

The logic of intention reports

Thomas Grano

Indiana University

ms. date: August 9, 2016

Please direct comments to: tgrano@indiana.edu

Abstract:

Unlike belief and desire reports, intention reports (e.g., *John intends to leave soon*) are not well studied in formal semantics. This paper aims to begin to fill this gap, focusing on empirical similarities and differences that intention reports bear in relation to other attitude reports and to other expressions that involve intentional action. These empirical properties are shown to follow from the view that an intention report *a intends p* denotes true iff *a* has a maximally ranked ACTION-RELEVANT or EFFECTIVE PREFERENCE (in the sense of Condoravdi and Lauer 2016) that *a* bears the RESPONSIBILITY relation (in the sense of Farkas 1988) to *p*. Also discussed are intention reports that do not involve syntactic control (e.g., *John intends for Bill to leave soon*), which in some previous literature have been argued to involve coercion. The proposed semantics for *intend* handles such sentences straightforwardly without coercion.

Contents

1	Introduction	3
2	The empirical landscape	5
2.1	Realism	6
2.2	Conjunction Introduction	7
2.3	Monotonicity	8
2.4	Non-gradability	10
2.5	Mood choice	11
2.6	Anankastic conditionals	12
2.7	Responsibility	13
2.7.1	The causation effect	13
2.7.2	Interaction with <i>(un)intentionally</i>	17
2.7.3	Uncontrollable outcomes	17
2.7.4	Conditions of satisfaction	18
3	Analysis part 1: <i>intend</i> as an effective preference predicate	19
3.1	Background on Condoravdi and Lauer 2016	19
3.2	Back to <i>intend</i>	22
3.2.1	Realism	22
3.2.2	Conjunction Introduction	23
3.2.3	Monotonicity	25
3.2.4	Non-gradability	26
3.2.5	Mood choice	28
3.2.6	Anankastic conditionals	29
3.2.7	Responsibility	30
4	Analysis part 2: <i>intend</i> and the RESP-relation	31
4.1	Background on Farkas 1988	31
4.2	Back to <i>intend</i>	33
4.3	More on the RESP-relation	35
4.3.1	Probing the semantic type of the RESP-relation	36
4.3.2	Probing the content of the RESP-relation	37
5	Compositionality: To coerce or not to coerce?	38
6	Comparison to previous approaches	41
6.1	Grano 2015	41
6.2	Pearson 2016	44
7	Conclusion	46

1 Introduction

The context for this paper is set by two empirical areas that have been central to many investigations in formal semantics. The first area is mental attitude reports, including especially belief reports like (1a) and desire reports like (1b).¹ The second area is a family of morphemes and structures that encode in some way the idea of intentional action, and these include inherently intentional verbs (2a), agent-oriented adverbs (2b), rationale clauses (2c), and imperatives (2d).²

- (1) a. John thinks it's raining.
b. John wants to go home.
- (2) a. John murdered Bill.
b. John intentionally broke the window.
c. John broke the window in order to get inside.
d. Break the window!

At the intersection of these two areas lie intention reports like (3). Speaking informally, intention reports are used to describe a kind of internal commitment or resolve on the part of the attitude holder to carry out some action.³

- (3) John intends to leave soon.

But despite the wealth of literature surrounding mental attitudes on the one hand and intentional action on the other hand, the formal semantics literature is rather sparse when it comes to intention reports. A broad goal of this paper is to begin filling this gap, in a way that elucidates the empirical similarities and differences that intention reports bear in relation both to other mental attitude reports and to other phenomena that involve intentional action.⁴

¹The literature on mental attitude reports is much too vast to do justice to here, but the particular corner of that literature most relevant to this paper is the one that focuses on the nuances of desire reports. Important work in this vein includes Stalnaker 1984; Asher 1987; Heim 1992; Giannakidou 1999; von Stechow 1999; Levinson 2003; Villalta 2008; Crnič 2011; Lassiter 2011; Rubinsteyn 2012; Anand and Hacquard 2013; Condoravdi and Lauer 2016.

²See, for example, Dowty 1972; Zwicky and Sadock 1975; Kamp 1999–2007; Martin and Schäfer 2014 on inherently intentional verbs, Egré 2014 on *intentional(ly)*, and Farkas 1988 on rationale clauses. Imperatives, like mental attitudes, have inspired a literature much too vast to survey here, but see Zanuttini, Pak, and Portner 2012:1260–1263 for an illuminating overview of some of that literature.

³Looking to the philosophy literature, see Bratman 1987 for a characterization of intention in which commitment plays an important role. See also Setiya 2014 for a useful overview of the philosophical literature on intention.

⁴The few previous treatments of *intend* in the semantics literature that I am aware of are: a recent proceedings paper (Grano 2015), as well as brief remarks made in passing by Heim (1992:199–

The central thesis of this paper is the marriage of two ideas, one drawn from the mental attitude literature and one drawn from the intentional action literature. On the mental attitude side, I follow Condoravdi and Lauer’s (2016) proposal that in modeling an attitude holder’s preferences, some linguistic phenomena are sensitive to a privileged subclass of preferences called ACTION-RELEVANT or EFFECTIVE PREFERENCES, defined as preferences that an agent uses to guide his or her actions. As a subclass of preferences, they bear certain similarities to ordinary preferences such as those routinely expressed via desire reports using *want*. But since they are used to guide action, they are also subject to certain rationality constraints that make them similar in other ways to beliefs. On the intentional action side, I follow Farkas’s (1988) proposal that some linguistic phenomena invoke the RESP(ONSIBILITY)-relation. On the version of this relation pursued here, RESP is true of an individual *a* and a proposition *p* iff *a* intentionally brings it about that *p*.

Putting these two ideas together sets the stage for:

(4) **Central thesis of this paper:**

$\llbracket a \text{ intends } p \rrbracket^w = 1$ iff $\text{RESP}(a,p) \in \max[\text{Effective-Preference}(a,w)]$

(A sentence of the form *a intends p* denotes true iff the proposition that *a* stands in the RESP-relation with *p* is a member of *a*’s maximally ranked effective preferences in the evaluation world.)

Once properly unpacked, (4) makes sense of a host of empirical properties associated with intention reports to a greater degree of accuracy than any other approach I am aware of. The bulk of this paper is concerned with laying out these empirical properties, unpacking (4), and showing how the latter explains the former.

Against the backdrop of (4), a secondary goal of this paper is to engage with a corner of the literature in which intention reports have interacted with the study of syntactic control: a number of linguists including Perlmutter (1968); Jackendoff (1996); Jackendoff and Culicover (2003); Culicover and Jackendoff (2005); Grano (2015) have observed that when *intend* combines with a complement that is not syntactically controlled, as in (5a), the resulting sentence is intuitively paraphrasable as a minimally different sentence in which the complement to *intend* is embedded under a causative predicate syntactically controlled by the attitude holder, as in (5b).

- (5) a. John intends for Bill to leave soon.
 b. \approx John intends to bring it about that Bill leave soon.

200); Ninan (2005:16–17); Crnič (2011:83–84); Pearson (2016:734); and Condoravdi and Lauer (2016:23–27). See also Jackendoff 1996 and remarks by Jackendoff and Culicover (2003:537) for an approach to *intend* couched in conceptual semantics. These works will be cited below as appropriate; some of Condoravdi and Lauer’s (2016) proposals in particular will play a rather important role in the analysis to be presented.

The authors cited above take this paraphrase relation to suggest that (5a) involves coercion (or something close to it: Perlmutter argues for a silent causative predicate). I will show instead that this paper's central thesis makes sense of the observed paraphrase relation in a way that does not rely on coercion, thereby leading to a kind of conceptual argument against coercion: coercion is superfluous once we properly understand the meaning of *intend*. And I will also present a possible empirical argument against the coercion analysis based on an observation about how intention reports interact with VP ellipsis. The argumentation deals with an issue central to argument structure and thematic relations concerning how and to what extent intentional agency is encoded into the semantics of ordinary event descriptions. If the reasoning is on the right track, it implies that sentences like *John broke the window* are underspecified with respect to whether the action was undertaken intentionally rather than being ambiguous between an intentional and a non-intentional reading.

The organization of the rest of this paper is as follows. Section 2 surveys the empirical landscape of intention reports, highlighting their similarities and differences with respect to belief and desire reports. Section 3 provides background on Condoravdi and Lauer's (2016) proposal regarding effective preferences and shows how a treatment of *intend* as an effective preference predicate makes sense of a number of its key properties. Section 4 provides background on Farkas's (1988) RESP-relation and shows how including it in the denotation for *intend* explains the rest of the key properties of *intend*. Section 5 discusses intention reports that do not involve syntactic control and argues against the view that such sentences involve coercion. Section 6 compares this paper's analysis of intention reports to two recent alternatives. Finally, section 7 concludes.

2 The empirical landscape

In this section, I lay out seven key semantic properties of intention reports. The first four properties (realism, conjunction introduction, monotonicity, and non-gradability) all highlight ways in which *intend* behaves like *believe* and unlike *want*, while the next two properties (mood choice and anankastic conditionals) show that in other ways, *intend* behaves like *want* and unlike *believe*. Finally, the last property (responsibility) highlights a way in which *intend* behaves unlike both *believe* and *want*. Taken together, these seven properties will constitute the empirical foundation for the analysis of intention reports to be presented in sections 3–4 below.

Before we begin, a note on the scope of the investigation: I limit myself throughout this paper to the verb *intend*. But it bears noting that *intend* is closely related to at least three classes of verbs. First, there are verbs that, like *intend*, involve a kind of internal or private commitment to take action, such as those in

(6a). Then there are those verbs that involve a kind of public or externally-directed commitment to take action, such as those in (6b). (See also Grano 2015, and the discussion thereof in section 6.1 below, for more on the relationship between private and public commitments.) Finally, there are those verbs that have to do with one agent bringing it about or trying to bring it about that another agent come to have an obligation and/or intention to act, as in (6c). Along with desire predicates like *want*, these constitute the core classes of predicates that typically instantiate syntactic control. Consequently, much of the previous work on these verbs is found within the control literature, especially the literature on controller choice: see especially Farkas 1988; Sag and Pollard 1991; Jackendoff and Culicover 2003; Rooryck 2008. (See also Sharvit 2003; Grano 2011 for formal semantic approaches to *try* and Condoravdi and Lauer 2011 on verbs like *promise* and *order*.)

- (6) a. verbs of private commitment: aim, choose, decide, endeavor, intend, plan, try
- b. verbs of public commitment: agree, offer, pledge, promise, swear, threaten
- c. verbs of influence: advise, beg, command, order, persuade, urge

It is a very interesting question — and one that I hope will inspire future research — whether and to what extent these verbs pattern like *intend* with respect to the properties outlined in this section. But for reasons of space and tractability, it is not a question that I will have anything more to say about in this paper.

2.1 Realism

The first property of intention reports to be considered here is what I will call REALISM: ordinarily, an intention report comes along with the expectation that the attitude holder believes the named outcome to be within the realm of possibility. This is evidenced by examples like those in (7), which at worst sound contradictory and at best portray John as being irrational.

- (7) a. #John intends to fly to the moon, even though he knows this is impossible.
- b. #John intends to turn into a unicorn, even though he knows this is impossible.

Desire reports make for a useful comparison here. Heim (1992) proposed that desire reports also carry such a presupposition. But as Heim herself observed, this proposal is challenged by examples like (8), which seem perfectly felicitous.

- (8) I want this weekend to last forever. (But I know, of course, that it will be over in a few hours.) (Heim 1992:199)

Consistent with this observation, there is a fairly clear contrast between (7) and the minimal variants in (9) that substitute in *want* in place of *intend*. The examples in (9) are perhaps a bit awkward, especially when judged in comparison with the explicitly counterfactual and impeccably felicitous formulations in (10). But they still do not descend to the level of infelicity associated with (7).⁵

- (9) a. John wants to fly to the moon, even though he knows this is impossible.
b. John wants to turn into a unicorn, even though he knows this is impossible.
- (10) a. John wishes he could fly to the moon, even though he knows this is impossible.
b. John wishes he could turn into a unicorn, even though he knows this is impossible.

2.2 Conjunction Introduction

A second property of intention reports, noted by Condoravdi and Lauer (2016), is that they validate CONJUNCTION INTRODUCTION: if *a* intends *p* and *a* intends *q*, then we would ordinarily expect that if *a* is rational, *a* should also (at least upon reflection) intend $p \wedge q$.

This property is detectable in two ways. First, suppose *p* and *q* contradict each other. Then, if *a* intends *p* and *a* intends *q*, and intention reports obey conjunction introduction, we can conclude that *a* intends $p \wedge q$. But then, given the realism property of intention reports discussed in section 2.1 above, this should mean that *a* believes (irrationally) that a logical contradiction is possible. Consistent with this expectation, we find that conjoined intention reports of the form *a* intends *p* and *a* intends *q* do indeed sound odd — or at the very least portray *a* as irrational — when *p* and *q* contradict each other ((11b) is taken from Condoravdi and Lauer 2016:25):

- (11) a. #John intends to stay at home tonight, but he also intends to go out tonight.
b. #John intends to move in with his girlfriend, but he also intends to keep living alone.

⁵One might wonder: is it strong enough to say that an intention report merely requires that the attitude holder believe in the *possibility* of the outcome, as opposed to belief that the outcome will certainly come about? (Cf. in this connection the discussion in Heim 1992:199.) See section 6.2 below, where I revisit this question in the context of assessing Pearson's (2016) proposal to analyze intention reports as a particular kind of belief report.

The second way in which conjunction introduction manifests itself is in cases where p and q do not contradict each other, but where a *intends* p and a *intends* q is paired with a denial of a *intends* $p \wedge q$. As expected, doing so has an odd result ((12) is adapted from Condoravdi and Lauer 2016:26–27, building on a similar example due to Levinson 2003):

- (12) #John intends to visit Paris this summer. And John intends to visit Rome this summer. But he doesn't intend to visit both Paris and Rome this summer.

