

Detaching *If*-Clauses from *Should*¹

Ana Arregui
University of Ottawa

1. Introduction

My objective in this paper is to investigate some aspects of the semantics of deontic *should*-conditionals. My particular interest is to understand what are the actual world facts that make deontic statements true. The starting point in my investigation is a famous puzzle presented by Chisholm in a 1963 paper (*Chisholm's Paradox*). Chisholm put forward a set of English sentences that he claimed were consistent and independent of each other. He then went on to show that the resources made available by a standard system of deontic logic were insufficient for capturing the intuitions regarding the sentences. Chisholm's Paradox has been much discussed in the logical and philosophical literature dealing with deontic modality (for an overview, see Åqvist 2002, Carmó and Jones 2002). It will be important in this paper because reconciling the data in Chisholm-style examples involves arriving to some conclusion regarding the interaction between what we consider ideal and what is actually true. My proposal in this paper will be an account of how facts affect the evaluation of *should* formulated in a way that does not predict that the divergence between ideals and facts leads to contradictions.

Chisholm's examples made use of the modal *ought to*. However, in this paper we will make use of *should* instead. The relevant behaviour of *ought to* can be reproduced with *should*, and examples with *should* can at times sound more familiar.² A *should*-version of Chisholm-sentences is provided in (1):

- (1)
- a. She should return the library book on time.
 - b. If she returns the library book late, she should pay a fine.
 - c. It should be the case that, if she returns the library book on time, she does not pay a fine.
 - d. She returns the library book late.

Chisholm showed that it was not possible to give the sentences in (1) a consistent and independent characterization with the resources of standard deontic logic (a modal logic with a monadic necessity operator: SDL). Chisholm suggested the following formalization (where *O* is a monadic necessity operator): (1a) *O book-on-time*, (1b) $\neg \textit{book-on-time} \rightarrow O \textit{pay-a-fine}$, (1c) $O(\textit{book-on-time} \rightarrow \neg \textit{pay-a-fine})$, (1d) $\neg \textit{book-on-time}$. With this characterization, the

¹ This paper is undergoing revisions. I am grateful for comments I have received from an anonymous reviewer.

² A comparison of differences between deontic modals lies outside the scope of this paper. The limitation of the study to the case of *should* is a disadvantage of the current work that must be addressed in future research. A discussion of differences regarding the relative strength of deontic modals can be found in Copley (2006) and von Fintel and Iatridou (2008) (who also presents cross-linguistic data).

sentences are predicted to be inconsistent with each other: from (1b) and (1d) we will derive $O \text{ pay-a-fine}$, and from (1a) and (1c) we will derive $O \neg \text{pay-a-fine}$ (since O will ‘distribute’ in (1c), giving rise to the inference that $O \text{ book-on-time} \rightarrow O \neg \text{pay-a-fine}$, and so with (1a) we will be able to derive $O \neg \text{pay-a-fine}$). Since O is a universal quantifier over possible worlds, $O \text{ pay-a-fine} \ \& \ O \neg \text{pay-a-fine}$ is a contradiction. (For alternative characterizations and discussions of consistency and independence, see Carmó and Jones (2002)).

The key player in debates of Chisholm’s Paradox has been a type of deontic conditional that Chisholm called *Contrary to Duty Imperatives* (CTDs). The conditional in (1b) above is a CTD. It is a conditional that spells out a ‘secondary duty’ (to pay a fine), a duty that arises in less-than-ideal circumstances in which a primary duty has been overlooked (to return the library book on time). The CTD in (1b) has a crucial role in deriving the contradiction: the SDL-characterization of (1b) together with the truth of the antecedent (1d) predicts the inference that *she should pay a fine* ($O \text{ pay-a-fine}$). If we allow contingent truths about the world (like (1d)) to trigger inferences in CTDs in this manner, we will derive obligations corresponding to less-than-ideal circumstances and this will contribute to the SDL-contradiction discussed above.

It will be useful to introduce some terminology. We will use the term *detachment* as a cover term for inference patterns in conditionals. We will distinguish between *factual detachment* and *deontic detachment*.³ *Factual detachment* is the pattern that holds when it is possible to draw an inference from a conditional of the form *if α , obligatory β* and the premise α to the conclusion *obligatory β* . *Deontic detachment* is the pattern that holds when it is possible to draw an inference from a conditional of the form *if α , obligatory β* and the premise *obligatory α* to the conclusion *obligatory β* . Here, *obligatory* stands in for a deontic necessity modal (in this paper, *should*). If we examine the SDL-characterization of (1) provided by Chisholm, we see that we have made use of both factual detachment (between (1b) and (1d)) and deontic detachment (between (1a) and (1c)) in deriving the contradiction. In an SDL-characterization of deontic conditionals, both are predicted to be valid.

There is something intuitive about the idea that *should*-conditionals license factual detachment.⁴ This seems particularly clear in the case of CTDs. Chisholm singled out CTDs as very important because *most of us need a way of deciding, not only what we ought to do, but also what we ought to do after we fail to do some of the things we ought to do* (Chisholm 1963: 35-36). It is hard to see how CTDs could fulfil this role without factual detachment. If the truth of the antecedent does not lead to unconditioned *shoulds*, how could deontic conditionals ever tell us what we should do when things have gone wrong? Consider the example in (2):

- (2) If she returns the library book late, she should pay a fine.

The CTD conditional in (2) makes a hypothesis that corresponds to less-than-ideal circumstances in which she has violated a library rule, and tells us what is best given those circumstances. Our intuitions seem to favor an interpretation that validates factual detachment. If we accept that (2) is true and we find out that she returned the book late, we will conclude that she should pay a fine.

The intuitions regarding deontic detachment, on the other hand, seem to go the other way. Imagine that library regulations not only spell out the duties of library patrons,

³ See Feldman (1986): 74 for terminology

⁴ This is briefly mentioned in Åqvist (2002), who also points to other references.

but also indicate that good behaviour should be rewarded. If this were the case, the sentences in (3) could both be true:

- (3) a. She should return the library book on time.
 b. If she returns the library book on time, she should receive a reward.

But even if we believe both sentences, we'll hesitate to claim outright that she should receive a reward. The truth of *she should receive a reward* depends on what actually happens with the book. It may be that *ideally* she returns the book on time and receives a reward, but the truth of (3a) is not enough to detach an unconditional *should*-statement from (3b).⁵

Following up on the brief discussion of (2) and (3), it may be tempting to conclude that in characterizing deontic statements, we ought to make a proposal that allows *should*-statements and conditionals to validate factual detachment and block deontic detachment. This would amount to requiring that the truth-conditions of *should* pay close attention to what is actually going on, so as to make sure that facts can trigger factual detachment and that facts can block deontic detachment. Yet, as we will see below (briefly), some of the most influential accounts of deontic conditionals in philosophy and logic bypass Chisholm's Paradox with dyadic deontic operators that fail to validate factual detachment and do validate deontic detachment (see Carmó and Jones 2002).⁶

Dyadic operators are operators that take two propositions as arguments. An example of a dyadic operator analysis of deontic conditionals can be found in Lewis (1973). Lewis (1973) was mainly dedicated to an account of *counterfactual* conditionals that made use of a system of spheres to encode contextually determined similarity. But the proposal for counterfactuals was also re-interpreted to provide an account of deontic conditionals. The system of spheres could be used to encode the relative value or 'goodness' of worlds, with worlds in the inner spheres being considered better than worlds in spheres further away. With a deontic interpretation, a conditional of the form *if α , Cond-Op β* came out true in a world iff the best worlds in which α was true were also worlds in which β was true (where context could define the measure of goodness and set up the system of spheres).

Dyadic operators such as Lewis's do not give rise to Chisholm's Paradox because they do not validate factual detachment and so do not license the problematic inference we saw in (1) between (1b) and (1d). With a dyadic-operator analysis, a conditional *if α , Cond-Op β* claims that the best α -worlds are also β -worlds. This means that the actual truth of α will not allow us to draw any inferences regarding the consequent. Factual detachment is not valid because the actual world, while an α -world, may not be amongst the best α -worlds, and so may be excluded from the domain of quantification of the conditional operator. Interestingly, dyadic operators do validate deontic detachment. If it is the case that in the

⁵ Intuitions regarding deontic detachment could be confounded by the possibilities made available by *modal subordination*. That is, if we have two modal statements following each other, at times it is possible to understand the second one as taking as a starting point the possibilities predicated over by the first. Here this would amount to an interpretation of *She should receive a reward* evaluated in worlds in which she has returned the book on time. A modal subordination interpretation is not the interpretation that interests us here.

⁶ For reasons of space, I will not attempt a full comparison between the proposal in this paper and other proposals available in the literature. I will briefly mention Lewis (1973), but I will leave other proposals aside. In particular, I will not be able to offer a comparison between the ideas presented here and those in Kratzer (1981, etc.).

best α -worlds, β is true and the best worlds are α -worlds, then the best worlds will be β -worlds (I present the patterns informally here, but a proof can be found in Acqvist 2002).

The debate regarding detachment patterns is a debate about the meaning of deontic modals and conditionals (which we have restricted here to the case of *should*). Our preliminary discussion of examples (2) and (3) supported a view that validated factual detachment and blocked deontic detachment. Such a view would tie the meaning of the modal and conditional very closely to what is going on. However, as we have just noted, an influential line of thought has resolved Chisholm's Paradox by blocking factual detachment (while allowing deontic detachment). Ultimately, the goal of this paper will be to look for an account that validates factual detachment (and blocks deontic detachment). However, to understand the problem properly, it is important to give the dyadic-operator approach its due and consider more complex examples that have been used to argue against factual detachment. We will examine an example from Feldman (1986):

Feldman (1986) presents a study in moral obligation, a type of deontic necessity that corresponds to our obligation to 'do the best we can'. According to Feldman, there are examples appealing to moral obligation that show that factual detachment is not valid. Consider the following scenario: A doctor is in charge of administering medicine to a patient. The best results would be obtained if the doctor gave medicine A on Monday and medicine A on Tuesday. However, if the doctor gave medicine B on Monday, the best follow up would be to give B on Tuesday (the mixture of medicines would be fatal). Suppose moreover that any other possibility regarding medication would be fatal, and suppose that, as a matter of fact the doctor will give medicine B on Monday. It is now Sunday, and we are faced with the sentences below:

- (4)
- a. The doctor should give medicine A on Monday and A on Tuesday.
 - b. If the doctor gives medicine B on Monday, he should give B on Tuesday.
 - c. The doctor will give medicine B on Monday.

According to Feldman, all sentences in (4) are true on Sunday. Even if, as a matter of fact, the doctor will give medicine B on Monday (maybe on purpose, maybe by mistake), on Sunday it is true that the doctor should give A on Monday and A on Tuesday (the doctor is morally obliged to do so). By characterizing the conditional in (4b) with a dyadic deontic operator, Feldman blocks factual detachment and the sentences in (4) do not lead to a contradiction. An account of deontic conditionals that validates factual detachment will have to provide a different perspective on this data.

This brief discussion sets the main parameters of the debate. The paper will investigate the meaning of *should*-statements and conditionals, setting up an account that allows for factual detachment without deriving contradictions. The view of deontic modality adopted here will be a 'bare bones' view, that will examine a general kind of 'ought-to-be' modality without distinguishing subtypes. As the examples in (1) and (4) illustrate, the challenge is to correctly explain the interaction between what is going on (the facts) and what is best (the ideals). The account I will present has three main ingredients: a proposal for the meaning of *should* (presented in §3); a proposal for the interpretation of clauses embedded under *should* (presented in §4); and an account of the strategies available for attaching *if*-clauses and detachment patterns (presented in §5). The next section (§2) will set the stage for what follows by presenting and discussing an account of deontic modality based on

manipulations of the temporal parameter of an accessibility relation. We will return to Chisholm's Paradox in the conclusion (§6).

The remainder of the introduction will be dedicated to clarifying the scope of the investigation and pointing to limitations of the current work. The semantics of deontic modals is usually stated in terms of quantification over possibilities, and the proposal here will follow this line. As with all such proposals, the key lies in figuring out what gets into the quantificational domain of the modal. The case of deontic modals is complex because on the one hand, deontic statements are not meant to describe what is going on, and so it is quite expected to find in the quantificational domain possibilities that are known to be different from the actual one. On the other hand, the truth of deontic statements is clearly dependent on what is actually going on, and so what is going on affects the quantificational domain of the modal. To correctly capture deontic modality, judgments about what is best (our values) have to be reconciled with what is going on (the facts). In this paper it will be argued that the key to making deontic modals appropriately sensitive to what is going on lies in a framework that allows us to make reference to parts of what is going on in the actual world (situations smaller than worlds), and not only to worlds themselves. We will show that a semantics for deontic modals that makes use of a situations framework as developed by Kratzer (1989, 2006) allows us to make distinctions that are fine-grained enough to correctly capture the quantificational domain of deontic modals, making them sensitive to facts in a manner that does not lead to contradictions.

A correct account of modality needs to explain the interaction of information from various sources. This is well known since Kratzer's 'bi-dimensional' proposal for the semantics of modality, in which a modal base interacts with an ordering source (Kratzer 1981, 1991). The proposal in this paper derives variations in the interpretation of *should*-statements from the interaction of a unified meaning for *should* and the presuppositions carried by aspectual heads in clauses embedded under *should*. The goal is to articulate the different restrictions on the truth conditions of modal claims compositionally, teasing apart the contributions made by aspect from those made by the modal itself. One of the claims made here is that aspect plays an important role in relating modal claims to what is going on. The claim regarding the 'modal relevance' of aspect follows Arregui (2004, 2007), and is in line with other recent research that has investigated the interaction between aspect and modality (e.g. Bhatt 2006 Ippolito 2008, Haquard 2006). By separating the contribution made by aspect from that made by the modal, the proposal presented here takes some steps towards providing an account of the syntactic articulation of the different components of deontic modality.

Two clarifications are in order. I have adopted a simplified view of the syntactic structure of *should*-statements and conditionals, making use of tripartite structures for conditionals without spelling-out details regarding the dynamic composition of meaning. An investigation into the interpretation of conditional statements must (eventually) address the context-change effect of conditionals. However, the research in this paper can be evaluated independently of any particular implementation of a dynamic system, so the simplification is hopefully justified. I would also like to note that the paper concerns itself exclusively with the truth conditions of deontic statements. Deontic statements clearly have a performative aspect, but nothing will be said about that here. The tension between the performative and the truth-conditional is a classic problem in the domain of deontic modality. In a paper dealing with the logic of normative statements and law, Alchourrón (1993) noted the 'systematic ambiguity' of deontic statements: *Sometimes they are used to express normative propositions designed to describe the consequences of the existence or non-existence of norms. Normative*

propositions do have truth values, and their truth value depends on normative actions (promulgation, derogation, etc.) of the authorities competent to create or modify the law. But on many other occasions, particularly when the sentences are used (for example by the authorities) to rule other people's behavior, they express norms which lack truth values (...). The proposal in this paper follows Alchourrón in distinguishing the truth-conditional content of deontic claims from their performative aspect (though we would prefer to remain neutral regarding the matter of ambiguity). At times, I will speculate that some of our intuitions regarding the oddness of certain statements should be linked to the performative instead of the truth-conditional. However, without a theory of the performative impact of *should*-statements, such claims can only be preliminary speculations.

2. A *time*-sensitive analysis

Deontic *should*-claims encode our judgments regarding what is best, and these appear to change through time. Consider the examples in (5):

- (5) a. She forgot to return her library book on time. #But she should return it on time.
 b. She forgot to return her library book on time. But she should have returned it on time.

A simple *should*-statement is not really possible in the context provided by (5a). This could be taken to indicate that after the book's due date has passed, it is no longer true that she should return the library book on time. At this point, what is true is that she should have returned the library book on time (5b).

