPLE VIEW formulated in terms of all-out intentions. After all, such formulations would require that Millie have an all-out intention to kill Mr. Smith. She clearly lacks such an intention. However, the contrastivist has a response: Millie intends to kill Mr. Smith rather than neither Mr. nor Ms. Smith (call this intention 'MR'), and she intends to kill Ms. Smith rather than neither (call this intention 'MS'). But she does not intend to kill Mr. Smith rather than Ms. Smith. Nor does she intend to kill Ms. Smith rather than Mr. Smith.

This takes us some way toward dissolving Bronner and Goldstein's challenge. However, it does not get us all the way. After all, Bronner and Goldstein point out

that Millie's intention would seemingly be satisfied by killing either Mr. Smith or Ms. Smith. When Mr. Smith opens the package and is blown to smithereens, MR is satisfied. But what about MS? To resolve this issue, we must recall that the satisfaction of one contrastive intention can render another such intention inert if they are both relativized to the same alternative (which is precluded when one intention is satisfied) and if they are both formed relative to the same ends. This was the lesson of SUIT. It applies equally here. Although MS goes unsatisfied, it is rendered inert by the satisfaction of MR because the satisfaction of MR precludes the contrasting possibility (that neither Mr. nor Ms. Smith perish) and it is formed relative to the same end (the desire to harm at least one Smith). In addition, Millie's all-out intention to kill either Mr. or Ms. Smith is satisfied. The proponent of the SIMPLE VIEW should, thus, deny that all of Millie's intentions are satisfied. However, they can capture the intuition that her intentions are satisfied by pointing out that her ultimate end is achieved, and those intentions that go unsatisfied are rendered inert—they are in no way frustrated.

This concludes my application of contrastivism in defense of the SIMPLE VIEW together with HOLISTIC DECISION and DECISION INTENTION. The potential explanatory power of contrastivism provides substantial support for the view that at least some intentions are contrastive. Moreover, as already indicated, considerations of theoretical unity suggest that if some intentions are contrastive, then all intentions are contrastive.

ANDREW PEET

LEEDS UNIVERSITY

andrewpeet123@gmail.com

References

- Anscombe, Elizabeth. (1963) *Intention*. 2d ed. Oxford: Blackwell.
- Aune, Bruce. (1978) 'Hall on Intention and Decision'. *Journal of Philosophy*, 75, 564.
- Bennett, Jonathan. (1980) 'Morality and Consequences'. In S. M. McMurrin (ed.), *The Tanner Lectures on Human Values*. Cambridge: Cambridge University Press.
- Bratman, Michael. (1984) 'Two Faces of Intention'. *The Philosophical Review*, 93, 375–405.
- Bratman, Michael. (1987) *Intentions, Plans, and Practical Reason*. Cambridge, MA: MIT Press.
- Bronner, Ben, and Simon Goldstein. (2018) 'A Stronger Doctrine of Double Effect'. *Australasian Journal of Philosophy*, 96, 793–805.
- Davidson, Donald. (1980) 'Intending'. In *Essays on Action and Events*. (Oxford: Oxford University Press), 83–102.
- Ferrero, Luca. (2009) 'Conditional Intentions'. *Noûs*, 43, 700–41.
- Frankfurt, Harry. (1988) *The Importance of What We Care About*. Cambridge: Cambridge University Press.
- Goetz, Stewart. (1995) 'The Choice-Intention Principle'. *American Philosophical Quarterly*, 32, 177–85.
- Goldman, Alvin. (1970) *A Theory of Human Action*. Princeton University Press.
- Hall, James. (1978) 'Deciding as a Way of Intending'. *Journal of Philosophy*, 75, 553–64.
- Harmatn, Gilbert. (1986) *Change in View*. Cambridge, MA: MIT Press.
- Hieronymi, Pamela. (2005) 'The Wrong Kind of Reasons'. *Journal of Philosophy*, 102, 37–57.
- Hieronymi, Pamela. (2006) 'Controlling Attitudes'. *Pacific Philosophical Quarterly*, 87, 45–74.
- Holton, Richard. (1999) 'Intention and Weakness of the Will'. *Journal of Philosophy*, 96, 241–62.

- Holton, Richard. (2003) 'How is Strength of the Will Possible?' In S. Stroud and C. Tappolet (eds.), *Weakness of the Will and Practical Irrationality* (Oxford: Clarendon Press), 39–67.
- Holton, Richard. (2004) 'Rational Resolve'. *Philosophical Review*, 113, 507–35.
- Holton, Richard. (2009) *Willing, Wanting, Waiting*. Oxford: Oxford University Press.
- Kane, Robert. (1996) *The Significance of Free Will*. Oxford: Oxford University Press.
- Kaufman, Arnold. (1966) 'Practical Decision'. *Mind*, 75, 25–44.
- Kavka, Gregory. (1983) 'The Toxin Puzzle'. *Analysis*, 43, 33–36.
- Knobe, Joshua. (2004) 'Intention, Intentional Action, and Moral Considerations'. *Analysis*, 64, 181–87.
- McCann, Hugh. (1986) 'Intrinsic Intentionality'. *Theory and Decision*, 20, 247–73.
- McGuire, John. (2016) 'Can One Decide to do Something Without Forming an Intention to do it?' *Analysis*, 76, 269–78.
- McHugh, C. (2013) 'Normativism and Doxastic Deliberation'. *Analytic Philosophy*, 54, 447–65.
- Meiland, Jack. (1970) *The Nature of Intention*. London: Methuen.
- Mele, Alfred. (1992). *Springs of Action: Understanding Intentional Behavior*. Oxford University Press.
- Mele, Alfred. (2003) *Motivation and Agency*. Oxford: Oxford University Press.
- O'Shaughnessy, Brian. (1980) *The Will: A Dual Aspect Theory*. Cambridge University Press.
- Paul, Sarah. (2012) 'How Do We Know What We Intend?' *Philosophical Studies*, 161, 327–46.
- Pink, Thomas. (1996) *The Psychology of Freedom*. Cambridge: Cambridge University Press.
- Raz, Joseph. (1975) *Practical Reason and Norms*. Hutchinson.
- Rozeboom, Grant. (2015) 'Side Effects and the Structure of Deliberation'. *Journal of Ethics and Social Philosophy*, 9, 1–19.
- Schaffer, Jonathan. (2007) 'Knowing the Answer'. *Philosophy and Phenomenological Research*, 75, 383–403.
- Searle, John. (2001) *Rationality in Action*. MIT Press.
- Sennet, Adam. (2016) 'Ambiguity'. *Stanford Encyclopaedia of Philosophy*. <https://plato.stanford.edu/entries/ambiguity/>.
- Setiya, Kieran. (2022) 'Intention'. In Edward N. Zalta (ed.), *Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/entries/intention/>.
- Shah, Nishi. (2008) 'How Action Governs Intention'. *Philosopher's Imprint*, 8, 1–19.
- Shah, Nishi. (2013) 'Why We Reason the Way We Do'. *Philosophical Perspectives*, 23, 311–25.
- Shepherd, Joshua. (2013) 'The Apparent Illusion of Conscious Deciding'. *Philosophical Explorations*, 16, 18–30.
- Snedegar, Justin. (2013) 'Reason Claims and Contrastivism about Reasons'. *Philosophical Studies*, 166, 231–42.
- Snedegar, Justin. (2017) *Contrastive Reasons*. Oxford University Press.