

This is a (very) long version of my ‘Contrastive Intentions’ paper – too long to send out to journals. The basic argument is contained within the shorter version. Moreover, the sections of this paper that do not appear in the shorter version are a lot more rough. Please don’t cite without permission, the shorter version is currently in progress/under review.

Contrastive Intentions – Director’s Cut

Andrew Peet

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.

1. A Puzzle

Consider the following case:

POISON: Ben is struggling in poverty when an eccentric billionaire offers to pay him a huge sum of money to manually pump a large amount of poisonous gas into a house. Ben agrees, and starts pumping. However, he quickly discovers that there is a family inside the house, and realizes that they will die if he keeps pumping. He does not wish to kill anyone. However, he

also doesn't wish to stop pumping, as this would prevent him from receiving payment. So he begrudgingly continues, and ends up killing the family.¹

Cases like this – so called 'side effect cases' - are familiar from the literature on intentional action and the doctrine of double effect.² They are also deeply puzzling. Consider the following sentences:

(1) Ben intentionally killed the family.

(2) Ben intended to kill the family.

(1) seems true. After all, Ben knowingly killed the family. He was able to do otherwise, yet he chose to bring about their demise. In light of this, it would make perfect sense to ask him “why did you kill the family?” and expect a reasonable answer.³ Moreover, there are strong reasons to reject (2). Firstly, it is

-
- 1 C.f. Anscombe (1963: 41-42) & Setiya (2009/2018). Anscombe discusses a number of variations of this case, about which she has varying judgments. None of her cases correspond exactly to this case (see footnote 3 for discussion of one of Anscombe's variants). Similar cases are discussed widely in the literature on intentional action and the doctrine of double effect. For example, Bennett's strategic bomber case (Bennett 1980), or any other such well known case of apparently unintended side effects, can be substituted for POISON. They are also central to the literature on the Knobe effect (see (Knobe 2003)). Indeed, this literature has shown judgments about side effect cases to be shifty and potentially untrustworthy in a number of ways. Thankfully, this literature has also granted us a reasonably good understanding of the factors which lead to the variable judgments regarding side effects (even if the exact mechanisms remain mysterious), and there is no reason to suspect that any such factors are at play with respect to the central intuitions appealed to in this paper. See footnote 37 for a brief discussion of the Knobe effect in light of the contrastivist view of intentions developed here.
- 2 What I label cases of 'known but apparently unintended side effects' should be taken to include cases in which the agent is merely justified in believing that the side effect will occur, and cases in which the side effect has a high probability of occurring (in footnote 35 I also discuss cases where the side effect has a low expected probability but remains salient due to its high cost). Moreover, I will be assuming a fallibilist conception of knowledge, such that an agent can know that their ϕ ing will bring about effect ψ even when they cannot rule out all possibilities in which they ϕ without bringing about effect ψ (fallibilism is, at this point, fairly standard within epistemology, but for a conception of fallibilism particularly suited to the position developed here see Worsnip (2015)). This will be important later on, as infallibilists will have a harder time adopting my positive proposal (see footnote 28). I do not take this to be a major cost as infallibilism already faces a litany of problems.
- 3 The closest we get to a denial of (1) is Anscombe (1957) who, in response to a similar case, denies that the occupants are poisoned intentionally. Anscombe's case differs from the version presented here in two important ways: Firstly, in Anscombe's case the subject's usual job is to pump the house full of water. In this one instance, he is aware that the water is poisoned. However, he is (according to Anscombe) utterly indifferent to this. This is an important difference. As I point out in footnote 37, the account I present later predicts that when the subject is utterly indifferent to the foreseen effect they won't count as bringing about that effect intentionally. Secondly, in Anscombe's case the poisoned individuals are nazi rulers the elimination of which will bring about a new morally praiseworthy regime. This difference

widely taken as a datum that agents in such cases do not intend the known but undesirable consequences of their chosen actions. This judgement is central to the framing of such cases both in the literature on intentional action, and the related literature on the doctrine of double effect.⁴ The most natural explanation for the almost universal judgement that (2) is false is that it *is* false.⁵

Furthermore, there are seemingly compelling reasons to reject (2) beyond its mere appearance of falsity.⁶ The reasons are as follows:

(A) Intention is a pro-attitude. Intending to ϕ involves being positively disposed toward

ϕ ing. Yet Ben does not seem positively disposed toward killing the family. Indeed, he might

is important since, plausibly, in this case killing the occupants has a different moral valence, and it is well known at this point that the moral valence of a foreseen effect impacts people's judgments as to whether the action is intentional. That is, we tend to judge the effect to be intentional when it has negative moral valence but not when it has positive moral valence. This is known as the Knobe effect (see Knobe 2004). I discuss the Knobe effect further in footnote 37 where I suggest a tentative diagnosis in light of the approach to intention developed here (in line with Scaife and Webber (2013)).

- 4 This intuition is not only common among philosophers. 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 . This said, a minority do attribute an intention in such cases. As I note in footnote 37, I believe that contrastivism is actually able to explain why some participants continue to attribute an intention in such cases. This is because contrastivism captures an important sense in which the apparent side effect is intended.
- 5 The hypothesis that (2) is false isn't the *only* potential explanation for its seeming falsity. It may be that Ben does not intend to kill the family when he starts pumping, but forms the intention to do so upon discovering their presence. In this case, the intuition that (2) is false may be explained by the fact that Ben did not initially intend to kill the family. However, it is easy to produce versions of the case where the verdict that (2) is correct in an unqualified form stretches credulity. We might, for instance, imagine that Ben, driven by his desire not to kill the family, frantically shouts warnings, trying to get the family to leave. We might even imagine that he has the means with which to kill the family were they to leave (a gun perhaps), but that he is in no way disposed to utilize said means in the event that the family leaves the house. (1) still seems true in this version of the case. After all, Ben still knowingly kills the family. He is able to do otherwise, yet he continues pumping, fully aware of the likely consequences of his action. Yet, (2), in its unqualified form, seems highly implausible here, even after Ben's discovery. Furthermore, as noted in footnote 1, this is one instance of a species of cases which are well studied in ethics and the philosophy of action. In many of these cases (for example, Bennett's (1980) strategic bomber case, and Knobe's (2003, 2004) CEO cases) the agent is aware of the likely side effect before they make their decision. Given that these cases are all parallel in structure, and evoke parallel intuitions (i.e. the judgment that the agent intentionally ϕ s, but does not have an intention to ϕ) it is likely that these intuitions have a common source. Nonetheless, if one remains unconvinced by POISON then it is possible to substitute in any of these other cases. The dilemmas, and the proposed response, will remain the same. Thanks to XXXX for pushing me on this.
- 6 I am skeptical as to the force of these considerations, I will return to this issue in §4 when I critically discuss these considerations as they apply to the claim that Ben intends to kill the family *rather than stop pumping*.

even see their death in a negative light.⁷

Moreover, Ben's attitude toward the family's demise seems to lack 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, it might be thought, will not see himself as being posed the problem of how to kill the family. There is no pressure on him to identify a means which will guarantee their death.

(C) Intentions place rational constraints on further intentions. If one intends to ϕ , and an option opens up which is incompatible with ϕ ing, one cannot rationally take that option whilst retaining the intention to ϕ . Yet Ben is not obligated to avoid courses of action simply because they are inconsistent with killing the family. If he was presented with the option to keep pumping without killing the family he would likely take it. He would be rational to do so, and his doing so would not require a revision of his prior intentions.

(D) Intending to ϕ usually leads to one's endeavouring to ϕ . That is, intentions have volitional force. Yet it does not seem that Ben endeavours to kill the family. This is not something he sets out to achieve.

Thus, there are strong reasons to hold that (2) is false.

However, there are also compelling reasons to accept (2). First, consider the Simple View of intentional action: The Simple View holds that if one intentionally ϕ s one must intend to ϕ .⁸ The Simple View is

⁷ An argument along these lines is presented by Nadelhoffer (2006), who points out that the agents in Knobe effect cases report indifference to the occurrence of the foreseen negative outcome. He takes this as evidence that the foreseen outcomes are not intended since, according to Nadelhoffer, an agent does not intend to ϕ if they are indifferent as to whether or not they ϕ .

