

COUNTERFACTUAL-STYLE REVISIONS IN THE SEMANTICS OF DEONTIC MODALS¹

Ana Arregui
University of Ottawa

1. INTRODUCTION

This paper investigates some aspects of the interpretation of deontic modals. According to standard analyses, deontic modals make quantificational claims regarding what happens in the best possible circumstances. What counts as ‘best’ usually depends on a complex interaction between facts and ideals. Context has been argued to play an important role, setting up a salient measure of goodness that is used to compare the relevant options. The range of options under comparison depends on what is going on (the facts).² In evaluating whether a deontic statement is true or not, we usually take certain facts for granted and set others aside. This paper investigates the factual grounding of deontic statements. How do we pay attention to facts in the evaluation of deontic modals? Do we consider facts in isolation, or do we take into account the web of dependencies that holds between facts in the actual world? In this paper I will show that to get the interpretation of deontic modals right, we need to track dependencies between facts. If we isolate facts from each other, we will get the truth conditions of deontic statements wrong.³

The claim that factual dependencies matter leads to an interesting parallelism between the interpretation of deontic modals and the interpretation of counterfactuals. In the literature on counterfactuals, it has long been observed that it is necessary to pay attention to dependencies between facts. As far as counterfactuals go, ‘facts stand and fall together’ (e.g. Kratzer 1981b, Veltman 2005): if we don’t pay attention to dependencies

¹ I am very grateful to three anonymous reviewers and to Philippe Schlenker for observations and suggestions. It was only with their help that I was able to see the shape this work should take. I am also very grateful to María Biezma for comments. Earlier versions of this work were presented at Sinn und Bedeutung 14 held at the University of Vienna on September 28-30 2009 and in the semantics workshop held in honor of Angelika Kratzer at MIT on December 6 2008. I am very grateful to the audiences for comments. Remaining errors are my own.

² I will use the word *fact* informally.

³ In Arregui (2010) I presented independent arguments for a fact-based semantics for deontic modals on the basis of detachment patterns for conditionals. In this paper, my concern will be with *dependencies* between facts.

between facts, we get the semantics of counterfactuals wrong. In this paper I claim that deontics are like counterfactuals in this respect: if we don't pay attention to dependencies between facts, we will get the semantics of deontics wrong. I argue that to account for the similarity we need an analysis of deontic modals that appeals to a notion of 'revision' like the one found in counterfactuals, one that pays attention to dependencies between facts. I will spell out such a proposal on the basis of the analysis of modals and counterfactuals in the work of Kratzer and Veltman (e.g. Kratzer 1981a, Veltman 2005).

Throughout the paper, *should* will be used as a stand in for a deontic necessity modal.⁴ The interaction between primary and secondary duties will provide important empirical evidence in favor of the view sketched above. To clarify terminology, I will begin with an example. Consider the deontic statements in (1), which spell out household regulations regarding the storage of wine:

- (1) a. The wine should be stored in the fridge.
b. If it is not stored in the fridge, it should be stored in the cellar.

The pattern in (1) is fairly common. We often identify a best option (1a), as well as a back-up plan (1b). Such back-up conditionals are known in the literature as 'Contrary to Duty Imperatives' (CTDs).⁵ In the example above, (1a) counts as the 'primary' duty. It indicates the best option in an absolute/unconditioned way. The CTD in (1b) spells out the 'secondary' duty: it indicates the best option in less-than-ideal circumstances in which the primary duty is not satisfied. In (1) the primary duty is that the wine should be stored in the fridge, and the secondary duty is that the wine should be stored in the cellar.

Examples like (1) are curious because the conditions in which the CTD instructs us to move to the secondary duty are circumstances in which the primary duty still holds. Imagine that you come home and find the wine sitting on top of the fridge. What will you say? *The wine should be stored in the fridge* or *the wine should be stored in the cellar*? It is true that the

⁴ There are important differences between deontic modals (e.g. with respect to epistemic possibilities, quantificational strength) that will not be addressed in this paper. For some discussion, the reader is referred to von Fintel and Iatridou (2008) and Copley (2006).

⁵ The term was first used by Chisholm (1963). I will use it for conditionals that spell out secondary duties. There is no present claim that such conditionals are imperatives. Alternative terminology can be found in the vast literature on the topic. For an overview, see Åqvist (2002) and Carmó and Jones (2002).

wine should be stored in the fridge. But, given that it isn't, is it also true that it should be stored in the cellar?

The answer to this question, interestingly, is 'it depends on the facts'. Sometimes, when the wine is not stored in the fridge, it is straightforwardly true that it should be stored in the cellar, and sometimes it isn't. The way that the world brings about that the wine is not stored in the fridge will affect our intuitions regarding whether or not it should be stored in the cellar. Dependence relations between the fact that the wine is not in the fridge and other facts will play a role in the interpretation of the deontic statement. To see this, consider the exchanges below:

- (2) a. You: The wine should be stored in the fridge.
b. Me: Yes, you are right. But the fridge door got stuck and I can't open it.
c. You: Well, it should be stored in the cellar (then).
- (3) a. You: The wine should be stored in the fridge.
b. Me: You are right, but I forgot to put it in the fridge.
c. You: #Well, it should be stored in the cellar (then).

We have no problem falling back on the secondary duty in (2c), but would not naturally do so in the case of (3c). This is puzzling. The circumstances in (3) appear to fit the condition for the secondary duty: the wine is not in the fridge. Why are our judgments regarding the secondary duty different in (2) and (3)? I will defend the hypothesis that to explain this difference, we need a semantics for *should* that pays maximal attention to facts and to dependencies between facts (we will consider alternatives to this view in section 3). I will propose an analysis for *should* that will predict that the secondary duty is true in (2c) but false in (3c). I have marked (3c) with # to indicate that there is infelicity but notice that the claim to be defended here is that in this scenario the sentence is false. Infelicity in this context follows from its obvious falsehood⁶ (I assume that there are pragmatic constraints like Gricean Maxims that disfavor making assertions that are obviously false. While it is true that

⁶ A reviewer has pointed out that not all false statements are infelicitous, but I will not discuss other kinds of cases here.

not all false assertions lead to infelicity, e.g. irony, in this case we do not recover any information that could justify the utterance.)

Let me briefly sketch out the idea, which will be presented in detail in later sections. Suppose that in the evaluation of *the wine should be in the cellar*, *should* pays attention to all the facts except those that are incompatible with the wine being in the cellar. In the scenario corresponding to (2c), *should* will pay attention to the fact that the fridge door was stuck and that the wine was not stored in the fridge (notice that given the way things are in the actual world, there is a dependence between those two facts – the wine enters the fridge through the door, so *if the fridge door is stuck, the wine is not in the fridge*). This means that *should* in (2c) will only compare alternatives in which the wine is not in the fridge. In some, the wine will be in the cellar, and in some, it won't. (2c) will be true iff in the best alternatives, the wine is in the cellar. Things are different in the scenario corresponding to (3c). In this scenario we assume that I am the person who puts the wine away. We take my reply to indicate that I forgot about the wine and did not put it away. And, if I forget to put the wine away, it will not be in the fridge nor in the cellar (in the actual world, there is a --possibly mediated-- dependence between those facts: *if I forget to put the wine in the fridge, it is not in the fridge and it is not in the cellar*).⁷ Since we are paying maximal attention to facts in the actual world, we will not entertain alternatives which deviate from actuality in that another person randomly puts the wine in the fridge/cellar. To evaluate the truth of (3c), we'll need to consider alternatives in which I didn't forget to put the wine away (if we include that fact, it will bring along the fact that the wine is not in the cellar). Once we bring in alternatives in which I did not forget to put the wine away, we open the door to various possibilities: there will be cases in which the wine is in the cellar and cases in which it is in the fridge (and possibly some other places!). (3c) will then be false: in the best alternatives, the wine is in the fridge (not in the cellar).

To understand the difference between (2c) and (3c), we need to pay attention to dependencies between facts. If the fact that the wine is not in the fridge counted as an independent fact when evaluating *should*, then it could be taken for granted both in the case of (2c) and (3c) (it is compatible with the wine being in the cellar), mistakenly predicting that

⁷ Notice that the example is naturally biased to signal that there is such a dependence: we usually address *should*-statements to the people we consider responsible for making things happen. Our intuitions in (3c) take such a dependence for granted (an alternative scenario will be considered in section 4).

(3c) should be true. It is only if we pay attention to the relation that links the fact that the wine is not in the fridge to the fact that I forgot to put the wine away, and also realize that this fact is linked to the fact that the wine is not in the cellar (and must therefore be set aside in evaluating (3c)) that we can understand why the secondary duty comes out false in (3) (and this is so even though the circumstances appear to fit the proposition in the antecedent of the CTD in (1b), i.e. the wine is not in the fridge⁸).