Consider, by contrast, *want*. As observed by Levinson (2003) and further discussed by Condoravdi and Lauer (2016), *want* does not always validate conjunction introduction. Substituting in *want* for *intend* in the above examples results in perfectly felicitous sentences. In (13), John is not portrayed as being irrational but rather simply has desires pulling him in conflicting directions. Similarly, as Condoravdi and Lauer (2016) note, (14) is felicitous in a scenario where (for example) John independently considers both Paris and Rome desirable travel destinations but prefers not to expend the time or money that would be needed to visit both places.

- (13) a. John wants to stay at home tonight, but he also wants to go out tonight.
b. John wants to move in with his girlfriend, but he also wants to keep living alone.
- (14) John wants to visit Paris this summer. And John wants to visit Rome this summer. But he doesn't want to visit both Paris and Rome this summer.

Finally, to round out the picture, we observe that *believe* patterns together with *intend* to the exclusion of *want* in validating conjunction introduction. The sentences in (15)–(16) are at worst contradictory and at best portray John as being irrational.

- (15) a. #John believes he'll stay at home tonight, but he also believes he'll go out tonight.
b. #John believes he'll move in with his girlfriend, but he also believes he'll keep living alone.
- (16) #John believes he'll visit Paris this summer. And John also believes he'll visit Rome this summer. But he doesn't believe he'll visit both Paris and Rome this summer.

2.3 Monotonicity

The third property of intention reports to be considered is (UPWARD) MONOTONICITY: if a *intends* p , then we would ordinarily expect that if a is rational, a should

also, at least upon reflection, intend logical consequences of p . Consistent with this expectation, we see in examples like (17) that if p entails q and we pair a denial of a *intends* q with an assertion of a *intends* p , the result is at worst contradictory and at best portrays a as being irrational. ((17a) and (17b) are fashioned after Heim 1992 and Anand and Hacquard 2013, respectively, based on example sentences that these authors construct to probe the monotonicity of *want*.)

- (17) a. #John doesn't intend to teach next semester, but he intends to teach Tuesdays and Thursdays next semester.
 b. #John doesn't intend to die, but he intends to die quickly.
 c. #John doesn't intend to use his voice, but he intends to sing.

Again, a comparison to *want* is instructive. It is a matter of some controversy whether a *wants* p is upward-monotonic on p . See, among others, Stalnaker 1984; Asher 1987; Heim 1992; Giannakidou 1999; von Stechow 1999; Levinson 2003; Villalta 2008; Crnić 2011; Anand and Hacquard 2013; Condoravdi and Lauer 2016. At this stage, I do not want to take a stance on this controversy, but rather wish to use some facts that bear on the controversy to illustrate the contrast between *intend* and *want*. As we see in (18), when we substitute in *want* for *intend* in the examples above, the sentences become perfectly felicitous: they are neither contradictory nor do they portray John as being irrational.

- (18) a. John doesn't want to teach next semester, (but given that he has to,) he wants to teach Tuesdays and Thursdays next semester.
 b. John doesn't want to die, (but given that he has to,) he wants to die quickly.
 c. John doesn't want to use his voice, (but given that he has to,) he wants to sing.

In the examples in (18), I infix *... but given that he has to...* between the relevant pieces because this helps reinforce the felicity of such sentences, a fact which is no doubt an important clue into how these sentences work. Crucially, this infixation produces no comparable improvement in the felicity of intention reports:

- (19) a. #John doesn't intend to teach next semester, (but given that he has to,) he intends to teach Tuesdays and Thursdays next semester.
 b. #John doesn't intend to die, (but given that he has to,) he intends to die quickly.
 c. #John doesn't intend to use his voice, (but given that he has to,) he intends to sing.

Finally, to round out the picture, we note that *believe* again here patterns like *intend* and unlike *want* in that it displays upward monotonicity:

- (20) a. #John doesn't believe he'll teach next semester, but (given that he has to,) he believes he'll teach Tuesdays and Thursdays next semester.
 b. #John doesn't believe he'll die, but (given that he has to,) he believes he'll die quickly.
 c. #John doesn't believe he'll use his voice, but (given that he has to,) he wants to sing.

2.4 Non-gradability

As discussed by Villalta (2008); Lassiter (2011); Anand and Hacquard (2013), *want* displays properties characteristic of gradable predicates, such as being able to occur in comparatives (21a) and superlatives (21b) and with degree adverbs (21c). As pointed out by Crnič (2011:83, note 22), however, when we turn to *intend*, we find that it does not display such properties, as seen in (22). In this respect, *intend* behaves like *believe*, which is also uncomfortable in gradable environments, as shown in (23).

- (21) a. John wants to go to Paris more than he wants to go to London.
 b. What John wants the most is to be happy.
 c. John wants very much to leave.
 (22) a. ?John intends to go to Paris more than he intends to go to London.
 b. ?What John intends the most is to be happy.
 c. ?John intends very much to leave.
 (23) a. ?John believes he'll go to Paris more than he believes he'll go to London.
 b. ?What John believes the most is that he'll be happy.
 c. ?John believes very much that he'll leave.

I also note in passing one other phenomenon that has *intend* and *believe* patterning together to the exclusion of *want*. As observed by Crnič (2011) and illustrated in (24) (based on Crnič 2011:81), *want* licenses what Crnič calls “weak *even*” in its complement, where *even* is “weak” in Crnič’s sense just in case the minimal clause containing it “denotes a proposition that is most likely (least noteworthy) among its alternatives” (Crnič 2011:11). By contrast, *intend* and *believe* do not license weak *even*.

- (24) a. Mary wants to find even ONE party that cares for the people.

- b. #Mary intends to find even ONE party that cares for the people.
- c. #Mary believes she'll find even ONE party that cares for the people.

While a full investigation of weak *even* would take us too far afield, what Crnić (2011) suggests in a nutshell is that the asymmetry in (24) has to do with the fact that *want* has a semantics that involves comparison over alternatives in a way that *intend* and *believe* do not. If this is right, then this is yet another manifestation of the non-gradability of *intend*.

2.5 Mood choice

Whereas all of the properties considered so far involve *intend* patterning together with *believe* to the exclusion of *want*, we turn in this and the next subsection to properties that *intend* and *want* share to the exclusion of *believe*. The first has to do with mood choice. It is well documented that in languages that mark an indicative/subjunctive mood distinction, ‘believe’ usually selects for an indicative complement (with some important exceptions) whereas ‘want’ selects for a subjunctive complement. Which way does ‘intend’ pattern? A difficulty in investigating this question is that ‘intend’ usually combines with a control complement (see section 5 below), which means for many languages that the complement takes the form of an infinitive that does not mark an indicative/subjunctive contrast. But when we turn to Greek, which marks mood and lacks nonfinite complementation, the observation is that *skopevo* ‘intend’/‘plan’ patterns with *thelo* ‘want’ in combining with complements introduced by the subjunctive mood particle *na*, whereas *pistevo* ‘believe’ ordinarily combines with complements introduced by the indicative complementizer *oti*. (See e.g. Giannakidou 2009 and references therein; the following examples are adapted from Spyropoulos 2008.)

- (25) Skopevo na ertho.
intend.1SG SUBJ come.1SG
‘I intend to come.’ (Spyropoulos 2008:166)
- (26) Thelo na ertho.
want.1SG SUBJ come.1SG
‘I want to come.’ (adapted from Spyropoulos 2008:166)
- (27) Pistevo oti tha ertho.
believe.1SG that will come.1SG
‘I believe that I will come.’ (adapted from Spyropoulos 2008:162)

2.6 Anankastic conditionals

A second way in which *intend* patterns like *want* and unlike *believe* is with respect to anankastic conditionals. Anankastic conditionals were first brought to the attention of linguists by Sæbø (1985, 2001) and have since been investigated by many formal semanticists such as von Stechow and Iatridou (2005); Huitink (2005, 2008); Nissenbaum (2005); von Stechow, Krasikova, and Penka (2006); Condoravdi and Lauer (2016). Informally speaking, an anankastic conditional is a sentence of the general form in (28) that expresses a necessary condition q for achieving a specified goal or desire p .

(28) If a [appropriate embedding predicate] p , [necessity modal] q .

For example, a salient interpretation of (29a) is as indicated by the paraphrase in (29b). This is the anankastic interpretation.

- (29) a. If you want to get good grades, you must study.
b. \approx You must study in order to get good grades.

Notably, not all conditionals that match the pattern in (28) are ordinarily read as anankastic conditionals. For example, the most natural interpretation of (30a) is NOT one that can be paraphrased by (30b).

- (30) a. If you want to get bad grades, you must reexamine your priorities.
b. $\not\approx$ You must reexamine your priorities in order to get bad grades.

What matters for our purposes is not the full range of issues raised by these sentences but rather just the observation due to Sæbø that the embedding predicate in the antecedent of an anankastic conditional need not be *want* but can also be a verb like *intend*, *hope*, *plan*, or *aim*. Sæbø suggests that what these predicates share is that they name an “intentional attitude”. Consistent with this generalization is the observation that *intend* indeed patterns like *want* in supporting anankastic conditionals, as shown in (31), whereas *believe* does not, as shown in (32).

- (31) a. If you intend to get good grades, you must study.
b. \approx You must study in order to get good grades.
(32) a. If you believe you’ll get good grades, you must study.
b. $\not\approx$ You must study in order to get good grades.

Taken together with the mood choice facts discussed above, this means that we have now identified two ways in which *intend* and *want* pattern together to the exclusion of *believe*.

2.7 Responsibility

In this final subsection of our empirical overview, I document a number of phenomena whereby *intend* patterns one way and *believe* and *want* pattern another way. In anticipation of the analysis to be presented later, let me say right at the outset that I believe all of these phenomena have the same underlying source. And the facts to be presented are perhaps more easily digested if they are considered with this underlying source in mind. So, let me now attempt to preview the analysis in a semi-informal way: in calculating the truth conditions of an intention report of the form *a intends p*, we consider only that subset of *p*-worlds in which *a* intentionally brings it about that *p*. Other *p*-worlds are ignored. By contrast, in calculating the truth conditions of a belief or desire report of the form *a believes/wants p*, we consider all *p*-worlds. I borrow Farkas's (1988) term RESPONSIBILITY to pick out this property of intention reports. All of the facts to be presented in this subsection can be read as trying to substantiate this core idea.

2.7.1 The causation effect

It has been observed before both in the linguistics literature (Perlmutter 1968; Jackendoff 1996; Jackendoff and Culicover 2003; Culicover and Jackendoff 2005; Grano 2015) and in the philosophy literature (Brand 1984; Ludwig 2007) that when *intend* combines with a complement that is not syntactically controlled, as in (33a), the resulting sentence is intuitively paraphrasable as a minimally different sentence in which the complement to *intend* is embedded under a causative predicate controlled by the attitude holder, as in (33b).

- (33) a. John intends (for) Bill to leave.
b. \approx John intends to bring it about that Bill leave.

Notably, no such paraphrase relation holds for belief or desire reports: (34a) cannot be sensibly paraphrased as (34b) nor can (35a) be sensibly paraphrased as (35b) (a point made by Jackendoff 1996). I will call this special property of intention reports the CAUSATION EFFECT.

- (34) a. John believes Bill will leave.
b. $\not\approx$ John₁ believes he₁'ll bring it about that Bill leave.
- (35) a. John wants Bill to leave.
b. $\not\approx$ John wants to bring it about that Bill leave.

Although I believe the causation effect is quite intuitive, it has not gone unchallenged, and it will be instructive to work through a potential objection that

will lead to an important qualification. In particular, the robustness of the causation effect has been called into question by Boeckx, Hornstein, and Nunes (2010) (cf. also Vermazen 1993). Responding to Culicover and Jackendoff's (2005) claim that (36a) can be paraphrased approximately as (36b), Boeckx et al. say that "it is not clear to us that [(37)] is a contradiction, which it should be if [(36a)] *meant* [(36b)]" (p. 233).

- (36) a. Hilary intended for Ben to come to the party.
- b. \approx Hilary intended to bring it about that Ben come to the party.
- (37) Hilary intended for Ben to come to the party, though being lazy and complacent, she intended to do nothing whatsoever to bring this about.

I agree with Boeckx et al. that (37) is not clearly contradictory. But does this really tell against the putative synonymy of (36a) and (36b)? (36b) asserts that Hilary intends to bring something about. By contrast, what the relevant portion of (37) denies is not that Hilary intends to bring something about but rather that Hilary intends to *do something* to bring something about. So what we have to ask is whether bringing something about entails doing something. And linguistic intuition tells us that it does not, at least not necessarily. For example, consider the scenario described in (38). Here, there is a sense of *do* and a sense of *bring about* that render sentences like (38a–b) true and non-contradictory in the relevant context. So the facts are consistent with the view that *bring about* involve a rather weak kind of causation that need not involve action in the everyday sense. And the fact that this same scenario cannot be accurately described monoclausally with causative *drown*, as seen in (38c), suggests that this weak notion may be INDIRECT CAUSATION in the sense of Wolff (2003).

- (38) CONTEXT: Kim and Sandy are on a boat. Sandy accidentally falls overboard and Kim, although perfectly capable of rescuing her, chooses not to.
 - a. (By not doing anything,) Kim brought it about that Sandy drowned.
 - b. (Kim did not do anything, and) Kim (thereby) brought it about that Sandy drowned.
 - c. #Kim drowned Sandy.

Returning to the Hilary/Ben scenario, then, the suggestion is that (37) is compatible with a scenario in which Hilary intends to bring it about that Ben come to the party, insofar as she (a) thinks Ben will come, (b) thinks that she could stop him if she wanted to, but (c) is committed to not taking any such obstructive measures. On this view, it is this commitment not to act that the truth of the intention

report hangs on. (In fact, the philosophy literature has a special name for such intentions: NEGATIVE INTENTIONS. See e.g. Harman 1976.)