In evaluating Chisholm's examples or Feldman's examples, we might be tempted to conclude that there is a temporal key to the puzzles. Take, for example, Feldman's sentences in (4). We may want to claim that on Sunday, it is true that she should take medicine A on Tuesday, but on Tuesday it is true that she should take medicine B. One way to explain why our judgments regarding what is best appear to change through time would be to propose that what is best is established relative to part of the history of the world (maybe the history of the world up to a certain time). Temporal variations in our judgements can then be explained as variations with respect to the pieces of the history of the world that are being taken into account. What is best relative to a particular stretch may be different from what is best relative to a larger stretch.

In this section we will investigate a notion of deontic necessity that takes time seriously and pays attention to temporal anchoring. In §2.1 we will spell out the idea that in evaluating *should*-statements, we take into account the history of the world up to the present moment, and in evaluating *should have* statements we take into account the history of the world up to some past moment. The proposal will be discussed in §2.2, where it will be argued that the difference between *should* and *should have* is not purely temporal. This will set the stage for the situations-based proposal in §3.

2.1 A proposal

Feldman (1986) characterized a general notion of deontic necessity that he termed *ought-to-be* (OB). Feldman (like Lewis) set up a system that compared the relative goodness of worlds. He made use of a function IV that delivered the intrinsic value of each world (for whatever measure of value was relevant). Feldman differed from Lewis in being explicit about the temporal anchoring of deontic necessity. In figuring out what ought to be at a time t , we consider all the worlds that are still physically possible at t , and see what happens in the best of them (in a sense, t functions as the ‘branching time’ in the world’s history). If α occurs in one of those worlds, and $\neg\alpha$ doesn’t occur in any as good or better, then at t it ought to be that α . The world accessibility relation (WA) is roughly characterized as follows: for all worlds w , w' and times t , $WA(w')(w)(t) = 1$ iff as of t , w' has not been ruled out by what has physically happened in w . The relation of world accessibility (WA) is temporally sensitive. As time goes by, the set of worlds accessible from the actual world becomes smaller. Less and less is possible.

Let us begin by adapting Feldman’s ideas regarding OB to propose a semantics for *should* as in (6):⁷

$$(6) \quad \llbracket \text{should} \rrbracket = \\ \lambda p_{\langle s, t \rangle} \lambda t_i \lambda w_s \exists w'_s (WA(w')(w)(t) = 1 \ \& \ p(w') = 1 \\ \& \ \neg \exists w'' (WA(w'')(w)(t) = 1 \ \& \ IV(w'') \geq IV(w') \ \& \ p(w'') = 0))$$

According to (6), *should* combines with a proposition p and the output is a property of times true in a time t and a world w iff there is a world w' accessible to w at t and p is true in w' and there isn’t a world w'' accessible to w at t in which p is false that is better than w' . In a simple *should* α statement, in which there is no past marker, the temporal parameter associated with *should* could be provided by a silent speech time pronoun s^* ($\llbracket s^* \rrbracket = s^* =$ the speech-time), a sort of silent present tense morpheme. The logical form of simple *shoulds* could look like (7a), with the interpretation in (7b):

$$(7) \quad \begin{array}{l} \text{a.} \quad \llbracket s^* [\text{should } \alpha] \rrbracket \\ \text{b.} \quad \lambda w \exists w' (WA(w')(w)(s^*) = 1 \ \& \ \alpha(w') = 1 \\ \quad \& \ \neg \exists w'' (WA(w'')(w)(s^*) = 1 \ \& \ IV(w'') \geq IV(w') \ \& \ \alpha(w'') = 0)) \end{array}$$

According to (7), *should* α is true iff there is a world w accessible to the actual world at the speech time in which α is true and there isn’t a world accessible at the speech time in which α is false that is better than w (notice that if α is ‘already’ true in the past of the actual world, then the statement *should* α will be trivially true – Feldman notes this as a property of his proposal).

Having an explicit temporal parameter for world accessibility, it is straightforward to characterize *should have* as a past form of *should*. The idea that in modal constructions with *have*, *have* functions as a temporal operator taking scope over the modal has a long tradition in literature dealing with modality (a.o. Condoravdi (2002), Ippolito (2003), etc.). We end up with a view according to which the contrast between the two forms is a tense contrast. A concrete proposal for *have* as a past operator is provided in (8) (for simplicity, I have anchored it directly to the speech time s^*):

⁷ For other proposals that have relativized world accessibility to a temporal parameter in a different context, see a.o. Condoravdi (2002), Ippolito (2003).

$$(8) \quad \llbracket \text{have} \rrbracket = \lambda P_{\langle i, t \rangle} \lambda w_s \exists t_i (t \leq_{\text{precedes}} s^* \& P(t)(w) = 1)$$

According to (8), *have* is a past operator that combines with a property of times and shifts its evaluation to some time before the speech time. A *should have* α statement will have the LF and denotation in (9):

$$(9) \quad \begin{array}{l} \text{a.} \quad \llbracket \text{have} [\text{should } \alpha] \rrbracket \\ \text{b.} \quad \lambda w_s \exists t_i (t \leq_{\text{precedes}} s^* \& \exists w' (WA(w')(w)(t) = 1 \& \alpha(w') = 1 \\ \& \neg \exists w'' (WA(w'')(w)(t) = 1 \& IV(w'') \geq IV(w') \& \alpha(w'') = 0)) \end{array}$$

Given (9), *should have* α will be true in the actual world if there is a world w' and there is a time t that precedes the speech time and w' is accessible to the actual world at t and α is true in w' and there isn't a world w'' that is better than w' that is accessible to the actual world at t in which α is false.

2.2 Discussion

The proposal presented in §2.1 follows Feldman in providing a temporal anchor for deontic necessity. The temporal argument associated with the modal corresponds to the temporal parameter of the accessibility relation and can be manipulated by either a silent present tense pronoun or by past *have*. Simple *shoulds* and *should haves* differ with respect to the stretch of the history of the world that affects their truth-value. However, as we will see in this section, there are examples that cast doubt on a purely temporal characterization of the contrast between simple *shoulds* and *should haves*. We will show that times do not provide adequate anchoring points for deontic necessity.

2.2.1 Some examples

The proposal in (6/8) seems to make sense given the contrast between examples like (10a) and (10b):

- (10) a. She forgot to return her library book on time. But she should have returned it on time.
 b. She forgot to return her library book on time. #But she should return it on time.

We can explain the oddness of (10b) by saying that, in the absence of *have*, the temporal parameter for the accessibility relation is the speech time. At the speech time, there aren't physically accessible worlds in which she returns the book on time, the domain of quantification is empty, and the claim is infelicitous. In (10a), on the other hand, *have* shifts the temporal parameter of the accessibility relation to a past time, and there is a point in the past such that in the best worlds physically accessible at that time, she returns the library book on time (such a time will be found, presumably, before the book's due date).

The proposal in (6/8) will also be successful with more complex examples. Suppose library regulations require that she return the book on a specific date, next Monday, but she already did it yesterday. The time at which she was supposed to return the book lies in the future, but her actions in the past have already made such worlds inaccessible. The proposal in (6/8) correctly predicts that a *should have* form will be chosen (even though the time at which she should have fulfilled her obligation lies in the future):

- (11) a. She returned her library book yesterday. #She should return it next Monday.
 b. She returned her library book yesterday. She should have returned it next Monday.

The proposal in (6/8) predicts that (11a) is odd because there isn't a world physically accessible from the speech time in which she returns the library book next Monday. In the past, however, there were such worlds, and indeed they were the best ones. So (11b) is correctly predicted to be true.

2.2.2 Accessibility is not shifted to the past

I will present two examples to illustrate the difficulties faced by a times-based analysis of the difference between *should* and *should have*. Consider a situation in which Jack, one of three candidates, has been elected as mayor after tough elections. He was elected by a large majority, but some small procedural rules regarding the timing of the elections were ignored:

- (12) Jack should have been elected three days after the nominations were announced.

We could judge (12) true even if there isn't any time in the past in which it was true that that Jack should be elected three days after the nominations are announced. It might well be that Jack is not a better candidate than Jim or John, and the worlds in which he is elected are not better than the worlds in which one of the other candidates is elected. When we judge (12) true, we take for granted that Jack was elected, and worry only about what would have been best regarding the date. We won't obtain this result by treating *should have* as a past *should*.

Another example making the same point. There is a military parade in front of the governor's house, and it is being shown on TV. A coin is tossed and one of the soldiers is randomly chosen to be shown in a close-up on TV.⁸ Looking at him, someone could utter (13):

- (13) He should have shaved.

Suppose we judge (13) true. This will be the case even if there is no point in the past at which it is the case that he should shave. For imagine we travel backwards in time till early morning, when the soldier was deciding whether to shave or not. It is not the case that in the best worlds accessible at that time he shaves. He was randomly chosen to be on TV. Worlds in which he does not shave and is not chosen will be just as good as worlds in which he shaves and is chosen. As in the previous example, in our judgment of (13) we take for

⁸ This example is inspired by Tichy's tossing of the coin examples in counterfactuals.

granted something that happens later (in this case, the outcome of the tossing of the coin), and this is not predicted by an analysis of *should have* as past *should*.

The examples above illustrate that the role of *have* in *should have* is not to shift the temporal parameter of the accessibility relation to the past. The examples show that when evaluating *should have*, we consider worlds that are like the evaluation world at times that follow the presumed past temporal parameter for the accessibility relation. This is not predicted by a purely temporal view of the contrast.

2.2.3 The importance of the context set

In discussing the difference between indicative and subjunctive conditionals, Stalnaker (1975) suggested that the role of subjunctive morphology in English was to indicate that presuppositions are being suspended, and that quantification is taking place over worlds that may be outside the context set. I would like to follow this intuition to establish a difference in the case of deontic conditionals, and claim that the contrast between *should* and *should have* is best understood by including a reference to the context set.

There is clearly a temporal component to the contrast between simple *should* and *should have* (which, I will argue, is due to differences in the interpretation of aspect). But the contrast cannot be reduced to a temporal matter. When a speaker chooses a simple *should*, two types of information are brought together: (a) that the temporal location of the embedded clause eventuality is non-past, and (b) that we are choosing amongst worlds in the context set (i.e. when a speaker claims *should* α , she is claiming that in the best worlds in the context set, α is true). The proposal follows the spirit of Stalnaker's observation by arguing that in some cases we choose one form over the other to indicate either that we are remaining in or moving out of the context set.

Examples that show that the choice between *should* and *should have* is not purely temporal can be constructed following the pattern in (14). Imagine two software developers discussing their new iPhone design:

- (14) John: I've had a great idea. We should include a breakfast application in the new iPhone!
 Jane: Haven't you heard? We'll be laid off at the end of the month and the whole project will be cancelled!!
 John: Oh no! That is dreadful! ... But isn't it a great idea? We should have included a breakfast application in the new iPhone!

In (14), John switches from a simple *should* to *should have* as an acknowledgment of the information provided by Jane. In switching, John acts as if he accepts what Jane says as true. Notice that the temporal semantics in (6/8) does not predict that a switch should be necessary. In the best worlds accessible from the actual world at the speech time, the developers are not laid-off, the project goes through (and there is a breakfast appliance in the new iPhone).


3. Building *should*

In making statements about what is best, deontic modals compare the intrinsic values of different possibilities. To understand the meanings of such modals, it is necessary to know which possibilities enter the domain of comparison. In this section it will be argued that *should* identifies the possibilities to be compared by paying attention to what has happened in the actual (evaluation) world. The actual world facts⁹ that *make should*-statements true are the facts that affect the domain of comparison of the modal. In this section we will develop a proposal for the meaning of *should* that follows this intuition. The proposal will be built in steps, focusing on the interaction between the facts in the actual world and the possibilities compared by the modal.

3.1 Whole slices?

The proposal in §2.1 puts constraints on the temporal parameters in the accessibility relation, and the outcome is to give the modal access to a temporal part (slice) of the history of the world. We can think of the temporal parameter in the accessibility relation as giving us the ‘rightmost’ edge of the relevant slice. In the case of *should*, it is the slice corresponding to history up to the speech time; in the case of *should have*, it is a past slice. In both cases, a whole slice is brought into play (we cannot ignore parts of it) and in both cases the history of the world that follows the relevant slice is not expected to affect truth conditions.

In this section we will examine both of these assumptions, with the goal of motivating the proposal to be presented later on. We have already seen examples that suggest that the assumptions are at least partly wrong. In discussing examples (12) and (13), we noted that our judgements about the truth of *should* statements was affected by what had happened at times later than the ‘accessibility’ time. Let us focus on (13). If we had to picture what happened (simplistically!) we might come up with a schema like (15):



I have used the diagrams above to illustrate the informal intuition that different ‘sequences of events’ interact with each other in the history of the world. When we make the judgment that the soldier should have shaved, we are claiming that given the fact that he will appear on TV, he should have shaved. We consider that a departure from the actual facts regarding his shaving history would be better given the actual facts regarding the outcome of the tossing of the coin (that sequence is taken for granted). A view that evaluates modal claims with

⁹ I am using the world *fact* informally to talk about what happens in the world, the things that ‘occur’.

respect to whole slices will have difficulties with examples like this, in which we seem to ‘carve up’ the world in a manner that does not respect the partitioning corresponding to times.

Our intuitions regarding detachment patterns also suggest that in evaluating modal claims we take into account parts of the history of the world that do not correspond to temporal slices. To see this, let us begin with a ‘temporal’ example involving CTDs, and examine it in light of the proposal in (6/8). Consider the following conversation between library assistants:

- (16) A: Sara should return the book on time.
 B: And if she doesn’t?
 A: If she doesn’t return the book on time, she should pay a fine.

Suppose that Sara has borrowed a book, and returned it late, and refused to pay a fine. Given this history, the proposal in (6/8) predicts that both statements in (17) are true:

- (17) a. She should have returned the book on time.
 b. She should have paid a fine.

There is a time in the past, such that in the best worlds (physically) accessible to the actual world at that time, she returns the book on time (this will be a time before the book’s due date). This makes (17a) true. There is also a time in the past such that in the best worlds (physically) accessible to the actual world at that time, she pays a fine (a time after the book’s due date and before the due date of the fine). This makes (17b) true (we will come back to these examples in §5.4).

But consider now the following scenario. A parent leaves the baby at home with the baby-sitter and strict instructions (including a CTD back-up plan):¹⁰

- (18) a. The baby should be asleep by 9pm.
 b. If she is awake (*at 9pm*), she should be watching *Baby Einstein*.

The baby was asleep at 8.30 but woke up with the sound of the TV. The baby-sitter was watching *The Simpsons*.¹¹ The baby was instantaneously wide awake, and stayed up watching it with her. In this scenario, both of the following could be judged true:

- (19) a. The baby should have been asleep (*at 9pm*).
 b. The baby should have been watching *Baby Einstein* (*at 9pm*).

Imagine an exasperated parent who the following day finds out that at 9pm, the baby had been wide awake and watching *The Simpsons*. She could burst out with *She should have been watching Baby Einstein!* She would feel this was true. However, we don’t have access to this intuition if we only evaluate *should*-claims with respect to slices of time. The time at which the baby woke up and the time at which she began watching *The Simpsons* is the same. We won’t be able to find a past branching time that precedes the watching of *The Simpsons*

¹⁰ Arguments against a temporal account of CTD conditionals can also be found in Prakken and Sergot (1996), who discuss many examples.

¹¹ *Baby Einstein* and *The Simpsons* are TV programs.

without also preceding the time of waking up. If we go back to a time before she watches *The Simpsons*, the best worlds from then on are worlds in which she remains asleep. We cannot account for the intuition that in certain circumstances, the parent could consider that (19b) is true. To account for that intuition, we would need to keep separate the facts regarding the baby being awake from the facts of the baby watching *The Simpsons*. But these parts are not spread out in time. We cannot ‘carve out’ these parts following temporal lines. A proposal that only gives the modal access to temporal slices, and not to smaller parts, won’t be sensitive to the actual world features that make (19b) true. The conclusion is that if we have access only to temporal parts, we will be able to explain our intuitions regarding detachment in examples like (17), where the relevant features differ in terms of their temporal location, but not in examples like (19), where we recognize different features and follow up on their consequences, but there is temporal overlap.