The relevance of dependencies between facts has been pointed out in the domain of counterfactual modality, where it has long been recognized that ‘facts stand and fall together’ (Kratzer 1981b, Veltman 2005). Standardly, it is claimed that in evaluating counterfactuals, we consider alternatives that are maximally similar to the actual world: a counterfactual of the form *if A, would B* is true if B is true in the most similar worlds in which A is true (in a *Lewis-Stalnaker* style analysis, Stalnaker 1968, Lewis 1973). The domain of quantification is identified on the basis of facts in the actual world. But if we worry about facts in isolation and do not take into account dependencies between them, we’ll make wrong predictions (a.o. Kratzer 1981b, 1989, 1990, 2002, Veltman 2005). The relevance of dependencies between facts in the case of counterfactuals can be appreciated by examining an example from Tichy (discussed by Veltman 2005):

- (4) ‘Consider a man, call him Jones, who is possessed of the following dispositions as regarding wearing his hat. Bad weather invariably induces him to wear a hat. Fine weather, on the other hand, affects him neither way: on fine days he puts his hat on or leaves it on the peg, completely at random. Suppose moreover that actually the weather is bad, so Jones is wearing his hat.’ (Tichy 1976)

The issue at stake is the truth value of the counterfactual *If the weather had been fine, Jones would have been wearing his hat*. In the scenario described above, we would probably judge this counterfactual false. But if we quantify over worlds like the actual world except for the fact that the weather was fine, we predict that the counterfactual will be true: if the difference lies in the weather, then Jones will be wearing a hat in the most similar worlds in which the weather is fine, just as he is actually doing. The problem, as discussed by Veltman (2005), is that in identifying the worlds quantified over, we have maintained similarity with respect to

⁸ For a discussion of the related topic of detachment in CTDs, the reader is referred to Arregui (2008, 2010).

the wearing of the hat. But, according to (4), Jones is wearing his hat *because* the weather is bad. If we give up on the fact that the weather is bad – ‘the very *reason* why Jones is wearing his hat’ (Veltman 2005: 164)- there is no need to assume that he was wearing his hat. This is an example of facts that ‘stand and fall together’. There is a dependence between the state of the weather and the wearing of the hat, and this dependence must be taken into account to obtain the right truth conditions for the counterfactual. If we ignore this dependence, we will make wrong predictions.

The claim at the core of the paper is that the semantics of deontic modals is similar to the semantics of counterfactuals in paying attention to dependencies between facts. In the case of deontics too, facts stand and fall together. The similarity arises because a similar notion of ‘revision’ is at stake in both types of modalities. The paper will cash out this proposal on the basis of a semantics for deontics that builds on the analysis of counterfactuals found in Kratzer (1981, 1991) and Veltman (2005). Both Kratzer and Veltman have been concerned with the role of relations between facts in the semantics of counterfactuals. We will make use of some of their proposals to probe similar issues in the semantics of deontics.

The paper is structured as follows: §2 examines the interaction between primary and secondary duties, investigating the difference between cases in which secondary duties are true and cases in which they are false; §3 presents some background for the semantics of deontic modals, giving a brief overview of some previous proposals and discussing various challenges; §4 presents a semantics for deontic modals that builds on a semantics of counterfactuals; §5 presents additional empirical support for the proposed semantics on the basis of ‘counterfactual style’ puzzles arising with deontic modals; §6 spells out some brief conclusions.

2. SECONDARY DUTIES

Our intuitions regarding secondary duties provide an illuminating perspective from which to investigate the interpretation of deontic modals. Prakken and Sergot (1996) observed that we sometimes have the intuition that primary and secondary duties ‘hold together’. They offer the following examples:

- (5) a. The children should not be cycling on the street.
b. If the children are cycling on the street, they should be cycling on the left side of the street.
- (6) a. There should be no fence.
b. If there is a fence, it should be a white fence.
(reported as a genuine example taken from a set of cottage regulations)

Consider (5). We come across the children cycling on the right side of the street. In some sense, it is both the case that they should not be cycling on the street and that they should be cycling on the left side of the street. The same holds for (6). If there is a blue fence, it may be the case that it should not be there and that it should be white.

Examples like (5) and (6) indicate that primary and secondary duties can be compatible, and we will need a semantics that makes sense of that (see also Arregui 2010). Do compatible primary and secondary duties always hold together? As we have seen in (2) and (3), the answer appears to be ‘no’. Another example making this point is presented below. Imagine that you are given the instructions in (7) regarding how to leave your house in a -20C morning:

- (7) a. You should be wearing your winter coat.
b. If you are not wearing your winter coat, you should be wearing your raincoat (it blocks the wind).

Imagine now that a friend sees you at the bus stop, wearing only a shirt and a thin cardigan, shivering:

- (8) a. Your friend: You should be wearing your winter coat.
b. You: I know, but I left it behind at my sister’s house.
c. Your friend: Well, you should be wearing your raincoat (then).
- (9) a. Your friend: You should be wearing your winter coat.

- b. You: I know, but I was late this morning and just ran out of the house without my winter coat.
- c. Your friend: #Well, you should be wearing your raincoat (then).

As in (2) vs. (3), there is a contrast in our willingness to fall back on the secondary duty. This seems straightforward in the scenario in (8), but not in the scenario in (9). Why not? After all, the circumstances do appear to satisfy the conditions for the secondary duty as described in (7b). And arguably, the best worlds in which I ran out of the house without my winter coat in a very cold morning are worlds in which I take my raincoat with me (which will at least block the wind). Why don't we fall back on the secondary duty then? It will not do to claim that our judgments in (9c) arise because it is funny/odd/uncooperative to mention a coat after being told that you ran out of the house. In the circumstances described in (9), your friend could insist on the winter coat, even though the raincoat is out:

- (10) a. Your friend: You should be wearing your winter coat.
- b. You: I know, but I was late this morning and just ran out of the house without my winter coat.
- c. Your friend: That's too bad. You really should be wearing your winter coat.

In contrast with (9c), the statement in (10c) is perfectly fine. This indicates that what goes wrong in (9c) cannot be characterized as a purely pragmatic effect arising because of lack of cooperation/ understanding/ good will on part of a speaker who insists on talking about a coat. (As a preview, note that according to the fact-based proposal to be defended in this paper, (10c) will be predicted to be true in this scenario.)

As with the examples in (2) and (3), differences in dependencies between facts underlie the differences in intuitions in (8) and (9). Suppose, again, that in evaluating *should* we pay attention to all the facts compatible with the embedded proposition (which will be made more precise in section 4). When we evaluate (8c), *should* will pay attention to the fact that you forgot your winter coat at your sister's house (compatible with the fact that you wear a raincoat). All the alternatives compared by *should* will match the actual world with respect to this fact. And they will also match the actual world with respect to the fact that

you are not wearing your winter coat (notice that there is a dependence between the facts: *if you forgot your coat in your sister's house, you don't put it on when you leave your home*). In the case of (8c), *should* will only compare alternatives in which you don't wear your winter coat. The sentence will be true if in the best ones, you wear your raincoat. Since this is the case, (8c) is true. Things are different in the case of (9c). In the scenario in (9), you were late that morning and ran out of the house without your winter coat. There is a relation between being late and running out and not wearing your winter coat. But we understand that this relation actually extends to all coats: *if you are late and run out of the house, you don't put on any coat*. If this generalization holds, then when we evaluate *should* in (9c), we will set aside the fact that you were late and ran out (paying attention to that fact would bring along with it the fact that you are not wearing your raincoat, incompatible with wearing a raincoat). The alternatives evaluated by *should* in (9c) will therefore include worlds in which you did not run out of the house in a hurry. In some of these you were wearing your raincoat and in others you were wearing your winter coat. The result is that the secondary duty will not hold in this scenario: in the best worlds, you were wearing your winter coat. The sentence in (9c) is false.

To understand the difference between (8) and (9) we need to pay attention to dependencies between facts: running out of the house brings along with it not putting on any coat (not your winter coat nor your raincoat) (i.e. these facts 'stand and fall together'). In evaluating the secondary duty, we have to set aside the fact that you ran out of the house (in order to set aside the fact that you were not wearing your raincoat). Doing so also sets aside that you were not wearing your winter coat, making worlds where you are wearing your winter coat available for comparison. The result is that the secondary duty is false: worlds in which you wear your winter coat are better. This link between facts is responsible for our intuitions in (9c).

Here is one final example to argue for the relevance of dependencies between facts. Suppose you work in the customer service department at the post office and there are regulations regarding your work clothes when you deal with clients:

- (11) a. You should be wearing the postal service shirt.
b. If you are not wearing the postal service shirt, you should be wearing a dark blue shirt.

You come in to work with a red shirt and the manager is angry. Consider now the following scenarios:

- (12) a. The manager: You should be wearing the postal service shirt.
b. You: I know, but it got ripped in the washing machine.
c. The manager: You should be wearing a dark blue shirt (then).
- (13) a. The manager: You should be wearing the postal service shirt.
b. You: I know, but the baby was still asleep this morning and I got dressed in the dark.
c. The manager: #You should be wearing a dark blue shirt (then).

In (13c) we see another example in which we don't fall back on the secondary duty. Even though in the best worlds in which you don't wear the postal service shirt and get dressed in the dark it is true that you wear a dark blue shirt, the secondary duty does not hold. It would be possible (though insensitive) for your boss to insist on the postal service shirt, but not on the dark blue shirt (indicating that we are not dealing with a purely pragmatic problem arising because of stubborn insistence with the regulations):

- (14) a. The manager: You should be wearing the postal service shirt.
b. You: I know, but the baby was still asleep this morning and I got dressed in the dark.
c. The manager: (I don't care!) You should be wearing the postal service shirt/ #You should be wearing a dark blue shirt (then).