⁸ The Simple View does not claim that ϕ ing with the intention to ϕ is sufficient for intentional action. At the very least the agent's ϕ ing must be related to their intention to ϕ in the right way (for example, through skilled execution of the

attractive largely due to its simplicity. It initially strikes many 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,⁹ and the claim that choosing to ϕ involves forming an intention to ϕ . However, according to the simple view (1) entails (2). So, if (1) is true, and (2) is false, then we must reject the Simple View.

Secondly, (2) seems to be 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 .¹⁰ This makes a great deal of sense: typically, in deciding to ϕ we settle the

intention to ϕ). It may be that even this is insufficient for intentional action (this is one lesson which could be drawn from the Knobe effect, see footnotes 31 and 37 for discussion).

9 If an armed robber demands my wallet then, in one sense, I have little choice but to hand it over. But this is just to say that one of the options available to me (that of handing my wallet over) is so much better than the alternative (getting stabbed) that the alternative is not worth considering. There is still a clear sense in which I do have a choice: I can choose to hand my wallet over, or to hold onto it and suffer the repercussions.

10 See, for example, Kaufman, (1966), Hall (1978), Davidson (1980), McCann (1986), Frankfurt (1988), Mele (1992, 2000, 2003), Kane, (1996), Pink (1996), Searle (2001), Hieronymi (2005, 2006) Clarke (2008, 2010), Holton (2009), and Shepherd (2013). Some theorists reject (or do not clearly endorse) 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 which 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 *ceritus paribus* version of DECISION INTENTION which generates the same problem. The view that deciding to ϕ involves intending to ϕ has been challenged by Bratman (1987), Rozeboom (2015) and McGuire (2016). Each point to side effect cases (like POISON) as counter

practical question of whether or not to ϕ in favour of ϕ ing. Intention is naturally thought of as the state that results from deliberation about what to do, much as belief is the state that results from deliberation about what is the case. Just as believing that p consists in settling on p , and being committed to its truth, intending to ϕ involves settling on ϕ , and thereby being committed to ϕ ing (Hieronymi (2005, 2006)).¹¹ Moreover, even in cases where the question of whether to ϕ and whether to intend to ϕ have sometimes been thought to come apart (for example, in Kavka's toxin puzzle (Kavka (1987))), reasons in favour 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 .

The basic thought behind HOLISTIC DECISION is that our decisions encompass more than just our ends abstracted from their contexts and considered consequences. We are not able to, and should not, screen off the known consequences of our decisions. Rather, we choose between the situations which we have the power to bring about (at a level of detail corresponding to our considered knowledge of how those situations will play out). Like DECISION INTENTION, HOLISTIC DECISION has a great deal of intuitive support. As Bratman says of his own version of the principle:

examples. 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. Firstly, as Bratman himself maintains, policies admit of exceptions. One might have a policy to wear a seat belt and nonetheless rationally retain that policy whilst violating it in special circumstances. Likewise, policy violations in side effect cases may constitute exceptions. Secondly, if one genuinely has a policy against ψ ing, one envisages that in ϕ ing one will thereby ψ , and one chooses to ϕ anyway, it's not clear that one can avoid the charge of irrationality or inconsistency simply by filtering ϕ ing from one's intentions. Rozeboom 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 tough pill to swallow. If we wish to hold onto 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 don't place many new and distinctive demands on our practical reason. As I discuss in §4, they sometimes will place new demands on our practical reasoning, but when they do this will generally be for good reason).

11 See also Shah (2008, 2013) and McHugh (2013).

“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, the rejection of HOLISTIC DECISION is unpalatable for a number of reasons.¹² Firstly, as Bratman points out, if we do not compare and decide upon whole scenarios in our decision making, and 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. Moreover, as Rozeboom (2015) emphasizes, decisions conclude practical reasoning, and thus resolve practical unsettledness. Thus, if we know that ϕ ing will also involve ψ ing, and we resolve to ϕ whilst remaining unsettled on ψ (for example, though not practically accepting that we will ψ) then we are behaving irrationally. We are treating ϕ as settled and ψ as unsettled whilst recognising that they come together.

Unfortunately, HOLISTIC DECISION and DECISION-INTENTION, together with the facts of POISON, appear to entail (2). Ben knew that he would, by pumping the house full of gas, kill the

¹² Nonetheless, some have rejected principles similar to HOLISTIC DECISION. For example, Harman (1986) considers and rejects a variant on this principle according to which rational agents must consider and accept all of the foreseeable consequences of ϕ ing when deciding to ϕ . He rejects this principle on the grounds that it is too cognitively demanding. As Bratman (1987) points out, Harman's objections do not apply to principles such as HOLISTIC DECISION which only take into account effects which are seriously considered. HOLISTIC DECISION might also be thought to have the implausible consequence that, for example, when I decide to go for a walk I decide to contract my muscles. After all, I know that walking involves the contracting of my muscles. However, unless this realization is occurrent and factors in the subject's reasoning HOLISTIC DECISION does not have this consequence.

family. This is something he seriously considers, yet he chooses to keep pumping any way. So, by HOLISTIC DECISION, he decides to kill the family. This seems correct: there is a clear sense in which he does decide 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.¹³ However, rejecting these principles comes at a severe cost, and forces us to complicate our pictures of intentional action and practical reason. I propose that we can keep the Simple View, HOLISTIC DECISION and DECISION-INTENTION, whilst 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.

The rest of the paper will be structured as follows: §2 sketches the basic idea behind contrastive intentions, it also provides a basic argument for contrastivism, and presents a suggestive argument for the hypothesis that all intentions may be fundamentally contrastive. §3 illustrates how contrastivism solves the side effect puzzle. This discussion informs §4, which puts some flesh on the contrastivist bones by giving the beginnings of a unified account of intentions within the contrastivist framework. §5 considers the question of whether contrastive intentions can be incorporated into the standard Bratmanian framework with the use of contrastive policies. It is argued not only that contrastive policies fail to capture our data, but that they actually illustrate an additional role played by contrastive

¹³ HOLISTIC DECISION is rejected by Goetz (1995), and DECISION INTENTION is rejected by Bratman (1987) and Rozeboom (2015).

intentions for limited planning agent such as ourselves. Finally, §6 further illustrates the structure and explanatory power of contrastivism by showing how it solves a number of other problems for the Simple View of intentional action.

2. Contrastive Intentions

2.1 Contrastive Intention ascriptions

Not all intention ascriptions have the surface form “I intend to ϕ ”. For example, we can also make explicitly contrastive intention claims, such as “I intend to ϕ rather than ψ ”. I might say “I intend to eat salad rather than chips”, or “I intend to go to the library rather than the party”.¹⁴

There are several ways we could potentially treat such sentences, some more promising than others. On the one hand we could treat them as meaning “I intend to ϕ and not ψ ”. So, 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).¹⁵ On the other hand, we could take them to be genuinely contrastive. That is, we could take intention to at least sometimes constitute 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 relative to which it is intended.

14 As observed by Anscombe (1963), and more recently by Setiya (2008), and Marušić and Schwenkler (2018), intention ascriptions also frequently take the form “I will ϕ ”. Indeed, this is perhaps the paradigmatic means of intention expression. This might be thought to lend further support to the hypothesis that Ben, at least in some sense, intends to kill the family in POISON. After all, it would not seem out of the ordinary for him to state or think, perhaps with regret, “I will kill the family” after realizing that the family is in the house. We also get contrastive ‘will’ (and related) statements, such as ‘I will/have decided to/am going to/have resolved to kill the family rather than stop pumping’. I do not wish to place much weight on such considerations. I worry that statements of the form “I will ϕ ”, or ‘I am going to ϕ ’, although often expressions of intention, might also sometimes serve as mere expressions of belief or knowledge. Statements of decision or resolution are obviously closer to statements of outright intention, but they also seem less common.