3.2 Quantifying over situations

So far I have mentioned parts of the history of the world informally, and argued that in order to capture out intuitions regarding *should* statements, we need a framework that gives us access to parts in a more fine-grained manner than is possible with a framework dealing only with temporal slices. In talking about parts in this paper, I have in mind the situations framework developed in Kratzer (1989), which I will present here and throughout the following sections (informally). Kratzer uses the term *situation* to refer to parts of the histories of worlds. The parts can be very small (according to Kratzer, situations can consist, e.g., of merely a thin particulate and a property) or very big. Situations can be parts of other situations. The ‘maximal situations’ are called *worlds* (they are situations that are not proper parts of any other situations).

The proposal in (6) built on the idea that the truth conditions for *should* evaluated what was best given a part of the history of the world that was a temporal slice. What would happen if instead of taking into account whole slices, we allowed the modal to see all situations (big and small)? To try to see clearly what is at stake, we will build the proposal up in steps. Let us begin by imagining that she has returned the library book late. We could say:

(20) She should pay a fine.

In a sense, her tardiness in returning the book is responsible for the truth of (20). The situation corresponding to the late return of the book appears to be what makes (20) true (at first sight). Following this idea, we might reason as follows: (20) is true because there exists a situation in the world (a ‘trigger’ situation: her returning the book late) such that, in its best follow-ups, she pays a fine. We’ll call this the *Existential Hypothesis*.

In order to spell out the Existential Hypothesis, we need to clarify what it means to talk about ‘the best follow ups’. To do this, I will introduce an auxiliary definition to allow us to make use of ‘best extensions’. Roughly speaking, the best-extensions of a situation will be the worlds that contain that situation which are amongst the best worlds to do so.¹² But there is a hitch. The situations framework developed by Kratzer adopts the Lewisian view that individuals are part of at most one world (we will not go into the motivations for this


¹² I will talk as if it is possible to identify the best worlds, given a particular criteria for goodness. In talking, I will set aside issues related to the limit assumption in the domain of deontic conditionals.

move here). For any situation s , there is at most one world w such that s is part of w (the part-of relation is a reflexive relation indicated with \leq). If we are talking about an actual world situation (the actual world situation of her returning the book late), we cannot go on to say simply that that situation is part of (\leq) another world. That situation exists only in the actual world. To talk about situations across worlds, we will follow the standard Lewisian strategy of appealing to counterparts. Talk about an actual world situation being part of another world should be understood as an informal claim that the actual world situation has a counterpart in another world, where counterparts are established via contextually determined similarity relations (there is a vast literature on counterparts, for some discussion of counterparts in frameworks with situations, see e.g. Percus (2000), Kratzer (2002), Arregui (2008)).

With these clarifications in place, we can now introduce ‘best-extensions’ in the following way:

- (21) *modal part of*: a situation s is a modal part of (\leq_m) a situation s' iff there exists a situation s'' such that s'' is a counterpart of s and $s'' \leq s'$
- (22) *best-extension*: a world w is a best-extension of a situation s iff $s \leq_m w$ and there isn't another world w' such that $s \leq_m w'$ and w' is better than w (for a contextually given measure of goodness).

We will understand the ‘better than’ relation in terms of Feldman’s proposal to use a function IV to measure the intrinsic value of a possible world for some relevant measure of goodness (Feldman 1986). Given two worlds w and w' , w' is better than w iff $IV(w') > IV(w)$. To see how (22) works, consider the schema below:



Suppose that s_2 is a situation in the actual world (w_0). The world w_1 will be a *best extension* of s_2 iff it includes a situation (let's say s_3) that is a counterpart of s_2 , and there isn't any other world that includes a situation that is a counterpart of s_2 that has a higher intrinsic value than w_1 . Notice that the definition in (22) allows for ties in goodness, and a single situation may have several *best-extensions*. Notice also that the definition in (22) does not have temporal directionality built in (more on this later).

With this machinery in place, we can now tackle the Existential Hypothesis more precisely. Here it is:

- (24) Existential Hypothesis (preliminary)
 $\llbracket \text{should } \alpha \rrbracket$ is true in a world w iff
 $\exists s. s \leq w$ and $\forall w': w'$ is a best-extension of $s, \alpha(w') = 1$

The big difference between (24) and (6) (if we were to rephrase it in terms of situations) is that (6) gives us a specific situation (in (7), for example, the history of the world up to the

speech time) and only temporal slices are under consideration, whereas (24) can see much smaller pieces. With (24), we could arrive at the following truth conditions for (20):

- (25) **||should [she pay a fine]||** is true in the actual world (w_0) iff
 $\exists s. s \leq w_0$ and $\forall w'$: w' is a best extension of s , she pays a fine in w'

Given the circumstances described above, (25) predicts that (20) is true in the actual world: there is a situation in the actual world (her returning the book late) such that in its best extensions, she pays a fine. (What *would* the best extensions look like? Maybe they are almost perfect worlds where everybody is very happy and the only ‘bad’ thing that has ever happened is that she returned the book late. But she paid her fine and was appropriately sorry.)

The Existential Hypothesis is interesting in that it gives the modal access to parts of what happens that are not temporal slices, but it also goes seriously wrong. Our judgments regarding deontic statements do not depend on isolated details of what has happened. To see that the Existential Hypothesis is too weak, consider the following. Suppose she is a careless, wealthy supermodel, who constantly returns books late and simply does not care. We might say ‘ok, she should pay a fine’. But suppose instead that she is a poor, hardworking grad student, who slipped up once, is very sorry, and would be completely ruined by a fine and have to leave school (she has no rich relatives, will not win the lottery, etc.). Our intuitions would probably waver. Taking into account only library regulations we would say: ‘go ahead, fine her’. But our judgements are often more complex, the measure of goodness can be an ‘overall’ measure, sensitive to various factors. Library regulations can be overridden by more important considerations. In judging (20), we do not only pay attention to the fact that she returned the book late, we also worry about the other things going on in the world. The problem with the Existential Hypothesis is that it only pays attention to identifying a ‘trigger’ situation, whereas our intuitions indicate that in evaluating *should* we take into account the other things that are going on too.

In evaluating whether she should pay a fine or not in (20), we have to give everything that is going on its due weight. We should worry about the best outcome *given that she returned the book late and everything else*. Suppose that in the actual world she is a wealthy unrepentant supermodel. Let’s call this situation s_1 . It is not the case that in the best-extensions of s_1 she pays a fine (in itself, being a supermodel is not illegal). However, given that in the actual world she has returned the book late, in the actual world there will be a bigger situation that will include s_1 and the situation of her returning the book late (this situation is an extension of s_1 , let’s call it s_2). And, in the best-extensions of s_2 , she does pay a fine. Suppose now that in the actual world she is a poor repentant grad student (let’s call *this* situation s_1). This situation will not have best extensions in which she pays a fine (it is not illegal to be a grad student). Given that in the actual world she did return the book late, there will again be a bigger situation in the actual world that will include s_1 and her returning the book late (s_2). But in this case, in the best extensions of s_2 , it may well be false that she paid a fine. If the relevant measure of goodness is sensitive to things like moral character, people’s overall well-being, the importance of being sorry for one’s wrong-doings, etc., it may be that in the best extensions of s_2 , she got away with a warning.

In evaluating *should*-statements, we seem to be quite methodical regarding the need to take into account everything that happens. Imagine that there are two patients with a serious illness, which can only be cured by medicine A. There is only one pill left (and only a whole

pill is effective). The disease is not as advanced on both patients. Sara is suffering a lot of pain, but her life is not at risk. Rosa has a very severe case, and will die without the pill. In a context in which human life was of the utmost importance, we would probably consider that Rosa should get the last pill, not Sara. However, if we evaluated the situations independently, we would not arrive at this result. As far as the situation of Sara's illness is concerned, in the best extensions it is true that she gets the pill. And as far as the situation of Rosa's illness is concerned, it is also true that in the best extensions she gets the pill. The examples show that in evaluating *should* we compare amongst a domain of possibilities that reflects all the things that happen, not only some selected parts.

To capture this insight, we will formulate an alternative hypothesis for *should*, that we will call the *Universal Unrestricted Hypothesis*:

- (26) Universal Unrestricted Hypothesis (preliminary)
 $\llbracket \text{should } \alpha \rrbracket$ is true in a world w iff
 $\forall s: s \leq w. \exists s'. s \leq s'$ and $\forall w': w'$ is a best-extension of $s, \alpha(w') = 1$

The proposal in (26) claims that a *should* α statement will be true in the actual world iff every situation in the actual world is part of a situation in the actual world such that α is true in its best extensions.

Notice that it is important to make a universal claim about the situations that make up the world. Let us go back to the scenario with the poor repentant grad student who returned the book late. She will actually have bad luck all her life, will not inherit money, will not win the lottery, and would be ruined by a fine. But suppose that in evaluating (20) we select amongst all the parts of the world the situation corresponding to her current lack of wealth and grad-student life (let's call this situation s_3) and ignore the fact that she will not inherit money, etc. There is an extension of s_3 in the actual world (let's call it s_4) in which she returned the book late (the situation corresponding to s_3 extended to include the facts of the late return of the book). As far as s_4 is concerned, it may be true that in all the best extensions, she pays a fine. This is because in the best extensions of s_4 she could win the lottery or be given millions of dollars by a friend, and would thus be able to satisfy library regulations to their full extent and save herself from ruin. Starting out from s_3 , we could move on to s_4 and its best extensions and find out that she does pay a fine. But this will not justify the truth of (20). The proposal in (26) captures this intuition. When evaluating (20), we'll have to consider all situations in the world. So, we will be forced to deal also with the actual situation of her being a poor grad student who will not find herself in possession of a large amount of money.

3.3 Counterfactuals?

There is an obvious problem with the Universal Unrestricted Hypothesis in (26). Consider again the case of the parent who finds out that the baby has been watching *The Simpsons*. Exasperated, she asserts (27):

- (27) She should have been watching *Baby Einstein*.

According to (26), (27) will be true only if every situation in the actual world has an extension such that in its best extensions, the baby was watching *Baby Einstein* at 9pm. But the baby was actually watching *The Simpsons*. The baby-watching-*The Simpsons*-situation will not have an actual world extension with best extensions in which the baby watches *Baby Einstein*: the proposition that the baby watches *Baby Einstein* is incompatible with the actual world situation of the baby watching *The Simpsons*.

There is a counterfactual air to *should*-statements that we have ignored so far. With *should*-statements, we can describe worlds that are different, and known to be different, from the actual world. The proposal in (26) does not allow this. It forces us to consider every situation in the world. But in evaluating *should* we seem to be able to ignore the situations that are incompatible with the proposition embedded by the modal. Those situations don't bother us. We only pay attention to the situations that can extend into situations in which the embedded proposition is true. This means that the unrestricted view of universal quantification in (26) won't work as stated. To be able to formulate an appropriate restriction, I will introduce an auxiliary notion to talk about the compatibility of a situation and a proposition:

- (28) (preliminary)
modal compatibility: a situation s is compatible with a proposition α
 iff there exists a situation s' such that $s \leq_m s'$ and α is true in s' .

With the help of *modal compatibility*, we can refine the proposal in (26) to the proposal in (29):

- (29) Universal Hypothesis (preliminary)
 $\llbracket \text{should } \alpha \rrbracket$ is true in a world w iff
 $\forall s: s \leq w$ that are α -compatible,
 $\exists s': s \leq s'$ and $\forall w': w'$ is a best-extension of s' , $\alpha(w') = 1$

With (29) in hand, examples like (27) appear less problematic. In checking whether situations in the actual world have extensions that ideally lead to the baby watching *Baby Einstein*, we will set aside the actual world situation of the baby watching *The Simpsons* together with other situations that are incompatible with watching *Baby Einstein* (and thus capture the counterfactual air of the statement). (Note that (29) leaves room for context-driven flexibility: in figuring out whether a situation is compatible with the embedded proposition or not, we will appeal to context-driven similarity criteria that will allow us to decide whether a situation in another world counts as a counterpart of an actual world situation).

However, another modification will be needed to the proposal in (29) to get things right. To see this, consider again the example with the doctor administering medicines in (4). Suppose that as a matter of fact the doctor administered medicine B on Monday. Medicine B had the side-effect of raising the patient's temperature (on Tuesday morning, the patient had a fever). Suppose also that it would be fatal for the patient to have a fever after being given medicine A. Our intuitions tell us that in this scenario, we still think that the doctor should have given the patient medicine A on Monday. This means that we exclude from the domain of quantification of the modal not only the situation of the doctor actually giving the patient medicine B on Monday, we also exclude from the domain of quantification the situation of the patient having a fever on Tuesday. This indicates that we exclude from the domain of quantification of the modal the situations that are incompatible with the embedded

proposition AND the situations causally dependent on such incompatible situations. From now on I will use the term α -compatible in a broad sense to include causal dependence:

- (30) *modal compatibility (broadly)*: a situation s is compatible with a proposition α iff there exists a situation s' such that $s \leq_m s'$ and α is true in s' and s is not causally-dependent on a situation s'' incompatible with α (where the notion of causal-dependence holds only between situations in the same world and is defined in terms of the laws of the evaluation world)¹³

If we understand compatibility broadly, then the proposal in (29) will ensure that when we evaluate whether the doctor should have given the patient medicine A on Monday, we exclude from the domain of quantification the situation of the patient having a fever on Tuesday.

The proposal in (29) has a problematic side-effect that is already familiar from the Feldman-style proposal in (6). If α is true, *should* α will be trivially true. In a sense, the truth of α collapses the semantic machinery built around *should*. This is because our proposal for *should* is closely tied to the facts. It looks for the best alternatives given the facts. For any particular thing that has happened, its best alternatives will include it. Thus, if it is true that she returned the book late, then for all the situations in the world, there will be an extension (to a situation that includes the late return of the book), such that in its best extensions, she returned the book late (for that situation, it will be true in all of its extensions that she returned the book late, therefore, it will also be true in its best extensions). If she returned the book late, we predict that it is true that she should have returned the book late.

The fact that the truth of the complement makes the truth of the deontic statement trivial is something that is problematic, and I will not be able to offer a general solution here. I will propose the stop-gap solution of excluding from the domain of quantification of *should* situations that themselves make the embedded statement true.¹⁴ The actual world situations that need to be considered when evaluating *should* α are the situations compatible with α that do not themselves make α true. I will refer to this set as the set of *relevant** situations, and propose a last revision to the meaning of *should* (where α is the proposition argument of *should*, the *relevant** situations = α -compatible situations that do not themselves make α true)

- (31) Universal Hypothesis (final)
 $\llbracket \text{should } \alpha \rrbracket$ is true in a world w iff
 $\forall s \leq w$ that are *relevant**,
 $[\exists s' [s \leq s' \ \& \ \forall w' \text{ that are best extensions of } s', \ \alpha(w') = 1]]$

¹³ The problem discussed in this paragraph is a problem often addressed in the literature on counterfactual conditionals, and is sometimes called the problem of ‘retraction’. When we evaluate modal claims we often need to set aside facts incompatible with the modal claims and those causally related to facts incompatible with the modal claims. What exactly determines which things go together and how to handle this is a matter of debate. Here I have informally appealed to the notion of a causal chain of events, similar to what we find in Bennett (2003) for the case of counterfactuals. The problem has been addressed in different forms by means of ‘lumping’ (Kratzer 1989), and ‘retraction’ (Veltman 2005).

¹⁴ The fact that in proposing a semantics for modals we may need to ensure diversity in the domain of quantification in order to avoid triviality has been observed in the literature with respect to various kinds of modality, and has been dealt with, e.g. by the *Diversity Condition* in Condoravdi 2002, and has also been discussed by Frank 1997, amongst others.