It would be difficult to explain our intuitions under an approach that considered that facts are independent of each other. If facts were independent, we could pay attention to the fact that you are not wearing the postal service shirt and the secondary duties in (13c) and (14c) would come out true (notice that in evaluating the secondary duty, it is definitely relevant whether you were wearing your postal service shirt). Again, I will argue that the key to the difference between (12c) and (13c) lies in the dependencies between facts. There is a dependence between dressing in the dark (i.e. without knowing what you are putting on) and

not wearing regulation clothing (*if you dress in the dark, you don't end up wearing your work clothes*). In evaluating (13c), we must set aside the fact that you dressed in the dark, making available alternatives in which you wore the postal service shirt and alternatives in which you wore a dark blue shirt. The postal service shirt alternatives are better, and so (13c) is false. In the case of (12c), we only compare alternatives in which you did not wear the postal service shirt, and in the best ones, you wore a dark blue shirt.

The various examples discussed above illustrate that the truth of the proposition in the antecedent of a CTD conditional does not automatically make the secondary duty true. *Should* pays attention to the ways in which the world brings about the truth of that proposition. Our intuitions regarding secondary duties are sensitive to the dependence relations holding between facts. In section 4 I will present an analysis of *should* that cashes out these remarks. But first I would like to strengthen the claim that in the relevant scenarios (3c, 9c, and 13c), the secondary duty is false (not merely independently infelicitous). The view that the statements of secondary duties are false in the relevant scenarios is supported by the exchanges presented below. Consider the example in (1) (repeated here):

- (1) a. The wine should be stored in the fridge.
b. If it is not stored in the fridge, it should be stored in the cellar.

Imagine now that you walk into the kitchen and find the wine on top of the (fully functional) fridge:

- (15) a. You: The wine should be stored in the cellar.
b. Me: No, that's not true. It should be in the fridge. I will put it there right away!

In the scenario above, in which the facts do not rule out the primary duty, the secondary duty does not hold and this is explicitly stated. On the other hand, if the facts (compatible with the secondary duty) rule out the primary duty, the secondary duty will be true:

- (16) a. You: The wine should be stored in the cellar.

- b. Me: You are right. The fridge door is stuck and we cannot put it in the fridge. I will put it in the cellar.

We see similar intuitions regarding truth with respect to example (7), repeated below:

- (7) a. You should be wearing your winter coat.
- b. If you are not wearing your winter coat, you should be wearing your raincoat (it blocks the wind).

Imagine that you come across your friend, violently shivering at the bus stop, wearing only a shirt and a thin cardigan:

- (17) a. You: You should be wearing your raincoat.
- b. Your friend: No, that's not true. I should be wearing my winter coat. But I was late this morning and just ran out of the house.

Or imagine that a colleague sees you in a red shirt at the costumer counter, where the guidelines for work clothes are as in (11), repeated below:

- (11) a. You should be wearing the postal service shirt.
- b. If you are not wearing the postal service shirt, you should be wearing a dark blue shirt.

Worried about potential problems with the manager, your colleague might warn you:

- (18) a. Your colleague: You should be wearing a dark blue shirt.
- b. You: That's not true. I should be wearing the postal service shirt. But I got dressed in the dark this morning!

In the exchanges above we see that it is possible to explicitly claim that the secondary duty is false in scenarios that appear to fit the CTD conditions, supporting the view that it is truth that is at stake (this will be taken up again in section 3.1).

To sum up. In this section we have examined examples that show that sometimes, primary and secondary duties hold together. We will need a semantics for *should* that allows for that. But we have also seen cases in which they don't. Even though the circumstances appear to fit the CTD antecedent, the secondary duty is false. In presenting these examples, we have noted the relevance of dependencies between facts. Such dependencies (facts 'standing and falling together') affect our intuitions regarding secondary duties and this tells us that they have to be taken into account in the semantics of *should*.

3. TOWARDS A COUNTERFACTUAL ANALYSIS

In section 2 we have informally appealed to an analysis of *should* according to which *should* exhibits 'maximal' dependency on facts (i.e. it pays attention to all compatible facts and to dependencies between facts). This view of *should* is very much reminiscent of standard analyses of counterfactual conditionals, and in section 4 I will present a 'counterfactual style' analysis of *should* building on work by Kratzer and Veltman. The goal in this section is to build motivation for that proposal and to briefly consider alternatives.

3.1 KRATZER (1981A, 1991)

I will begin with a brief discussion of the proposal for deontic modals found in Kratzer (1981a, 1991) (some familiarity with Kratzer's influential proposal will be presupposed). Kratzer (1981a, 1991) presents a context-dependent account of the semantics of modals according to which the interpretation of a modal is established in relation to two context-dependent functions: a modal base (f) and an ordering source (g). They operate together to fix the truth value of modal statements. The modal base identifies a set of possible worlds and the ordering source imposes an ordering amongst the set. Depending on the quantificational strength of the modal, the truth of the modal statement depends on what happens in all of the best possible worlds, or in some, etc. The different flavors of modality result from different combinations of modal bases and ordering sources.

In my informal description above, I have taken for granted the ‘limit assumption’ (Lewis 1973), according to which it is always possible to identify a set of best possible worlds. Kratzer does not make this assumption (for reasons that will not be discussed here), and her definition of modal necessity needs to be presented in a more complex manner. For reasons of simplicity, I will make the limit assumption here, and, simplifying Kratzer’s original proposal, propose the definition of modal necessity in (19):⁹

(19) A proposition p is a *necessity* in a world w with respect to a modal base f and an ordering source g iff the following condition is satisfied:

$\forall w' \in \bigcap f(w)$: if w' is a $g(w)$ -best world, then $w' \in p$. [adapted from Kratzer 1991]

The modal base f applied to w will deliver a set of propositions and quantification will take place over worlds within the set characterized by the set of propositions. The ordering source g applied to w will also deliver a set of propositions. These ones will be used to order the worlds within the modal base. The definition of what it means to be a $g(w)$ -best world is provided in (20):

(20) Given an ordering source g and possible world w :

w' is a $g(w)$ -best world iff there isn’t a world w'' such that $w'' <_{g(w)} w'$. (i.e. w' is a $g(w)$ -best iff there isn’t a world that is $g(w)$ -better).

With these definitions in hand, we can now turn our attention to deontic modal necessity. Kratzer actually considers two options for interpreting deontic modals. *Option 1*: an empty modal base and a normative ordering source. With this interpretation, *should* ϕ will claim that in the best of all possible circumstances, given the norms, ϕ is true. The truth of *should* ϕ will then be independent of the facts in the evaluation world. *Option 2*: a circumstantial modal base and a normative ordering source. This modal base will assign to the evaluation world (some) true propositions. With a circumstantial modal base, the worlds quantified over will

⁹ Kratzer’s original definition of modal necessity is given in (i):

(i) A proposition p is a *necessity* in a world w with respect to a modal base f and an ordering source g iff the following condition is satisfied:
For all $u \in \bigcap f(w)$ there is a $v \in \bigcap f(w)$ such that $v \leq_{g(w)} u$ and for all $z \in \bigcap f(w)$: if $z \leq_{g(w)} v$, then $z \in p$. (Kratzer 1991)

match the evaluation world with respect to certain facts (i.e. those described by the context dependent modal base). The statement *should* ϕ will be true iff ϕ is true in the best worlds (given the norms/contextually established measure of goodness) that match the actual world with respect to the modal base propositions.

We can illustrate how the two options work with respect to example (21) (recall that in the relevant scenario the wine was found sitting on top of the fridge):

(21) The wine should be stored in the cellar.

It is clear that a purely normative interpretation will not work here (*Option 1*). In the best of all possible worlds, the wine is stored in the fridge, and (21) is false. *Option 1* cannot account for those cases in which we judge (21) true (e.g. 2c).

The outcome with a circumstantial modal base (*Option 2*) is more complex. A circumstantial modal base assigns to the evaluation world a set of propositions true in the evaluation world. Which propositions get into the set depends on the context of utterance. If the set of propositions corresponding to the circumstantial modal base includes the proposition that the wine is not stored in the fridge, and does not include any proposition that is inconsistent with the proposition that the wine is stored in the cellar, then in the best worlds the wine will be stored in the cellar, and (21) will be true.

The process that decides which propositions get into the circumstantial modal base in Kratzer's account is a pragmatic process. It makes sense from this point of view that the modal base should not include the proposition that the wine is not stored in the cellar (or propositions that entail it). That would trivialize the truth conditions for (21) and it is reasonable to assume that there is a pragmatic constraint against triviality. But what about the proposition that the wine is not stored in the fridge? If that proposition is missing from the modal base, then the set of worlds quantified over will include worlds in which the wine is stored in the fridge and worlds in which it is stored elsewhere. In the best worlds, the wine is stored in the fridge, predicting that (21) is false.