In light of these considerations, a better way of probing the synonymy (36a) and (36b) is to try pairing (36a) with the negation of (36b), as in (39). Whereas (37) is not clearly contradictory, (39) is somewhat more difficult to judge — so difficult, in fact, that I am reluctant to hang any conclusion on it.

(39) ?#Hilary intended for Ben to come to the party, though being lazy and complacent, she did not intend to bring it about that Ben come to the party.

Given the subtleties involved in the Hilary/Ben scenario, it will be useful to turn to a clearer (and as far as I know, fresh) challenge to the robustness of the causation effect, namely the existence of naturally occurring (and to my ear, fairly natural sounding) examples like those in (40). Sentences like these do not invite the causative paraphrase, as indicated in (41).

- (40) a. I didn't intend for it to rain while we were planting today.
(http://www.mlive.com/news/jackson/index.ssf/2014/04/in_the_classroom_da_vinci_inst.1.html; retrieved 7/29/16)
- b. Although Rizwan & Darsana didn't intend for it to rain on their wedding day, they did intend on having an amazing party and that is just what they did!
(<http://undercoverlive.com/category/weddings/orange-county-weddings/page/3>; retrieved 7/29/16)
- (41) a. (40a) $\not\approx$ I didn't intend to bring it about that it rain while we were planting today.
- b. relevant portion of (40b) $\not\approx$ Rizwan and Darsana didn't intend to bring it about that it rain on their wedding day.

There are at least three tacks one could take in responding to the challenge posed by examples like this. The first would be to give up on the idea that intention reports encode responsibility in the way I claim it does. But this would sit awkwardly both with the other sources of evidence for the responsibility effect to be considered in the remainder of this section, and with the fact that most non-control intention reports do invite the causative paraphrase.

The second tack would be to maintain the idea that intention reports encode responsibility, but to weaken or bleach the notion of responsibility (a “generalize to the worst case” strategy) so that it does not explicitly encode causation, hence suppressing the prediction that intention reports should always be amenable to causative paraphrases. See in this connection the semantics for *intentional* proposed by Egré (2014), where the basic proposal is that an outcome counts as intentional iff the agent sufficiently desires the outcome and sufficiently foresees how it

will come about. (See also section 4.3.2 below for additional discussion.) We might then say that one particularly salient way in which one foresees how an outcome will come about is by causing it: deliberately taking actions so as to bring the outcome about. But at the fringe are cases where the intentional status of an outcome is verified by pure foresight, stripped of causation, as in the rain examples.

Finally, the third tack we could take is to entertain the idea that responsibility does necessarily involve causation, but that under some conditions, the verb *intend* is used in a separate but related sense whereby it does not encode responsibility. Here, I will take this third tack. This reason is twofold. First, it allows us to maintain the strong thesis that responsibility has to do with causation (which seems best in keeping with how Farkas 1988 understood it: see section 4.1 below), and thereby not lose the insight that *intend* (in the usual, responsibility-laden sense) normally sounds odd in the context of weather outcomes just like is the case for other expressions having to do with intentional action to be considered in section 4.1 below such as imperatives (42b) and rationale clauses (42c).

- (42) a. #John intended for it to rain.
b. #Make it rain!
c. #It rained in order to water the plants.

The other reason for taking the third tack relates to the observation that the challenging examples can all be comfortably paraphrased using *count on*, as in (43).

- (43) a. Hilary counted on Ben coming to the party (, though being lazy and complacent, she did not intend to bring it about that Ben come to the party).
b. I didn't count on it raining while we were planting today.
c. Rizwan and Darsana didn't count on it raining on their wedding day.

As will be discussed in a bit more detail in the conclusion in section 7 below, *count on* is quite similar to *intend* in that it names a preference that the attitude holder thinks has a good chance of being realized and that the attitude holder is using to guide his or her actions. But unlike *intend*, it is responsibility-free; i.e., it strips away any notion that the satisfaction conditions depend on the attitude holder achieving the outcome him/herself. (See section 7 below for the actual linguistic diagnostics confirming these claims: here I just sketch the basic intuition.) This invites the hypothesis that although *intend* ordinarily encodes responsibility, there are some contexts in which responsibility is stripped away and it takes on a meaning like *count on*. This seems plausible, given that such slippage in the relationship between an attitude lexeme and a particular attitude expressed is reminiscent of the way *want* sometimes takes on a meaning like *intend* (see section 3 below). It is also reminis-

cent of the cross-linguistically attested existence of belief/desire-hybrid predicates: see Nguyen 2013 on Mandarin *xiǎng* and Bogal-Allbritten 2016 on Navajo *nízin*.

In summary, none of the three tacks outlined here are particularly attractive because they all entail a weakening of the theory. But the facts compel us to choose one. And tack three has the relative virtue of making the challenge a familiar and inevitable one: it comes as no surprise that words can have multiple related senses. That being said, a full investigation of the factors conditioning the variation between the senses is something that will unfortunately have to await further research.

2.7.2 *Interaction with (un)intentionally*

A second way in which *intend* parts company from *believe* and *want* has to do with the effect of inserting the adverbs *unintentionally* and *intentionally* into the complement of the attitude verb. Turning first to *unintentionally*, the observation is that with *believe*, such an insertion is felicitous (as long as the predicate names an action that can in principle be undertaken unintentionally) as in (44a). Similarly for *want*: (44b) reports a coherent albeit somewhat unusual desire. (44c), by contrast, requires much more work to make sense of and when decontextualized sounds quite odd. (Though see section 4.2 below for a suggestion about a context that might make (44c) true.)

- (44) a. John believes he'll break the window unintentionally.
b. John wants to break the window unintentionally.
c. #John intends to break the window unintentionally.

An analogous effect is witnessed with *intentionally*. With a belief or desire report, the insertion of *intentionally* into the complement of the attitude verb is a coherent and informative addition to the content of the attitude being expressed, as seen in (45a–b). By contrast, (45c) sounds redundant.

- (45) a. John believes he'll break the window intentionally.
b. John wants to break the window intentionally.
c. #John intends to break the window intentionally.

2.7.3 *Uncontrollable outcomes*

The next effect to be considered has to do with what happens when we embed under the attitude verb a predicate that names an outcome which is ordinarily not under the control of the individual named by the subject. The observation is that in the case of *intend*, something special has to happen in order for the resulting sentences to be felicitous. Consider for example the sentences in (46), which are awkward

insofar as they require that John believe that it is within the realm of possibility that he has control over the named outcomes (cf. also Heim 1992:199–200; Jackendoff 1996; Jackendoff and Culicover 2003:542 for relevant discussion).

- (46) a. #John intends to go through puberty next year.
- b. #John intends to snore while he is asleep tonight.
- c. #John intends to resemble his father.
- d. #John intends to be tall.

Witness in this connection the corresponding infelicity of the variants in (47), which provide additional linguistic support for the idea that the embedded predicates in (46) name outcomes that cannot ordinarily be construed as being brought about intentionally by the individual named by the subject. If anything, the sentences in (47) are even more infelicitous than those in (46), because instead of just requiring us to entertain the idea that John *thinks* he might be able to have control over the outcome, these sentences are used to assert that John actually did have control over the outcome.

- (47) a. #John intentionally went through puberty last year.
- b. #John intentionally snored while he was asleep last night.
- c. #John intentionally resembled his father.
- d. #John was intentionally tall.

By contrast, when we turn to *want* and *believe*, there is no corresponding infelicity, as seen in (48)–(49). To make sense of (48)–(49), we need not think that John believes himself to have control over the named outcomes.

- (48) a. John wants to go through puberty next year.
- b. John wants to snore while he is asleep tonight.
- c. John wants to resemble his father.
- d. John wants to be tall.
- (49) a. John believes he'll go through puberty next year.
- b. John believes he'll snore while he is asleep tonight.
- c. John believes he'll resemble his father.
- d. John believes he'll be tall.

2.7.4 *Conditions of satisfaction*

Finally, yet a fourth way of approaching the special nature of *intend* comes from the work of Searle (1983). For Searle, attitudes have conditions of satisfaction: a belief is either true or false, a desire is either fulfilled or unfulfilled, and an intention is

either carried out or not carried out. What determines, for a given attitude, its conditions of satisfaction? According to Searle, this is rather transparent in the case of belief and desire: a belief that p is satisfied iff p is true and a desire that p is satisfied iff p is true.⁶ But crucially, according to Searle, the conditions of satisfaction for intention are more complicated. For an intention that p to be satisfied, it must be the case not only that p is true but more specifically that p is true in virtue of standing in an appropriate causal relation with the very intention it represents.⁷

For example, suppose John intends to break the window. Before he has a chance to carry out his intention, he is outside throwing a ball around. He accidentally throws the ball at the window and breaks it. Our linguistic intuition tells us that in this scenario, it is not the case that John broke the window intentionally. So this shows us that it is possible for someone to intend to do something and to do it, but not intentionally. And according to Searle, it is false in this scenario to say that John's intention was satisfied, because the intention was not an antecedent in an appropriate kind of causal chain to the outcome of breaking the window. To the extent that it is possible to have intuitions about conditions of satisfaction, this seems right. By contrast, if John wants to break the window or believes he'll break the window, then his breaking the window satisfies the attitude regardless of how his breaking the window comes about (though see note 6).

3 Analysis part 1: *intend* as an effective preference predicate

3.1 Background on Condoravdi and Lauer 2016

Building on Davis (1984) and Levinson (2003), Condoravdi and Lauer (2016) develop a semantics for *want* that takes as its starting point the view that *want* has two kinds of readings, which Condoravdi and Lauer illustrate using the examples in (50a–b) respectively. Of interest is the fact that (50a) and (50b) are both truthful and appropriate answers in the same situation, despite the fact that they seem to contradict each other.

- (50) Do you want to play tennis?
a. I want to, but I have to teach.

⁶An anonymous reviewer points out that this may be an oversimplification: Kratzer (1998) considers cases where a belief report is construed *de re* with respect to a specific event. In such cases, a belief that p can fail to be satisfied even if p is true, provided the belief holder is mistaken about the event that verifies p .

⁷For this reason, Searle says that intentions are CAUSALLY SELF-REFERENTIAL. According to Searle (2009), this term was coined by Harman (1976). Another important forebear to the idea is Davidson 1973. For criticism, see Mele 1987.

- b. No [= I don't want to], I have to teach.

Following Levinson (2003), Condoravdi and Lauer suggest that whereas the reading of *want* in (50a) simply reports a desire “as a matter of psychological fact” (Condoravdi and Lauer 2016:22), the reading of *want* in (50b) reports “the kind of desire accompanying intentional action” (Levinson 2003:223). Fleshing out this distinction, Condoravdi and Lauer (2016:22–23) write as follows:

For an agent who has to decide between alternative courses of action, the decision is driven by two factors. On the one hand, he has certain beliefs, including beliefs about which actions are available, and what their consequences are. On the other hand, he has certain preferences for how the world turns out to be, relative to outcomes over which he might have some influence. But not all desires or preferences that the agent has as a matter of psychological fact need to count among the preferences that guide action choice. He might simply fail to take some of his desires into account, or a more important preference might defeat a less important one. We call the preferences that the agent takes into account when choosing actions his *effective preferences*.

To model effective preferences⁸, Condoravdi and Lauer draw on their earlier work (Condoravdi and Lauer 2011, 2012) to propose that, at any given world, an agent has a set of PREFERENCE STRUCTURES corresponding to “the various sources of his preferences” (Condoravdi and Lauer 2016:28), such as “desires, inclinations, personal moral codes, and obligations” (Condoravdi and Lauer 2012:45). As defined in (51), a preference structure is a pair $\langle \mathbf{P}, \prec \rangle$, where \mathbf{P} is a set of propositions (a set of sets of possible worlds) and \prec is a strict partial order on \mathbf{P} .

- (51) Given a set of worlds W , a *preference structure* is a pair $\langle \mathbf{P}, \prec \rangle$, where $\mathbf{P} \subseteq \mathcal{P}(W)$ and \prec is a strict partial order on \mathbf{P} .
(Condoravdi and Lauer 2016:29)

Condoravdi and Lauer propose that among an agent's preference structures is his effective preference structure, which serves to “integrate his various preferences into an over-arching set of preferences that can guide action” (p. 29). They furthermore propose that an agent's effective preference structure is subject to at least two special requirements. The first is consistency, as defined in (52): if any of the propositions in the preference structure are inconsistent with each other relative to the agent's beliefs, then these propositions must be asymmetrically ranked.

⁸Although Condoravdi and Lauer (2016) do not explicitly make this connection, their concept of effective preferences seems to be similar to Hansson's (2001) concept of CHOICE-GUIDING PREFERENCES in his formal philosophical work on values and norms.

The intuitive basis for this requirement is that if an agent has two or more preferences that cannot all be realized together given what the agent believes about the world, then it would be irrational for the agent to use all of these preferences to guide his actions; instead, the agent must rank the conflicting preferences so that the outranked ones do not guide action.

(52) **Consistency**

A preference structure $\langle \mathbf{P}, \prec \rangle$ is *consistent* with respect to an information state B iff for any $X \subseteq \mathbf{P}$, if $B \cap \bigcap X = \emptyset$, there are $p, q \in X$ such that $p \prec q$. (Condoravdi and Lauer 2016:29)

The second requirement that Condoravdi and Lauer propose for effective preference structures is realism, as defined in (53): each of the propositions in the preference structure must be compatible with the agent's beliefs. The intuitive basis for this requirement is that it would be irrational for an agent to let his actions be guided by a preference that would be impossible to realize.