With this addition, the actual truth of α will not have consequences regarding whether *should* α is true. Both the situations that are incompatible with α and the situations that make α true will be ignored. The truth value of the statement will only depend on everything else that is going on.

By excluding from the domain of quantification the actual world situations that are incompatible with the embedded proposition, we arrive at an interpretation that is, in a sense, counterfactual. The question then arises as to how counterfactual exactly it is. Both in the case of counterfactual statements and deontic statements, we pay attention to only parts of the world, and set others aside. The difference lies in the criteria we use to ‘fill in’ the blanks. In the case of counterfactuals, we identify the worlds quantified over by considering what would have happened given the laws that actually govern the world. In the case of deontic statements, we identify the worlds quantified over by considering what would have happened ideally given some measure of value.

Notice that the fact that *should* doesn’t (necessarily) care about the laws that govern the world does not mean that any absurd *should* statement will come out true. Consider (32):

(32) There should have been a pink giraffe in my living room.

There are two ways in which we could imagine (32) coming out true. Either a pink giraffe in my living room is so intrinsically ideal that the embedded proposition would be true in all ideal worlds, no matter what (which is not the case). Or there is some fact in the actual world that favors a pink giraffe in my living room, and for every relevant* situation in the world, there is an extension (to include that fact) such that in the best extensions there is a pink giraffe in my living room (which is not the case either). So (32) will be false. Not because it is incompatible with the laws that govern the world, but because it doesn’t improve things given the actual facts.

Given the proposal in (31), counterfactual and deontic statements differ not only in terms of how we fill in the blanks, but also in terms of what I will call loosely ‘temporal directionality’. The issue of temporal directionality and counterfactuals has a long history in philosophy and theories of causation, and I will not attempt to address it here (for discussion of ‘backtracking’ see a.o. Lewis 1979, Frank 1997, Bennett 2003, Arregui 2004, 2005, Schulz (forthcoming)). Let us simply observe that in evaluating counterfactual statements, we tend to think in terms of causes and consequences, where the causes temporally precede the consequences. The proposal in (31), however, makes no claims regarding temporal directionality. In evaluating *should*-statements we take into account all parts of the world, and consider situations that are to our past or future (the notion of a best-extension does not encode a temporal direction). The facts that we feel ‘are responsible’ for the truth of the deontic claim may follow the eventuality described by the clause embedded under *should*. In (13), we already saw an example in which the truth of a deontic claim was ‘justified’ by something that happened later. Another example is given in (33):

(33) A friend of mine gives me a blue scarf for my birthday. I exclaim:
 a. I should have bought blue shoes last week (instead of yellow ones)

The sentence in (33) seems true even though last week I did not have anything to wear with my blue shoes, nor any idea that I would get a blue scarf.

We might worry about the absence of temporal directionality in the characterization of best-extensions. Imagine a group of doctors treating a patient with a terrible disease whose quality of life is deteriorating rapidly. They get word of a medical breakthrough and a new drug that works only if vitamin D has been administered a week before. Is one of them likely to utter (34)?

(34) We should have given the patient vitamin D a week ago.

An utterance of (34) in this context could be considered strange (very unhelpful). The drug didn't even exist a week ago. But I would like to suggest that the oddness is not due to temporal directionality, but is linked instead to the performative effects that are associated with the utterance. (34) would be odd uttered as a (self)reproach. I would like to speculate that some performative effects arise only with deontic claims that satisfy constraints regarding temporal directionality. Notice that if we remove the reproachful note, (34) seems better:

(35) Oh no! We should have given the patient vitamin D a week ago. But we couldn't possibly have known that!

3.4 Examples

In this section we will examine some examples to see (31) at work. According to (31), both sentences in (36) are true in the scenario described above:

- (36) a. The baby should have been asleep by 9pm.
b. The baby should have been watching *Baby Einstein*.

The truth conditions for the sentences are spelled out in (37):

- (37) a. **[[should [the baby have been asleep by 9pm]]]** is true in a world w iff
 $\forall s \leq w$ that are relevant*,
 $[\exists s' [s \leq s' \ \& \ \forall w' \text{ that are best extensions of } s',$
the-baby-has-been-asleep-by-9pm in w']]
- b. **[[should [the baby have been watching B.E.]]]** is true in a world w iff
 $\forall s \leq w$ that are relevant*,
 $[\exists s' [s \leq s' \ \& \ \forall w' \text{ that are best extensions of } s',$
the-baby-has-been-watching-B.E. in w']]

The key to understanding how the two sentences can be true is to notice that the actual world situations quantified over in each case are different. In (37a) we quantify over situations compatible with the baby being asleep. We'll set aside the actual situation of the baby being awake at 9pm and the baby watching *The Simpsons* at 9pm (and whatever else is broadly incompatible with the baby sleeping). For all situations quantified over, there are extensions (to situations that include the tired baby) such that in their best extensions, the baby was asleep by 9pm. In (37b) we quantify over situations compatible with the baby

watching *Baby Einstein* at 9pm. This time, we'll set aside the actual situation of the baby watching *The Simpsons*, but we will consider the situation of the baby being wide-awake till very late. It will be the case that for all actual world situations there are extensions (to situations that include the wide awake baby) such that in their best extensions, the baby is watching *Baby Einstein*.

The result is that the proposal in (31) predicts that both sentences in (36) can be true (they are not inconsistent). And, in a similar way, we could predict that sentences like (38) could both be true (note that this was a result already available to the temporal proposal in (6/8)):

- (38) a. She should have returned the library book on time.
b. She should have paid a fine.

A difficulty that arises with these results is that we wouldn't (normally?) conjoin these sentences:

- (39) a. The baby should have been asleep and the baby should have been watching *Baby Einstein*.
b. She should have returned the library book on time and she should have paid a fine.

I would like to propose that the answer to this puzzle lies not in the semantics of deontic statements, but in their pragmatics. Suppose that both sentences in (36) are true. We can still wonder about why somebody would choose to utter one or the other. Imagine that the parent leaves the baby at home with two baby-sitters. One is in charge of putting the baby to sleep by 9pm, and the other is in charge of making sure that if the baby is not asleep, she is watching *Baby Einstein*. The parent finds out that the baby has been up till very late watching *The Simpsons*. We could imagine she would say (36a) to the baby-sitter in charge of putting the baby to sleep by 9pm, and (36b) to the baby-sitter in charge of making sure the baby watched *Baby Einstein*. The parent's choice about what to say to whom would then be based on the goal of reproaching each baby-sitter. Both statements are true but they only count as a reproach when uttered in the right context. I would like to speculate that we should not ask the semantics of deontic statements to predict these facts, but should leave it instead to a theory of pragmatics that spells out the performative effects associated with deontic statements.¹⁵

3.5 A small summary

In this section we have proposed a fact-sensitive meaning for *should* that balances what is best given everything that happens (but leaving out some parts):

¹⁵ An alternative would be to encode 'hearer sensitivity' in the semantics of *should* (where *should* would mean 'best given what is/was in your control', or something similar). However, it seems to me that the range of performative effects associated with *should* is broader than what we would achieve by simply encoding hearer sensitivity. However, without an overt pragmatic story, it is not possible to truly evaluate that part of the proposal.

- (31) Universal hypothesis (final)
 $\llbracket \text{should } \alpha \rrbracket$ is true in a world w iff
 $\forall s \leq w$ that are relevant*,
 $[\exists s' [s \leq s' \ \& \ \forall w' \text{ that are best-extensions of } s', \alpha(w') = 1]]$

In this proposal, there are two layers of quantification. One is related to ‘what is best’, characterized in terms of quantification over best extensions (possible worlds). The other has to do with making sure that all (relevant*) actual-world facts are taken into account. This is the quantifier over situations that make up the evaluation world.

In this proposal, the evaluation world serves as the ‘evaluation frame’ for the deontic claim: it provides us with the totality of things we take into account when figuring out what is best and so decides what possibilities will be compared when evaluating what is best. Actual world facts affect the truth-value of *should*-statements by determining what possibilities will be compared in evaluating the modal claim. In a sense, this layer of quantification has to do with ensuring sufficient similarity with the actual world amongst the possibilities compared by the modal.¹⁶

According to (31), we evaluate *should*-statements in worlds. We could wonder what would happen if we evaluated them in situations smaller than worlds. Doing so would amount to allowing *should*-statements to denote non-persistent propositions (in the sense of Kratzer (1989)). If a *should*-claim were to be evaluated with respect to a situation that was smaller than a world, the situation would then serve as the evaluation frame. A *should*-

¹⁶ I would like to briefly examine a referential alternative to (31). In (31) there is *quantification* over actual world situations compatible with the proposition embedded under the modal. In a *referential* alternative to (31), the modal statement would be evaluated relative to a single situation in the actual world. A referential semantics could be given as in (*) (where the modal statement is made true in a world by virtue of what happens in one (proper) part of the world):

(*) $\llbracket \text{should } \alpha \rrbracket$ is true in a world w iff $\forall w'$ that are best-extensions of s' , $\alpha(w') = 1$, where s' is the maximal relevant* situation in w .

With a proposal like (*), we target directly the ‘biggest’ situation in the evaluation world compatible with the embedded proposition, and then examine what happens in its best extensions. It might seem that this achieves the same results as (31), but in a more straightforward manner. However, the implicit uniqueness presupposition in (*) is not correct: there may be a variety of ways of picking-out the facts compatible with the embedded proposition (α). Intuitions regarding examples like (**) below indicate that we need to consider all options when evaluating *should* (which leads us back to a quantificational account):

(**) I am at a party with two friends, and we each pick-up a glass from a tray holding three glasses. I am not very satisfied with my drink, and tell my friends: *I should have taken one of the other two glasses* (any of those two, this is meant with a narrow scope reading for the object DP).

It turns out that in one of those glasses there is a drink I like very much, and in the other there is a drink that is harmful. In such circumstances, we would judge the *should*-statement in (**) false (I would never have said what I said if I had thought there was a harmful drink in one of the other two glasses!). Notice that there are (at least) two ways of ‘carving out’ of the actual world a maximal situation compatible with me taking one of the other two glasses: we either set aside the situation in which one of my friends takes the glass with the pleasant drink or we set aside the situation in which one of my friends takes the glass with the harmful drink. If we take into account both such situations (as is required by the quantificational account in (31)), we will get the truth-value of the sentence right: starting out from the situation in which the glass with the harmful drink has not been taken, it is not the case that in its best extensions, I take it (in this discussion I am obviously assuming that my well-being is the measure of goodness). (Notice that the fact that in evaluating modals we may need to consider different combinations of actual world facts and embedded propositions is already addressed in Kratzer’s semantics for counterfactuals (a.o. Kratzer 1981), where different maximal consistent combinations of propositions are taken into account in evaluating *would*. See also the discussion in Veltman 2005.) (I am grateful to a reviewer for bringing up a referential alternative.)

statement could be true in a situation s smaller than a world and false in the world of the situation s (i.e. false when more things were taken into account).

Finally, let us note that the proposal in (31) will be obviously helpful in explaining the semantics of *should*-conditionals and CTDs. As we saw in §3.4, (31) predicts that primary obligations may be compatible with secondary obligations. If we assign *should* the interpretation in (31), factual detachment in CTDs will not give rise to Chisholm's Paradox.

4. *Should vs. should have*

In constructing the meaning of *should*, I have made use indiscriminately of examples with simple *shoulds* and *should have*. It is now time to clarify the difference between the two. In §2, I argued that the difference was not purely temporal, and appealed to a Stalnaker-style intuition that there was a relation with the context set. In this section I will spell out a proposal for simple *should*-statements that builds on this view. I will argue that the difference between simple *shoulds* and *should haves* arises because of differences in the propositions embedded under *should*, not because of differences in the interpretation of the modal itself. And I will argue that the propositions embedded under *should* vary because of differences in the embedded aspectual heads. Once the discussion of simple *shoulds* is complete, we will turn to *should*-conditionals in §5.

The obvious contrast between simple *shoulds* and *should have* statements has to do with the nature of embedded aspect. In the case of *should have*, the modal *should* embeds a clause headed by the perfect aspectual head *have*. This is not the case with simple *shoulds*, where there is no perfect. Following up on Arregui (2004, 2007), I propose that in the absence of *have* morphology, the clause embedded under *should* will either contain a default silent perfective aspectual head (in the case of eventive predicates) or will lack an aspectual head all together (in the case of stative predicates). If we look in more detail at the structure of *should* statements, we will find the following options:

- | | | | |
|------|----|--|--------------------------|
| (40) | a. | She should pay a fine. | <i>Simple perfective</i> |
| | b. | [should [\emptyset _{perfective} [she pay a fine]]] | |
| (41) | a. | She should have paid a fine. | <i>Perfect</i> |
| | b. | [should [have [she pay a fine]]] | |
| (42) | a. | She should be here. | <i>Simple stative</i> |
| | b. | [should [she be here]] | |

Even though there is no overt aspectual morphology in (40a), there is a silent aspectual head with the semantics of perfective aspect. This is not the case for simple *shoulds* that are stative, since states do not receive a perfective interpretation.

My claim in this section is that the interpretation of *should* remains invariant across these three cases. It is the interpretation of aspect that is responsible for the differences we observe between the types of clauses. I will begin this section by presenting a view of aspect that predicts interaction with modality (§4.1). This proposal for aspect has been argued to be helpful in understanding variations in the case of counterfactual statements (Arregui 2004, 2007). I will then spell out and evaluate the predictions made for the interpretation of *should*-

statements (§4.2), and discuss the special case of statives (§4.3). Finally, I will make a small aside to discuss a difference between stative *shoulds* and perfect *shoulds* with respect to the behaviour of some presupposition triggers (§4.4) (this section can be skipped without continuity problems).

4.1 On the interpretation of aspectual heads

Arregui (2004, 2007) presented a proposal for the interpretation of aspectual heads that had the objective of explaining variations in the interpretation of counterfactual conditionals. The choice of aspectual heads was argued to be the key for understanding variation in the interpretations of clauses embedded under counterfactual modals. In particular, it was claimed that perfective aspect resulted in an anchoring of propositions to the worlds in the context set. The proposal provided an analysis for the Stalnakerian intuition that the choice of morphology could be driven by the desire to indicate that claims were being made regarding worlds within/outside the context set. (For an alternative analysis, see Ippolito (2003). The idea that aspect may be linked to epistemic issues or have a modal dimension has also shown up recently in work outside counterfactuals (Bhatt 1999, 2006, Hacquard 2006)).

The proposal in Arregui (2004, 2007) makes use of the basic syntactic architecture and insights in Kratzer (1998). Following Kratzer, VPs headed by eventive verbs are characterized as denoting properties of events. Aspectual heads project above VP and relate properties of events to properties of times. In the case of perfective aspect, the aspectual head locates the running time of the event within the reference time. In Arregui (2007) I argued for a referential view of perfective aspect, and used that characterization, together with a pragmatic diagonalization strategy (à la Stalnaker) to derive the presupposition that a perfective clause denoted a proposition true only in the context set. Here I will simplify the presentation, and ignore the internal mechanisms that gave rise to the presupposition.