In a standard Kratzer-style account, whether the proposition that the wine is not stored in the fridge gets into the circumstantial modal base in evaluating (21) depends on context. The fact that we easily fall back on the secondary duty in the scenario in (2) tells us that in that case the proposition did get into the circumstantial modal base. The fact that we

did not fall back on the secondary duty in the scenario in (3) tells us that in that case, it did not get into the circumstantial modal base. The judgments are quite robust. However, it is difficult to see what notion of contextual relevance could predict this difference. Given the CTD in (1b), it seems clear that it should always be relevant whether the wine was stored in the fridge or not (this is the condition for triggering the secondary duty). What would be the reason to exclude that proposition from the contextually determined circumstantial modal base in the context corresponding to (3c) (and not in the context corresponding to (2c))? This is a challenge faced by a purely context-dependent account.

Support for the view that the evaluation of *should* statements depends on what the facts are and their relations, as opposed to what might be contextually salient/relevant, comes from the observation that in evaluating *should* statements, we are very much interested in the facts. Recall the preferences regarding the storage of the wine spelled out in (1), and consider now the reasoning in (22):

- (1) a. The wine should be stored in the fridge.
 b. If it is not stored in the fridge, it should be stored in the cellar.
- (22) The wine is not stored in the fridge and it is not stored in the cellar. It should be stored in the cellar.

We do not judge the *should*-statement in (22) automatically true. Before agreeing that the wine should be stored in the cellar, we would probably ask: *Why? What's wrong with the fridge?* However, the first sentence has arguably made salient the proposition that the wine is not stored in the fridge. Given (1b), it is clear that that proposition is relevant when evaluating whether the wine should be stored in the cellar. But we don't automatically put it into the circumstantial modal base of the secondary duty. We ask about the facts instead. How did the world bring about that the wine is not in the fridge? This is what matters to the evaluation of the secondary duty.¹⁰

¹⁰ The first sentence has arguably also made salient the proposition that the wine is not stored in the cellar. But we set aside that proposition in evaluating the secondary duty to avoid triviality. Notice that the salience of a proposition is not enough to rule out a *should*-statement:

(i) The wine is not stored in the fridge. It should be stored in the fridge.

It is clearly very difficult to conclusively argue against a purely contextual determination of the circumstantial modal base given the absence of a theory of the relevance/salience of propositions that makes predictions with relevant precision and detail. When judging such matters, we usually (and reasonably) appeal to common sense (from a common-sense point of view, I have argued that it is hard to see the difference between (2) and (3) regarding the relevance/salience of the proposition that the wine is not stored in the fridge). I hope, however, to have shown that a purely contextual account of the interpretation of *should* faces challenges (the arguments with (2) vs. (3) can be reproduced with the other examples). The proposal made in this paper simply places the bets in a different way. The claim is that whether *should* pays attention to a fact or not does not depend on context-based interests or relevance, rather, it depends on the web of dependencies between facts in the world. The truth of a *should* statement will be determined by what is going on in the world, not by what we conversationally choose to pay attention to or what is conversationally relevant.¹¹ (Though whether we are willing to assert a *should* statement will undoubtedly be affected by what we are paying attention to and what we know.) Of course, this proposal faces challenges too, which will be taken up in later sections.

3.2 REVISING THE PREMISE SET

In our informal discussion of the interpretation of *should* in section 2, we appealed to the idea that *should* pays attention to all the facts (minus those incompatible with the embedded proposition). In this section we will develop a simple proposal to capture this intuition, and use it to motivate a more complex account in section 4.

Given the idea that *should* pays attention to all the facts, one possible starting point would be to claim that its modal base is built on the basis of a totally realistic modal base (for Kratzer, a modal base is *totally realistic* if it uniquely characterizes a world). We cannot use a totally realistic modal base itself as the modal base of *should* (it would only deliver the evaluation world, and trivialize the semantics), but we can start off from a totally realistic

¹¹ This statement must be immediately qualified since context-dependent interests will have a huge influence on the interpretation of *should*-statements by affecting the ordering source.

modal base f and make minimal modifications so as to make it compatible with the embedded proposition.¹² In the case of *should* ϕ , the revised modal base will be f_ϕ (a totally realistic modal base minimally revised so as to obtain consistency with the proposition embedded under the modal, $[[\phi]]$ ¹³). The set of worlds characterized by the revised modal base f_ϕ can then be ordered by the contextually salient ordering source, and a *should* ϕ statement will be true iff ϕ is true in the best worlds in the resulting set. The outcome will be a maximally fact-dependent interpretation for *should*. The proposal is presented below (where f is a totally realistic modal base and g is a contextually given ordering source):

(23) Semantics for *should* (preliminary)

For all worlds w , $[[\textit{should } \phi]]^{f,g}(w) = 1$ iff

$\forall w' \in \cap f_\phi(w)$: if w' is a $g(w)$ -best world, $w' \in [[\phi]]$

Given (23), the statement in (24a) will receive the truth conditions in (24b):

(24) a. **[should [the wine be stored in the fridge]]**

b. For all worlds w , $[[24a]]^{f,g}(w) = 1$ iff

$\forall w' \in \cap f_\phi(w)$: if w' is a $g(w)$ -best world, then the wine is stored in the fridge in w'

In revising the totally realistic modal base, we will remove from the set of propositions assigned to the evaluation world the proposition that the wine isn't stored in the fridge (and other propositions inconsistent with the proposition that the wine is stored in the fridge). The worlds corresponding to the remaining propositions will be ordered by the ordering source, and the modal statement will be true iff in the best worlds the wine is stored in the fridge.

Given (23), the statement in (25a) will receive the truth conditions in (25b):

¹² This idea is obviously inspired by Kratzer's treatment of counterfactuals (Kratzer 1981a, 1991), but this is a simplified version.

¹³ For the sake of simplicity, I will assume that given a world w and a proposition ϕ false in w , there is just one totally realistic modal base and there is just one way to revise it to make it compatible with ϕ . This is not the case (as has been pointed out by Kratzer 1981, etc.). Technically speaking, it would be necessary to consider all the ways of revising a totally realistic modal base, but for the sake of simplicity, I set this aside here.

- (25) a. [should [the wine should be stored in the cellar]]
b. For all worlds w , $[[25a]]^{f,g}(w) = 1$ iff
 $\forall w' \in \cap f_\phi(w)$: if w' is a $g(w)$ -best world, then the wine is stored in the cellar
in w' .

In revising the totally realistic modal base this time, we remove from the set of propositions assigned to the evaluation world the proposition that the wine is not stored in the cellar and other propositions inconsistent with the proposition that the wine is stored in the cellar. We will keep the proposition that the wine is not stored in the fridge. We will order the worlds corresponding to the remaining propositions and the modal statement will be true iff in the best worlds the wine is stored in the cellar.

According to the proposal in (23), primary and secondary duties can be compatible with each other. Given (23), we quantify over different worlds when evaluating (24a) and (25a). The revision of the totally realistic modal base takes place with respect to different propositions and both (24a) and (25a) can be true. The proposal in (23) thus allows us to make sense of Prakken and Sergot (1996)'s observation that primary and secondary duties can hold together.¹⁴

3.3 FACTS STAND AND FALL TOGETHER

The proposal in (23) gets us partway where we need to go (primary and secondary duties can be compatible) but not all the way there. It ignores dependencies between facts. The literature on counterfactuals has observed that failure to take into account dependencies between facts leads to mistakes in truth conditions (e.g. Kratzer 1981b, Veltman 2005). We'll briefly review that point before discussing (23).

Let us reconsider our earlier counterfactual example. When evaluating the truth of (26) in the scenario in (4), it is necessary to keep track of the dependence between the state of the weather and the wearing of the hat (recall that we judge (26) false in that scenario):

¹⁴ For a brief discussion of the outcome of conjoining primary and secondary duties, the reader is referred to Arregui (2010).

(26) If the weather had been fine, Jones would have been wearing his hat.

Suppose we adopted a naïve approach, and ignored dependencies between facts. To evaluate (26), we would look at worlds in which the proposition that the weather is fine is true, as well as all the propositions true in the actual world consistent with that proposition.¹⁵ The true proposition that Jones is wearing his hat would be in that set. So, the naïve approach would mistakenly predict that in the scenario in (4), (26) is true: we would quantify over worlds in which Jones is wearing his hat.

In discussing Tichy's example, Veltman noted that there is a dependence between the state of the weather and the wearing of the hat. The link between those two facts is not accidental, but responds to a law/generalization operative in the actual world: *if the weather is bad, Jones wears his hat*. When we evaluate counterfactuals, we should not consider these facts independently of each other. When we give up the proposition that the weather is bad, we should also give up the proposition that Jones is wearing his hat. Then we will make correct predictions for (26).

In section 2 we have seen that to understand what happens in the evaluation of secondary duties (2c vs. 3c), (8c vs. 9c) and (12c vs. 13c), we need to keep track of dependencies between facts, very much like we do in the evaluation of counterfactuals. This indicates that the semantics of *should* needs to be sensitive to such dependencies. The proposal in (23) is not. Let us go back again to the wine storage regulations, taking into account (1) in the scenarios described in (2) and (3):

- (2) a. You: The wine should be stored in the fridge.
b. Me: Yes, you are right. But the fridge door got stuck and I can't open it.
c. You: Well, it should be stored in the cellar (then).
- (3) a. You: The wine should be stored in the fridge.
b. Me: You are right, but I forgot to put it in the fridge.
c. You: #Well, it should be stored in the cellar (then).