15 Strictly speaking there are two potential ‘and-not’ readings: ‘I intend to ϕ and I intend not to ψ ’, and ‘I intend to ϕ and it is not the case that I intend to ψ ’. This ambiguity will not matter here as both readings entail that ‘I intend to ϕ rather than ψ ’ entail ‘I intend to ϕ ’. And it is this entailment that is problematic.

There is a boring version of contrastivism according to which whenever we intend to ϕ we always intend to ϕ rather than not- ϕ (or, we intend to ϕ rather than any alternative ψ which we know to be inconsistent with ϕ ing). This approach gains intuitive support from the thought that when one forms an intention one is committed to not failing. And it would capture the sample sentences I have just presented insofar as library and party attendance/salad and chips eating are inconsistent. This contrastivism is very much in the spirit of standard approaches to intention as involving all out commitments.

A more ambitious contrastivism would hold that one can intend to ϕ rather than ψ , whilst not intending to ϕ rather than χ , for some known alternative χ which is incompatible with ϕ . This would allow that one could intend to ϕ rather than ψ , whilst not intending to ϕ rather than not- ϕ . 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 p387) 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 these alternatives. Likewise, Ben is only able to partially 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 would involve a preference-like choice between alternatives viewed as possible outcomes of action.¹⁶ The intention to ϕ rather than ψ involves a ranking of alternatives which

16 There are multiple senses of ‘option’ or ‘alternative’ relevant to action and practical reasoning. So it is worth clarifying

place ϕ ahead of ψ .¹⁷ 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. So it would commit the agent to ϕ over ψ . But it would not commit the agent to ϕ over χ .¹⁸

To be clear, the analogy with preferences here is purely structural. Traditional conceptions of intention can be seen as treating intention as desire like in the following sense: it is possible to desire some outcome simpliciter. One needn't 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

what I mean by 'options' or 'alternatives'. In one sense, our options/alternatives are possible effects we are able to bring about in the world. This is the relevant sense of option/alternative here: when we form an intention to ϕ we commit to bringing about some concrete effect. When we are choosing between the alternatives ϕ and ψ we are deciding which concrete effects we are going to commit to bringing about. We won't always know what our options are in this sense (this will become relevant in §5). And choosing the option of ϕ ing will not always result in our successfully ϕ ing. The second sense of 'option' concerns the different intentions it is available to us to form. For example, one's options might be to intend to ϕ or to intend to ψ . These options are arguably more relevant to decision theory and the subjective ought, as the expected utility of intending to ϕ may differ from the expected utility of actually ϕ ing due to the possibility of failure to successfully implement one's intention (see Hedden (2012), and Koon (forthcoming) for a response). Both notions of 'option' are closely related, and have important roles to play in explaining rational action. The former options are effects which the agent may commit themselves to, and the latter are commitments the agent may undertake.

17 This is merely an outline of robust contrastivism about intentions. I believe there are multiple ways in which contrastivism can be developed, and they will all solve the core problems addressed in this paper in roughly the same way. I will outline my own preferred view in §4. However, I consider it more important to establish the plausibility of some form of contrastivism (fitting the rough outline above) than I do to establish the correctness of my preferred contrastivist view of intentions. Hence separating my discussion of the basic contrastivist framework and my preferred contrastivist theory of intentions.

18 Thanks to XXXXX for helping me with this formulation of the view.

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 . Rather, it merely involves a commitment to ϕ ing rather than ψ ing. Moreover, preferences and desires are, unlike intentions, do not require a commitment to the desired or preferred outcomes.

The first two approaches (“and-not” and boring contrastivism) are consonant with standard assumptions about intentions as involving all out commitment. Thus, they have some claim to being the default approaches. However, they encounter problems when we consider cases like POISON. Consider the following:

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

Unlike (2), (3) does not seem obviously false. Indeed, this would be a natural way for Ben to describe his intention.¹⁹ This provides some inconclusive support for the truth of (3). But this 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 reconcile plausible versions of HOLISTIC DECISION, DECISION INTENTION, and the Simple View with the standard intuitions about side effect cases, it is able to resolve a number of other challenges to the Simple View. Moreover, I argue that contrastive intentions play several important roles for limited planning agents. I believe these considerations provide weighty independent support for contrastivism.

¹⁹ This is, of course, not the only appropriate way to describe his behavior. It is, for example, also appropriate to say that he would rather kill the family than stop pumping, or that he will kill the family rather than stop pumping. Neither of these descriptions strictly imply that he intends to kill the family rather than stop pumping. However, as noted in footnote 14, to say of one’s future or current behavior ‘I will ϕ ’, is typically to express an intention to ϕ . Moreover, the description ‘he would rather ϕ than ψ ’ describes the agent’s reason for ϕ ing. It is, thus, perfectly consistent with his intending to ϕ , or his intending to ϕ rather than ψ . It is also appropriate to say he chose to ϕ rather than ψ , or that he decided to ϕ rather than ψ . These more plausibly entail some form of intention to ϕ (as discussed in §1).

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 killing the family. If the option arose to keep pumping whilst allowing the family to live he would likely take it. And his doing so would involve no revision of his prior commitments.

Thus, if (3) is true then we must reject the assumption that intending (in a non-conditional way) necessarily involves an all-out commitment. We must allow that Ben intends to kill the family relative to some alternatives (such as the cessation of pumping), but not relative to all alternatives. For example, we should allow that that (3) is consistent with the falsity of (4):

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

This makes POISON a case in which an agent rationally intends to ϕ (kill the family) rather than ψ (stop pumping) for one alternative, and does not intend to ϕ (kill the family) rather than not- ϕ (allow the family to live). Thus, if (3) is true, then we must embrace robust contrastivism.²⁰

²⁰ Importantly, this does not entail that ‘I intend to ϕ rather than ψ ’ *never* expresses (or, at least, is used to communicate) ‘I intend to ϕ and not ψ ’. There are some cases where such a reading seems clearly appropriate. For example, it seems appropriate for ‘I intend to eat the salad rather than the chips’. Other uses ‘rather than’ also seem to admit of different readings. For example ‘The police officer wants to catch criminals rather than allow them to commit crimes’ is naturally read as claiming that the police officer desires to catch criminals, and either desires not to allow them to commit crime or, at least, does not desire that they be allowed to commit crime. On the other hand, a contrastive reading seems more

This argument not only suggests that some intentions must be contrastive, it also places important constraints on the nature of contrastive intentions. If contrastivism is to work, it should be thought of, as outlined above, as a thesis about the structure of intentions rather than their contents. The reason for this is that it is not clear what genuinely 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 (or lack thereof)).²¹ Yet it is not clear what the satisfaction conditions could be for the content $[\phi \text{ rather than } \psi]$. What is it for $[\phi \text{ rather than } \psi]$ to obtain? I can only see two obvious options. The first is for reality itself to be contrastive, and to be such that $[\phi \text{ rather than } \psi]$. The alternative is for $[\phi \text{ rather than } \psi]$ to demand that ϕ obtain and ψ not obtain. The first option is hard to swallow. Yet the second would render $[\phi \text{ rather than } \psi]$ 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 ’. Yet we have seen that the ‘and-not’ reading is unable to capture the apparent truth of (3), and its apparent consistency with the falsity of (4). We rejected it for this reason. Of course, nothing I have said rules out the possibility of analyzing contrastive intentions as, for example, some complex of conditional and/or disjunctive intentions. It would be impossible to rule this out a priori. But no such analysis seems forthcoming, and a simpler alternative is available.²²

appropriate for ‘The police officer wants to quit rather than be held accountable for his actions’. This seems only to express a preference. Presumably the police officer would rather remain on the force and not be held accountable for his actions. It is possible that ‘rather than’ constructions are ambiguous. It is also possible that the contrastive reading is fundamental, and that ‘and-not’ readings emerge through some process of pragmatic enrichment.

21 Arguably, there are exceptions to this. For example, as Grzankowski (2016, 2018) observes, we can have attitudes to objects. For example, if I am asked to think of a number, and I think of the number seven, then I have an attitude with the number seven as its content. This attitude does not represent the world as being any particular way. It is not the sort of attitude that can be satisfied or realized etc. However, intentions are very much not like this. Intentions can be satisfied – we typically strive for their satisfaction. The content of an intention needs to be the sort of thing that can obtain. Friedman (2013) has argued there is a species of question directed attitudes. Examples include wondering, curiosity, and suspension of judgement. However, once again these attitudes are very much unlike intentions in that, although some of them (e.g. curiosity) can be satisfied, their satisfaction does not place demands on the world.