Adapting the proposal in Arregui (2007) to a situations-based framework, we obtain the semantics for perfective aspect in (43). According to (43), the perfective head combines with a property of events and the output is a partial function that takes situations as inputs and delivers truth-values. The restriction indicates that the function is defined only for situations that are parts of worlds within the context set:

$$(43) \quad \llbracket \emptyset_{\text{perfective}} \rrbracket^c (P) = \lambda t. \lambda s: \exists w. w \in c \ \& \ s \leq w. \exists e. [P(e)(s) = 1 \ \& \ \tau(e) \subset t]$$

Where P is a property of events, and the restriction on the output function indicates that the function is defined only for situations that are parts of worlds in the context set (c) (in the future, this will be abbreviated as follows: $\lambda s: s \leq w \in c$).¹⁷

According to the proposal in (43), perfective aspect is responsible for situating an event in a situation at a time. The aspectual head claims that the time of the event is included within

¹⁷ Some terminology: I take it that if P is a property of events, $P(e)(s)$ is true iff e is a P -event that occurs in s . I will follow Kratzer (2006-8) in characterizing events as self-connected situations that *exemplify* properties of situations (Kratzer 2006-8: 29). This means that $P(e)(s)$ is true iff e is situation that exemplifies P and is part of s . (The reader is referred to Kratzer 2006-8 for examples and discussions regarding Davidsonian events in a situations framework.)

the reference time (it is a subinterval of the reference time) and carries the presupposition that the event is found in a world in the context set. An example is given in (44):

- (44) Example:
 $\llbracket [\emptyset]_{\text{perfective}} [\text{she returns her book on time}] \rrbracket^c =$
 $\lambda t. \lambda s: s \leq w \in c. \exists e. [\text{she-returns-her-book-on-time}(e)(s) = 1 \ \& \ \tau(e) \subset t]$

The denotation of the aspectual phrase is a property of times. If this predicate were to be shifted to a non-past time t (as it happens in the scope of a modal), we would obtain the property of situations in (45):

- (45) $\lambda s: s \leq w \in c. \exists e. [\text{she-returns-her-book-on-time}(e)(s) = 1 \ \& \ \tau(e) \subset t]$,
 where t is a non-past time and the proposition is defined only for (situations in) worlds in the context set

The predicate in (45) will be true of a situation if there exists an event of her returning the book on time in that situation, that event has a running time in the future (not in the past), and the situation is part of a world in the context set.

Let us turn now to the case of perfects. According to Arregui (2007), perfect aspect does not encode information regarding the context set. Following Kratzer (1998), we will adopt a ‘result state’ view of the perfect according to which the perfect combines with a property of events and gives rise to a property of times true of times that follow an event with the relevant property (see also Parsons 1994). We will modify the proposal to adapt it to the situations framework we are using here:

- (46) $\llbracket \text{have}_{\text{perfect}} \rrbracket(P) = \lambda t. \lambda s. \exists e. [P(e)(s) = 1 \ \& \ \tau(e) <_{\text{precedes}} t \ \& \ t \subset \tau(s)]$

According to (46), perfect combines with a property of events and the result is a property of times true of a time t and a situation s iff the temporal location of the situation s extends (at least) until the time t and there is an event of the appropriate type in the situation that precedes the time t .

- (47) Example
 $\llbracket [\text{have}_{\text{perfect}}] [\text{she returned the book on time}] \rrbracket^c =$
 $\lambda t. \lambda s. \exists e. [\text{she-return-the-book-on-time}(e)(s) = 1 \ \& \ \tau(e) <_{\text{precedes}} t \ \& \ t \subset \tau(s)]$

The denotation of the aspectual phrase is again a property of times. If this predicate were to be shifted to a non-past time t , we would obtain the property of situations in (48):

- (48) $\lambda s. \exists e. [\text{she-return-the-book-on-time}(e)(s) = 1 \ \& \ \tau(e) <_{\text{precedes}} t \ \& \ t \subset \tau(s)]$
 where t is a non-past time

The predicate in (48) would be true of a situation if there exists an event of her returning the book on time in the situation and the running time of the event precedes a non-past time t and the situation extends at least as far as t . Notice that the claim is that the state that follows the event should extend as far as non-past t . This allows the event to be past, present or future (with respect to the speech time). Given that the claim is made with perfect aspect,

there isn't any presupposition regarding the context set. This proposition can be true of a situation regardless of whether the situation is part of a world in the context set or not.

4.2 Aspect and *should* statements

With a semantics of aspect in hand, we can now turn to the problem of explaining how the choice of aspect interacts with the interpretation of *should* and results in the differences we have observed. In §3 we ended up with a proposal for the meaning of *should* that evaluated what happened in the best extensions of situations in the actual world (repeated in (49)):

- (49) Universal hypothesis (final)
 $\llbracket \text{should } \alpha \rrbracket^c$ is true in a world w iff
 $\forall s \leq w$ that are relevant*,
 $[\exists s' [s \leq s' \ \& \ \forall w' \text{ that are best extensions of } s', \alpha(w') = 1]]$

Before becoming explicit about the meaning of aspectual heads, we discussed the interpretations of *should*-statements in a simplified manner, and appealed to truth conditions presented as in (50):

- (50) $\llbracket [\text{should } [\text{she return the library book on time}]] \rrbracket^c$ is true in a world w iff
 $\forall s \leq w$ that are relevant*,
 $\exists s' [s \leq s' \ \& \ \forall w' \text{ that are best-extensions of } s',$
she-returns-her-book-on-time (w') = 1]

In (50), we are using *she-returns-her-book-on-time* as shorthand for the embedded proposition. There are two pieces of information that do not show up in this shorthand: (1) that the modal will set the time of returning at a non-past time, and (2) the proposition is true only in worlds in the context set (the embedded clause has perfective aspect). We will take the temporal information for granted, and set aside the mechanisms that allow the modal to shift the evaluation time of the embedded clause (but see a.o. Enç (1996), Iatridou (2000), Condoravdi (2002), Gennari (2003) for discussion). Our interest lies in the aspect-triggered context-set presuppositions of the embedded clause. Taking into account the meaning contribution of a perfective aspectual head as in (45), we propose (51):

- (51) $\llbracket \text{should } [\emptyset_{\text{perfective}} [\text{she return the library book on time}]] \rrbracket^c (w) = 1$ iff
 $\forall s \leq w$ that are relevant*,
 $\exists s' [s \leq s' \ \& \ \forall w' \text{ that are best-extensions of } s',$
 $[\lambda s: s \leq w \in c. \exists e. [\text{she-returns-her-book-on-time}(e)(s) = 1 \ \& \ \tau(e) \subset t]] (w') = 1]$
 (where t is a non-past time)

Given the truth conditions in (51), the proposition embedded under *should* is claimed to be true in all situations (worlds) that are best extensions of relevant actual world situations. My claim is that the presuppositions associated with the embedded proposition are accommodated within the restriction of the quantifier over possible worlds that quantifies over best extensions, and in this way put a restriction on the possibilities compared by the modal. Accommodating the presuppositions of the 'nuclear scope' proposition has the effect

of restricting quantification to worlds in the context set. (Notice that the idea that presuppositions can be accommodated in quantification domains is familiar from the literature. An early discussion can be found in Heim (1982), who noted that in sentences like *Every boy hates his father*, we quantify only over boys who have fathers – the presuppositions of the possessive in the nuclear scope are accommodated in the domain of quantification of the universal *every*). If we take into account presupposition accommodation and simplify the notation a little, we end up with (52):

- (52) $\llbracket \text{should } [\emptyset_{\text{perfective}} \text{ [she return the library book on time]}] \rrbracket^c(w) = 1$ iff
 $\forall s \leq w$ that are relevant*
 $\exists s' [s \leq s' \ \& \ \forall w' \text{ that are best-extensions of } s' \text{ and members of } c,$
 $\exists e. [\text{she-returns-her-book-on-time}(e)(w') \ \& \ \tau(e) \subset t]$
 (where t is a non-past time)

As (52) shows, the choice of perfective aspect will convey the information that quantification is taking place over worlds in the context set. If the proposition embedded under *should* is known to be false, the choice of perfective aspect will be infelicitous.

Things will be different when the modal embeds a clause headed by perfect aspect. If we look into the details of the structure and embedded proposition (48), we will find (53):

- (53) $\llbracket \text{should } [\text{have}_{\text{perfect}} \text{ [she return the library book on time]}] \rrbracket(w) = 1$ iff
 $\forall s \leq w$ that are relevant*,
 $\exists s' [s \leq s' \ \& \ \forall w' \text{ that are best-extensions of } s',$
 $\exists e. [\text{she-returns-her-book-on-time}(e)(w') = 1 \ \& \ \tau(e) <_{\text{precedes}} t \ \& \ t \subset \tau(w')]$ ¹⁸
 (where t is a non-past (future) time)

According to (53) we begin by identifying the situations quantified over in the actual world. The *should*-statement will be true iff for all of them there is an extension in the actual world such that in its best extensions, she is in the state of having returned the book on time at some non-past time (which could mean that she returned the book in the past or future). Contrary to what we saw in the case of the perfective, there are no context-set presuppositions associated with the perfect, and no restrictions are added to the quantificational domain of the universal quantifier over best extensions. With perfect aspect, we are free to look for best extensions outside the context set.

The contrast between perfective and perfect in terms of context-set presuppositions has many interesting consequences. It allows us to predict that *should*-claims regarding future events should be made with *should have* instead of simple *should* if they are incompatible with the common ground. In example (14) (repeated below) we choose *should have* because we have accepted as true that the project will be cancelled:

- (54) John: I've had a great idea. We should include a breakfast application in the new iPhone!
 Jane: Haven't you heard? We'll be laid off at the end of the month and the whole project will be cancelled!!

¹⁸ It is unusual to end up talking about the running time of a world, but it does not seem problematic.

John: Oh no! That is dreadful! ... But isn't it a great idea? We should have included a breakfast application in the new iPhone!

The choice is conditioned by what we accept as knowledge / are willing to presuppose in the conversational context, not by what has happened in the world so far.

According to the semantics for *should* in (31), *should* and *should have* statements don't differ in terms of the parts of the world that affect their truth-values. In both cases, quantification takes place over all situations appropriately compatible with the embedded proposition. Moreover, the truth-conditions are not sensitive to temporal directionality. *Should* statements can be supported by situations that happen in the future. Consider (55):

- (55) A: The firing squad will execute a man tomorrow.
 B: They should offer him a cigarette first.
 (following Prakken and Sergot (1996))

The truth of the claim that they should offer the man a cigarette doesn't depend on what has happened in the world so far (we should not be fooled into thinking that (55B) is true because they have already decided/planned to execute a man – (55A/B) could be uttered by time travellers who know that a man will be executed, even if none of the machinery to make that happen has been set in motion yet and nobody else knows the man will be executed). Neither is it the case that it is true in the best worlds accessible from the speech time (ideally, the man would not be executed). Intuitively, the things that make (55B) true are in the future. It is only when we take for granted something that will happen in the future (the execution) that we would utter (55B).

The idea that *should*-statements are made true by everything that happens may appear worrying, since we don't know everything that happens and yet we are happy to utter *should*-statements. Suppose that I tell you that my sister will arrive at the bus station late tomorrow, and that I can't go pick her up. The bus station is in a dodgy neighbourhood and she is afraid of walking home alone in the dark. You might utter (56):

- (56) She should take a cab.

This seems a reasonable thing to say. But suppose that I were able to show you that tomorrow evening, all cab drivers waiting in the station will be mean and evil people. With that knowledge, it is likely that you will think that (56) is false. I would like to claim that, once we have agreed upon a measure of value, there is a fact of the matter regarding the truth of (56). It is decided by everything (relevant*) that happens in the world. At times, we are at an epistemic disadvantage, and do not have access to all the relevant knowledge (for example, we can't know what will happen). But we take an 'epistemic gamble', and use *should* all the same. And epistemic gambles are not exclusive to future-oriented statements, they are also taken with respect to what has happened in the past. Suppose the same story of my sister's arrival had been set in the past, and I explained to you that she had walked home alone and had been very afraid. You might go on to claim (57):

- (57) She should have taken a cab.

And again, I could have filled in the information you don't have (the cabdrivers at the station that night were all mean and evil), and you would have changed your mind.

4.3 Stative clauses

In the discussion of aspect and counterfactual modals in Arregui (2004, 2007), the case of simple statives was assimilated to the case of perfect states. I showed that in counterfactual conditionals, both could be used to describe worlds that are outside the context set. The key was the claim that stative clauses did not project perfective aspect and lacked context-set presuppositions (simple states could combine with modals directly). In the context of *should* statements, this view predicts that simple stative *shoulds* can, in the relevant sense, be ‘counterfactual’ deontics (they do not carry context set presuppositions, and so the proposition embedded under the modal can be true in worlds outside the context set).

The classification of simple stative *shoulds* as counterfactual is a delicate matter. On the one hand, simple stative *shoulds* seem perfectly fine in cases in which we know the embedded clause is false:

(58) Your mother isn’t here. She should be here.

On the other hand, they can trigger the performance of actions in a way that might not be really expected in the case of true counterfactual deontics. Consider (59):

(59) a. Parent (looking at the baby playing on the floor): The baby should be in bed.
b. The baby-sitter (picks up the baby, puts her to bed): There!

To understand what happens in this kind of example, it is important to consider the temporal span associated with clauses embedded under *should*. When stative clauses are embedded under *should* as in (58) and (59), we understand that there is a temporal span corresponding to the embedded clause. In the case of (59a), for example, we understand the parent to claim that the baby should be in bed at a time that includes the speech time and extends some time into the future (presumably, the temporal span associated with states is pragmatically identified, notice the difference between *The baby should be in bed* and *The president should live in The White House*). States are homogeneous, and when they hold at an interval, they hold throughout that interval. That is, they are true at all the subintervals. Faced with an annoyed parent who is claiming that the baby should be in bed throughout an interval that includes now and some future times, the baby-sitter may feel that the best thing to do is to ensure that at least for the times that are left in that interval, the baby IS in bed.

The temporal flexibility of states can make it difficult to see that they can be used to make deontic statements that are counterfactual. This becomes easier to see if we use adverbials to manipulate the temporal span associated with the state:

(60) a. Parent: The baby should be in bed.
b. The baby-sitter (picking up the baby): I will put her to bed right away!
c. Parent: What’s the use? The baby should be in bed NOW!/
The baby should already be in bed!

Of course, if the parent uses a perfect, the baby-sitter will know immediately that there is nothing to be done:

(61) Parent: The baby should have been in bed.

There is nothing the baby-sitter can do now that will make it be the case that the baby was in bed in the past.

I will conclude that it is possible to treat simple statives on a par with perfect statives, and allow both to make claims that are not anchored to the context set. The result is that with simple statives, as with perfects, the proposition embedded under the modal may be true in worlds that lie outside the context set, and the resulting deontic claim may, in this sense, be counterfactual.

4.4 Presupposition triggers in *should*-statements

In this section I will make a brief aside to discuss the behaviour of some presupposition triggers in *should*-statements (readers may skip this section without missing material that will be referenced later). The discussion follows up on part of the observations and discussions in Ippolito (2003, etc.) who has investigated presuppositions in the context of counterfactual conditionals. The purpose of this section is to highlight some data discussed by Ippolito which could be thought of as potentially problematic for the approach presented here. My goal is to point to a venue for future research that could account for the data within the current framework.

In §4.3 I proposed to group simple stative and perfect clauses as clause-types compatible with ‘counterfactual’ deontics. However, working with counterfactual conditionals, Ippolito (2003, etc.) has noted that simple statives and perfects can differ with respect to the behaviour of presupposition triggers. Parallel observations can be found in the domain of *should*-statements. We will be particularly interested in the presuppositions triggered by pronouns, proper names and definite descriptions (referential expressions). The key point is that as far as (existence, uniqueness) presuppositions are concerned, simple statives differ from perfects in that perfects can shift presuppositions away from the evaluation world, while simple statives cannot.¹⁹ We will provide illustrations with definite descriptions, but the general observations and proposals can be extended to the other cases. Consider the examples in (62):

- (62) a. It is a pity your parents are in Atlanta. Your mother should be here.
 b. #It is a pity your parents have passed away. Your mother should be here.
 c. It is a pity your parents have passed away. Your mother should have been here.²⁰

¹⁹ I would like to note on the basis of very informal observations that there seems to be some speaker variation in the evaluation of this data. It also seems to be the case that there is variation depending on the types of referring expressions (some people get the effect more easily with proper names and pronouns). Experimental work would probably be required to obtain a truly accurate picture, but this goes beyond the scope of this paper.

²⁰ Parallel examples with proper names and pronouns could be provided as in (i):

- (i) a. It is a pity Sarah is in Atlanta. She should be here.
 b. #It is a pity Sarah passed away. She should be here.
 c. It is a pity Sarah passed away. She should have been here.