¹⁵ There may be more than one way of achieving consistency, but I set this aside here.

The secondary duty *The wine should be stored in the cellar* has the truth conditions noted in (25) (repeated below):

- (25) a. [should [the wine should be stored in the cellar]]
b. For all worlds w , $[[25a]]^{f,g}(w) = 1$ iff
 $\forall w' \in \cap f_\phi(w)$: if w' is a $g(w)$ -best world, then the wine is stored in the cellar
in w' .

The truth conditions in (25b) don't predict the contrast between (2c) and (3c). In both cases, we quantify over worlds in which all the propositions true in the actual world compatible with the proposition that the wine is stored in the cellar are true. In both cases, this will include the proposition that the wine is not in the fridge. The prediction is that the secondary duty will be true in both scenarios, contrary to fact.

To understand what happens in (3), we need a semantics for *should* that can track the relation between me forgetting to put the wine away and the wine not being in the fridge and neither being in the cellar. When we evaluate the secondary duty (3c), we need to consider worlds in which I don't forget to put the wine away (otherwise, it would not be in the cellar!). But if we give up on me forgetting to put the wine away, we should also give up on the wine not being in the fridge (there is a dependence between those two facts in the actual world, much like the dependence between the bad weather and the fact that Jones was wearing his hat). This will correctly predict that we don't fall back on the secondary duty: in the best alternatives, I put the wine in the fridge and the primary duty holds. As in the case of counterfactuals, we will get the right result with a semantics that tracks dependencies between facts.

4. A COUNTERFACTUAL-STYLE ANALYSIS OF DEONTIC *SHOULD*

In my discussion of the parallelism between counterfactuals and deontics so far, I have talked informally about dependencies between facts. The intuition, which I think we can recognize, is that given that something has happened (or that something is a certain way), some other thing has happened too (or is a certain way). Dependencies between facts are,

themselves, (accidental) facts about the evaluation world. Consider the wine example in (1) again. The intuitions reported for (3c) respond to the fact that we are taking for granted that I am the person who puts the wine away. This means that if I forget to put the wine away in the fridge, it will not be in the fridge *and* it will not be in the cellar. We take this dependence to be a fact about the world and evaluate (3c) accordingly. But things could have been different. We could have been in a world in which I was not the person who puts the wine away. Maybe I was only the person in charge of putting the wine in the fridge, and somebody else was in charge of putting it in the cellar if I forgot to store it in the fridge. If we knew this to be the case, it would make perfect sense for you to claim that the wine should be in the cellar after I tell you that I forgot to put it in the fridge (I would probably agree with you!). In that world, there is no dependence between me forgetting to put away the wine and the wine not being in the cellar. As this alternative scenario makes clear, the truth of examples like (3c) really hinges on the dependencies between facts in the world of evaluation.

There are many different types of generalizations that can be considered examples of dependence relations between facts, e.g. natural laws (if I drop the apple, it falls), habits (if the weather is bad, Jones wears his hat), part-whole relations (if I come along, my arms and legs come too – after an example by Kratzer). The literature on counterfactuals has used various kinds of technology to address dependencies. Lewis (1979), for example, talked about laws and the relevance of laws in measuring similarity. In early work, Kratzer appealed to differences in the totally realistic modal base characterizing the evaluation world (Kratzer 1981b). In later work, Kratzer appealed to lumping (Kratzer 1989) and to natural propositions (Kratzer 1990, 2002). Veltman (2005) appeals, informally, to laws, which include generalizations arising from conventions, habits, etc., as well as natural laws. Schulz (2007) makes a proposal that addresses dependencies between facts in terms of the framework of Pearl (2000). In this paper I will follow Veltman and use the word ‘law’ informally to cover the kinds of generalizations that will matter. I will not however, have anything substantial to say about the nature of such laws.¹⁶

The main idea that I am arguing for is that there is a parallelism between counterfactuals and deontics in terms of factual dependence. I will cash out this idea with a semantics for *should* that builds on Veltman’s analysis of counterfactuals (Veltman 2005). I

¹⁶ For a critical appraisal of Veltman’s proposal, the reader is referred to Schulz (2007).

have chosen this framework because it will provide us with an intuitive and graphic account that will be easy to transpose to the case of deontics (where we will also use Kratzer-style ordering sources). In principle, it would be possible to formulate an analysis for deontic *should* taking as a starting point a different analysis of counterfactuals. Kratzer (1990, 2002), for example, deals with dependencies between facts in the semantics of counterfactuals by appealing to ‘natural propositions’. These are propositions constructed in a way that pays attention to facts in the evaluation world and to dependencies between facts. Since the domain of quantification of counterfactuals is identified in terms of natural propositions, the worlds quantified over will be worlds that match the evaluation world with respect to facts and dependencies. By setting up the semantics of counterfactuals with natural propositions, Kratzer packs into the propositions themselves the relations that Veltman (2005) captures in terms of laws. There are many interesting differences between Kratzer’s and Veltman’s proposals, meriting further investigation. However, for the purposes at hand, those differences do not appear relevant. I will build a semantics for *should* making use of Veltman’s proposal. It would in principle also be possible to develop an appropriately fact-sensitive semantics for *should* in terms of Kratzer-style natural propositions. However, that project remains for future work.

4.1 VELTMAN (2005)

I’ll begin the process of setting up the semantics of *should* by presenting Veltman’s analysis of counterfactuals. Veltman works within a premise-set style semantics, defined in a way that pays attention to laws (dependencies between facts). “In making counterfactual assumptions, we are not prepared to give up propositions that we consider to be general laws.” (Veltman 2005: 166). Veltman’s objective in the 2005 paper is to make a dynamic proposal for the interpretation of counterfactuals that deals both with the need to revise the information in the premise set on the basis of counterfactual assumptions and the need to track information growth throughout a conversational exchange. I will not make use of the full machinery of Veltman’s proposal, as the issues dealt with in this paper can be handled without the full dynamic framework. I will simplify Veltman’s proposal and focus only on the aspects relevant to the revision of premise sets.

Veltman's goal is to offer a sophisticated account of the revision of premise sets that is able to handle relationships between facts. One of Veltman's key claims is that to characterize revision in a way that does not lead to mistakes it is necessary to carry out revisions with respect to facts that are independent of each other, and allow the 'laws' to fill in the gaps by bringing along dependent facts. This will make the revision of premise sets sensitive to the law-like relationships between facts. In Veltman's analysis, the premise set for a world is identified as the set of *independent* facts true in the world. Revision takes place with respect to the independent facts in a manner that satisfies logical consistency and the laws. I will present the relevant aspects of Veltman's proposal briefly in this section. Readers are referred to Veltman's work for full details.

In his presentation, Veltman characterizes worlds as a valuation function on a finite set of atomic sentences: given a set of atomic sentences A , a world is a function with domain A and range $\{1, 0\}$. A world is thus identified with the valuation function that assigns 1 to the atomic sentences true in the world and 0 to the atomic sentences false in the world. Set notation is introduced to facilitate the exposition. Given sentences p and q and a world w , Veltman makes use of the notation $\langle p, 1 \rangle \in w$ to indicate that $w(p) = 1$ and $\langle q, 0 \rangle$ to indicate that $w(q) = 0$. Situations are characterized as partial functions from A to $\{1, 0\}$. Given sentences p and q and a situation s , the notation $\langle p, 1 \rangle \in s$ is used to indicate that $s(p) = 1$ and $\langle q, 0 \rangle \in s$ is used to indicate that $s(q) = 0$. Veltman calls pairs like $\langle p, 1 \rangle$ and $\langle q, 0 \rangle$ *facts*. Given this terminology, a situation s is a subset of a world w ($s \subseteq w$) iff all the facts members of s are also members of w . Sentences are taken to express propositions, which are sets of possible worlds. Given a sentence p , the proposition expressed by p will be $[[p]]$ (where $[[p]] = \{w \in W : w(p) = 1\}$ and $[[\neg p]] = W \sim [[p]]$).

In presenting his analysis, Veltman appeals to the notion of the 'universe' of a cognitive state S (U_S): the set of possible worlds in which all the propositions that the agent of the state S considers to be general laws hold. In setting up a non-dynamic proposal, I will appeal to a different, though related, notion: the 'law-horizon' of a world (U_w). For any world w , U_w is the set of possible worlds that obeys the laws (regularities) operative in w . The law-horizon of the actual world ($w_{@}$) will be $U_{@}$.

In order to explain how to carry out revisions, we will need some of the auxiliary notions defined by Veltman. They have been adapted here to fit our present goals (see Veltman 2005: 167-168 for the original proposals):

- (27) a. A situation s *determines* a world w in $U_{@}$ iff for all w' in $U_{@}$ such that $s \subseteq w'$, $w' = w$.
- b. A situation s is a *basis* for a world w iff s is a minimal situation that determines w in $U_{@}$.
- c. A situation s *forces* a proposition P within $U_{@}$ iff for every world w in $U_{@}$ such that $s \subseteq w$, $w \in P$.