(53) **Realism**

A preference structure $\langle \mathbf{P}, \prec \rangle$ is *realistic*, relative to an information state B , iff for all $p \in \mathbf{P}$: $p \cap B \neq \emptyset$. (Condoravdi and Lauer 2016:29)

Against this backdrop, Condoravdi and Lauer (2016) propose an underspecified semantics for *want* as in (54)–(55) whereby *want* is interpreted relative to a preference structure P . Specifically, on this approach, a *wants* ϕ denotes true relative to a world w and a preference structure P iff $\llbracket \phi \rrbracket$ is a member of the highest ranked subset of P relative to a in w .

(54) $want_P(a, \phi)$ is true in w iff $\llbracket \phi \rrbracket \in \max[P(a, w)]$

(55) $\max[\langle \mathbf{P}, \prec \rangle] := \{p \in \mathbf{P} \mid \neg \exists q \in \mathbf{P}: p \prec q\}$
(Condoravdi and Lauer 2016:30)

Thus, a sentence like *John wants to play tennis* might be true relative to one of John's preference structure (because *John plays tennis* is among the most highly ranked propositions in that preference structure) but false relative to John's effective preference structure (for example, it could be that consistency requires that *John plays tennis* be asymmetrically ordered with respect to *John teaches*, and the latter proposition wins out).

3.2 Back to intend

Condoravdi and Lauer (2016) acknowledge the close relationship between effective preferences and intentions, though they are explicit about not conflating the two. Instead, they suggest (p. 23):

One way to fit intentions into our set-up is to assume that they are a particular kind of effective preference with special properties. Specifically, intentions could be those effective preferences that the agent has decided, for the time being, to maintain indefinitely into the future, until they are realized or consciously reconsidered. This would be close to the conception of intention that Bratman (1987) articulates.

They may well be right about this. But since we have to start somewhere, what I want to do in this section is adopt the working hypothesis that *intend* is just like *want* as analyzed by Condoravdi and Lauer except that it exclusively targets effective preferences.⁹ Then we can see how far this hypothesis goes toward accounting for the observations surrounding *intend* laid out in section 2 above.

Specifically, the hypothesis is that *intend* has the meaning indicated in (56), where $\text{Effective-Preference}(a,w)$ is a 's effective preference structure at w and \max is as defined in (55) above.

$$(56) \quad \llbracket a \text{ intends } p \rrbracket^w = 1 \text{ iff } p \in \max[\text{Effective-Preference}(a,w)] \quad (\text{to be revised})$$

With this hypothesis in mind, we can now revisit each of the seven properties of *intend* discussed in section 2 above.

3.2.1 Realism

The first property of *intend* we identified is that an intention report ordinarily comes along with the expectation that the attitude holder believes the named outcome to be within the realm of possibility, so that sentences like (57) sound odd.

$$(57) \quad \# \text{John intends to fly to the moon, even though he knows this is impossible.}$$

This property follows straightforwardly from the hypothesis that *intend* exclusively targets effective preferences, provided we follow Condoravdi and Lauer in imposing the realism requirement on effective preference structures. Realism, repeated in (58), is explicitly denied in (57), thereby accounting for its odd status.

⁹In this connection, Heim (1992) in fact suggests that “*want* has a reading more or less equivalent to *intend* ... Probably this is not really an ambiguity but indicates a broader sort of vagueness” (p. 200).

(58) **Realism**

A preference structure $\langle \mathbf{P}, \prec \rangle$ is *realistic*, relative to an information state B , iff for all $p \in \mathbf{P}$: $p \cap B \neq \emptyset$. (Condoravdi and Lauer 2016:29)

3.2.2 *Conjunction Introduction*

The second property of intention reports we identified was conjunction introduction, as witnessed in two distinct ways. First, we saw that it is odd to coordinate two intention reports that involve mutually inconsistent outcomes, as in (59).

(59) #John intends to stay home tonight, but he also intends to go out tonight.

If *intend* is an effective preference predicate, then the oddity of examples like this is captured straightforwardly by Condoravdi and Lauer's requirement that effective preferences obey the consistency property, repeated in (60). In particular, what goes wrong in (59) is that it asserts that *John stays home tonight* and *John goes out tonight* are both among John's maximally ranked effective preferences. But if John's belief worlds exclude the possibility of these two propositions both being true simultaneously, then by (60), they cannot be equally ranked. One must be demoted and thereby not count among John's maximally ranked effective preferences (and thereby not be something that John intends).

(60) **Consistency**

A preference structure $\langle \mathbf{P}, \prec \rangle$ is *consistent* with respect to an information state B iff for any $X \subseteq \mathbf{P}$, if $B \cap \bigcap X = \emptyset$, there are $p, q \in X$ such that $p \prec q$. (Condoravdi and Lauer 2016:29)

The second way conjunction introduction manifests itself is in the observation that it is odd to follow a pair of intention reports with a denial that the attitude holder intends the conjunction of the two outcomes, as in (61).

(61) #John intends to visit Paris this summer. And John intends to visit Rome this summer. But he doesn't intend to visit both Paris and Rome this summer.

As it happens, the analysis as presented so far does not account for the oddity of (61). Let *John visits Paris this summer* be p and let *John visits Rome this summer* be r . Then (61) asserts that: p is among John's maximally ranked effective preferences, and r is among John's maximally ranked effective preferences, but $p \wedge r$ is *not* among John's maximally ranked effective preferences. There is no violation of consistency here as defined in (60); i.e., there are no equally ranked, mutually inconsistent propositions in John's effective preference structure.

Condoravdi and Lauer consider a possible way out of this shortcoming. In particular, they are concerned with the minimal variant of (61) in which *intend* is replaced by *want*. This is given in (62).

- (62) John wants to visit Paris this summer. And John wants to visit Rome this summer. But he doesn't want to visit both Paris and Rome this summer.

Condoravdi and Lauer suggest that the key to understanding examples like (62) is the existence of a third unstated preference on John's part, namely the preference not to expend the resources needed to visit both Paris and Rome. Returning to the crucial example in (61), then, the suggestion would be that its oddity stems from the fact that (63a–c) are mutually inconsistent.

- (63) a. John visits Paris.
 b. John visits Rome.
 c. John does not expend the resources needed to visit both Paris and Rome.

But this does not actually solve the problem, because even on the assumption that in interpreting (61) we have to posit the third unstated preference in (63c), consistency would then merely require that these preferences not all be equally ranked. For example, consistency would be satisfied if (63a) and (63b) are both maximally ranked and (63c) has a lower ranking.

Consequently, what I would like to propose is that there is another restriction imposed on effective preferences structures, namely that maximally ranked preferences are closed under conjunction:

- (64) **Closure under conjunction:**
 A preference structure $\langle \mathbf{P}, \prec \rangle$ is *closed under conjunction* iff for all p, q , if $p \in \max[\langle \mathbf{P}, \prec \rangle]$ and $q \in \max[\langle \mathbf{P}, \prec \rangle]$, then $p \wedge q \in \max[\langle \mathbf{P}, \prec \rangle]$

By (64), if *John visits Paris* and *John visits Rome* are both among John's maximally ranked effective preferences, then so is *John visits Paris and John visits Rome*. This straightforwardly accounts for the oddity of (61). And while (64) may look at *ad hoc* at first blush, I would like to suggest that it is independently quite plausible and natural as a rationality constraint on effective preferences: if an agent is using two preferences to guide his actions, then it is hard to imagine how that agent could not also be using the conjunction of those two preferences to guide his actions.

Furthermore, once we adopt (64), it is not clear that we need consistency as an independent constraint. Consider again the example involving mutually inconsistent outcomes, repeated here in (65).

(65) #John intends to stay home tonight, but he also intends to go out tonight.

(64) requires that if (65) is true, then among John’s maximally ranked effective preferences is the proposition that John stays home tonight and John goes out tonight. But since this is a contradiction, it runs afoul the realism requirement. So at least as far as this particular kind of example goes, consistency is superfluous as an independent constraint. Since closure under conjunction as defined in (64) only targets maximally ranked effective preferences, one thing consistency does that closure under conjunction does not do is ensure that there are no inconsistencies even among equally but non-maximally ranked preferences. But unless we can identify natural language expressions that target non-maximally ranked effective preferences, it is not clear how we would test whether this is needed.

3.2.3 Monotonicity

The third property of intention reports to revisit is upward monotonicity, as witnessed by the oddity of examples like (66).

(66) #John doesn’t intend to teach next semester, but he intends to teach Tuesdays and Thursdays next semester.

As it stands, the hypothesized semantics for *intend* does not predict the infelicity of these kinds of examples. It does not follow from anything presented so far that if p entails q and p is a maximally ranked effective preference, q is a maximally ranked effective preference as well.

There are at least two tacks we could take toward addressing this shortcoming. One option is to modify the denotation for *intend* from (67) to (68). Something like (68) is in fact explicitly entertained (though ultimately not adopted) by Condoravdi and Lauer (2016) in connection with their semantics for *want*. Condoravdi and Lauer use the term ‘exact-match’ for the kind of attitude semantics instantiated by (67), and they contrast this with what they call the ‘Quine-Hintikka’ kind of attitude semantics instantiated by (68). The latter denotation ensures upward monotonicity and would thereby account for the data in (66).¹⁰

¹⁰A more explicitly Hintikkan formulation would look like (i), or equivalently, (ii). Like (68), (i)/(ii) ensures upward monotonicity. If we took this kind of approach, then the closure under conjunction requirement on effective preference structures that I proposed in section 3.2.2 above could be naturally understood as the requirement that $\bigcap \text{max}[\text{Effective-Preference}(a,w)]$ be non-empty.

(i) $\llbracket a \text{ intends } p \rrbracket^w = 1$ iff $\forall w' [w' \in \bigcap \text{max}[\text{Effective-Preference}(a,w)] \rightarrow p(w')]$

(ii) $\llbracket a \text{ intends } p \rrbracket^w = 1$ iff $\bigcap \text{max}[\text{Effective-Preference}(a,w)] \subseteq \{w' \mid p(w') = 1\}$

- (67) *Exact-match* (cf. Condoravdi and Lauer 2016:21):
 $\llbracket a \text{ intends } p \rrbracket^w = 1$ iff $p \in \max[\text{Effective-Preference}(a,w)]$
- (68) *Quine-Hintikka* (cf. Condoravdi and Lauer 2016:21):
 $\llbracket a \text{ intends } p \rrbracket^w = 1$ iff $\exists q [q \subseteq p \text{ and } q \in \max[\text{Effective-Preference}(a,w)]]$

The second option would be to maintain the ‘exact-match’ analysis of *intend* as in (67), but to impose yet another requirement on effective preference structures, along the following lines:

- (69) **Monotonicity**
 A preference structure $\langle \mathbf{P}, \prec \rangle$ is *monotonic* iff for all p, q , if $p \in \max[\langle \mathbf{P}, \prec \rangle]$ and $p \rightarrow q$, then $q \in \max[\langle \mathbf{P}, \prec \rangle]$

By (69), if *John teaches Tuesdays and Thursdays next semester* is among John’s maximally ranked effective preferences, then so is *John teaches next semester*. As with closure under conjunction, this constraint may seem *ad hoc*, but a case could be made that it is actually an independently plausible rationality constraint on effective preferences: if an agent is to use some preference to guide his actions, then it seems natural to assume that his actions would also have to be guided by logical consequences of that preference.

The choice between adopting a Quine-Hintikka analysis of attitude predicates and adding a monotonicity constraint on effective preference structures will not be crucial to anything that follows, but for concreteness, I will adopt the second option. The reason is that the data that establish the monotonicity property of intention reports seems to be on a par with the data that bear on realism and conjunction introduction, in the sense that they have to do with assumptions we ordinarily make about rational agents. Consequently, it seems appropriate to treat them in parallel fashion in the analysis, rather than singling out one for special treatment.

3.2.4 Non-gradability

The next property of intention reports to be dealt with is their non-gradability, as evidenced by the oddness of examples like (70).

- (70) a. ?John intends to go to Paris more than he intends to go to London.
 b. ?What John intends the most is to be happy.
 c. ?John intends very much to leave.

This non-gradability follows straightforwardly from the analysis of *intend* under consideration, repeated in (71). The reason is that although effective preference structures, like all preference structures, come with an ordering, the semantics for

intend in (71) is defined in such a way that only *maximally ranked* effective preferences are considered. By definition, maximally ranked preferences are not asymmetrically ranked with respect to each other, and so there is nothing on which to hang gradability.

(71) $\llbracket a \text{ intends } p \rrbracket^w = 1$ iff $p \in \max[\text{Effective-Preference}(a,w)]$ (to be revised)

In fact, what is of more interest is the observation that *want*, unlike *intend*, is gradable, as we've already seen in examples like those repeated in (72). If *want* is like *intend* in exclusively targeting *maximally ranked* preferences, it is not obvious how examples like this could be accommodated.

- (72) a. John wants to go to Paris more than he wants to go to London.
 b. What John wants the most is to be happy.
 c. John wants very much to leave.

Since this paper's focus is *intend* rather than *want*, a full investigation of this issue would take us too far afield. But a few remarks will help clarify how *want* and *intend* differ and thereby contribute to a better understanding of *intend*. To a first approximation, what seems to be needed for (72a), to take one example, is a semantics whereby John has some (contextually resolved) preference structure, according to which the proposition *John goes to Paris* is more highly ranked than *John goes to London*.¹¹ In other words, the sentence tells us something about the ranking of the two preferences with respect to each other without saying anything about their absolute standing in the preference structure. Consistent with such an analysis, it is noteworthy that (73a) can be true even in a situation where (73b) is false, analogous to the observation that (74a) does not entail (74b).

- (73) a. John wants to go to Paris more than he wants to go to London.
 b. \nrightarrow John wants to go to Paris.

- (74) a. John is taller than Bill.
 b. \nrightarrow John is tall.

According to an influential approach to the semantics of gradable predicates (see e.g. Kennedy 2007), the so-called positive form of a gradable predicate such as

¹¹More complicated are cases like (i), which require us to make a comparison across distinct preference structures (in this case, one of John's preference structures and one of Bill's preference structures). Such sentences seem to call for more than what we get for free with Condoravdi and Lauer's ordering relation, which only applies *within* a given preference structure.