Since proper names and pronouns can be analyzed as denoting situation-sensitive functions, the proposal for definite descriptions can be extended to these cases too.

In (62a) we see a felicitous example of a definite description in a stative *should*-statement. In (62b), however, the definite description is less felicitous. Both (62a) and (62b) would be considered examples of ‘counterfactual’ deontic statements (given the context, we know that the embedded clause does not describe the actual world). The difference in felicity arises with respect to the anchoring of the existence (and uniqueness) presuppositions associated with the description. The examples illustrate that even if the *should*-statement is ‘counterfactual’, with simple stative *shoulds* we expect to find the presuppositions associated with the definite satisfied in the actual (evaluation) world. This is not the case for perfects, as is illustrated by the (felicitous) (62c).


In the remainder of this section I would like to speculate on this point, and sketch a proposal that would allow us to get a handle on this data. The gist of the proposal will be as follows: I will argue that the important difference between simple statives and perfect *should*-statements is that with perfects there is an intensional operator embedded under the modal (the perfect *have*), while this is not the case with simple statives. According to (46), the perfect is capable of binding intensional arguments. A perfect will be able to bind the intensional argument of a definite description, and shift the evaluation index, affecting the calculation of the existence and uniqueness presuppositions. This will result in an ‘opaque’ reading for the definite. With simple statives, on the other hand, there is no intensional operator, and the intensional argument of a definite description will remain free. This will trigger a context-dependent valuation which will anchor the description to the actual world. The result will be a transparent reading. This means that the uniqueness and existence presuppositions must be met in the actual world.

For the sake of explicitness, I will present an implementation of this proposal, taking advantage of some of the ideas in Percus (2000) (for the sake of simplicity, I will ignore times in this discussion). Percus (2000) spelled out a theory regarding the syntactic distribution of world (situation) variables that was in part motivated by the contrast between transparent and opaque readings of noun phrases in modal contexts. Percus pointed out that while the situation argument associated with verbal projections always received an opaque reading, the noun phrase could either be opaque or transparent. I would like to suggest that one way of capturing this result would be to set up a system in which only noun phrases introduce syntactically-visible situation arguments and in which furthermore, the situation variable associated with the noun phrase can only be bound in the presence of an intensional operator that introduces a binder.²¹ This view would give us a handle on the contrast between the simple *should* examples and the perfect *should* examples.

In the case of a simple stative clause, like (63a), the clause embedded under *should* will be as in (63b):

- (63) a. should [your mother be here]

²¹ The proposal has been inspired by Kusumoto (2005)’s treatment of temporal opacity in the interpretation of noun phrases. Temporal opacity in noun phrases has been investigated a.o. by Enç (1986), Musan (1995), Kusumoto (1999, 2005).




With the assumption that a free situation pronoun is interpreted deictically as making reference to the utterance situation (s_0), the modal in (63a) would take as argument the proposition in (64):

(64) $\lambda s. [\text{your-unique-mother-in-}s_0 \text{ is here in } s]$

This proposition will be true in situations in which (a counterpart of) your actual mother is here. Nothing ties this proposition to our epistemic state, and it can be used to describe a situation that is known to be outside the context set. The presuppositions of the definite, however, remain anchored in the actual world, and the definite description will be infelicitous if you do not have a unique mother.

To capture the intensional effects of the perfect, I propose to allow the perfect head itself to introduce a syntactically visible intensional parameter. The proposal is spelled out in (64):

- (64) a. Your mother should have been here.
 b.



To implement this proposal, it would be necessary to modify the denotation of perfect aspect given in (46) slightly so that the aspectual head can combine with a situation argument first. This could be achieved with the modification in (65):

(65) $\llbracket \text{have}_{\text{perfect}} \rrbracket = \lambda s. \lambda P. \exists e_{\text{state}}. [P(e_{\text{state}})(s) = 1]$

The lexical entry in (46) has been modified so that the situation argument is introduced first. (I have also annotated the event variable to show that we are quantifying over eventualities that are states (e_{state}), just for ease of reference.) According to (65), the perfect is responsible for existentially quantifying over eventualities and locating them within situations (again, we have ignored times).

Given the set-up in (64), the perfect will manipulate the intensional parameter of a description in the same way other ‘intensional operators’ do. *Have* will be able to ‘bind’ the situation argument in the DP in those cases in which that situation pronoun is coindexed with the situation pronoun associated with *have*. When this happens, abstraction over the *have*-situation-variable will lead to abstraction over the DP-situation-variable, and the result will be that both variables will be co-bound. The proposition that will serve as argument of the modal will be (66):

$$(66) \quad \lambda s. \exists e_{\text{state}}. [\text{your-unique-mother-in-s-is-here}(e_{\text{state}}) \text{ in } s]$$

The proposition in (66) will be true in a situation *s* iff it includes an eventuality (state) of your mother in *s* being here. Given the presence of an operator able to bind the intensional parameter of the definite in the embedded clause, the presuppositions of the definite will be shifted to the situations/ worlds described by the verbal predicate and need not be satisfied in the actual world. Since there are no epistemic presuppositions, such situations may well be in worlds outside the context set.

5. Conditioning *should*

With a proposal for the meaning of *should* in hand, and an understanding of how the interpretation of aspectual heads may interact with modal claims, we can now turn to structures in which a *should* statement is conditioned by an *if*-clause. I will establish that there are (at least) two possible strategies when conditioning *should*. There are *epistemic if*-clauses, that put constraints on the set of epistemically available possibilities, and there are *deontic if*-clauses, that are interpreted within the scope of the deontic modal. I will argue in favor of this distinction on the basis of contrasts in the interpretation of tense morphology in the *if*-clause itself.

The structure of this section is as follows: In §5.1 we will examine the interpretation of *if*-clauses, noting variations in the interpretation of tense and the correlation with the interpretation of the *if*-clause. In §5.2 we will make a proposal for the structures that result in the different interpretations for tense. (These two sections argue for the view that there are two different kinds of *if*-clauses, and may be skipped by readers willing to take this for granted.) In §5.3 we will investigate the semantic contribution of the two types of *if*-clauses. Finally, in §5.4 I will examine the detachment patterns associated with the *if*-clauses.

5.1 The interpretation of tense in *if*-clauses

The interpretation of tense morphology in *if*-clauses associated with *should* can vary. In some cases, tense morphology appears to receive the standard, speech-time oriented, past/present interpretation. In others, it does not. We find examples with standard interpretation in (67):

- (67) a. If she returned the book late (*yesterday*), she should pay a fine.
 b. If she returned the book late (*yesterday*), she should have paid a fine.

Here we see a simple past perfective form in the antecedent, and a consequent clause whose shape varies depending on whether paying the fine is still an (epistemic) option or not (i.e. there is perfective or perfect morphology in the consequent). As far as the antecedent clause is concerned, past tense seems to receive a standard past interpretation, and the epistemic constraints associated with the perfective hold (we would only utter the sentences in (67) if the proposition that she returned the book late was compatible with the common ground).

When the antecedent clause has present tense morphology, the interpretation of tense can be shifted towards the future (we'll consider this a 'non-standard' interpretation of present tense morphology), or it can be deictically anchored to the speech time:

- (68) a. If she returns the library book late (*tomorrow*), she should pay a fine.
 b. If she still has library books at home (*right now*), she should have been given a warning weeks ago.

In these cases too, we see that aspect morphology in the consequent varies depending on the epistemic status of the proposition. Of interest to us here is that, in both cases, the antecedent clause is compatible with our knowledge (even though there is a state in the antecedent, which carries no context set presuppositions, we could not interpret (68b) as a 'counterfactual' deontic conditional, where the antecedent is not compatible with the common ground). When tense is present, the pattern is the same even in the presence of a perfect in the antecedent clause:

- (69) a. If she has returned all her library books late, she should receive a stiff fine.
 b. If she has kept pets in her apartment for several years, she should have been given a warning a long time ago.

A present perfect in the antecedent leads to a hypothesis compatible with our knowledge. We could not felicitously utter either (69a) or (69b) in a context in which it was known that the antecedent was false.

It is also possible to find *if*-clauses associated with *should* in which a past tense does not receive a standard interpretation. In such cases, we find past tense morphology with future shifted readings. Examples are provided in (70):

- (70) a. The goldfish died on Saturday, so she buried it in the garden. But if it had died next Monday (*instead*), she should have given it to the vet.
 b. She returned the library books on time, so everything was fine. But if she had returned them next week, she should have received a stiff fine.

As these examples show, even if we want to talk about the future, we use a past perfect in the antecedent clause if the antecedent is incompatible with the common ground. A present perfect in the antecedent clause would not achieve the same effects:

- (71) a. The goldfish died on Saturday. #If it has died next Monday (*instead*), she should (have) bury it in the garden.
 b. She returned the library books on time. #But if she has returned the library books next week, she should (have) be given a fine.

As well as future shifted past perfects, it seems possible to find future shifted past perfectives (they seem a bit more awkward):

- (72) a. The goldfish is fine so far. But if it died tomorrow, she should feed it to the cats.
 b. She still has the library books at home. If she returned them tomorrow, she should get a reward.

In these examples we see a past perfective antecedent with an interpretation shifted towards the future. The presuppositions associated with the perfective are evident in the usual way, and we would only use examples like these in circumstances in which the antecedent was compatible with our knowledge. There is, however, a ‘distancing’ effect achieved by the choice of a past perfective instead of a present perfective. We seem to use a past when the antecedent, though compatible with the common ground, seems less likely to be true (this effect has been dubbed by Iatridou (2000) ‘future-less-vivid’. We will return to this below.)

The data discussed in this section show that there are two parameters that affect the shape of *if*-clauses given their relation to the common ground. When a perfective aspectual head shows up in an *if*-clause, the context set presupposition that we discussed in §4 is activated, and the *if*-clause is felicitous only if it is compatible with the common ground. This is irrespective of whether tense morphology in the antecedent is present or past, or whether tense morphology receives a standard interpretation or not. When dealing with statives however, whether it be simple statives or perfects, tense is crucial. When tense morphology is present, the *if*-clause can only be used felicitously when it is compatible with the common ground (this is the case even though simple statives and perfects do not carry context set presuppositions). To stray outside the common ground, past morphology must be used with a non-standard interpretation. In the next two sections I will provide an account of this variation, beginning with the structural conditions that affect the distribution of tense features (§5.2) and continuing with the interpretation of *if*-clauses (§5.3).

5.2 Attaching *if*-clauses

In examining the conditionals above, we came across examples in which tense morphology in the antecedent receives its standard interpretation, and examples in which it appears ‘bleached’ of its standard interpretation, and shifted towards the future (this is the case for both present and past tense morphology). This pattern of data can be accounted for if we allow for the possibility that the *if*-clauses that appear associated with *should* are not all actually attached in the same place. In this section I will present two strategies for conditioning *should* with two possible attachment sites for *if*-clauses.

Before turning to the attachment of *if*, it is worth considering the possibility that in examples like *If she had returned the library book late, she should have paid a fine* we are dealing with a *should* statement embedded in the consequent clause of a counterfactual conditional.²² That is, the possibility that in examples like this there is a silent counterfactual modal (SM_{would}) that

²² See Frank (1997) for an analysis of deontic conditionals of the form ‘*if* α , *would have to have* β ’ as involving a deontic operator in the scope of a counterfactual. Arguments against a uniform view of ‘counterfactual’ deontics as involving a deontic operator in the scope of a counterfactual can be found in von Stechow and Iatridou (2008).

takes the antecedent as a restrictor and has a *should*-statement as the nuclear scope (73a). Could the logical form for this kind of conditionals look like (73b)?

- (73) a. If she had returned the library book late, SM_{would} she should have paid a fine.
 b. SM_{would} [*if she had returned the library book late*] [should [she have paid a fine]]

There is something intuitive about this idea. The intuition could be described as follows: in this kind of examples we seem to start out by imagining circumstances just like the actual world except for the fact that she returned the library book late, and then claim that in those circumstances, the best thing would have been for her to have paid a fine. However, while this intuition may indeed resonate with us, a proposal built on the logical form in (73b) would not deliver the correct truth conditions (at least, with the standard similarity-based analysis of counterfactuals following Lewis and Stalnaker). According to (73), the main operator in the conditional is a counterfactual modal. The semantics of counterfactuals will ensure that the domain of quantification of the modal consist of the most similar worlds to the actual world in which the antecedent is true. Given (73), we will then claim that in those worlds, the consequent is true (*she should have paid a fine*). But such an analysis allows similarity with the actual world to play a role that is bigger than what we actually observe in the interpretation of deontic conditionals. Consider the example in (74):

- (74) If she had built that building, she should have paid a fine.

Suppose that the woman in question is a reckless architect with a construction company of her own, who never bothers to carry out proper soil surveys before putting up high buildings. Given her nature, we know that if she had built that building (which is high), she would not have carried out the proper soil surveys. And if she had not carried out the proper soil surveys, she should have been fined. However, in spite of all this knowledge, we are not willing to judge (74) true in the circumstances described. This shows that the consequent clause is not evaluated in the most similar worlds in which she has built that building, and thus that *should* is not in the scope of a counterfactual modal.


5.2.1 *If*-clauses in the scope of *should*

In the examples discussed in §5.1 we have come across conditionals in which tense morphology in the antecedent clause does not receive its standard deictic interpretation. I will present a ‘sequence of tense’ analysis of this morphology (the morphology simply reflects the tense features in the immediate environment and does not feed into the interpretation). In my proposal, tense morphology in *if*-clauses in the scope of *should* copies a past morphology feature associated with the modal *should* (in the account presented here, this feature is not interpreted – it may be a fossilized piece of morphology left over from the days in which *should* was transparently the past form of *shall*).²³

²³ Recent literature has paid a lot of attention to the role of past tense in counterfactual constructions (a.o. Ogihara 2000, Iatridou 2000, Condoravdi 2002, Ippolito 2003, etc, Arregui 2004, 2008.) The matter has not received much discussion in the domain of deontic statements (see von Stechow and Iatridou 2008 for an exception). The proposal made here does not assign past tense any role in the interpretation of *should*.

I will illustrate a sequence-of-tense account of the distribution of tense morphology in *if*-clauses appealing to the referential theory of tense in Kratzer (1998), and the account of sequence of tense phenomena that is spelled out there (for other accounts of sequence of tense, some similar, some very different, the reader is referred to von Stechow 1995, Ogiwara 1996, Kusumoto 1998, 2005, etc.). Let us begin with the cases in which past tense morphology does not receive a standard interpretation. I will argue that in these cases, tense in the antecedent clause is a variable tense that receives an interpretation from a variable assignment, and does not carry deictic constraints. Following Kratzer (1998) (a.o.) I will call this tense a *zero* tense (\emptyset). Zero tenses lack morphology of their own. The tense morphology associated with a zero tense is agreement morphology that copies tense features higher up in the clause. My proposal is that when we observe past tense morphology in the antecedent clause with an interpretation that is shifted towards the future, we are dealing with a zero tense interpreted within the scope of *should* that copies past tense features associated with the modal. In such cases, the structure could look like (75):

- (75) a. If the goldfish had died next Monday (instead), she should have given it to the vet.
b.



Since zero tense is a variable tense, its interpretation will depend on the contextually salient variable assignment. As a variable tense, it will give rise to bound variable readings in syntactic configurations in which it is bound by a higher operator. In the structure in (75b), the zero tense is bound by an index higher up in the structure (represented informally as λ_i), and as a result, the *if*-clause denotes a property of times (for a proposal regarding indices as abstractors, the reader is referred to Heim and Kratzer (1998)) (for a more detailed discussion of a proposal along these lines, see Arregui (2004, 2008)).


When the *if*-clause is interpreted in the scope of the modal, the modal is responsible for shifting its interpretation to a non-past time. Both the *if*-clause and the consequent clause are interpreted as future-shifted propositions, with a reference time that is either present or future. Even though there is past morphology in the antecedent clause, the antecedent proposition is not affected by the semantics of a deictic past tense.