22 In §5 I consider and reject the possibility that contrastive intentions are a form of conditional intention: i.e. the intention

This problem does not arise if we take contrastivism as a thesis about the structure of intentions. It is not at all puzzling that attitudes could be contrastive in structure. After all, preferences constitute a clear and uncontroversial example of an attitude which is contrastive in structure. Thus, in order to capture this data (without extreme metaphysical commitments) we should take contrastive intentions to be contrastive in structure, rather than a structurally non-contrastive attitude which can take contrastive contents.

2.3. Are all intentions contrastive?

The argument presented in §2.1 and §2.2 suggests that at least some intentions are genuinely contrastive. This is a significant result. As should become clear in §3, §5, and §6, even this limited contrastivism is able to do substantive work: it is able to save the Simple View, along with plausible versions of HOLISTIC DECISION, and DECISION INTENTION. Moreover, merely contrastive intentions seemingly play a number of important roles for limited planning agents. However, before considering the explanatory work contrastive intentions can do, I believe it is worth considering a stronger view: that *all* intentions are contrastive. The basic 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 fundamentally contrastive is unified. It sees all intentions as structurally alike. Intentions can, thus, be seen as a unified attitude. In the remainder of this section I will present a linguistic argument that picks up on this theme.

There's a well-known strategy for generalising contrastivist analyses beyond explicitly contrastive cases: If "intend" isn't always at least implicitly contrastive, then "intend" is ambiguous. Yet there's

to φ if the options are φ and ψ .

little independent reason to believe that “intend” is ambiguous in this sense (Velleman (1989) suggests that ‘intend’ is ambiguous in another sense: between goal and decided upon outcome), and if we can avoid attributing ambiguity then we should.

Furthermore, “intend” fails standard tests for ambiguity, such as the coordination across conjunction test (employed by Schaffer (2007) and Snedegar (2013, 2017) in defence 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, whilst only explicitly using the term once. If the term is ambiguous then the sentence should sound infelicitous or, at the very least, require some interpretative effort. Consider the following examples:

1a. The colours are light.

1b. The feathers are light.

1.c. The colours and the feathers are light.²³

1c is not uninterpretable. However, it is semantically anomalous. It requires some interpretative effort.

We clearly cannot take the same sense of “light” to apply to both the colours and the feathers.

Conjunctions involving explicitly contrastive intention claims and outright intention claims do not seem similarly infelicitous: Consider the following:

(4) 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 special interpretative effort.

Yet the single occurrence of “intend” attaches to both conjuncts. And the second conjunct is explicitly contrastive. So “intend”, in this case, is contrastive. However, the first conjunct is not explicitly

²³ This example is taken from Sennett (2016).

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.²⁴ However, it 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.²⁵ This suggests that the default contrast for “I intend to ϕ ” is “not- ϕ ”. That is, “I intend to ϕ ” will be interpreted as meaning “I intend to ϕ rather than not- ϕ ”. I will, from here on out, 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.²⁶

24 In footnote 20 I noted the possibility that ‘rather than’ constructions are ambiguous – sometimes they admit of an ‘and-not’ reading. If the ‘and-not’ reading was appropriate for the ‘rather than’ in (4) this argument would fail. However since (4), like (3), seems an appropriate way for Ben to describe his action, the ‘and-not’ reading is not appropriate in this setting.

25 Such an all-out commitment could still be conditional on background assumptions (Ferrero (2009)). Moreover, it is important to remember that we are talking about possibilities viewed as potential outcomes of action. Not all possibilities will count. For example, I know that nothing that I can do will have any impact on whether or not the world will end tomorrow. So, when I undertake an all-out commitment to finishing this paper I am not choosing the finishing of this paper over the ending or the saving of the world. I know that I could work on a different paper however, or that I could leave this paper unfinished. I could also enjoy the weather or go home early to see my son. When I undertake an all-out commitment to finishing the paper I am choosing the completion of the paper over any of these alternatives insofar as they are inconsistent with the completion of the paper (not all these options necessarily are inconsistent with the completion of the paper. I could, for example, sit outside and finish the paper, thereby enjoying the weather).

26 It is important to note that for this argument to work (i.e. for it to show that all intentions are structurally contrastive) it is crucial that we have already established a contrastive reading for ‘rather than’ in ‘I intend to ϕ rather than ψ ’. Otherwise the argument would show too much. After all, we can construct a similar sentences with respect to belief:

(6) I believe I will get paid, and that I will kill the family rather than stop pumping.

The felicity of this sentence does not, by itself, suggest that beliefs are contrastive. The reason for this is that we have no reason to hold that “I believe I will kill the family rather than stop pumping” should be assigned a genuinely contrastive reading. The ‘and-not’ reading seems perfectly reasonable here. So coordination across conjunction does not require that belief is structurally contrastive.

This is not to say that beliefs are not contrastive. For example, Yalcin (2018), and Hoek (forthcoming) have argued for the related hypothesis that belief is question sensitive (see Elga and Rayo (MS) for a similar view). Anybody who accepts contrastivism about intention and also wishes to retain some form of cognitivism about intention will have to maintain a view of belief along these lines. Although I am sympathetic to the view that beliefs are question sensitive I hope to stay neutral on the issue here.

The argument would also over-generate if contrastivism was a thesis about propositional contents. Consider the following:

These considerations suggest that not only some, but all intentions may be structurally contrastive. This is, of course, a strong hypothesis. And, as we will see in §3, §5, and §6, 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 it seriously. We will return to this issue in §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, as the solution informs the positive view I present in §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 widespread 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 not- ϕ ”. 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, nor would he attempt to ensure their demise in any other way (even if he was easily able to do so). Thus, it is clear that he is committed to killing the family only insofar as this is a consequence of his continued pumping. His commitment goes no further. So (2), on its most natural reading, is false (if it is unclear at this point in

(7) I intend to pay my debts, and to kill the family or stop pumping.

The felicity of this sentence does not show that all intentions are disjunctive. Intention, like other attitudes, can take contents with many different logical forms. The felicity of (7), thus, does not require that ‘intend’ acquire a different reading in the first and second conjuncts. This stands in contrast to (5) which would, if the generalization of contrastivism is denied, require a single use of ‘intend’ to pick out two structurally distinct attitudes. Thanks to XXXX for pushing me on this.

what sense Ben is *committed* to killing the family, this will be clarified in the next section).²⁷²⁸

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, in making his decision 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*, not as an all-out decision to kill the family.

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

27 Of course, there are also distant possibilities in which he would cease his pumping. For example, he might abandon his pump upon learning that his employer was unable to pay him. However, this situation is quite different: in this situation he revises his intentions upon acquiring new information. This stands in contrast to the situations discussed above: his failure to, for example, murder the family as they pass him, or to prevent them from leaving were they to try, would involve no such revision of his intentions. Indeed, in the version of the case we are considering, he actually tries to bring about one of these possibilities by attempting to warn the family.

28 It is here that infalibilists about knowledge may run into problems. If we assume that Ben knows that his continued pumping will result in the family’s demise, then he must be so positioned that he can rule out any possibility in which his continued pumping fails to do so. Thus, it is less clear that we can deny that his commitment is all-out. After all, he is committed to killing the family in every open possibility. The infalibilist is not without options here. For example, they might deny that, in the more plausible versions of our case, Ben really *knows* that the family will die (after all, they might leave). However, it will certainly be less straightforward for the infalibilist to adopt the contrastivist solution to the puzzles I have outlined.

HOLISTIC DECISION:²⁹

HOLISTIC ALL-OUT DECISION (HAOD): 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 make an all-out decision to ψ .