- (i) John wants to go to Paris more than Bill wants to go to London.

tall as used in (74b) incorporates a silent morpheme (the so-called *pos* morpheme) that interacts with the rest of the material in the sentence to yield truth conditions to the effect that John stands out along the dimension named by the gradable predicate relative to a threshold determined in part by the context. In (74a), by contrast, *pos* is absent, and instead, the comparative morpheme interacts with the other material to yield truth conditions to the effect that John's height exceeds Bill's height.

A reasonable null hypothesis then is that (73a–b) is analytically parallel to (74a–b): in the semantics for (73b), there is a silent *pos* morpheme inserted in the absence of overt degree expressions such as the comparative morpheme in (73a). Such an analysis moves us away from the 'maximality' based semantics for *want* proposed by von Stechow (1999) (and later adopted by Condoravdi and Lauer (2016)) on analogy with the role that ordering sources play in a Kratzer-style modal semantics, and instead moves us in the direction of approaches to *want* taken by authors such as Villalta (2008); Lassiter (2011); Rubinstein (2012); Anand and Hacquard (2013).

While these remarks only scratch the surface, the important point is that on the view suggested here, *want* and *intend* are both built on preference structures, but *intend* deals in maximally ranked preferences only and hence does not tolerate gradability, whereas *want* allows us to access non-maximal portions of the preference structures it invokes.

3.2.5 Mood choice

We turn now to the question of why *intend* should pattern like *want* and unlike *believe* in taking subjunctive complements. Portner (2015) generalizes that most theories of mood choice fall into one of two categories. The first is what Portner calls the comparison-based theory, according to which "The subjunctive marks a clause which functions as the argument of a predicate with a comparative semantics" (Portner 2015:4). The other is what Portner calls the truth-based theory, according to which "The indicative marks a clause which is entailed to be true throughout a designated set of worlds" (Portner 2015:4). Let's consider each in turn.

At first blush, the comparison-based theory would seem to stand at odds with the observation that *intend* behaves like a non-gradable predicate. And in fact, this is a familiar problem: Villalta (2008) argues for a version of the comparison-based theory, based mostly on data from Spanish, but also points out that directive predicates like *mandar* 'order' and causative predicates like *hacer* 'make' select for subjunctive complements even though these predicates do not pass standard tests for gradability. As a way out of this problem, Villalta suggests that for these predicates, "the comparative meaning component is embedded more deeply within

the meaning of the predicate” (Villalta 2008:484). A similar approach could be entertained for *intend*: it has a comparative meaning in the sense that it is built on the logic of preference structures, which encode an ordering relation, and this is why it selects for the subjunctive. But because non-maximally ranked preferences are ignored, this gradability does not manifest itself in the ordinary ways.

As for the truth-based theory of mood choice, let’s consider Giannakidou 2009 as a concrete example. Building on her previous work, Giannakidou suggests that an attitude verb selects for the indicative when the attitude holder “is committed to the truth of the complement sentence” (Giannakidou 2009:1887). This is interesting insofar as *intend* has to do intuitively with an attitude holder’s commitment to undertake some action. But commitment to the truth of a proposition and commitment to undertaking an action are not necessarily the same thing. And crucially, on the analysis proposed for *intend* in this paper, an intention report of the form *a intends p* does not entail that *a* is committed to the truth of *p*; rather, it comes only with the weaker realism constraint that *a* believes *p* to be within the realm of possibility. (And although it is an admittedly delicate matter, see section 6.2 below for potential evidence that intention reports involve this relatively weak kind of realism.) Also of relevance is the observation that *intend*, like *want* and unlike *believe*, prototypically takes future-oriented complements and does not allow past-oriented complements.¹² Futurity has been argued to be a nonveridical category (Giannakidou and Mari 2016) and hence also consistent with the appearance of subjunctive mood on the truth-based theory.

3.2.6 Anankastic conditionals

Now we return to the observation that *want* and *intend* support anankastic readings of conditionals whereas *believe* does not:

- (75) a. If you want to get good grades, you must study.
 b. \approx You must study in order to get good grades.
- (76) a. If you intend to get good grades, you must study.
 b. \approx You must study in order to get good grades.

¹²As with *want* as depicted in (ia), present-oriented interpretations of the complement to *intend* are atypical but possible when the context is set up in the right way: see (ib).

- (i) a. John is at home. And that is good because that’s exactly where John wants to be right now.
 b. John is tapping the table with his pen. But don’t stop him because that’s exactly what he intends to be doing right now.

- (77) a. If you believe you'll get good grades, you must study.
 b. $\not\approx$ You must study in order to get good grades.

Since Condoravdi and Lauer's (2016) semantics for *want* (and by extension my proposed semantics for *intend*) is designed with data from anankastic conditionals in mind, there is not much to be said here beyond briefly reviewing Condoravdi and Lauer's approach to anankastic conditionals. In a nutshell, Condoravdi and Lauer propose that when the modal in the consequent of the conditional is interpreted teleologically, its ordering source is constituted by the relevant agent's maximally ranked effective preferences. The anankastic reading arises when the embedding verb in the antecedent of the conditional is construed as an effective preference predicate, thereby potentially affecting the ordering source of the teleological modal as we move from one world to the next (the link between the antecedent and the consequent being mediated by a covert epistemic necessity modal). Crucially, *want* can be read as an effective preference predicate when its underspecified preference structure variable is resolved in the appropriate way. And on the analysis entertained here, *intend* is always read as an effective preference predicate. By contrast, *believe* does not deal in preferences (effective or otherwise), so the lack of an anankastic reading in (77) is no mystery.

3.2.7 Responsibility

Finally, we return to the observation that in an intention report, the outcome is necessarily construed as being intentionally brought about by the attitude holder, as witnessed by the causation effect (78), the oddness of denying intentionality in the complement (79), the redundancy of asserting intentionality in the complement (80), and the oddness of complements that name outcomes ordinarily not under one's control (81).

- (78) a. John intends (for) Bill to leave.
 b. \approx John intends to bring it about that Bill leave.
- (79) #John intends to break the window unintentionally.
- (80) #John intends to break the window intentionally.
- (81) #John intends to resemble his father.

We have proposed that an intention report is used to describe an effective preference, i.e., a preference that guides action. On the one hand, it is tempting to think that such a proposal makes the effects witnessed here seem somewhat unsurprising: more specifically, it is tempting to think that if a preference is to guide action, then it is only natural that the realization of that preference would have to be

something under the deliberate control of the attitude holder. (Though see section 7 below, where some facts surrounding the predicate *count on* constitute a possible reason for resisting this temptation.) But on the other hand, nothing in the theory guarantees that *intend* encodes responsibility. The goal of the next section is to correct this.

4 Analysis part 2: *intend* and the RESP-relation

4.1 Background on Farkas 1988

Farkas (1988) argues that in order to understand certain natural language phenomena, we need to recognize the existence of a relation she calls RESP(ONSIBILITY). For Farkas, RESP is a two-place relation between an individual *i* and a situation *s* that holds true “just in case *s* is the result of some act performed by *i* with the intention of bringing *s* about” (Farkas 1988:36).

As evidence for the RESP-relation, Farkas points to four phenomena wherein a verb phrase or clause is felicitous only if the situation it describes can be construed as being intentionally brought about by some individual. The first phenomenon is rationale clauses. As seen in (82a–b), rationale clauses are felicitous in combination with event or state descriptions that can be construed as being intentionally brought about, such as John reading a book or a shopwindow having a sale sign in it. By contrast, as seen in (82c–d), rationale clauses are infelicitous in the context of situations not ordinarily under one’s control such as a son resembling his father or the weather being good lately.

- (82) a. John read ‘Anna Karenina’ in order to impress Mary.
b. The shopwindow has a big sale sign in it in order to attract customers.
c. #John resembles his father in order to annoy his grandmother.
d. #The weather has been good lately in order to please the tourists.
(Farkas 1988:36)

The second phenomenon is imperatives. Based on the data in (83), Farkas suggests that imperatives are felicitous only when the situations they describe can be construed as standing in the RESP-relation with the addressee.

- (83) a. Be polite!
b. #Be tall!
c. #Resemble your father!
(Farkas 1988:39)

The third phenomenon is the adverb *intentionally*. Based on the data in (84), Farkas suggests that *intentionally* can be used only when the RESP-relation holds

between the subject and the situation described by the verb phrase.

- (84) a. John hurt Mary intentionally.
b. John fell off the ladder intentionally.
c. #John resembles his father intentionally.
d. #John is tall intentionally. (Farkas 1988:39)

Finally, Farkas's fourth piece of evidence for the RESP-relation is a class of control predicates. Farkas generalizes that in the case of object-control predicates, if x persuades / convinces / forces / urges / requires y to VP, then the satisfaction conditions for the state of affairs described are such that y is responsible for bringing about the situation named by the VP. This is supported by the observation that when the VP names a situation not under one's control, the result is odd:

- (85) #John {convinced/persuaded/requested/ordered/required} Pete to
{resemble Bill/be tall/be blue-eyed/bleed}

And, Farkas points out, there are some subject-control predicates whose conditions of satisfaction involve the subject standing in the RESP-relation with the VP. Although Farkas mostly focuses on *promise*, as in (86), the same observation holds for other predicates of public commitment (87a), as well as for predicates of private commitment such as *intend* (87b).¹³

- (86) #John {promised} to {resemble Bill/be tall/be blue-eyed/bleed}
(87) a. #John {agreed/offered} to {resemble Bill/be tall/be blue-eyed/bleed}
b. #John {chose/decided/intended/agreed/offered} to {resemble Bill/be tall/be blue-eyed/bleed}

Farkas is careful to distinguish the RESP-relation from the conceptually very similar thematic relation AGENT. Farkas points to two important differences. First, whereas an individual standing in the agent relation to a situation must be syntactically an argument of the predicate associated with the relevant situation, this need not hold for the RESP-relation. For example, in (88), the situation of the shopwindow having a big sale sign in it stands in the RESP-relation with some individual

¹³In the interest of space, I will not discuss in detail here other phenomena not considered by Farkas that may also involve the RESP-relation. But I would like to mention a few candidates: in later work, Farkas and Giannakidou (1996) put the RESP-relation to use in making sense of the distribution of extra-wide scope. See also Barker 2002 for an application of the RESP-relation in understanding a particular class of evaluative adjectives. And although not explicitly connected to Farkas's work, two other empirical phenomena that may be fruitful to explore from this angle include futarates (Copley 2008) and *have*-causatives (Copley and Harley 2009), which the cited authors argue involve commitment to bringing about an action.

not explicitly represented in the argument structure. Similarly in (89), according to Farkas, John is understood as standing in the RESP-relation with the situation named by the complement to *promise* even though *John* binds no argument in this complement.

(88) The shopwindow has a big sale sign in it in order to attract customers.

(89) John promised Mary that the children will be in bed by 8.

To (88)–(89) we may add a couple of other kinds of examples of this phenomenon not considered by Farkas. First is the use of *intentional* in (90), which works in relevant respects just like (89). Second is the existence in some languages of imperatives with non-second-person subjects, as in (91) (see Zanuttini et al. 2012 for discussion, from which both of these examples are taken). In these imperatives, the addressee is understood as bearing the RESP-relation to the situation described by the imperative even though there is no second-person argument in the imperative.

(90) It was intentional on John’s part that the children were in bed by 8.

- (91) a. Tebulwa: sa:ph rahe!
 table-nom clean-nom be-imp.3s
 ‘The table be clean!’ BHOJPURI
- b. Che venga anche lui!
 that come also he
 ‘(See to it that) he comes as well.’ ITALIAN

The second difference Farkas identifies between the RESP-relation and the agent relation is that the agent relation need not encode any notion of intention. For example, it would be relatively uncontroversial to say that John bears the agent relation in sentences like (92), even though this sentence is compatible with a situation in which John broke the window unintentionally (see section 5 for additional discussion).

(92) John broke the window.

4.2 Back to intend

Combining the effective preference analysis of *intend* with the insights from Farkas (1988) leads naturally to the proposal that we revise our preliminary semantics for *intend* in (93) to that in (94).¹⁴

¹⁴Farkas (1988) treats RESP as a relation between individuals and situations, whereas I assume here that it is a relation between individuals and propositions. See section 4.3.1 for discussion.

(93) $\llbracket a \text{ intends } p \rrbracket^w = 1$ iff $p \in \max[\text{Effective-Preference}(a,w)]$ (*to be revised*)

(94) $\llbracket a \text{ intends } p \rrbracket^w = 1$ iff $\text{RESP}(a,p) \in \max[\text{Effective-Preference}(a,w)]$ (*final*)

Let us now consider how the revised proposal accounts for the responsibility effects we observed. Since the RESP-relation is a meta-language predicate, we cannot directly access linguistic intuitions about the conditions under which it holds. But the following rough natural-language equivalences seem in keeping with Farkas's proposal (though see section 4.3.2 below for further discussion):

(95) $\text{RESP}(a,p)$
 \approx it is intentional on a 's part that p
 $\approx a$ intentionally brings it about that p

For the special case where the subject of p happens to be a (i.e., where we have $\text{RESP}(a,P(a))$), the following also seems appropriate:

(96) $\text{RESP}(a,P(a)) \approx a$ intentionally P s

Armed with these rough equivalences, let's now reconsider the causation effect:

(97) a. John intends (for) Bill to leave.
b. \approx John intends to bring it about that Bill leave.

On the current proposal, we get the following semantics for (97a) (here and in what follows, let $\llbracket \text{John} \rrbracket = j$):

(98) $\llbracket \text{John intends (for) Bill to leave} \rrbracket^w =$
1 iff $\text{RESP}(j, \llbracket \text{Bill leave} \rrbracket) \in \max[\text{Effective-Preference}(j,w)]$
'John's maximally ranked effective preferences include the proposition that John intentionally brings it about that Bill leave.'