To identify the worlds quantified over when making a contrary-to-fact hypothesis (e.g. *the weather is fine* in (26)), we first identify a basis of $w_{@}$ (there may be more than one) and then we identify, within each situation that constitutes a basis for $w_{@}$ the maximal subsituation that does not ‘force’ the inconsistent proposition (e.g. the proposition that the weather is not fine/ the weather is bad).

To illustrate how the definitions work, let us go back to Tichy’s example:

- (28) $p =$ *the weather is bad*
 $q =$ *Jones is wearing his hat*
 $r =$ *Ottawa is the capital of Canada* (a random sentence for illustration)

Suppose that the actual world is subject to a law that states that *if p is true, q is true* (*if the weather is bad, Jones is wearing his hat*). Illustrative examples of possible worlds are provided below (not all possibilities, for reasons of space), where $w_{@}$ is the actual world and worlds that do not belong to $U_{@}$ have been crossed out (double strike-through):

(29)	p	q	r	
	$w_{@}$	1	1	1
	w_1	0	1	1
	w_2	1	1	0
	w_3	0	0	1
	w_4	1	0	1
	w_5	1	0	0

The actual world has a single basis: $s = \{ \langle p, 1 \rangle, \langle r, 1 \rangle \}$ (for all $w' \in U_{@}$, if $s \subseteq w'$, $w' = w_{@}$).¹⁷

Veltman's two-step procedure to make a counterfactual hypothesis works as follows: first, we identify a set of worlds on the basis of the facts of the actual world that are compatible with the counterfactual assumption taking into account the relevant laws. I will call this set of worlds the 'revision set' built on the basis of the counterfactual hypothesis. The second step is to update the revision set with the proposition corresponding to the counterfactual hypothesis.

Since here we are interested mainly in how to carry out revisions to premise sets, we will focus on the first step (the construction of a revision set). Informally, the revision set will be the set of law-like worlds (members of $U_{@}$) that maximally match $w_{@}$ with respect to independent facts compatible with the embedded proposition. In the example above, the revision set for the proposition that the weather was fine will consist of the set $\{w_{@}, w_1, w_3\}$. These are the worlds in $U_{@}$ in which Ottawa is the capital of Canada ($\langle r, 1 \rangle$ is the only actual world independent fact compatible with $[[\neg p]]$, the proposition that the weather is fine).

In our analysis of deontic *should*, the revision set for the embedded proposition will serve as the modal base for *should*.¹⁸ To identify the revision set for a proposition, we begin by defining an auxiliary set:

- (30) Where P is a proposition and w is a possible world,
 $w \downarrow P$ (the *retraction* of P from w) is the set of situations s such that $s \subseteq w$ and there is a basis s' for w such that s is a maximal subset of s' not forcing P within U_w .¹⁹

What this set does is identify the maximal sets of independent facts in w that do not force ('bring along') P . With this auxiliary set in hand, we can now define the *revision set* for a world w and a proposition P (abbreviated as $Rev_{w/P}$):

¹⁷ The situation s ($\{ \langle p, 1 \rangle, \langle r, 1 \rangle \}$) determines $w_{@}$ in $U_{@}$ (for all $w' \in U_{@}$, if $s \subseteq w'$, $w' = w_{@}$) and it is the *minimal* situation to do so: there is no situation $s' \subset s$ such that s' determines $w_{@}$ in $U_{@}$. This means that s is a basis for $w_{@}$. There is no other situation in $w_{@}$ satisfying these requirements, so s is the only basis for $w_{@}$.

¹⁸ I use the term 'modal base' informally here to talk about the set of worlds that will be quantified over, not in Kratzer's technical sense (a function from worlds to sets of propositions).

¹⁹ Veltman's definition is different because of differences in frameworks, see Veltman (2005: 168).

- (31) Where P is a proposition and w, w' are possible worlds,
 $w' \in Rev_{w/P}$ iff $w' \in U_w$ and there is some $s \in w \downarrow P$ such that $s \subseteq w'$.

The revision set (given a world w and a proposition P) will be those members of the law-horizon of w that ‘extend’ maximal sets of independent facts of w that do not force the proposition P .

To illustrate how the definitions work, let us go back to Tichy’s example. As we noted, there is a single basis for the actual world: $s = \{ \langle p, 1 \rangle, \langle r, 1 \rangle \}$. Let P in (31) be the relevant proposition ($[[p]]$, the proposition that the weather is bad). The set corresponding to $w_{@} \downarrow [[p]] = \{ \langle r, 1 \rangle \}$. This is the maximal set of independent facts in $w_{@}$ that does not force $[[p]]$. The revision set ($Rev_{w_{@}/[[p]]}$) will be the set of worlds $w \in U_{@}$ such that $\{ \langle r, 1 \rangle \} \subseteq w$. Below is an illustration of the worlds that will be members of $U_{@}$ and $Rev_{w_{@}/[[p]]}$. The worlds that are not members of $U_{@}$ have been crossed out with double lines (double strike-through), and the worlds that are not members of $Rev_{w_{@}/[[p]]}$ have been crossed out with a single line. The worlds that are members of $Rev_{w_{@}/[[p]]}$ have not been crossed out at all.

(32)

	p	q	r
$w_{@}$	1	1	1
w_1	0	1	1
w_2	1	1	0
w_3	0	0	1
w_4	1	0	1
w_5	1	0	0

$Rev_{w_{@}/[[p]]}$ will include the actual world, and worlds like w_1 and w_3 . It will exclude worlds like w_2 (which differs from the actual world with respect to the independent fact that r is false).

Having identified the revision set for $w_{@}$ given $[[p]]$, we can now follow Veltman in taking the second step, and make the counterfactual hypothesis that $\neg p$ (*the weather is fine*). This requires updating the revision set with $\neg p$. As we have seen, $Rev_{w_{@}/[[p]]} = \{w_{@}, w_1, w_3\}$. If we update this set with $\neg p$ we will end up with $\{w_1, w_3\}$. Sentence q (*Jones is wearing his hat*) is

not true in all the worlds in this set (see w_3). The prediction is that the counterfactual in (26) (repeated below) is false:

(26) If the weather had been fine, Jones would have been wearing his hat.

By allowing the revision of premise sets to take place only with respect to independent facts and letting the laws ‘fill in’ the others, Veltman’s account provides correct truth conditions for the counterfactual in (26).

Veltman’s proposal is interesting to us since it provides us with a method to identify the independent facts in the evaluation world (there may be more than one way of ‘carving up’ the independent facts) and allows us to characterize revision with respect to the independent facts. The laws are allowed to do their work and ‘fill in’ all the dependent facts. Lawlike relations between facts play a role in revision because laws are taken into account when identifying the independent facts. Veltman’s system spells out in a very explicit manner the effect of lawlike dependencies on the revision of the premise set. In the next section we will use revision sets as the modal bases of deontic modals.

4.2 REVISION SETS IN THE SEMANTICS OF DEONTIC *SHOULD*

We will follow Kratzer (1981, 1991) in differentiating between a modal base and an ordering source in the semantics of deontic *should*. We will maintain Kratzer’s views about the ordering source and revise the proposal regarding how to identify the modal base. In evaluating a sentence of the form *should* ϕ in a world w , we will use as a modal base the revision set for w given the embedded proposition ($[[\phi]]$). In looking for the revision set, our objective is to identify the maximal set(s) of independent facts in w compatible with $[[\phi]]$.²⁰

²⁰ The notion of deontic necessity that will be presented here could be thought of as a case of ‘ceteris paribus’ preference: *should* ϕ is true in a world w iff everything else being equal, in w the best worlds are ϕ -worlds. The role and characterization of the ‘ceteris paribus’ clause has been the object of much discussion in the study of preference outside the domain of linguistics (see von Wright 1963 for an early formulation) and has been of interest both in philosophy and artificial intelligence. Van Benthem et al. (2009) investigate ‘ceteris paribus’ in the evaluation of preference within the framework of a modal preference logic. They characterize several preference operators, as well as a modal logic that captures ‘ceteris paribus’ understood as ‘all other things

This means we have to look for the maximal sets of facts that do not force the proposition $[[\neg\phi]]$ (i.e. we need to *retract* $[[\neg\phi]]$). The proposal for deontic *should* is given in (33) (where g is a contextually provided ordering source and the modal base will be provided by $Rev_{w/[[\neg\phi]]}$):

(33) Semantics for *should*

For all worlds w , $[[should\ \phi]]^g(w) = 1$ iff

$\forall w' \in Rev_{w/[[\neg\phi]]}$: if w' is a $g(w)$ -best world, then $w' \in [[\phi]]$.

The auxiliary definitions for revision sets are given in (34) and (35) (repeated from above in a slightly different format to facilitate discussion of *should*):

(34) Where $[[\neg\phi]]$ is a proposition and w is a possible world, $w \downarrow [[\neg\phi]]$ is the set of situations s such that $s \subseteq w$ and there is a basis s' for w such that s is a maximal subset of s' not forcing $[[\neg\phi]]$ within U_w .