The examples in (72) showed that it is possible to shift the interpretation of *if*-clauses with past perfective morphology towards the future too. In such examples, the interpretation of tense follows the same lines as in (75), but a difference arises because of the context-set presuppositions associated with perfective aspect: future-shifted past perfective *if*-clauses can only be felicitous if the antecedent proposition is compatible with the common ground.

5.2.2 *If*-clauses above *should*

We have seen in §5.1 that when we are not dealing with future-shifted past morphology in the *if*-clause, the antecedent proposition can only be felicitous if it is compatible with the common ground. I would like to propose that in these cases the *if*-clause is not interpreted in the scope of *should*, but functions instead as an epistemic restrictor, adding conditions to the common ground. Within a dynamic framework, it would be possible to spell out the semantics of *if*-clauses in a manner that allowed them to interact with the common ground/context set directly (for some discussion of conditionals in a dynamic setting, the reader is referred to a.o. von Stechow (1999), Kaufmann (2005)). However, since we are working with a static framework, I will cash out the proposal by treating epistemic *if*-clauses as restrictors on a silent epistemic modal higher up in the clause. The fact that in these cases present tense morphology may appear bleached of its usual deictic meaning with an interpretation shifted towards the future indicates that the silent modal is associated with present tense features. In parallel with (75), I propose structures like in (76):

- (76) a. If she returned the library book late (yesterday), she should pay a fine.
- b. If she returns the library book late (tomorrow), she should pay a fine.
- c.



According to (76c), the antecedent clause in examples like (76a/b) restricts a silent epistemic modal high up in the clause, and the *should*-clause is actually the nuclear scope of that operator. If tense morphology in the antecedent is present (as in (76b)), the tense in the antecedent could well be a zero tense whose features copy the features of a higher present tense (i.e. *tense* = \emptyset_i). If tense morphology in the antecedent is past, it could only be a case of a true, deictic past tense (in this position, it could not be a case of agreement tense) (i.e. *tense* = *past*).

The prediction is when the *if*-clause has a zero tense with present morphology, or a deictic past tense, the antecedent proposition will be compatible with the common ground (since it is restricting an epistemic modal). With this interpretation of tense, we expect stative and perfect *if*-clauses to also require compatibility with the common ground. It won't be only

if-clauses with perfective aspect that are tied to the context set. This explains the distribution of data we saw in §5.2.

There is a strong ‘counterfactual’ flavour associated with antecedent clauses with past perfect morphology. Consider the contrast in (77):

- (77) a. If she returned the library book late (yesterday), she should pay a fine within a week.
 b. If she had returned the library book late (yesterday), she should have paid a fine within a week.

In (77a) we observe simple past tense morphology in the antecedent, and perfective aspect. We expect to find the *if*-clause attached above *should* and we expect to find context-set presuppositions (the *if*-clause is adding restrictions to the common ground). This is the form that is marked as compatible with what we know (are willing to take for granted). In (77b) we observe a past perfect antecedent. This form is not marked as compatible with what we know. There is strong counterfactual intuition associated with (77b). We would normally understand that she returned the book on time. But this is not required by the semantics of *should*. I will make the hypothesis that the counterfactual effect arises as an implicature because of the contrast with a form marked as ‘knowledge compatible’ (this type of account of counterfactual effects has also been proposed for past perfect *if*-clauses attached to *would*-conditionals, see a.o. Ippolito (2003)). As an implicature, the effect can be eliminated, leaving the assertion intact. This is what we observe in examples like (78):

- (78) A: If she had returned the library book late, she should have paid a fine.
 B: She did return the library book late!
 A: Well, then, she should have paid a fine!

By choosing a past perfect, A gives rise to the implicature that she didn’t return the library book late. The fact that the implicature can be rejected and we can still draw the inference shows that it is not really part of the content of the conditional.

It may be possible to understand the ‘future less vivid’ effects of past perfective antecedents with similar reasoning. Consider the contrast in (79):

- (79) a. The goldfish is fine so far. But if it died tomorrow, she should feed it to the cats.
 b. The goldfish is fine so far. But if it dies tomorrow, she should feed it to the cats.

Both conditionals in (79) would be uttered only in contexts in which it is compatible with what is known that the goldfish dies tomorrow. We can explain this because both examples have perfective aspect in the antecedent, and thus trigger context-set presuppositions. The interpretation of tense morphology indicates that in (79a) the antecedent clause is restricting *should* (it has future-shifted *past* morphology) whereas in (79b) the antecedent clause is an epistemic restrictor (it has future-shifted *present* morphology). It may be that the ‘vividness’ associated with (79b) comes about because of the choice to mark explicitly that the *if*-clause is modifying an epistemic modal/ interacting directly with the context set.

5.3 Interpreting *if*-clauses in the scope of *should*

Making use of proposals for sequence of tense phenomena, we were able to argue that differences in the interpretations of tense morphology in *if*-clauses can be correlated with differences in the attachment sites of *if*-clauses. *If*-clauses can either serve as restrictors on an epistemic modal with present tense features, or on a *should*-modal with past tense features. This section will be dedicated to the interpretation of *if*-clauses in the scope of *should*. I will not discuss the case of *if*-clauses as epistemic restrictors (a proposal on this topic lies outside the scope of my paper, the reader is referred to von Stechow and Gillies (2007) and von Stechow and Gillies (2008) for discussions of the complex issues related to epistemic modality). I will simply note that when *if*-clauses restrict epistemic modals, factual detachment is expected to be valid (after all, the actual world is one of the worlds compatible with what we know) and deontic detachment is not (this matter will be taken up in §5.4).

The main objective of this section is to argue that *if*-clauses are able to put restrictions on the extensions of actual world situations that enter the domain of comparison of the modal. The analysis will be built in steps, and along the way we will motivate both the need to appeal to minimal extensions and the need to appeal to facts. In the end, we will settle for an analysis in which *if*-clauses serve as restrictors on the fact-sensitive *should* in (49). The presentation begins with the example in (80):

(80) If he were in trouble, he should be given assistance.

Let us suppose that we are willing to judge (80) true simply on the basis of the fact that the law requires that people in trouble be given assistance. And imagine the world is such that he is not in trouble, and he is not given assistance. A natural way to set up a restrictor-analysis of *if*-clauses would be to suppose that in examples like (80), we go through all the situations in the world compatible with his being given assistance. The sentence will be true iff for all relevant* situations s that can be (modally) extended to situations s' in which he is in trouble, it is the case that the situations s' have best extensions in which he is given assistance. A concrete proposal is found in (81):

(81) **[[if α , should β]]** is true in w iff (preliminary)
 $\forall s \leq w$ that are relevant*,
 $\{s': s \leq_m s' \ \& \ \alpha \text{ is true in } s'\} \subseteq$
 $\{s': \forall w': w' \text{ is a best extension of } s', \beta(w') = 1\}$

Though seemingly reasonable, the proposal in (81) makes some incorrect predictions. In extending the actual-world situations to situations in which the antecedent is true, we consider ALL extensions in which the antecedent is true. However, not all extensions in which the antecedent is true can have extensions in which the consequent is true. For, amongst the situations in which the antecedent is true there will be quite large situations, in which we will find many things besides the original actual world situation and the features that make the antecedent true. In these situations, the consequent may be false. These situations won't have extensions in which the consequent is true, and therefore they won't have best extensions in which the consequent is true. Allowing quantification over the 'antecedent' situations to include situations in which the fate of the consequent has already

been decided will not allow us to see the work of the deontic machinery (the extension to best-extensions will be trivialized).

Intuitively, things go wrong with (81) because it allows into the set of antecedent-situations situations that are ‘bigger’ than we would like. We would like the set of antecedent situations to be made up of the original actual world situations with a ‘small’ extension to allow the antecedent proposition to be true. We don’t want to allow in ‘extra’ stuff that could cloud the working of the deontic modal. To capture this effect, we will make use of quantification over ‘minimal extensions’, defined as below:

(82) *α -minimal extension*

Given a situation s and a proposition α , s' is an α -minimal extension of s iff $s \leq_m s'$ and α is true in s' and there isn’t a situation s'' such that $s \leq_m s'' \leq_m s'$ and α is true in s'' and $s'' \neq s'$.

To see how this definition works, consider the following case: imagine that s is a situation consisting only of a stretch of sand in a beach, and α is the proposition *there is a man*. Consider now a situation s' consisting only of s and a man. s' will count as a α -minimal extension of s because it is true that there is a man in s' and there isn’t any s'' such that $s \leq_m s'' \leq_m s'$ and there is a man in s'' and $s'' \neq s'$. Consider now an extension of s to a situation s' in which there are two men. Then s' will not be a α -minimal extension of s because there will be a situation s'' such that $s \leq_m s'' \leq_m s'$ and there is a man in s'' and $s'' \neq s'$.

With the definition in (82) at hand, we can revise the proposal in (81) as follows:

(83) **[if α , should β]** is true in w iff (preliminary)
 $\forall s \leq w$ that are relevant*,
 $\{s': s' \text{ is a } \alpha\text{-minimal extension of } s\} \subseteq$
 $\{s': \forall w': w' \text{ is a best extension of } s', \beta(w') = 1\}$

Given (83), (80) will be true iff for every relevant* situation in the world it is the case that all minimal extensions in which he is in trouble are such that in their best extensions, he is given assistance.²⁴ By using minimal extensions, we make sure that only the ‘stuff’ needed to make the antecedent true is added.

The proposal in (83) is better than the proposal in (81), but another adjustment needs to be made. Consider the conditional in (84):

(84) If he had forgotten to pay his taxes (last May), he should have been fined \$200.

Suppose that according to the law, the fine one must pay for not paying one’s taxes is linked to one’s salary. In his case, it would have been \$200 because he earns \$20 000 per year. Suppose also that as a matter of fact he paid his taxes and was not fined \$200. What happens with (84)? The proposal in (83) tells us that (84) will be true iff for every relevant* situation

²⁴ As before, *relevant** is defined with respect to the main clause proposition (this is the nuclear scope of the conditional, not its restrictive clause/ antecedent). Notice that for the situations that are incompatible with the *antecedent* the claim that the extensions in which the antecedent is true have best extensions in which the consequent is true will be trivial. With the current set up, such situations count towards the set of situations that verify the truth of the conditional (we could also have explicitly stated that such situations should be excluded from the domain of quantification).

in the world, it is true that for all minimal extensions in which he forgets to pay his taxes, in the best extensions he is fined \$200 (last May). But this won't give us the right truth-conditions. In the actual world there is a situation in which he earns \$20 000, but there are also many situations in which this is not true (situations that do not include him at all!). For those situations, minimal extensions to situations in which he forgot to pay his taxes won't result in situations in whose best extensions he is fined \$200. Earning \$20 000 is not in itself very valuable, so it won't necessarily be part of the best extensions of the antecedent situations.

To get the truth conditions of (84) right, we must make sure that the antecedent situations include the actual world features that make it true that he earns \$20 000 (but exclude the features that makes it true that he paid his taxes and was not fined). The facts regarding his salary must be taken into account when figuring out what possibilities make it into the domain of comparison of the modal. We can obtain this result by going back to our original fact-sensitive proposal for the semantics of *should* (49), and allow *if*-clauses to function as restrictors in that analysis (85):

$$(49) \quad \llbracket \text{should } \alpha \rrbracket \text{ is true in } w \text{ iff} \\ \forall s \leq w \text{ that are relevant*}, \\ [\exists s': s \leq s' \ \& \ \forall w' \text{ that are best extensions of } s', \alpha(w') = 1]$$

$$(85) \quad \llbracket \text{if } \alpha, \text{ should } \beta \rrbracket \text{ is true in } w \text{ iff} \\ \forall s \leq w \text{ that are relevant*}, \\ \{s'': \exists s': s \leq s' \ \& \ s'' \text{ is an } \alpha\text{-minimal extension of } s'\} \subseteq \\ \{s'': \forall w': w' \text{ is a best extension of } s'', \beta(w') = 1\}$$

The proposal in (85) treats *if*-clauses as restrictors, and at the same time allows specific facts to have an effect on the semantics of the conditional. According to (85), a conditional *if* α , *should* β will be true iff for all the (relevant) situations in the actual world, the α -minimal extensions of some actual-world extension all have best-extensions in which β is true.

The proposal in (85) will provide better results than (83). Consider (84) again. For all the (relevant) situations in the actual world it will be true that the set of minimal extensions of some actual world extension (i.e. a situation that includes the situation quantified over and the situation that he earns \$20 000 per year) in which he forgot to pay his taxes will be a subset of the set of situations that have as best extensions worlds in which he was fined \$200.

5.4 Detachment patterns

We can finally turn our attention to detachment patterns and tackle the puzzles of factual detachment and deontic detachment discussed in §1. We will be interested both in the case of epistemic restrictors and restrictors in the scope of *should*. As was noted earlier, when *if*-clauses are interpreted epistemically (as a restrictor on a silent epistemic modal), we predict that factual detachment will be valid. Consider, for example, (86):

$$(86) \quad \text{If she returned the library book late (yesterday), she should pay a fine.}$$

The *if*-clause in (86) acts as an epistemic restrictor (notice the standard deictic interpretation of tense morphology in the antecedent clause). The conditional in (86) will be true iff it is the case that for every world in the context set in which she returned the library book late yesterday, it is true that she should pay a fine. If (86) is true and the actual world is a world in which she returned the library book late yesterday, then the actual world will be a world in which she should pay a fine. So, factual detachment is valid.

Deontic detachment, however, will not be supported by epistemic *if*-clauses. It would be difficult to access intuitions regarding deontic detachment with the example in (86) (we would have to consider that she should have returned the library book late!). So, we will switch to example (87), in which we also see an epistemic *if*-clause:

(87) If she remembered to phone her mother yesterday, she should feel proud of herself.

(87) tells us that in all the worlds in the context set in which she remembered to phone her mother yesterday, it is true that she should feel proud of herself. Now, imagine that it is true that she should have remembered to phone her mother yesterday. This means that for every relevant* situation *s* in the actual world (compatible with her remembering to phone her mother yesterday) there exists a situation *s'* such that $s \leq s'$ and in all the best extensions of *s'*, she remembered to phone her mother yesterday. This can be true whether she did phone her mother yesterday or not. Suppose she didn't. Then it will be false that she should feel proud of herself (for the sake of simplicity, let us just assume that nothing else would warrant her feeling proud of herself): it is not true that for every relevant* situation *s* in the actual world (compatible with her feeling proud of herself) there exists a situation *s'* such that $s \leq s'$ and in the best extensions of *s'*, she feels proud of herself. Given that it is possible for (87) to be true, for the *should*-statement *She should have remembered to phone her mother yesterday* to be true, and for the *should*-statement *She should feel proud of herself* to be false, deontic detachment is not valid.

We can now turn to *if*-clauses in the scope of *should*. We will examine factual detachment first. Consider (88):

- (88) a. You: She returned the library book on time. But if she had returned it late, she should have been fined.
 b. Me: Actually, she did return the library book late!
 c. You: Oh! Well..., then she should have been fined.

The conditional in (88a) is true iff for every relevant* situation *s* in the actual world (compatible with her being fined) there exists some situation *s'* such that $s \leq s'$ and for all minimal extensions of *s'* in which she returns the book late, the best extensions are such that she is fined. Suppose also that (88b) is true and there is a situation in which she returned the book late. Then it follows that (88c) will be true too. If (88a) and (88b) are both true, then some of the antecedent situations quantified over by (88a) will be actual world situations. For every situation in the actual world compatible with her being fined there is an extension to a situation in the actual world (in which she returns the book late) such that in its best extensions, she is fined. This makes (88c) true.

Finally, let us consider the pattern of deontic detachment and examine the sentences in (89) (let us suppose she did not give money to charity last May and did not receive a tax discount):

- (89) a. If she had given money to charity (*last May*), she should have received a tax discount.
 b. She should have given money to charity (*last May*).
 c. She should have received a tax discount.