The causation effect is thus accounted for by the presence of the RESP-relation.

Consider next the observation that it is odd to deny intentionality in the complement of *intend* (99a), redundant to assert intentionality (99b), and odd to use a complement that names an outcome not ordinarily under one's control (99c).

(99) a. #John intends to break the window unintentionally.
b. #John intends to break the window intentionally.
c. #John intends to resemble his father.

On the current proposal, the denotations for (99a–c) are (100a–c) respectively:

- (100) a. 1 iff $\text{RESP}(j, [\text{John break the window unintentionally}]) \in \max[\text{Effective-Preference}(j, w)]$
 ‘John’s maximally ranked effective preferences include the proposition that John intentionally breaks the window unintentionally.’
 b. 1 iff $\text{RESP}(j, [\text{John break the window intentionally}]) \in \max[\text{Effective-Preference}(j, w)]$
 ‘John’s maximally ranked effective preferences include the proposition that John intentionally breaks the window intentionally.’
 c. 1 iff $\text{RESP}(j, [\text{John resemble his father}]) \in \max[\text{Effective-Preference}(j, w)]$
 ‘John’s maximally ranked effective preferences include the proposition that John intentionally resembles his father.’

We thereby accurately predict that (99a–c) should have a status similar to (101a–c), respectively.

- (101) a. #John intentionally broke the window unintentionally.
 b. #John intentionally broke the window intentionally.
 c. #John intentionally resembled his father.

Notice that the proposal does not predict that (99)/(101) are infelicitous in all contexts but only that they should have a similar status. In this connection, Paul Egré (pers. comm.) suggests that (101a) could be possible in a context in which someone orders John to break the window without thinking about it. So John reasons that if he takes the appropriate drugs, this will adversely affect his motor skills, thereby increasing the odds of inadvertently breaking things around him. In this context, (99a) seems true and felicitous, and if John is successful in his plan, (101a) seems true and felicitous as well.

4.3 More on the RESP-relation

Having laid out the core analysis and shown how it accounts for the relevant data, there are still a couple of issues surrounding the RESP-relation that need to be addressed. The first is reconciling the fact that Farkas treats RESP as a relation between individuals and situations whereas I treat it as a relation between individuals and propositions. The second has to do with whether we can be more precise about the content of the RESP-relation, moving beyond the natural language paraphrases that I relied on above. In what follows, I take these two issues up in turn.

4.3.1 Probing the semantic type of the RESP-relation

The proposed semantics for intention reports, repeated in (102), assumes that the infinitival complement to *intend* denotes a proposition, and that RESP is a relation between individuals and propositions. Other options are conceivable. Farkas (1988) for example treats RESP as a relation between an individual and an eventuality (or in her terminology: situation). This option could be implemented by treating the complement to *intend* as a property of eventualities rather than a proposition and redefining its truth conditions as in (103).

(102) $\llbracket a \text{ intends } p \rrbracket^w = 1$ iff $\text{RESP}(a,p) \in \max[\text{Effective-Preference}(a,w)]$

(103) $\llbracket a \text{ intends } p \rrbracket^w = 1$ iff $[\exists e [p(e) \wedge \text{RESP}(a,e)]] \in \max[\text{Effective-Preference}(a,w)]$

Actually, the choice between (102) and (103) conflates two separate issues that need to be distinguished. The first issue has to do with the semantic type of the infinitival complement to *intend*: is it proposition-denoting? Or, as an anonymous reviewer wonders about, could it be that it is not large enough to include whatever functional projection gives rise to quantification over the eventuality argument of the predicate (say, Aspect), and therefore denotes a property of eventualities? The second issue is whether RESP takes as its argument the denotation of the entire complement (whether that be a proposition or a property of eventualities) or just the eventuality argument associated with that complement.

I do not know how to settle the first issue. Wurmbrand (2014) argues that infinitival complements to future-oriented verbs like *intend* do not contain Tense (though they do contain the future modal *woll*), but does not explicitly argue for or against the presence of Aspect in such complements. Keshet (2008), on the other hand, argues that all infinitival complements have an aspectual position. It can be observed that complements to *intend* admit both the progressive (104a) and the perfect auxiliary (104b), but this counts as evidence for the propositional view only insofar as it can be established that these auxiliaries do indeed bind the predicate's eventuality argument, which is a delicate question. Consequently, I will remain neutral on this issue.

- (104) a. John intends [to be studying when his parents get home].
b. John intends [to have lost five pounds by Christmas].

The second issue, however, may be more tractable. Farkas's treatment of RESP as a relation between individuals and eventualities is *a priori* attractive insofar as it makes RESP look like a thematic relation, to which it bears a close intuitive resemblance. But an argument can be made to support the view that RESP has to take as its argument something larger than this. The argument rests on the premise that

the expressions *intentionally* and *intentionally bring it about that* are close natural-language analogues of RESP and therefore should have similar logical properties. The relevant observation is that both of these expressions induce referential opacity on nominal arguments embedded in the constituents with which they combine, as seen in (105) (cf. Thomason and Stalnaker 1973). Moreover, both expressions need to be able to scope over negation, as seen in (106). What these data suggest is that *intentionally* and *intentionally bring it about that* (and by extension RESP) need to take as their argument more than just an eventuality and in particular something that includes both the direct object and a position for negation. This is consistent with the view that RESP is a relation to a proposition or to a property of eventualities, but it is not consistent with the view that RESP is a relation to an eventuality.

- (105) a. Oedipus intentionally married his mother.
 \neq Oedipus intentionally married Jocasta.
 b. Oedipus intentionally brought it about that he married his mother.
 \neq Oedipus intentionally brought it about that he married Jocasta.
- (106) a. John intentionally did not break the window.
 \neq John did not intentionally break the window.
 b. John intentionally brought it about that he did not break the window.
 \neq John did not intentionally bring it about that he broke the window.

4.3.2 Probing the content of the RESP-relation

Above, I have relied on the idea that we can intuit the content of the RESP-relation by translating $\text{RESP}(a,p)$ into its close natural language analogues *a intentionally brings it about that p* or *it is intentional on a's part that p*, and for the special case where the subject of *p* is *a* (i.e., $\text{RESP}(a,P(a))$), *a intentionally Ps*. But a reasonable question to ask is whether we can do better than relying on these natural language analogues.

Let me first address this question in the context of the narrow goals of this paper. In the context of this paper, invoking the RESP-relation in the semantics of *intend* helps us achieve two goals. First, it helps account for the responsibility effects associated with intention reports discussed in section 2.7 above. Second, it helps us make sense of the similarities among various phenomena that involve intentional action (rationale clauses, imperatives, some agent-oriented adverbs, and some control predicates), because we can say that these similarities are underpinned by the common involvement of RESP in all of these phenomena. With respect to these two particular goals, it is not clear to me that any value would be added in further limning the RESP-relation. A comparison might be made here to thematic relations: these are often treated as semantic primitives, and for many purposes

this is sufficient, despite the fact that there are difficult puzzles surrounding the question of exactly when the agent relation (for example) holds and when it does not (see e.g. the recent discussion in Williams 2015:141–144). And of course thematic relations are just the tip of the iceberg: even a mundane predicate such as that named by *cat* has vague boundaries (Chierchia 2010:117), though for most purposes we can treat it as a primitive.

That being said, I think it is likely that a more sophisticated understanding of the RESP-relation is possible and that such an understanding may turn out to be useful for some purposes beyond those of this paper. In this connection, I simply want to point to the work of Egré (2014) (cf. also Egré and Cova 2015). Drawing on relevant work in experimental philosophy concerning people’s intuitions about the conditions under which someone’s action counts as *intentional* (see especially Knobe 2003a,b; Pettit and Knobe 2009), Egré proposes a semantics for the predicate *intentional* according to which it is a vague, gradable, and multi-dimensional predicate sensitive to the degree to which the agent desires the relevant outcome and the degree to which the agent can foresee how his/her actions will lead to the outcome. Plausibly, RESP is amenable to the same kind of approach. And since even core thematic relations like agent have been suggested to be vague and context-sensitive (Pietroski 2004:184, cited by Williams (2015):144), it seems plausible that RESP would be as well.

5 Compositionality: To coerce or not to coerce?

On the proposal presented above, intention reports that instantiate syntactic control, like (107), have essentially the same status as intention reports that do not instantiate syntactic control, like (108). In both cases, the complement to *intend* instantiates *p* in the formula repeated in (109). And in both cases, John names both the individual whose effective preference structure is accessed as well as the individual who bears the RESP-relation with *p*. The only difference is that in (107), in virtue of being a control sentence, John also names the participant associated with the subject position of *break the window*, whereas in (108), Bill takes this place.¹⁵ Call this a COERCION-FREE semantics for intention reports.

(107) John intended to break the window.

(108) John intended for Bill to break the window.

¹⁵How exactly the control relation gets encoded by the syntax and interpreted by the semantics is a very interesting question but one that would take us too far afield here: for two recent approaches to the syntax-semantics of control consistent with the general architectural assumptions of this paper, see Stephenson 2010; Pearson 2016.

$$(109) \quad \llbracket a \text{ intends } p \rrbracket^w = 1 \text{ iff } \text{RESP}(a,p) \in \max[\text{Effective-Preference}(a,w)]$$

A COERCION-BASED semantics for intention reports, on the other hand, is built around the idea that whereas (107) is compositionally straightforward, something special has to happen in the case of (108). More specifically, a proponent of the coercion-based view might react to what I have proposed in this paper in something like the following way: in (107), what this paper identifies as the RESP-relation embedded into the semantics of *intend* is actually just part of the ordinary thematic structure associated with *break the window*, whose subject bears the RESP-relation to the described event. What the predicate *intend* does is SELECT FOR a control complement whose unexpressed subject bears the RESP-relation. This selectional restriction is satisfied in (107). But it is not satisfied in (108). Consequently, the grammar repairs this selectional restriction violation by superimposing on top of the complement a new RESP-relation whose individual argument can be controlled from the matrix clause. (Something like this is entertained by authors such as Perlmutter (1968); Jackendoff (1996); Jackendoff and Culicover (2003); Culicover and Jackendoff (2005); Grano (2015), though I have paraphrased the idea to make it consistent with the terminology used elsewhere in this paper; none of these authors, for example, use the term RESP.)

Is there any independent evidence for assigning (108) a special, coerced status? Jackendoff and Culicover (2003) cite the paraphrase relation that intuitively holds between sentences like (110a) and sentences like (110b) as support for a coercion analysis.

- (110) a. John intended for Bill to break the window.
 b. \approx John intended to bring it about that Bill break the window.

But a paraphrase relation is not in itself probative of coercion. For example, a paraphrase relation also holds between (111a) and (111b), and (111a) is indeed often analyzed as involving coercion (see e.g. Pustejovsky 1995), but in a recent paper, Piñango and Deo (2016) mount a convincing case that there is in fact no coercion here.

- (111) a. John began the book.
 b. \approx John began to read/write the book.

Suppose, though, that we adopt the coercion analysis for the sake of argument and see where it takes us. An immediate problem for the coercion analysis is that there seems to be nothing inherently intentional about the relation borne by the subject in *break the window*. Although there do exist some inherently intentional verbs like *murder* (see Kamp 1999–2007), an event report like (112a) seems

to be underspecified for intention in the sense that it could be truthfully uttered in a scenario where John unintentionally broke the window. It is also compatible with an explicit denial of intention, as in (112b). This suggests that it is not enough for *intend* to rely on a pre-existing thematic relation; RESP is needed to encode intentionality and thereby account for the full range of responsibility effects, even though there will often be a fair amount of overlap between what RESP contributes and what is already contributed by core thematic structure.

- (112) a. John broke the window.
b. John unintentionally broke the window.

But we could be a bit more generous. Suppose we were to entertain the view that the predicate *break the window* harbors an ambiguity, and that one of its readings is specified for intentionality. Then it would be conceivable that *intend* comes with a selectional restriction that lets it combine with the intentional reading of *break the window* only. Is such an ambiguity plausible? An early discussion of the question of whether event reports like (112a) are underspecified or ambiguous with respect to intentionality is taken up by Zwicky and Sadock (1975), who ultimately conclude that the facts do not clearly support either the ambiguity view or the underspecification view. One point Zwicky and Sadock make is that the so-called ‘identity tests’ for ambiguity are unilluminating in cases where the two readings being tested stand in a privative rather than a polar opposition. For example, (113a) illustrates a standard attachment ambiguity. No ‘crossed’ reading is available whereby the ambiguity is resolved in one way in the first conjunct but the other way in the elided VP in the second conjunct. (113b), on the other hand, can be used to report a situation wherein John ate a *ham* sandwich but Bill ate a *turkey* sandwich, thereby supporting the view that *sandwich* is underspecified with respect to these details rather than ambiguous. But even though a crossed reading is available in (113c) wherein John broke a mirror intentionally and Bill broke a mirror unintentionally, this does not rule out a version of the ambiguity view wherein one reading is underspecified with respect to intention and the other reading is intention-specific. Such an approach is compatible with (113c) being true in a ‘crossed’ scenario, where the ambiguity is resolved to the underspecified reading in both conjuncts.

- (113) a. John saw the man with the telescope, and so did Bill.
b. John ate a sandwich, and so did Bill.
c. John broke a mirror, and so did Bill.

In the context of the present investigation, though, VP ellipsis actually is illuminating and the evidence points toward underspecification rather than ambi-

guity. In particular, it is telling that there is nothing contradictory about (114). If *break the window* had an unambiguously intentional reading that *intend* selected for, then the expectation would be that if this VP were embedded under *intend* and then served as an antecedent for VP ellipsis, the elided VP should also be unambiguously intentional. But this is not the case, as (114) attests to: it is possible to explicitly deny intention even when the VP in question is anteceded by a VP embedded under *intend*. (I assume here that the licensing conditions for VP ellipsis include the requirement that there be a semantically identical antecedent, as in e.g. Merchant 2001.)