HOLISTIC QUALIFIED DECISION (HQD): 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 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 his killing of the family. 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 they are rational, then they will form an intention to ϕ . However, this principle is ambiguous between all-out and qualified decisions, and all-out and merely contrastive intentions. It is not especially relevant how we should read DECISION-INTENTION as it applies to all out decisions since there is no pressure to hold that Ben makes an all-out decision to kill the family. After all, the most plausible reading of HOLISTIC-DECISION only requires a qualified decision in Ben's circumstances. So, we can accept that all out decisions generate all out intentions without issue. But what about the analogues of DECISION-INTENTION for qualified decisions? Here are the possible readings:

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

²⁹ Where left unqualified I take "decision" to mean "all-out decision" in these principles.

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 make 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. Thus, QDQI is the more plausible rendering of DECISION-INTENTION as applied to qualified decisions.

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).³⁰

This 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 non-contrastive 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 not- ϕ) replaced by unqualified non-contrastive intentions.³¹

30 The idea here is that when a subject knows that ϕ would result in ψ , and they decide to ϕ anyway, then they intend to ϕ rather than not ψ . This is one situation in which merely contrastive intentions are appropriate. However, this is by no means the only type of situation in which merely contrastive intentions are appropriate. I discuss further cases elsewhere (especially §6).

31 It is worth saying a brief word on the Knobe effect before continuing, as some authors (such as Nadelhoffer (2006))

4. The Nature of Contrastive Intentions.

In the previous 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, I think 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.³² Doing so will also allow me to explain why the reasons we observed in §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).

The contrastivist can (and, I believe, should) largely follow Bratman (1987) as to the nature of

have taken the Knobe effect to be problematic for the Simple View. Knobe (2003) observed that participants are more likely to claim that a foreseen outcome was brought about intentionally if the action was judged to be morally problematic than if the outcome was morally positive. This generates a challenge for the Simple View insofar as there are cases in which, as observed by Knobe (2004), we are happy to say an action is intentional, even when we are not happy to say that it is intended. However, this is just the old problem of apparently unintended side effects, and the explanation just given extends straightforwardly to Knobe’s cases. Beyond this the Knobe effect does not pose a distinct challenge to the Simple View, where the Simple View is taken to provide a necessary, not a sufficient condition on intentional action. I do believe that contrastivism is suggestive of a particular explanation of the Knobe effect. However, that will have to wait until footnote 37, when we will have discussed the nature of intentions further.

32 I do not claim that this is the only possible way to develop a viable contrastivist view of intentions. I suspect that it is possible to develop contrastivism within other frameworks, and that many of the arguments presented here will carry over. The following should be seen as just one way of putting flesh on the contrastivist bones.

intentions. Bratman distinguishes intentions from beliefs and desires by reference to their functional role. They are he tells us:

“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 2009/2018).³³

I do not believe this is quite true of intentions as a species of mental state. However, I believe it is close. I depart from Bratman in two key ways. The first will be clear in light of my remarks in §2 and §3: 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 their volitional role, and their role as inputs to practical reasoning.³⁴ I agree that intentions are pro-attitudes of a sort, and that they involve commitment. However, I take them 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. There may be other alternatives (e.g. χ) such that one is not committed to ϕ ng rather than χ ing. This will be the case when one’s intention to ϕ is merely contrastive. However, it is possible (indeed, it is common) to intend to ϕ rather than not- ϕ . This

33 This is a slight modification of what Bratman actually says, found in Setiya (2009/2018). Bratman (1987) does not strictly speaking state that intentions are inputs to means end reasoning, but rather inputs to reasoning “to yet further intentions”. Nonetheless, the context and later discussion makes it clear that he has means-end reasoning in mind. Thus, I take Setiya’s modification of the quotation to be appropriate.

34 What does it mean to say that intentions involve commitments? The basic idea is intentions, unlike other pro-attitudes such as desires, place certain rational demands on us. For example, if I desire to eat cake, but have no intention to do so as a result of my diet, then I am under no rational pressure to determine a means by which to acquire cake. However, if I decide to give up on my diet, and form the intention to eat cake for breakfast, then I am rationally required to identify a means to this end. If I form and maintain the intention to eat cake for my breakfast, but make no effort to acquire any cake by breakfast time, then I have exhibited a form of practical irrationality (this is the sense in which intentions, unlike desires, pose problems for practical reasoning). As I will shortly argue, the commitments placed on us by merely contrastive intentions needn’t involve the rational demand that we identify means to our ends. Rather, they place a demand on us to secure certain facilitating conditions which are necessary for us to actually follow through on our plans.

will entail a commitment to ϕ ing over any alternative which is incompatible with ϕ ing. That is, it will involve an all-out commitment. Thus, in this respect, all-out intentions will be largely akin to intentions more generally as conceived of by Bratman (hence the ease with which the disunified theorist was able to replicate the reasoning in §3 by simply replacing talk of all-out intentions with unqualified non-contrastive intentions).

The second 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. Thus, all-out intentions are, in this second respect also, akin to intentions more generally as conceived of by Bratman. 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 which allows them to play the functional roles Bratman assigns them.

Moreover, I am sympathetic to Bratman's general claim that intentions frame our practical reasoning. And I think it is clear that some merely contrastive intentions pose problems for practical reason. Consider, for example, SUIT: my intention to wear my coat rather than get rained on was merely contrastive. Yet it demanded of me that I take my coat with me, and thus formulate a plan regarding how I will transport it (i.e. do I just carry it, do I take a bag etc.).

However, it is not clear that all contrastive intentions pose problems for means end reasoning. So I am skeptical of the claim that it is a core aspect of the functional role of intentions as a species of mental

state 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. This can be illustrated by returning to our core example.

In §1 it was 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. As noted, he is not faced with the problem of identifying a strategy which will guarantee the family’s death (that is, he is not posed a question for means end reasoning). He is, however, faced with the problem of overcoming his moral scruples and following through on a course of action which 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. Although he is not posed the problem of how he will execute his intention, he is posed the problem of how he will **get himself** to follow through on this aspect of his plan. Without a strategy for overcoming his moral impulses he will likely fail. Thus, Ben is required, if he is to follow through on his plan, to identify such a strategy. His strategy could be something as simple as gritting his teeth and putting all thoughts out of his mind. Alternatively 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 the 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.³⁵³⁶

More generally, this illustrates why it makes sense to at least form a qualified intention to ϕ upon deciding to ϕ (or, upon forming an intention to ψ , in the realization that ψ ing will involve ϕ ing).

Known effects of our intended actions can raise obstacles which must be planned around, especially effects which, in Harman's terms, 'raise alarm bells', meaning that they are factored into our practical

35 It might be thought that this functional role can be played just as well by beliefs about side effects: 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 to those imposed by his contrastive intention. However, we can modify the case to eliminate the possibility that beliefs about side effects are playing the same functional role as merely contrastive intentions. 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 thus die. Furthermore, suppose that he resolves to keep pumping in this eventuality. This 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. It merely involves the side effect being considered a live possibility. More generally, we can make the expected probability of the side effect arbitrarily low, but still have it set off 'alarm bells' for the agent by raising the moral cost of the side effect. It might be thought that the mere doxastic possibility that the side effect will be brought about could play the same functional role as contrastive intentions. However, the mere doxastic possibility of a side effect will not always register as sufficiently relevant to warrant being factored into the agent's practical reasoning (e.g. if the cost is trivial and the probability is low).