Given the proposals in (49) and (85), it will not be possible to infer (89c) from the truth of (89a) and (89b). Suppose that (89a) is true. Then, for all relevant* situations s in the actual world (compatible with her receiving a tax discount) there exists a situation s' such that $s \leq s'$ and for all s'' that are minimal extensions of s' in which she gives money to charity, the best extensions are such that she receives a tax discount. If (89b) is also true, then for all relevant* situations s in the actual world (compatible with her giving money to charity last May) there exists a situation s' such that $s \leq s'$ and in all the best extensions of s' , she gave money to charity last May. But (89c) will not be true in the circumstances described above. For (89c) to be true, it should be the case that for all relevant* situations s in the actual world (compatible with her receiving a tax discount) there exists a situation s' such that $s \leq s'$ and in the best extensions of s' , she received a tax discount. Arguably, in the best extensions of many actual world situations, she will have given money to charity last May, so in the best extensions, she will also receive a tax discount. But in the circumstances described above, this won't be true for all situations s compatible with her receiving a tax discount. Given that in the actual world she did not donate money to charity last May, then there will be at least one situation s'' in the actual world that won't have as best extension worlds in which she donated money to charity last May and gets a tax discount. That fact about the actual world will be visible and will prevent the modal statement in (89c) from being true. As a result, the pattern of deontic detachment will not be valid.

The detachment patterns described above and the fact-sensitive semantics for *should* have many interesting consequences. On the one hand, as we have already noted, we do not predict inconsistencies between the *should*-statements in (90):

- (90) a. She should have returned the library book on time.
 b. She should have paid a fine.

Given the existential quantifier over situations in *should*, and the differences regarding the actual world situations being quantified over, (90a) and (90b) could both be true in the actual world (as we saw, this was a result that was also predicted by the time-sensitive proposal we discussed in §2). But what about examples like (91) (based on the original examples by Feldman)? Notice that the *if*-clause in (91b) is an epistemic restrictor (it has a future-shifted present tense):

- (91) a. The doctor should give medicine A on Monday and the doctor should give medicine A on Tuesday.
 b. If the doctor gives medicine B on Monday, he should give medicine B on Tuesday.
 c. The doctor will give medicine B on Monday.

Do we predict that (92) is true?

- (92) The doctor should give medicine A on Tuesday and the doctor should give medicine B on Tuesday.

Even though the conditional in (91b) does validate factual detachment, the proposal does not predict that (92) is true. The problem will arise at the level of the presuppositions regarding the context set triggered by aspect in the embedded clause. Notice that if the proposition that she gives medicine B on Monday is accepted as true, then all the worlds in the context set will be such that the doctor gives the patient medicine B on Monday (indeed, armed with this knowledge, we are happy to detach the unconditioned statement *The doctor should give medicine B on Tuesday*). But then, the context set presuppositions associated with perfective aspect in (91a) will not be satisfied (the context set presuppositions triggered by (91a) is that it is possible, as far as the common ground is concerned, that the doctor give the patient medicine A on Monday and that he give the patient medicine A on Tuesday). The problem arises because the conditions needed for detachment in (91b) are incompatible with the context set presuppositions in (91a).

There are two slightly different predictions made here. One is that we would never utter the sentences in (91) as one assertion during a conversation.²⁵ This is because the first sentence carries presuppositions that are incompatible with the last one. The other is that it will not be possible to draw the inference that (92) is true from the sentences in (91). This is because the first conjunct carries a presupposition that is incompatible with the information we must take for granted in order to assert the second conjunct.

The proposal presented here thus predicts an asymmetry between simple *shoulds* and *should haves*. We predict that one could truthfully assert (93):

(93) She should have returned the book on time and she should have paid a fine.

The circumstances in which we might utter (93) require some thinking. Maybe we are making a list of all the things she did wrong in the past month. We note that she ignored her duties twice. There are two different things she should have done and didn't do. However, we could not truthfully assert (94):

(94) She should return the book on time and she should pay a fine.

We predict that we cannot felicitously assert (94) because the information we need to have in order to claim that she should pay a fine is inconsistent with the presuppositions triggered by the claim that she should return the book on time. The result is that (94) cannot be understood as a claim regarding two independent obligations. This prediction appears in line with the intuitions reported in the literature.

6. Conclusions

In this paper we have investigated the relation between the truth of deontic claims and the facts in the world. We have used the inference patterns licensed by *should*-conditionals as a window to peer into the semantics of the modal. The proposal presented in the paper has

²⁵ A small clarification is probably needed. The prediction is that we would not assert the sentences in (91a/b) and add the proposition that she is given medicine B on Monday. Whether the *will*-sentence in (91c) simply denotes the proposition true if she is given medicine B on Monday or whether it has modal flavours is a problem that lies outside the scope of my investigation.

several parts: a situation-based semantics for the modal *should*, a view of the propositions embedded under *should* that allows aspect to play a crucial role in anchoring propositions to the context set, and a proposal for *if*-clauses that distinguishes between epistemic *if*-clauses and *if*-clauses in the scope of *should*, treating the latter as restrictors on the quantificational domain of the modal.

An important part of the account is the idea that restrictions on the interpretation of a modal (like *should*) could arise from the interaction between the semantics of the modal and features of the embedded clause (in this case, aspect). It has been widely accepted, since the bi-dimensional approach to modality in Kratzer (1981, etc.), that the interpretation of modals may call upon information of different sorts (in Kratzer's work, modals are associated both with a *modal base* and an *ordering source*). In the proposal presented here, some of the information standardly associated with the modal base in a Kratzer (1981)-style account is captured by the choice of aspectual heads in the embedded clause. Aspectual head tells us whether we are quantifying over epistemically accessible worlds or not (with perfective aspect we quantify over epistemically accessible worlds, which would be the case if we said that a modal had an epistemic modal base in Kratzer's system, whereas with perfects and states, we don't). Further investigations would be required to find out to what extent it would be possible to capture the interaction between the information Kratzer places in the modal base and ordering source of the modal compositionally in the structure, through the interaction of the interpretation of the modal and restrictions that may come from tense and aspect in the embedded clause.

We have addressed some of the concerns regarding CTDs found in the philosophical literature. We have provided a semantics for *should* that allows primary duties to be compatible with secondary duties. This goes a long way towards defusing Chisholm's Paradox. We have also characterized the difference between simple *should*-statements and *should have*s in relation to the context set. As a result, our take on Chisholm's sentences (repeated below) turns out to be a little different from his own:

- (1)
 - a. She should return the library book on time.
 - b. If she returns the library book late, she should pay a fine.
 - c. It should be the case that, if she returns the library book on time, she does not pay a fine.
 - d. She returns the library book late.

In presenting sets of sentences like (1), Chisholm took as a starting point the fact that they could all be true and independent of each other. But, given the proposal made here, this needs to be reconsidered. Taking into account the context set presuppositions associated with (1a), we would not expect to find a felicitous utterance that included (1a) and (1d). The information in (1d) is incompatible with the presuppositions of (1a). The prediction that (1a-d) would not be made as a single utterance seems on the right track. But it is important to note that Chisholm did not consider this issue, so in providing an 'answer' to his puzzle, we have modified it slightly.

Finally, we have made a proposal for the meaning of *should* that claims that everything that happens in the world is important for the truth-conditions of a *should*-statement. The proposal makes the truth of *should*-statements depend entirely on what actually happens. The view allows us to understand the conditions in which we may be willing to detach an *if*-clause from a CTD-conditional and utter an unconditioned *should*-statement. In examining CTDs, Prakken and Sergot (1994, 1996) and Carmó and Jones

(2002) have pointed out that we seem willing to detach unconditioned deontic claims when we ‘take the antecedent for granted’ (Prakken and Sergot, for example, have proposed to model this in logic with a conditional operator appealing to contexts). Taking a version of a Prakken and Sergot example, suppose it is true that *If they execute the dissenter at dawn, they should offer him a cigarette*. We will be willing to claim that *They should offer him a cigarette* in a context in which we have taken for granted that they will execute him at dawn. And if it turns out that they execute him at dawn, we will have said something true. The fact that the semantics of *should* predicts that truth depends on what actually happens allows us to understand why our willingness to detach *if*-clauses from CTDs depends on the things we are ready to take for granted.

References

- Aqvist, Lennart (2002). Deontic Logic. In Gabbay, D. and Guenther, f. (eds.) *Handbook of Philosophical Logic*, Volume 8, Kluwer, The Netherlands, 147-264..
- Alchourrón, Carlos (1993). Philosophical Foundations of Deontic Logic and the Logic of Defeasible Conditionals. In Meyer, J.J. and Wieringa, R.J. (eds.) *Deontic Logics in Computer Science: normative System Specification*, Wiley and Sons, Cichester, 43-84.
- Arregui, Ana (2004). *On the Accessibility of Possible Worlds: the Role of Tense and Aspect*. PhD dissertation, University of Massachusetts, Amherst.
- Arregui, Ana (2005). Layering modalities. University of Ottawa, ms.
- Arregui, Ana (2007). When Aspect Matters: the Case of *Would*-conditionals. *Natural Language Semantics* 15, 221-264.
- Arregui, Ana (2008). On Similarity in Counterfactuals, University of Ottawa, ms.
- Bennett, Jonathan (2003). *A Philosophical Guide to Conditionals*, Clarendon Press, Oxford.
- Bhatt, Rajesh (1999). *Covert Modality in Non-Finite Contexts*, PhD disseration, University of Pennsylvania.
- Bhatt, Rajesh (2006). *Covert Modality in Non-finite Contexts*, Walter de Gruyter.
- Bhatt, Rajesh and Roumyana Pancheva (2001). Conditionals, Ms. University of Texas at Austin and USC.
- Carmó, José & Andrew Jones (2002). Deontic Logic and Contrary-to-Duties. In Gabbay, D. and Guenther, F. (eds.) *Handbook of Philosophical Logic*, Volume 8, Kluwer, The Netherlands, 265-343.
- Chisholm, Roderick M. (1963). Contrary-to-Duty Imperatives and Deontic Logic. *Analysis* 24, 33-36.

- Condoravdi, Cleo (2002). Temporal Interpretation of Modals: Modals for the Present and for the Past. In Beaver, D., Kaufmann, S. and Casillas, L. (eds.) *The Construction of Meaning*, CSLI Publications, 59-88.
- Copley, Bridget (2006). What Should *Should* mean? (ms. under review)
- Enç, Mürvet (1986). Towards a Referential Analysis of Temporal Expressions. *Linguistics and Philosophy* 9, 405-426.
- Enç, Mürvet (1996). Tense and Modality. In Lappin, S. (ed.) *The Handbook of Contemporary Semantic Theory*, Blackwell, Oxford, 345-358.
- Feldman, Fred (1986). *Doing the Best We Can*, Springer.
- von Fintel, Kai. (1995). A minimal theory of adverbial quantification, ms. MIT
- von Fintel, Kai (1997-2005). How to count situations, ms. MIT.
- von Fintel, Kai (1999). Counterfactuals in a Dynamic Context, in Kenstowicz, M. (ed.) *Ken Hale: A Life in Language*, MIT Press.
- von Fintel, Kai and Anthony Gillies (2007). An Opinionated Guide to Epistemic Modality, in Szabo Gendler, Tamar and Hawthorne, John (eds.) *Oxford Studies in Epistemology: Vol. 2*, Oxford University Press, Oxford, 32-62.
- von Fintel, Kai and Sabine Iatridou (2008). How to Say *Ought* in Foreign: The Composition of Weak Necessity Modals. In Guéron, J. and Lecarme, J. (eds.) *Time and Modality*, Springer, 115-141.
- von Fintel, Kai and Anthony Gillies (2008). CIA Leaks. *The Philosophical Review* 117, 77-98.
- Frank, Anette (1997). *Context Dependence and Modal Constructions*, PhD dissertation, University of Stuttgart.
- Gennari, Silvia (2003). Tense meanings and temporal interpretation. *Journal of Semantics* 20, 35-71.
- Hacquard, Valentine (2006). *Aspects of Modality*, Ph.D. Dissertation, MIT.
- Heim, Irene (1982). *On the Interpretation of Definite and Indefinite Noun Phrases*, PhD dissertation, University of Massachusetts.
- Heim, Irene and Angelika Kratzer (1998). *Semantics in Generative Grammar*. Blackwell.
- Iatridou, Sabine (2000). The Grammatical Ingredients of Counterfactuality. *Linguistic Inquiry* 31, 231-270.

- Ippolito, Michela (2003). Presuppositions and Implicatures in Counterfactuals. *Natural Language Semantics* 11, 145-186.
- Ippolito, Michela (2006). Semantic composition and presupposition projection in subjunctive conditionals. *Linguistics and Philosophy* 29, 631-672.
- Ippolito, Michela (2008). Subjunctive Conditionals. In Grønn, Atle (ed.) *Proceedings of Sinn und Bedeutung* 12, ILOS, Oslo.
- Kaufmann, Stefan (2005). Conditional truth and future reference. *Journal of Semantics* 22, 231-280.
- Kratzer, Angelika (1981). The Notional Category of Modality. In H.J. Eikmeyer and H. Rieser (eds.) *Words, Worlds, and Contexts*, de Gruyter, Berlin-New York, 38-74. (Reprinted in Paul Portner and Barbara Partee (eds.) (2002). *Formal Semantics. The Essential Readings*. Blackwell, Oxford.)
- Kratzer, Angelika (1989). An investigations of the lumps of thought. *Linguistics and Philosophy* 12, 607-653.
- Kratzer, Angelika (1991). Modality. In von Stechow, A. and Wunderlich, D. (eds.) *Handbuch Semantik/Handbook Semantics*. de Gruyter, Berlin and New York, 639-650
- Kratzer, A. (1998). More Structural Analogies Between Pronouns and Tenses. In Strolowitch, D. and Lawson, A. (eds.) *Proceedings of Salt VIII*. Ithaca: CLC Publications.
- Kratzer, A. (2002). Facts: Particulars of Information units? *Linguistics and Philosophy*, 25, 655-670.
- Kratzer, A. (2006-2008). Situations in natural language semantics. (In the Stanford Encyclopedia of Philosophy.)
- Kusumoto, K. (1998). *Tense in embedded contexts*. PhD dissertation, University of Massachusetts, Amherst.
- Kusumoto, K. (2005). On the quantification over times in natural language. *Natural Language Semantics* 13, 317-357.
- Lewis, D. (1973). *Counterfactuals*. Harvard University Press, Cambridge.
- Lewis, D. (1979). Counterfactual Dependence and Time's Arrow. *Noûs*, 13, 455-476. (Reprinted in D. Lewis (1986), *Philosophical Papers*, Vol. 2, Oxford: Oxford University Press.)
- Musan, Renate (1995). *On the Temporal Interpretation of Noun Phrases*, PhD dissertation, MIT.
- Ogihara, T. (1996). *Tense, Attitudes, and Scope*. (Dordrecht: Kluwer Academic Publishers)
- Parsons, Terence (1994). *Events in the semantics of English: a study in subatomic semantics*, MIT Press, Cambridge.

Percus, Orin. (2000). Constraints on Some Other Variables in Syntax. *Natural Language Semantics* 8, 173-229.

Prakken, Henry & Mark Sergott (1996). Contrary-to-Duty Obligations. *Studia Logica* 57, 91-115.

Schulz, Katrin (forthcoming). "If you wiggle A, the B will change": Causality and the Meaning of Conditional Sentences. (Synthese).

Stalnaker, Robert (1975). Indicative Conditionals. *Philosophia* 5: 269-286. (Reprinted in Stalnaker, Robert (1999). *Context and Content*, Oxford University Press, Oxford, 63-77

von Stechow, Arnim (1995). On the Proper Treatment of Tense. *Proceedings of Semantics and Linguistic Theory V*, 362-386.

Veltman, F. (2005). Making counterfactual assumptions. *Journal of Semantics* 22, 159-180.