(114) John intended to [break the window], and in the end he did ~~(break the window)~~, albeit not intentionally.

This suggests that *break the window* does not have an inherently intentional reading, thereby undermining the coercion-based approach to intention reports.

A lot more could be said here. For example, a proponent of the coercion analysis could try to make sense of (114) via a decompositional approach to thematic structure whereby the intentionality encoded into the complement of *intend* is low enough in the verbal projection to be selected by *intend* but too high to be targeted by VP ellipsis. One might also assess the plausibility of coercion based on non-semantic factors such as sentence processing or cross-linguistic considerations. But I hope here to have made the case that a coercion-free analysis is plausible and perhaps even preferable.

6 Comparison to previous approaches

As noted in the introduction, the formal semantics literature is sparse when it comes to intention reports. But there are two recent proposals that I want to discuss in this section, namely those found in Grano 2015 and Pearson 2016. This will serve both to better situate this paper's analysis with respect to the previous literature and to throw into relief certain aspects of this paper's analysis that have not yet been explicitly highlighted.

6.1 Grano 2015

The background for Grano's (2015) semantics for *intend* is Portner's (2004; 2007) proposal that clause types are each associated with a distinct semantic type and relate to a distinct discourse component: declarative clauses denote propositions and can consequently be used to update the Common Ground (a set of propositions), interrogative clauses denote sets of propositions and can consequently be used to

update the Question Set (a set of sets of propositions), and imperative clauses denote properties and can consequently be used to update the addressee’s To-Do List, modeled as a set of properties such that “the participants in the conversation mutually assume that [the addressee] will try to bring it about that he or she has each of these properties” (Portner 2007:352).

Against this backdrop, Grano observes that it is always infelicitous to accept an imperative while denying the corresponding intention, even as one considers the different kinds of imperative flavors identified by Portner (2007):

- (115) a. Sit down right now! (order)
- b. OK. #But I don’t intend to. (Cf.: OK. But I don’t want to.)
- (116) a. Have a piece of fruit! (invitation)
- b. OK. #But I don’t intend to. (Cf.: OK. But I don’t want to.)
- (117) a. Talk to your advisor more often! (suggestion)
- b. OK. #But I don’t intend to. (Cf.: OK. But I don’t want to.)

(Grano 2015:244)

Grano therefore concludes, following an earlier suggestion by Ninan (2005), that To-Do Lists come in public and private varieties (analogously to the way the Common Ground has as its private counterpart individual belief): imperatives target the addressee’s Public To-Do List whereas intention reports target an attitude holder’s Private To-Do List.

Building on this conclusion, Grano proposes that *intend* has the semantics in (118), where PrivTDL is a function from an individual to the set of properties constituting that individual’s Private To-Do List.

$$(118) \quad \llbracket \text{intend} \rrbracket = \lambda P_{\langle et \rangle} \lambda x. P \in \text{PrivTDL}(x)$$

Crucial to Grano’s proposal is that *intend* combines with a property-denoting complement. Taking the view that control complements like (119) are property-denoting whereas non-control complements like (120) are proposition-denoting (Dowty 1985), (120) incurs a type mismatch. Updating an idea due to Perlmutter (1968), Grano proposes that this mismatch is repaired via the silent causative predicate in (121). In this way, the causation effect — i.e., the observation that (120) can be paraphrased as *John intended to bring it about that Bill leave* — is made sense of.

- (119) John intended [to leave].
- (120) John intended [for Bill to leave].
- (121) $\llbracket \text{CAUSE} \rrbracket = \lambda p \lambda x. x \text{ brings it about that } p$

The primary strength of this approach lies in the deep analytical connection it identifies between intention reports and imperatives, which seems quite attractive. But instead of cashing out that connection via the RESP-relation, as I have attempted to do, Grano cashes it out type-theoretically, the animating idea being that natural language models objects of commitment (whether public or private) as properties. This raises many questions. Here, in the interest of space, I will raise just what I see as the most serious concern. Consider the intention reports in (122).

- (122) a. John intended to break the window.
b. John intended for Bill to break the window.
c. John intended to resemble his father.

On the coercion-free analysis of intention reports that I have argued for above, all three of the intention reports in (122) involve the attitude holder standing in the RESP-relation with the complement. (122a), we might say, is the easiest to interpret, because it is not hard to imagine how John could have responsibility over an eventuality in which he is the agent of a window breaking. (122b) by contrast requires a bit more work since it requires us to construe John as overseeing an eventuality in which he is not a participant and that has its own agent, Bill. Finally, (122c) is the most difficult because one has to imagine ways in which John could have control over his resembling his father. In summary, on this kind of approach, the perceived differences in interpretive ease do not reflect different grammatical mechanisms (such as coercion) but rather the ease with which we can imagine a scenario that would verify the relevant RESP-relation.

Contrast this with the analysis of Jackendoff and Culicover (2003). They would say that *intend* selects for an action, which is satisfied in (122a), whereas in (122b–c), the complements denote non-actional situations and hence coercion would be needed to interpret these sentences. Modulo the reservations about appealing to coercion discussed in section 5 above, there is something attractive about this idea: (122b–c) require extra work in a way that (122a) does not.

Grano's (2015) analysis, by contrast, has the somewhat bizarre consequence that (122a) and (122c) pattern together in being coercion-free, since both sentences involve a control complement, whereas (122b) requires coercion. And yet if anything, the work involved in interpreting (122c) (as measured, say, by the urge to paraphrase (122c) as *John intended to bring it about that he resemble his father*) is even greater than the work involved in interpreting (122b).

In summary, it seems sensible to treat (122a–c) in a uniform way, or to single out (122b–c) and assign them a special coerced status. But to treat (122a) and (122c) uniformly to the exclusion of (122b), as Grano (2015) would have it, runs contrary the intuition that motivates the coercion analysis in the first place.

6.2 Pearson 2016

Although not focused on intention reports in particular, Pearson’s (2016) paper on partial control in attitude reports contains an appendix with worked out denotations for a number of attitude predicates including *expect*, *promise*, *remember*, *regret*, *be glad*, *be sorry*, and *intend*. Abstracting away from some orthogonal details that need not concern us here, the semantics that Pearson proposes for *intend* essentially boils down to the following:

$$(123) \quad \llbracket \text{intend} \rrbracket^w = \lambda p \lambda x. \forall w' [w' \text{ is compatible with what } x \text{ believes in } w \rightarrow x \text{ brings it about that } p \text{ in } w']$$

Pearson says of her proposed denotation that it captures the intuition that “if Mary intends to go to the movies, then she has the belief, ‘I will bring it about that I go to the movies’” (p. 734).

So, let’s consider how well (123) fares against the empirical properties of intention reports laid out in section 2 above. The first thing to note is that by analyzing intentions as a species of belief, (123) gets “for free” all of those empirical properties of intention reports that mirror belief reports: realism, conjunction introduction, monotonicity, and non-gradability.

If we take “ x brings it about that p ” at face value, then a shortcoming of the analysis is that it predicts that (124) should be a contradiction, contrary to fact. (124) would for example be true in a scenario in which John believes himself to be clumsy and accident-prone and so thinks that he will accidentally bring it about that he break the window.

$$(124) \quad \text{John believes he’ll bring it about that he break the window, but John does not intend to break the window.}$$

One move we could make toward remedying this shortcoming would be to replace “ x brings it about that p ” with $\text{RESP}(x,p)$:

$$(125) \quad \llbracket \text{intend} \rrbracket^w = \lambda p \lambda x. \forall w' [w' \text{ is compatible with what } x \text{ believes in } w \rightarrow \text{RESP}(x,p) \text{ in } w']$$

Then the question that needs to be asked is whether we can identify intending to act with believing that one will (intentionally) act in such a way. Typically, these notions go hand in hand, but Bratman’s (1987) influential philosophical work on intention proffers counterexamples in both directions. First, Bratman entertains a scenario where one can have a belief without the corresponding intention:

Suppose I have a fleeting craving for a chocolate bar, one which induces

a fleetingly predominant desire to eat one for dessert. And suppose that just as fleetingly I notice this desire and judge (in a spirit of resignation, perhaps) that it will lead me to so act. But then I stop and reflect, recall my dieting plans, and resolve to skip dessert. [...] I am inclined to say that I had no [...] intention [to eat a chocolate bar for dessert], for I was never appropriately settled in favor of such a dessert. (Bratman 1987:20)

Second, Bratman entertains a scenario where one can have an intention without the corresponding belief:

Perhaps I intend to carry out a rescue operation, one that requires a series of difficult steps. I am confident that at each stage I will try my best. But if I were to reflect on the matter, I would have my doubts about success. I don't have other plans or beliefs which are inconsistent with such success; I do not actually believe I will fail. But neither do I believe I will succeed. (Bratman 1987:38)

These two scenarios seem consistent with the truthful, non-contradictory status of (126) and (127), respectively. This would be difficult to reconcile with the semantics for *intend* in (125).

(126) I believed I would have a chocolate bar, but I did not intend to have a chocolate bar.

(127) I intend to carry out a rescue mission, but I do not believe I will be successful in carrying it out.

A second shortcoming of treating intending as a kind of believing has to do with the facts surrounding anankastic conditionals. In particular, the analysis leaves unexplained the asymmetry between (128) and (129): if intention reports are just beliefs about what one will (intentionally) bring about, then we would not expect *intend* to be able to trigger anankastic readings of conditionals as it does in (128). In short, by not treating *intend* as being about preferences, it becomes mysterious how the antecedent of the conditional in (128a) could affect the ordering source associated with the teleological modal in the consequent in a way that validates the paraphrase in (128b).¹⁶

¹⁶At first blush, it also looks like Pearson's account will have trouble explaining why *intend* should part ways from *believe* in taking subjunctive complements. But this is actually not a problem, since Pearson builds a causative semantics into her denotation, and causative predicates are known to trigger subjunctive mood (see e.g. Villalta 2008).

- (128) a. If you intend to get good grades, you have to study.
 b. \approx You have to study to get good grades.
- (129) a. If you believe (you will bring it about that) you will get good grades,
 you have to study.
 b. $\not\approx$ You have to study to get good grades.

7 Conclusion

This paper started out with the goal of understanding the analytical relationship that intention reports bear to other attitude reports and to other expressions that have to do with intentional action. In a nutshell, the proposal that we ended up with says that intention reports are similar to desire reports in the sense that they are built on preference structures. But unlike ordinary desire reports, intention reports exclusively target effective preference structures, thereby subjecting them to certain rationality constraints that ally them in some ways to belief reports. Furthermore, intention reports make use of the RESP-relation, just like other expressions that involve intentional action.

I will close with what I see as an important question raised by this analysis, a question that I think could fruitfully guide further investigation. Focusing just on the relationship between *want* and *intend*, one way of summarizing the core proposals of this paper is as indicated in the table in (130). The two predicates differ along three dimensions: *want* has an underspecified preference structure whereas *intend* targets effective preferences; *want* does not involve the RESP-relation whereas *intend* does; and *want* is gradable whereas *intend* is not.

		<i>want</i>	<i>intend</i>
(130)	preference structure?	underspecified	effective
	RESP-inducing?	no	yes
	gradable?	yes	no

Notably, nothing in the analysis I have presented ties any of these properties together in a fundamental way; that is, the theory allows for all logically possible combinations of these three dimensions. This yields a typology of eight kinds of possible predicates (possibly more, if there are other kinds of preference structures aside from effective ones that can be exclusively targeted). It is tempting to see this as a shortcoming of the analysis and to try derive these particular clusterings of properties from deeper principles. (And in fact I already gestured at this in section 3.2.7 when I entertained the idea that there is indeed a natural intuitive link between effective preferences and responsibility.) But I think that such a theoretical exercise would be premature, because we do not actually know *a priori* whether the other

six kinds of predicates allowed by the theory exist or not.

According to Condoravdi and Lauer (2016), when *want* is used as the embedding verb in the antecedent of a conditional, an anankastic reading of the conditional is available only when *want* targets an effective preference structure. So one way we could approach the question of possible clusterings of properties is by investigating whether, when *want* is put in this context that forces an effective preference construal, it becomes RESP-inducing and non-gradable. If it did, this would be a compelling reason for thinking that the three dimensions in (130) are deeply correlated. The contrast between (131)–(132) suggests that *want* is indeed RESP-inducing when it has the effective preference construal: it gives rise to the causation effect when an anankastic reading is salient (131), but not when an anankastic reading is not salient. (The validity of the paraphrase in (131b) depends on a very weak reading of *bring it about that: help* would be much more natural here. But the crucial point is that there is a contrast with (132b), which is intuitively not a valid paraphrase regardless of how weakly the causative is construed.)

- (131) a. If you want your daughter to get good grades, you have to tutor her.
b. \approx If you want to bring it about that your daughter gets good grades, you have to tutor her.
- (132) a. If you want your daughter to get bad grades, you have to reexamine your approach to parenting.
b. $\not\approx$ If you want to bring it about that your daughter gets bad grades, you have to reexamine your approach to parenting.

As for gradability, the following data show that at least some gradability-inducing constructions sound odd with effective-preference *want*, though the judgments are subtle and it is difficult to abstract away from potentially interfering pragmatic factors.

- (133) a. ?If you want very much to get good grades, you have to study.
b. ?If you want to get good grades more than you want to have fun, you have to study.
c. If what you want the most is to get good grades, you have to study.

So the foregoing considerations suggest that there is indeed a non-trivial relationship between effective preferences, responsibility, and (possibly) non-gradability. But on the other hand, we also identified in section 2.7.1 above the predicate *count on*, which seems to name effective preferences without being RESP-inducing. (134) shows that *count on* supports anankastic readings of conditionals; this establishes that it at least has the option of targeting effective preferences. And it seems plausible to entertain the stronger view that like *intend*, it targets effective preferences

exclusively, given that it obeys realism (135a), conjunction introduction (135b), and upward monotonicity (135c). But as witnessed in (136)–(137), *count on* does not give rise to responsibility effects. So perhaps, the connection between effective preferences and the RESP-relation is not so deep after all.