36 Contrastive intentions might 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 predicted weakness of the will by actively resolving to ϕ . Thus, it might be thought that in addition to triggering further practical reasoning, our intentions also serve as triggers for the will. And, it is our faculty of the will that largely 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 has resolved to kill the family rather than stop pumping. And this would presumably be enough to trigger his will, and thus allow him overcome weakness of the will. So, in this sense contrastive intentions might often solve the very problem they pose. Of course, this depends on buying into Holton's theory of the mechanisms of willing and resolution. I don't wish to take a stand on that here.

reasoning (Harman, 1986)(i.e. the kinds of effects covered by HOLISTIC DECISION).³⁷ The resultant planning will not necessarily involve identifying means with which to bring about the relevant effect. However, it may involve identifying means for bringing about background or facilitating conditions which must be met in order for one to implement one's overall plan. These background conditions may, for example, include being of a mental constitution whereby one can overcome moral scruples (or emotional aversions) which may otherwise cause one to falter in one's implementation of the plan.³⁸ I believe the relevant reasoning, even if it is not means-end deserves to be classified as practical and instrumental. It is practical because it is essentially action guiding (it is clearly not theoretical). And it is instrumental since it concerns the practicalities of carrying out our plans, rather than identifying

37 This is also suggestive of an explanation of the Knobe effect: when an apparent side effect of an action violates a norm this sets off alarm bells for the agent. The norm violation is factored into the agent's practical reasoning, as it has the potential to constitute an obstacle to be confronted and overcome during the execution of the plan. Effects which are normatively neutral do not raise such alarm bells, so they do not factor into the agent's practical reasoning, and are not intended. A similar solution is advocated by Scaife and Webber (2013). Indeed, since contrastivism holds that, in side effect cases, the subjects typically intend to ϕ rather than ψ , it is also able to explain why a substantial minority of participants in Knobe's (2004) study do attribute an intention in side effect cases. After all, contrastivism captures an important sense in which the subjects of such cases do intend to ϕ . It merely holds that these intentions are not all-out. These remarks are, of course, merely suggestive. A full exploration of and response to the Knobe effect would take us beyond the scope of this paper.

38 It could be responded that we can construct versions of POISON in which the killing of innocents holds no significance for the protagonist (who we will call 'Psycho-Ben'). The killing of innocents does not raise a problem for Psycho-Ben's practical reasoning. Thus, he does not intend to kill the family. However, if the killing of the family does not register as significant to Psycho-Ben in any way, meaning it is not factored into his practical reasoning, then it is not clear that we should say that Psycho-Ben intentionally kills the family (as mentioned in footnote 3, Anscombe (1957) denies that the poisoning is intentional in an analogous version of the case). The situation is similar to that of Bratman's marathon runner: Bratman (1987) asks us to consider a marathon runner who realizes that, by completing the race, they will wear down their sneakers. He holds that, in certain circumstances, we may be willing to claim that the runner intentionally wears down the sneakers. For this to be the case, the sneakers must bear some significance to the runner. They may, for example, be a family heirloom. Otherwise, if the sneakers have no significance at all to the runner, it seems strongly counter intuitive to claim that the runner wears them down intentionally. I suggest that the reason for this is that, if the sneakers do have some significance to the runner, then the act of wearing them down is something which requires at least some small degree of mental fortitude and planning. The problem the sentimental runner is posed is similar to (though far less significant than) the problem Ben is posed by the prospect of killing the family. The killing of innocents is, to Psycho-Ben, equivalent to the wearing down of sneakers for the non-sentimental runner. Just as the non-sentimental runner does not intentionally wear their shoes down, Psycho-Ben does not intentionally kill the family. I concede that it is hard to accept the verdict that somebody could knowingly but unintentionally kill a family in the way we are imagining. However, it is also hard to imagine somebody for whom the killing of a family genuinely holds no significance whatsoever. I doubt there are any real life examples of such individuals (for even genuine psychopaths have to live with the risks associated with the killing of innocents, even if moral concerns do not move them).

suitable ends.³⁹

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 to simply accept this consequence, in much the same way that it would be irrational to simply 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 which pose problems for instrumental practical reasoning (broadly conceived). All-out intentions are a subspecies of intention which involve an all-out commitment (akin to a preference over all alternatives), and this feature allows them to play important functional roles not played by all forms of intention.

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

³⁹ It might be responded that the reasoning at play here is actually means-end after all: Ben's continued pumping was a means to his end of getting paid. His acts of looking away and visualizing a problem free future were not token causes of his continued pumping. However, they were still indirect causes of his continued pumping since his continued pumping is causally dependent on him adopting some such strategy. I believe this is a merely terminological issue. I think means to ends are naturally thought of in terms of token causation. But if they are thought of in terms of a more general form of causal relevance then we should simply agree with Bratman that it is essential to the functional role of intentions that they pose problems for means-end reasoning, and claim that Ben is, after all, posed a means end problem with respect to the killing of the family.

(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 (and, as noted earlier, I suspect that we can intend to ϕ rather than ψ without any such preference). 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 which 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 since his commitment is only partial. If he had an all-out intention to kill the family then he would be far more constrained: he would be irrational to adopt any intention which conflicted with the killing of the family.

(D) Endeavoring: All-out intentions, because they involve all-out commitments, give rise to endeavorings. It would be akratic to adopt an all-out commitment to ϕ ing without thereby endeavoring to ϕ . However, merely contrastive intentions do not involve all-out commitments. Thus they will not inevitably give rise to endeavorings. So, it is unsurprising that Ben does not endeavor to kill the family.

So, we see that the reasons provided in §1 for rejecting (2) do not apply to (3). In the remainder of the paper I will do two things. Firstly, in §5, I will consider an alternative, and far less radical, way of incorporating contrastive intentions into Bratman's framework via the use of contrastive policies. Following this, in §6, I will further illustrate the structure and explanatory power of contrastivism by showing how it allows us to resolve two more objections to the Simple View.

5. Contrastive Policies

I have just argued that a general contrastivist view of intentions can be developed in line with the general thrust of Bratman's approach to intentions. However, it may be thought that contrastive intentions can be incorporated into Bratman's framework in a far less radical way: through the inclusion of contrastive policies. In this section I will argue not only that contrastive policies are insufficient to capture our data, but that they actually give us an additional reason to embrace the more radical approach taken here, highlighting another important role such intentions play for limited planning agents such as ourselves.

Bratman sees intentions as taking many forms. We don't just have present directed intentions and plans, we also have policies. Policies involve a *ceteris paribus* commitment to act in a particular way in particular (usually reoccurring) circumstances (or, in the case of self-governing policies, to give pride of place to certain considerations in one's deliberation). Bratman provides the following examples:

My general policy is to turn down a second drink when I have to drive home; I see the host of the party approach me with a second drink; applying my policy I form the intention to turn him down when he gets to me. My general policy is to allow myself to chair only one committee at a time. I am told that Clara is on the phone. While walking to the phone I surmise that she will ask me to chair a second committee and, applying my policy, form the intention to turn her down. My intentions to turn down the host and Clara are, when I form them, policy-based.

Bratman, (1987): 57

Our policies can clearly be contrastive. Many policies have the form 'In circumstances C, I will ϕ '. The relevant circumstances may involve us being presented with certain alternatives. For example, I may have a policy to ϕ when the options are ϕ and ψ . This would result in me ϕ ing whenever my

alternatives are ϕ and ψ (all else being equal). If our data could be captured purely in terms of contrastive policies then we would not need to depart in any serious way from the standard Bratmanian approach to intentions.

Unfortunately, this is not enough to capture our data. That is, it is not sufficient to capture the falsity of (2):

(2) Ben intended to kill the family.

Compatibly with the truth of (3):

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

The reason for this is that policies generate actions by issuing either present directed or near future directed intentions. When Bratman saw the host approaching his policy caused him to form an intention to turn down a second drink. Likewise, his policy to only chair one committee per semester causes him to form an intention to turn down Clara's request. If policies are the only form of intention that can be contrastive, then the resultant present or near future directed intentions must be non-contrastive.

Suppose that Ben has a policy to always kill the family when the options are to kill the family or stop pumping.⁴⁰ This policy will issue a non-contrastive near future directed intention to kill the family. But if this is the case then (2) is true. This is the result we were attempting to avoid. So, contrastive policies are not enough. We need contrastive present and near future directed intentions in order to capture our data.

Nonetheless, the possibility of contrastive policies does suggest another reason for embracing contrastive present and near future directed intentions. Suppose you know that the payoffs for ψ , ϕ , and

⁴⁰ Of course, it is unlikely that Ben would have such a policy, but without such a policy it is even less clear how contrastive policies could resolve our puzzle.

their negations will always be as follows:

$$\psi = -4$$

$$\phi = -2$$

$$\sim \psi = 4$$

$$\sim \phi = 0$$

On the basis of this you form a policy to always choose ϕ over ψ (or, equivalently, to choose $\sim\psi$ over $\sim\phi$) when these are the available (and incompatible) alternatives.