(35) Where $[[\neg\phi]]$ is a proposition and w is a possible world,

$Rev_{w/[[\neg\phi]]} = \{w' \in W : w' \in U_w \text{ and there is some } s \in w \downarrow [[\neg\phi]] \text{ such that } s \subseteq w'\}$

As before, the way we identify the revision set will track not only compatibility amongst facts in the evaluation world, but also law-like dependencies between facts. To see how this works, let us go back again to the examples in (2) and (3). We'll begin with (2):

- (2) a. You: The wine should be stored in the fridge.
b. Me: Yes, you are right. But the fridge door got stuck and I can't open it.
c. You: Well, it should be stored in the cellar (then).

Let us make the following assumptions:

being equal' (the notion of 'ceteris paribus' relevant to our proposal), as opposed to 'ceteris paribus' understood as 'all other things being normal'. Reasoning about preferences has also proven important in the domain of artificial intelligence. Bienvenu et al. (2010), for example, investigate the relation between preference logics and preference representation languages such as CP nets (mainly studied by AI researchers). They point to the 'ceteris paribus' principle as a commonality in the interpretation of preferences statements. (I am grateful to an anonymous reviewer for pointing out the broader implication of the 'ceteris paribus' clause, and directing my attention to the AI literature.)

- (36) $p =$ *the fridge door got stuck*
 $q =$ *the wine is stored in the fridge*
 $t =$ *the wine is stored in the cellar*
 $r =$ *Ottawa is the capital of Canada* (a random sentence to illustrate the system)

Suppose that in the actual world there is a law (generalization) according to which if p is true, q is false (in the actual world the wine goes into the fridge through the door, so *if the fridge door got stuck, the wine is not in the fridge*). Examples of possible worlds are presented in (37), where $w_{@}$ is the actual world and the worlds crossed out (double line) do not belong to $U_{@}$:

(37)

	p	q	t	r
$w_{@}$	1	0	0	1
w_1	0	0	0	1
w_2	1	0	1	1
w_3	0	1	0	1
w_4	1	0	1	0
w_5	1	1	0	1

Given the laws and facts as described, only $\langle q, 0 \rangle$ is a dependent fact and the unique basis for the actual world is $s = \{ \langle p, 1 \rangle, \langle t, 0 \rangle, \langle r, 1 \rangle \}$ (s determines $w_{@}$ and it is the minimal situation to do so).

We are interested in (2c) in this scenario (*the wine should be stored in the cellar*):

- (38) a. [**should [the wine be stored in the cellar]**]
b. For all worlds w , $[[\text{(38a)}]]^g(w) = 1$ iff
 $\forall w' \in \text{Rev}_{w/[\neg \text{the wine is stored in the cellar}]}$:
if w' is a $g(w)$ -best world, then the wine is stored in the cellar in w' .

Which worlds do we find in $Rev_{w@/[\neg \text{the wine is stored in the cellar}]}$ ²¹? First we need to identify the set $w@ \downarrow [\neg \text{the wine is stored in the cellar}]$. Since there is a single basis for $w@$, this set will contain a single situation: $s' = \{ \langle p, 1 \rangle, \langle r, 1 \rangle \}$. $Rev_{w@/[\neg \text{the wine is stored in the cellar}]}$ will contain the worlds in $U@$ that extend this situation. Examples are given in (39), with cases of worlds not in the revision set crossed out (single line):

(39)

	<i>p</i>	<i>q</i>	<i>t</i>	<i>r</i>
$w@$	1	0	0	1
w_1	0	0	0	1
w_2	1	0	1	1
w_3	0	1	0	1
w_4	1	0	1	0

Given (33), (38a) is predicted to be true in $w@$ iff in the $g(w@)$ - best worlds in $Rev_{w@/[\neg \text{the wine is stored in the cellar}]}$ the wine is stored in the cellar. In the toy example we are discussing, this will be true since w_2 is better than $w@$ (within the set of worlds in which the wine is not in the fridge, the worlds in which it is in the cellar are better).

The proposal in (33) predicts that in the case of (2), when we evaluate whether the wine should be stored in the cellar, the modal base will consist of worlds that are like the actual world with respect to the fact that the fridge door got stuck and that Ottawa is the capital of Canada. Since the worlds obey the actual laws, they will also be worlds in which the wine is not in the fridge. In some of these worlds the wine is stored in the cellar, and in some it isn't. The sentence *the wine should be stored in the cellar* will be true iff the wine is stored in the cellar in the best worlds in that set (given an ordering source g that corresponds to the contextually given measure of goodness regarding the storage of wine). The proposal in (33) makes correct predictions for the circumstances in which we judge the secondary duty true.

Note that the proposal above allows us to make sense of the intuition that we fall back on a secondary duty when the primary duty in some sense 'can't' be satisfied (i.e. in some sense, the primary duty is not possible). The analysis of (38) predicts that in the circumstances described above, when we evaluate whether the wine should be in the cellar

²¹ It would also be possible to write $Rev_{w@/[\neg r]}$, but I will use English to make it easier to read.

or not, we will leave in place the fact that the fridge door got stuck (that fact will be part of the basis for the evaluation world, it will be taken for granted). Since there are no law-like links between this fact and the fact the wine isn't in the cellar, the fact will not be affected by the revision process (it will survive in the revision set). All the worlds in the revision set will be worlds in which the fridge door got stuck. There is, however, a law-like generalization linking the stuck door to the fact that the wine is not in the fridge: given the fact regarding the door, the wine can't be in the fridge. The result is that when evaluating whether the wine should be in the cellar in this scenario, we only compare worlds in which it isn't in the fridge, and in which, given that the fridge door got stuck and we care about the laws, it couldn't be in the fridge. We fall back on the secondary duty when the facts determine that the primary duty can't be satisfied given the laws.

Let us turn now to the scenario in (3), where we noted that our intuitions were different:

- (3) a. You: The wine should be stored in the fridge.
b. Me: You are right, but I forgot to put it in the fridge.
c. You: #Well, it should be stored in the cellar (then).

We'll make the assumptions below:

- (40) $p = I \text{ forgot to put the wine in the fridge}$
 $q = \text{the wine is stored in the fridge}$
 $t = \text{the wine is stored in the cellar}$
 $r = \text{Ottawa is the capital of Canada}$ (a random sentence for illustration)

As we mentioned earlier, (3c) is false in a world in which I am the person who puts the wine away and there is a link between p , q and t : *if I forgot to put the wine in the fridge, the wine is not in the fridge and the wine is not in the cellar*. In this example, the relevant generalization (law) is that if p is true, then I forgot about the wine and both q and t are false. We find examples of possible worlds in (41), where $w_{@}$ is the actual world and worlds not in $U_{@}$ have been crossed out (double line):

(41)

	p	q	t	r
$w_{@}$	1	0	0	1
w_1	0	0	0	1
w_2	1	0	0	0
w_3	0	1	0	1
w_4	0	0	1	1
w_5	0	1	0	0
w_6	1	1	0	1
w_7	1	0	1	1

Given the laws and facts as above, the basis for the actual world is $s = \{ \langle p, 1 \rangle, \langle r, 1 \rangle \}$ (for all $w' \in U_{@}$, if $s \subseteq w'$, $w' = w_{@}$) (informally, we can say that the truth of p decides what happens to both q and t so neither is in the basis for the actual world – they are not independent).

We are interested in (3c) (*the wine should be stored in the cellar*) in the circumstances described above, with the truth conditions we have already noted (repeated below):

- (42) a. **[should [the wine is stored in the cellar]]**
 b. For all worlds w , $[[42a]]^g(w) = 1$ iff
 $\forall w' \in Rev_{w/[[\neg \text{the wine is stored in the cellar}]]}$:
 if w' is a $g(w)$ -best world, then the wine is stored in the cellar in w' .

Which worlds do we find now in $Rev_{w_{@}/[[\neg \text{the wine is stored in the cellar}]]}$? First we need to identify the set $w_{@} \downarrow [[\neg \text{the wine is stored in the cellar}]]$. Since there is a single basis for $w_{@}$, this set will contain a single situation: $s' = \{ \langle r, 1 \rangle \}$. $Rev_{w_{@}/[[\neg \text{the wine is stored in the cellar}]]}$ will now consist of the worlds in $U_{@}$ that extend this situation. Examples of worlds in $U_{@}$ are given in (43), with worlds not in the revision set crossed out (single line):

(43)

	p	q	t	r
$w_{@}$	1	0	0	1
w_1	0	0	0	1

w_2	— 1 —	— 0 —	— 0 —	— 0
w_3	0	1	0	1
w_4	0	0	1	1
w_5	— 0 —	— 1 —	— 0 —	— 0

Notice that in constructing the revision set this time, we have maintained similarity with the actual world only with respect to the fact that Ottawa is the capital of Canada. The revision set includes worlds in which the wine is neither in the fridge nor in the cellar, worlds in which the wine is in the fridge, and worlds in which the wine is in the cellar. The sentence in (42a) will be true iff in the best worlds in the revision set, the wine is stored in the cellar. But this is not the case. In the best worlds in this revision set, the wine is stored in the fridge (w_3). The prediction made by the proposed analysis (in (33)) is that in this scenario (42a) is false. The revision needed to evaluate (42a) allows us to take into account worlds in which the wine is stored in the fridge, and in the best worlds, this is indeed the case.