- (134) a. If you count on getting good grades, you have to study.
b. \approx You have to study in order to get good grades.
- (135) a. #John counts on turning into a unicorn, even though he knows this is impossible.
b. #John counts on staying at home tonight but he also counts on going out tonight.
c. #John counts on teaching Tuesdays and Thursdays next semester, but he doesn't count on teaching next semester.
- (136) a. John counts on Bill getting good grades.
b. $\not\approx$ John counts on bringing it about that Bill get good grades.
- (137) a. John counts on breaking the window unintentionally.
b. John counts on breaking the window intentionally.
c. John counts on going through puberty next year.

Clearly, these remarks only scratch the surface of what could be a much more detailed investigation of which clusters of properties are attested in natural language attitude predicates and which are not. I merely hope to have provided a suitable foundation for such a line of inquiry.

References

- Anand, Pranav, and Valentine Hacquard. 2013. Epistemics and attitudes. *Semantics & Pragmatics* 6:1–59.
- Asher, Nicholas. 1987. A typology for attitude verbs and their anaphoric properties. *Linguistics and Philosophy* 10:125–197.
- Barker, Chris. 2002. The dynamics of vagueness. *Linguistics and Philosophy* 25:1–36.
- Boeckx, Cedric, Norbert Hornstein, and Jairo Nunes. 2010. *Control as movement*. Cambridge: Cambridge University Press.
- Bogal-Allbritten, Elizabeth. 2016. Building meaning in Navajo. Ph.D. Dissertation, University of Massachusetts.
- Brand, Myles. 1984. *Intending and acting: Toward a naturalized action theory*. Cambridge: MIT Press.

- Bratman, Michael E. 1987. *Intentions, plans, and practical reason*. Cambridge, MA: Harvard University Press.
- Chierchia, Gennaro. 2010. Mass nouns, vagueness and semantics variation. *Synthese* 174:99–149.
- Condoravdi, Cleo, and Sven Lauer. 2011. Performative verbs and performative acts. In *Sinn und Bedeutung 15*, ed. Ingo Reich, Eva Horch, and Dennis Pauly, 149–164. Saarbrücken: Universaar.
- Condoravdi, Cleo, and Sven Lauer. 2012. Imperatives: Meaning and illocutionary force. In *Empirical issues in syntax and semantics 9*, ed. Christopher Piñon, 37–58.
- Condoravdi, Cleo, and Sven Lauer. 2016. Anankastic conditionals are just conditionals. *Semantics & Pragmatics* 9:1–61.
- Copley, Bridget. 2008. The plan's the thing: Deconstructing futurate meanings. *Linguistic Inquiry* 39:261–274.
- Copley, Bridget, and Heidi Harley. 2009. Futurates, directors, and *have*-causatives. *Snippets* 19:5–6.
- Crnič, Luka. 2011. Getting *even*. Ph.D. Dissertation, MIT.
- Culicover, Peter W., and Ray Jackendoff. 2005. *Simpler syntax*. Oxford: Oxford University Press.
- Davidson, Donald. 1973. Freedom to act. In *Essays on freedom of action*, ed. Ted Honderich, 137–156. London: Routledge and Kegan Paul.
- Davis, Wayne A. 1984. The two senses of desire. *Philosophical Studies* 45:181–195.
- Dowty, David. 1972. On the syntax and semantics of the atomic predicate CAUSE. In *Papers from the eighth regional meeting of the Chicago Linguistic Society*, ed. Paul Peranteu, Judith Levi, and Gloria Phares, 62–74. Chicago: Chicago Linguistic Society.
- Dowty, David R. 1985. On recent analyses of the semantics of control. *Linguistics and Philosophy* 8:291–331.
- Egré, Paul. 2014. Intentional action and the semantics of gradable expressions (on the Knobe Effect). In *Causation in grammatical structures*, ed. Bridget Copley and Fabienne Martin, 176–205. Oxford: Oxford University Press.
- Egré, Paul, and Florian Cova. 2015. Moral asymmetries and the semantics of *many*. *Semantics & Pragmatics* 8:1–45.
- Farkas, Donka F. 1988. On obligatory control. *Linguistics and Philosophy* 11:27–58.

- Farkas, Donka F., and Anastasia Giannakidou. 1996. How clause-bounded is the scope of universals? In *SALT VI*, ed. Teresa Galloway and Justin Spence, 35–52. Ithaca, NY: Cornell University.
- von Fintel, Kai. 1999. NPI-licensing, Strawson-entailment, and context-dependency. *Journal of Semantics* 16:97–148.
- von Fintel, Kai, and Sabine Iatridou. 2005. What to do if you want to go to Harlem: Anankastic conditionals and related matters. Ms., Massachusetts Institute of Technology.
- Giannakidou, Anastasia. 1999. Affective dependencies. *Linguistics and Philosophy* 22:367–421.
- Giannakidou, Anastasia. 2009. The dependency of the subjunctive revisited: Temporal semantics and polarity. *Lingua* 119:1883–1908.
- Giannakidou, Anastasia, and Alda Mari. 2016. A unified analysis of the future as epistemic modality: the view from Greek and Italian. Ms., University of Chicago and Institute Jean Nicod, CNRS/ENS/EHESS.
- Grano, Thomas. 2011. Mental action and event structure in the semantics of *try*. In *Semantics and linguistic theory (SALT) 21*, ed. Neil Ashton, Anca Chereches, and David Lutz, 426–443. eLanguage.
- Grano, Thomas. 2015. Getting your to-do list under control: Imperative semantics and the grammar of intending. In *Proceedings of NELS 45, vol. 1*, ed. Thuy Bui and Deniz Özyildiz, 241–252. Amherst, MA: GLSA Publications.
- Hansson, Sven Ove. 2001. *Structures of values and norms*. Cambridge: Cambridge University Press.
- Harman, Gilbert. 1976. Practical reasoning. *The Review of Metaphysics* 29:431–463.
- Heim, Irene. 1992. Presupposition projection and the semantics of attitude verbs. *Journal of Semantics* 9:183–221.
- Huitink, Janneke. 2005. Analyzing anankastic conditionals and sufficiency modals. *Student organization of linguistics in Europe (ConSOLE)* 13:135–156.
- Huitink, Janneke. 2008. Modals, conditionals and compositionality. Nijmegen: Radboud University Dissertation.
- Jackendoff, Ray. 1996. The conceptual structure of intending and volitional action. In *Evolution and revolution in linguistic theory*, ed. Héctor Campos and Paula Kempchinsky, 198–236. Georgetown: Georgetown University Press.
- Jackendoff, Ray, and Peter W. Culicover. 2003. The semantic basis of control in English. *Language* 79:517–556.
- Kamp, Hans. 1999–2007. Intentions, plans and their execution: Turning objects of thought into entities of the external world. Ms., University of Stuttgart.

- Kennedy, Christopher. 2007. Vagueness and grammar: the semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy* 30:1–45.
- Keshet, Ezra. 2008. Infinitival complements. In *Proceedings of SuB12*, ed. Atle Grønn, 303–317. Oslo: ILOS.
- Knobe, Joshua. 2003a. Intentional action and side effects in ordinary language. *Analysis* 63:190–193.
- Knobe, Joshua. 2003b. Intentional action in folk psychology: an experimental investigation. *Philosophical Psychology* 16:309–324.
- Kratzer, Angelika. 1998. Scope or pseudoscope? Are there wide-scope indefinites? In *Events and grammar*, ed. Susan Rothstein, 163–196. Berlin: Springer.
- Lassiter, Daniel. 2011. Measurement and modality. PhD dissertation, New York University.
- Levinson, Dmitry. 2003. Probabilistic model-theoretic semantics for *want*. In *Proceedings of SALT 13*, ed. R. Young and Y. Zhou, 222–239. Ithaca, NY: Cornell University.
- Ludwig, Kirk. 2007. Foundations of social reality in collective intentional behavior. In *Intentional acts and institutional facts*, ed. Savas L. Tsohatzidis, 49–71. Dordrecht: Springer.
- Martin, Fabienne, and Florian Schäfer. 2014. Causation at the syntax-semantics interface. In *Causation in grammatical structures*, ed. Bridget Copley and Fabienne Martin, 209–244. Oxford: Oxford University Press.
- Mele, Alfred R. 1987. Are intentions self-referential? *Philosophical Studies* 52:309–329.
- Merchant, Jason. 2001. *The syntax of silence: Sluicing, islands, and the theory of ellipsis*. Oxford: Oxford University Press.
- Nguyen, Elizabeth. 2013. Observations on Mandarin *xiang*. M.s., University of Connecticut.
- Ninan, Dilip. 2005. Two puzzles about deontic necessity. In *New work on modality. MIT Working Papers in Linguistics (MITWPL) 51*, ed. J. Gajewski, V. Hacquard, B. Nickel, and S. Yalcin, 149–178. Cambridge, MA: MITWPL.
- Nissenbaum, Jon. 2005. Kissing Pedro Martinez: (existential) anankastic conditionals and rationale clauses. In *The proceedings of SALT XV*, ed. E. Georgala and J. Howell, 134–151. Ithaca, NY: Cornell University.
- Pearson, Hazel. 2016. The semantics of partial control. *Natural Language & Linguistic Theory* 34:691–738.
- Perlmutter, David. 1968. Deep and surface structure constraints in syntax. PhD Dissertation, MIT.

- Pettit, Dean, and Joshua Knobe. 2009. The pervasive impact of moral judgments. *Mind and Language* 24:586–604.
- Pietroski, Paul M. 2004. *Events and semantic architecture*. New York: Oxford University Press.
- Piñango, Maria Mercedes, and Ashwini Deo. 2016. Reanalyzing the complement coercion effect through a generalized lexical semantics for aspectual verbs. *Journal of Semantics* 33:359–408.
- Portner, Paul. 2004. The semantics of imperatives within a theory of clause types. In *Proceedings of SALT 14*, ed. Robert B. Young, 235–252. Ithaca, NY: Cornell University.
- Portner, Paul. 2007. Imperatives and modals. *Natural Language Semantics* 15:351–383.
- Portner, Paul. 2015. Imperative mood. Talk given at Workshop on Nonveridical Expressions and Subjectivity in Language, University of Chicago, 11 December 2015.
- Pustejovsky, James. 1995. *The generative lexicon*. Cambridge, MA: MIT Press.
- Rooryck, Johan. 2008. Control via selection. In *New horizons in the analysis of control and raising*, ed. William D. Davies and Stanley Dubinsky, 281–292. Dordrecht: Springer.
- Rubinstein, Aynat. 2012. Root modalities and attitude predicates. Ph.D. Dissertation, University of Massachusetts Amherst.
- Sæbø, Kjell Johan. 1985. Notwendige Bedingungen im Deutschen: Zur Semantik modalisierter Sätze. Tech. rep. Universität Konstanz, Sonderforschungsbereich 99: Linguistik.
- Sæbø, Kjell Johan. 2001. Necessary conditions in a natural language. In *Audiatur vox sapientiae: A festschrift for Arnim von Stechow*, ed. Caroline Féry and Wolfgang Sternefeld, 427–449. Berlin: Akademie-Verlag.
- Sag, Ivan, and Carl Pollard. 1991. An integrated theory of complement control. *Language* 67:63–113.
- Searle, John. 1983. *Intentionality: An Essay in the Philosophy of Mind*. New York: Cambridge University Press.
- Searle, John R. 2009. *Making the social world: The structure of human civilization*. Oxford: Oxford University Press.
- Setiya, Kieran. 2014. Intention. In *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta. Spring 2014 edition.
- Sharvit, Yael. 2003. Trying to be progressive: The extensionality of *try*. *Journal of Semantics* 20:403–445.

- Spyropoulos, Vassilios. 2008. Finiteness and control in Greek. In *New horizons in the analysis of control and raising*, ed. William D. Davies and Stanley Dubinsky, 159–183. Dordrecht: Springer.
- Stalnaker, Robert. 1984. *Inquiry*. Cambridge: MIT Press.
- von Stechow, Arnim, Sveta Krasikova, and Doris Penka. 2006. Anankastic conditionals again. In *A festschrift for Kjell Johan Sæbø: In partial fulfillment of the requirements for the celebration of his 50th birthday*, ed. Torgrim Solstand, Atle Grønn, and Dag Haug, 151–171. Oslo: Forfatterne.
- Stephenson, Tamina. 2010. Control in centred worlds. *Journal of Semantics* 27:409–436.
- Thomason, Richmond H., and Robert C. Stalnaker. 1973. A semantic theory of adverbs. *Linguistic Inquiry* 4:195–220.
- Vermazen, Bruce. 1993. Objections of intention. *Philosophical Studies* 71:223–265.
- Villalta, Elisabeth. 2008. Mood and gradability: An investigation of the subjunctive mood in Spanish. *Linguistics and Philosophy* 31:467–522.
- Williams, Alexander. 2015. *Arguments in syntax and semantics*. Cambridge: Cambridge University Press.
- Wolff, Phillip. 2003. Direct causation in the linguistic coding and individuation of causal events. *Cognition* 88:1–48.
- Wurmbrand, Susi. 2014. Tense and aspect in English infinitives. *Linguistic Inquiry* 45:403–447.
- Zanuttini, Raffaella, Miok Pak, and Paul Portner. 2012. A syntactic analysis of interpretive restrictions on imperative, promissive, and exhortative subjects. *Natural Language & Linguistic Theory* 30:1231–1274.
- Zwicky, Arnold M., and Jerrold M. Sadock. 1975. Ambiguity tests and how to fail them. In *Syntax and semantics 4*, ed. John Kimball, 1–36. New York: Academic Press.