When you know you have to choose between ϕ and ψ the policy will straightforwardly issue an unqualified intention to ϕ . Now suppose that you are uncertain of your options. You may be in a situation where your exclusive options are ϕ and ψ . This seems very likely. But you may also be in a situation where $\sim\psi$ is actually compatible with $\sim\phi$. This would be optimal if it was an available option, but it's being a possibility seems unlikely.

In these circumstances, does it make sense to form an all-out commitment to ϕ ? One's policy only commits one to ϕ ing if ϕ and ψ are the only options (or, the only options not excluded by other considerations). So, if you are not completely certain that ϕ and ψ are the only options you can't be certain what your policy commits you to in your present circumstances. Merely contrastive intentions are well suited to this kind of situation. One can form a merely contrastive intention to ϕ rather than ψ , without undertaking an unqualified commitment to ϕ ing.

It might be thought that conditional intentions are even better suited to this task than contrastive intentions. The thought here would be that you are best off forming a conditional intention to ϕ if your only options are ϕ and ψ , and an intention not to ϕ if it turns out that $\sim\phi$ and $\sim\psi$ can be rendered

compatible. Indeed, it may even be thought that this is all a contrastive intention really amounts to. That is, it might be thought that an intention to ϕ rather than ψ is nothing more than a conditional intention to ϕ if the options are ϕ and ψ .

This would be a mistake. Contrastive intentions differ from such conditional intentions in a number of key ways. Firstly, they differ in their satisfaction conditions. Suppose I intend to ϕ conditional on ϕ and ψ being my only options. Furthermore, suppose I mistakenly believe that ϕ and ψ are my only options, and so I go on to ϕ . In this situation my conditional intention is not satisfied as the antecedent has not been actualized. Suppose, however, that I had instead formed the merely contrastive intention to ϕ rather than ψ . As long as my successful ϕ ing precludes ψ my contrastive intention will be satisfied.

This serves to distinguish contrastive from conditional intentions. However, it does not show that contrastive intentions are superior in the circumstances we are discussing. For that we need to turn to a second distinction between contrastive and conditional intentions: the demands they place on our planning. Suppose I have a conditional intention to ϕ if my only options are ϕ and ψ . I may believe it to be overwhelmingly likely that these will be my only options. With this in mind, I may or may not plan for the contingencies in which I have alternative options. If I don't carry out such planning then I conserve cognitive resources, but I leave myself open to the risk that I will have to plan on the fly if my information about my options changes. After all, if I plan to ϕ if my options are ϕ and ψ , and I discover that ϕ and ψ are not my only options, then I will be left with a gap in my plan that needs to be immediately filled in. In the example above this contingency planning would be easy to carry out. However, in reality there may be a great many unlikely possibilities with respect to our future options. It would be impossible to plan for all such contingencies.⁴¹

41 In the above example such contingency planning was already built in, since it was specified that the agent also intended not to ϕ if it turns out that $\sim\phi$ and $\sim\psi$ can be rendered compatible. However, on top of the fact that, in reality, there will be a great many more contingencies to plan for, the strategy of always simply intending to ϕ if the options are ϕ and ψ , and to not ϕ otherwise, is hopeless. It fails to factor in the possibility of situations in which there is an option χ which is

This is where contrastive intentions have an advantage. If I merely intend to ϕ rather than ψ , and I discover that ϕ and ψ are not my only options, I am not thereby forced to engage in further practical reasoning. I can, if such reasoning is not convenient or economical, simply continue on as before. No gap would be left in my plan in these circumstances. At the same time, however, because of the merely contrastive nature of my intention further practical reasoning would be an option if it were convenient or economical. Doing so would not necessarily involve abandoning the intention to ϕ rather than ψ . So, for this reason, contrastive present and future directed intentions are better suited than conditional intentions for navigating circumstances where we attempt to execute a contrastive policy whilst in a state of uncertainty with respect to our options.

Merely contrastive intentions may not be necessary for creatures who could be infallibly certain of their alternatives at all times. However, for limited creatures like us contrastive intentions are apt to play an important role.

6. Contrastive Intentions and the Simple View

6.1. Contrastivism and Consistency

We have seen that there is reason to endorse contrastivism about intentions. We have also seen that this allows for a version of the Simple View which 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 better grasp the structure and explanatory power of contrastive intentions.

Probably 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:

worse than both ϕ and ψ . Thanks to XXXX or this point.

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 very 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 any way. He reasons that his chances of hitting either target are very low, so by playing both games he maximises his chance of winning a prize. Michael's strategy ultimately pays off. He hits target A and thereby wins a prize (cf. Bratman (1984: 381-383)).

It seems clear that 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, effort, and monitoring. However, as Bratman points out, it seems problematic to claim that he intended to hit target A. If we say that Michael intended to hit target A then we should surely 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, supposing that it is irrational to have inconsistent intentions, 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 moves too quickly. It 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. It is true that Michael would be irrational if he had an all-out intention to hit target A, and an all-out intention to hit target B. Since 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 clearly inconsistent. However, the Simple View needn't entail that he has such intentions. 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. This is a perfectly natural way to describe the case.

It might be worried that if one intends to ϕ rather than ψ , and to χ rather than ψ , then one must intend to ϕ and χ rather than ψ . This inference seems acceptable in some cases. For example, 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 quite 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 will not have resolved Bratman's challenge: Michael intends to hit target A rather than neither target, and intends to hit target B rather than neither target, 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. When characterising contrastive intentions I noted that they should be seen as involving 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 $(\phi$ and $\chi)$ is a possible outcome. There are cases in which ϕ and χ are both individually possible outcomes, but $(\phi$ and $\chi)$ is not a possible outcome due to the incompatibility of ϕ and χ . GAMES is such a case.

Moreover, the inference breaks down in some cases in which the outcomes are compatible. Consider the following case:

SUIT: I am going out in my fancy suit. I know there is a chance of wind and rain, and I know that the rain would ruin my suit. I have an umbrella and a coat. The coat does not suit me, so I

would prefer not to have to wear it. Still, I would rather wear the coat than get rained on. My best option is to use my umbrella, as my umbrella is both stylish and, when functional, able to protect me from the rain. However, I am also aware that the chance of wind may prevent my umbrella from functioning properly. With this in mind, I set out with both my umbrella and my coat in my possession. 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 only wear my coat if my umbrella is unable to function due to the wind. So I do not intend to both wear my coat and use my umbrella rather than get rained on.⁴²

Here we see that one can consistently intend ϕ rather than ψ , and χ rather than ψ , without intending (ϕ and χ) rather than ψ , even when (ϕ and χ) is an open possibility. This leaves us with the challenge of explaining why the inference sometimes sounds reasonable. My diagnosis is as follows: In SUIT and GAMES the relevant actions (hitting either target, or utilizing the coat or the umbrella) were ways of achieving the same ultimate end. In GAMES the agent aims to win a prize. Hitting either target is sufficient for achieving this goal. In SUIT the agent aims to stay dry. The coat and the umbrella are each individually sufficient to keep them dry. Moreover, in each case performing one of the actions precludes the obtaining of 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

42 This case bears similarities to Richard Holton's (2008) library books case, which he uses to motivate the existence of partial intentions (thought of as a distinct mental state separate from full intention). I suspect that contrastivism allows us to do away with partial intentions as a distinct mental state. The contrastivist is, as noted earlier, able to see intention as a unified kind, and see full (or, all-out) intention as a species of this more general kind distinguished by virtue of requiring all out commitment (analogues to a preference over all alternatives, rather than merely some). However, a full exploration of the relation between merely contrastive intentions and Holton's partial intentions would considerably lengthen this already lengthy paper, so it will have to wait for another occasion.

one's coat rather than get rained on, inert.⁴³⁴⁴

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 attends the library in order to become more learned. One attends the gym with the goal of improving one's physical health. One will not, by attending the gym, become more learned. And one will not, by attending the library, become more physically fit. Thus, if one succeeds in attending 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. As a result 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 the gym rather than 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 non-contrastive intentions. Thus, contrastive intentions allow for a straightforward response to Bratman's video game argument. As we'll see in the next subsection, they also allow us to respond to a more recent objection to the Simple View presented by Bronner and Goldstein (2018).

6.2. Disjunctive Intentions

Bronner and Goldstein (2018) have presented a new challenge to the Simple View. They offer the

43 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 .

44 Of course, my description of SUI is somewhat of a simplification: there are many more ends at play than I have mentioned, and these ends relate to my intentions in complex ways. For example, my end in using my umbrella is not simply to stay dry, but rather to stay dry and remain stylish. By wearing my coat I only satisfy one of these ends. I stay dry, but I don't remain stylish. However, this does not undermine my point. For, in the circumstances under which I would be forced to wear my coat, I have been precluded from looking stylish at all. There is no way to satisfy this end. Thus, insofar as my intention to use my umbrella rather than get rained on relates to the end of remaining stylish it no longer places any rational demands on me.

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, as Bronner and Goldstein point out, 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. The first thing to note is that Millie will intend to kill Mr Smith rather than neither Mr nor Ms Smith (call this intention "MR"), and she will also intend 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 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, as 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, her all-out intention to either 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 which 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. I believe that the potential explanatory power of contrastive intentions provides substantial support for the view that at least some intentions are structurally contrastive. Moreover, as already indicated, I believe that considerations of theoretical unity suggest that if some intentions are structurally contrastive then all intentions are structurally contrastive.

7. Conclusion

I have developed and defended a contrastivist view of intentions according to which intention is a ternary relation between an agent, an action, and an alternative. I have presented a number of considerations in favour of such a contrastivist view, some of which have featured more prominently than others, but all of which are worth noting:

1. Contrastivism solves the puzzle generated by side effect cases, and the principles HOLISTIC DECISION, and DECISION INTENTION.
2. Contrastivism allows us to mount a defence of a plausible version of the Simple View against a number of objections.
3. Contrastivism explains the appropriateness of 'I intended to ϕ rather than ψ ' in cases where 'and-not' readings seemed inappropriate.

4. In §4 and §5 we saw that contrastive intentions are able to play important roles for limited planning agents.

5. As discussed in several footnotes, contrastivism combines fits in nicely with a promising explanation of the Knobe effect. Moreover, it explains why some subjects continue to attribute intentions in Knobe effect cases.

The contrastivist view I have developed has a number of important features. Firstly, it allows that an agent can coherently intend to ϕ relative to one alternative, whilst not intending to ϕ relative to other known alternatives. Secondly, agglomeration is not a norm for contrastive intentions. One can coherently intend to ϕ rather than ψ , and to χ rather than ψ , without intending to $(\phi \& \chi)$ rather than ψ . Thirdly, it allowed us to see a way in which the satisfaction of one intention can render another intention inert without the second intention being either satisfied or frustrated.

The potential implications of contrastivism go far beyond those discussed here. And many of the implications I do discuss deserve more attention. However, this is already a long paper, so this further discussion will have to wait.

Bibliography

Anscombe, Elizabeth. 1957. *Intention*. Harvard University Press.

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 (3). 375-405.

Bratman, Michael. 1987. *Intentions, Plans, and Practical Reason*. Cambridge, Mas: MIT Press.

Bronner, Ben. Goldstein, Simon. 2018. A Stronger Doctrine of Double Effect. *Australasian Journal of Philosophy* 96 (4). 793-805.

Davidson, Donald. 1980. Intending. From *Essays on Action and Events*. Oxford: Oxford University Press, 83-102.

Elga, Adam., Rayo, Agustín. Unpublished Manuscript. Fragmentation and Information Access.

Ferrero, Lucía. 2009. Conditional Intentions. *Noûs*, 43 (4). 700-741.

Frankfurt, Harry. 1988. *The Importance of What We Care About*. Cambridge: Cambridge University Press.

Friedman, Jane. 2013. Question Directed Attitudes. *Philosophical Perspectives* 27. 145-174.

Goetz, Stewart. 1995. The Choice-Intention Principle. *American Philosophical Quarterly* 32 (2). 177-185.

Grzankowski, Alex. 2016. Attitudes Towards Objects. *Noûs* 50 (2). 314-328.

Grzankowski, Alex. 2018. Non-Propositional Contents and How to Find Them. *Journal of Consciousness Studies* 25 (3-4). 233-241.

Harman, Gilbert. 1986. *Change in View*. Cambridge, Mas: MIT Press.

Hedden, Brian. 2012. Options and the Subjective Ought. *Philosophical Studies*, 158 (2). 343-360.

Hieronymi, Pamela. 2005. The Wrong Kind of Reasons. *Journal of Philosophy* 102. 37-57.

Hoek, Danniël. Forthcoming. Minimal Rationality and the Web of Questions. In Dirk Kindermann, Peter van Elswyk, and Andy Egan (eds.) *Unstructured Content: Foundations and Applications of Possible Worlds Contents*. Oxford University Press.

- Holton, Richard. 1999. Intention and Weakness of the Will. *The Journal of Philosophy* 96 (5). 241-262.
- 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 (4). 507-535.
- Holton, Richard. 2008. Partial Belief, Partial Intention. *Mind*. 117 (465). 27-58.
- Holton, Richard. 2009. *Willing, Wanting, Waiting*. Oxford: Oxford University Press.
- Kane, Robert. 1998. *The Significance of Free Will*. Oxford: Oxford University Press.
- Kaufman, Arnold. 1966. Practical decision. *Mind* 75: 25-44.
- Knobe, Joshua. 2003. Intentional Action and Side Effects in ordinary Action. *Analysis*. 63 (3). 190-194.
- Knobe, Joshua. 2004. Intention, Intentional Action, and Moral Considerations. *Analysis*, 64. 181-187.
- Koon, Justis. Forthcoming. Options must be External. *Philosophical Studies*. 1-15.
- Marušić, Berislav., and Schwenkler, John. 2018. Intending is Believing: A Defense of Strong Cognitivism. *Analytic Philosophy* 59 (3). 309-340.
- 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. 269-278.
- Mele, Alfred. 2003. *Motivation and Agency*. Oxford: Oxford University Press.
- Nadelhoffer, Thomas. 2006. On Trying to Save the Simple View. *Mind and Language* 21 (5). 565-586.
- Paul, Sarah. 2012. How do we know what we intend? *Philosophical Studies*, 161. 327-346.
- Pink, Thomas. 1996. *The Psychology of Freedom*. Cambridge: Cambridge University Press.
- Rozeboom, Grant. 2015. Side Effects and the Structure of Deliberation. *Journal of Ethics and Social Philosophy* 9 (2). 1-19.

Scaife, Robin. Webber, Jonathan. 2013. Intentional side-effects of action. *Journal of Moral Philosophy*. 179-203.

Schaffer, Jonathan. 2007. Knowing the Answer. *Philosophy and Phenomenological Research*. 75 (2). 383-403.

Sennet, Adam. 2016. Ambiguity. *Stanford Encyclopaedia of Philosophy*.

<https://plato.stanford.edu/entries/ambiguity/>

Setiya, Kieran. 2008. Practical Knowledge. *Ethics* 118. 388-409.

Setiya, Kieran. 2009/2018. Intention. *Stanford Encyclopedia of philosophy*.

<https://plato.stanford.edu/entries/intention/>

Shah, Nishi. 2008. How Action Governs Intention. *Philosopher's Imprint* 8 (5). 1-19.

Snedegar, Justin. 2013. Reason Claims and Contrastivism about Reasons. *Philosophical Studies* 166 (2). 231-242.

Snedegar, Justin. 2017. *Contrastive Reasons*. Oxford: Oxford University Press.

Velleman, D. 1989. *Practical Reflection*. Princeton University Press.

Worsnip, Alex. 2015. Possibly False Knowledge. *Journal of Philosophy*, 112 (5). 225-246.

Yalcin, Seth. 2018. Belief as Question Sensitive. *Philosophy and Phenomenological Research*, 97 (1). 23-47.