As we see, the use of Veltman-style revision sets as modal bases allows us to capture our intuitions regarding the difference between (2) and (3). In evaluating whether a secondary duty *should* ϕ holds, we will remove from the premise set the facts that ‘bring along’ $[[\neg\phi]]$. If, in the absence of those facts, the primary duty becomes ‘possible’, then the primary duty will hold and the secondary duty will not (this is what happens in (3)).

The proposal in (33) gives us a handle on our intuitions regarding the interaction between primary and secondary duties in the examples discussed. However, as it stands, the proposal makes some potentially problematic predictions. So far, we have only considered examples in which *should* ϕ is evaluated in a world in which ϕ is false, and the predictions have been on target. But if we evaluate *should* ϕ in a world in which ϕ is true, the proposal in (33) as it stands predicts that the modal statement will be trivially true, and this is a mistake. The prediction arises because in (33) the revision process that identifies the modal base starts out from the independent facts in the evaluation world and looks for the maximal set(s) that fit with $[[\phi]]$ (i.e. the facts that do not force $[[\neg\phi]]$). If it is actually the case that $\neg\phi$ (as in the examples so far), then we will quantify over worlds that match the actual world only with respect to a subset of facts. But if it is actually the case that ϕ , then no revision will be necessary. In this case the modal base will be the singleton set containing the actual world (in Kratzer’s term, a ‘totally realistic’ modal base) and *should* ϕ is predicted to be trivially true.

The prediction that given the truth of ϕ , the truth of *should* ϕ is trivial is problematic. It means, for example, that if it is true that Jones murdered Smith, it will also be true that it should be the case that Jones murdered Smith.²² It would be possible to revise (33) in a way that blocks triviality predictions, as in (44) (where g is the contextually provided ordering source):

- (44) Semantics for *should* (modified)
 For all worlds w , $[[\textit{should } \phi]]^g(w) = 1$ iff
 $\forall w' \in \text{Rev}_{w/A}$: if w' is a $g(w)$ -best world, then $w' \in [[\phi]]$,
 where $A = [[\phi]]$ if ϕ is true in w and $A = [[\neg\phi]]$ if ϕ is false in w .

The modification in (44) guarantees that in evaluating *should* ϕ , neither the truth of ϕ nor its falsity lead to trivial results. The revised proposal allows us to take advantage of the insights of the analysis proposed in (33) and at the same time avoid problematic triviality predictions. Moreover, the revised proposal is in line with research that has argued that in evaluating modal statements, it is necessary to guarantee ‘diversity’ in the domain of quantification (see e.g. Condoravdi’s 2001 ‘diversity condition’).²³ Such arguments have been presented by e.g. Frank (1987), Condoravdi (2001), see also von Stechow (1999, 2001), and are in line with the proposal in Arregui (2010); views against such accounts have been presented in Zvolenszky (2002). It is worth noting that while the proposal in (44) blocks triviality, the arguments in favor of counterfactual-style revision presented in this paper have only involved examples in which *should*-statements were not trivial (further research would be needed to fully back up (44)).

5. REVISION PUZZLES

²² The observation that ‘maximally’ fact-dependent accounts of deontic modals predict that true statements are deontically necessary/possible can be found in the literature (see e.g. Feldman 1986). For a proposal that blocks this result for deontic modals within a Kratzer-style situations framework, see Arregui (2010).

²³ I am grateful to an anonymous reviewer for pointing to the broader issues regarding diversity in the domain of quantification.

The semantics for *should* presented in section 4 is maximally sensitive to facts and to relations between facts. Deontics are treated in a manner that is similar to counterfactuals. In the case of counterfactuals, it has long been accepted that truth depends on what is going on in the evaluation world: once we identify the relevant measure of similarity (contextually determined), the truth of a counterfactual is decided by the facts. Dependencies between facts are themselves facts about the evaluation world, and must be taken into account in the semantics. We have argued that in the case of *should* too, dependencies between facts are amongst the facts that must be taken into account in the semantics. Otherwise, we'll get the truth conditions wrong.

Kratzer has presented a series of puzzles (e.g. Kratzer 1981b) to argue that we need to take into account dependencies between facts in the semantics of counterfactuals. In this section we will investigate the explanatory power of the analysis of *should* proposed in (33)/(44) by reproducing in the domain of deontic statements the kinds of puzzles that Kratzer noted for counterfactuals. This will provide additional evidence in support of the revision-based proposal in (33). As we will see, it is possible to reconstruct in deontic statements puzzles parallel to those noted by Kratzer for counterfactual conditionals. As in the case of counterfactuals, the solution to those puzzles can be found in a semantics that pays attention to factual dependencies. The revision-based account of *should* that we have developed in section 4 [(33)] will provide good insights into the data.

5.1 REVISION-PUZZLES IN COUNTERFACTUALS

We will examine three challenges noted by Kratzer (1981b) for a naive 'revision style' analysis of counterfactuals (I will simply give a brief overview of the puzzles and will not attempt to present Kratzer's solution). We will dub the puzzles *the puzzle of the priority of certain facts*, *the puzzle of the grouping of certain facts* and *the puzzle of the irrelevance of certain facts*. We will examine them in turn below.

We will begin with the *puzzle of the priority of certain facts*. Here is Kratzer's example illustrating the issue:

- (45) *Hans and Babette spend the evening together. They go to a restaurant called “Dutchman’s Delight”, sit down, order, eat and talk. Suppose now, counterfactually, that Babette had gone to a bistro called “Frenchman’s Horror” instead. Where would Hans have gone? (I have to add that Hans rather likes this bistro).* (Kratzer 1981b: 206)

Kratzer’s point is that in the scenario above, we would judge the counterfactual in (46) true:

- (46) If Babette had gone to “Frenchman’s Horror” instead, Hans would have gone to “Frenchman’s Horror”.

In evaluating the counterfactual in (46), we will set aside some of the facts in the actual world. We do not want to quantify over worlds like the actual world with respect to the restaurant that Babette actually went to, since this is incompatible with the antecedent. What is interesting about our intuition that (46) is true is that it shows that we quantify over worlds that are also different from the actual world with respect to the restaurant Hans went to. As Kratzer notes, there is no logical (or natural!) incompatibility between the proposition that Babette went to Frenchman’s Horror and Hans went to Dutchman’s Delight. Yet, when we hypothesize that Babette went to a different restaurant, we also imagine that Hans went to a different restaurant. In evaluating (46), we seem to give priority to the fact that they should be in the restaurant together over the fact that Hans actually went to Dutchman’s Delight. Why should one fact have priority over the other?

Kratzer illustrates the *puzzle of the grouping of certain facts* with the example in (47):

- (47) *Regina and I go on a walk in the bush. We have to pass a hanging bridge. I pass first. Regina is waiting. I am in the middle of the bridge. Suppose now, counterfactually, that I had passed a bit faster and had just left the bridge. Where would Regina be? Would she still be waiting?*
(Kratzer 1981b: 206)

Kratzer points out that we do not judge the counterfactual in (48) true. Our intuitions tell us that if I had passed a bit faster, Regina might still be waiting but she might also have started to cross:

(48) If I had passed a bit faster, Regina would still be waiting.

The puzzle is that if we quantify over worlds that only differ from the actual world in that I passed a bit faster, we would be stuck with worlds in which Regina is still waiting, and the prediction is that (48) should be true in the scenario above. The mystery is to explain why, when we set aside the fact that I am in the middle of the bridge, we also set aside the fact that Regina is waiting (and in this way make available both worlds in which she is waiting and worlds in which she is not). The two facts appear grouped together (they stand and fall together) in the semantics of the counterfactual.

The last puzzle to be discussed is the *puzzle of the irrelevance of certain facts*. It is illustrated with Kratzer's example in (49):

(49) *I am taking dancing lessons at Wander's dancing school. Last Saturday, there were five men to dance with: John, James, Jack, Joseph and Jeremy. I danced with the latter three only. Suppose now, counterfactually, that I had danced with at least four of the men. With whom might I have danced?*
(Kratzer 1981b: 206)

Kratzer's point with this example is that we do not judge the counterfactual in (50) true:

(50) If I had danced with four men, I would have danced with Jack, Joseph, Jeremy and some other man.

We don't automatically judge (50) true. Intuitively, any combination of four men would have been possible. This means that in identifying the worlds quantified over, we do not pay attention to the facts regarding the three men who I have actually danced with. For some reason, the fact that I did dance with those specific three men can be ignored when we search for worlds in which I danced with four men.

Kratzer's puzzles present a challenge for a theory of counterfactuals that seeks to model the semantics of counterfactuals in terms of revision to a set of premises. Since this has been our strategy for *should* as well, our proposal potentially inherits these problems. My goal in the next section will be to construct examples with *should*-statements that parallel Kratzer's revision puzzles in the domain of deontic modals. We will evaluate the Veltman-

inspired revision proposal for *should* in (33) in the light of these new examples. As we will see, our proposal is sufficiently sensitive to the dependencies between facts to handle these challenges.

5.2 REVISION PUZZLES IN DEONTIC STATEMENTS

Let us begin by considering the *puzzle of the priority of certain facts*. Suppose that Babette and Hans have arranged to go to the cinema together. Babette arrives first and goes in, sitting down towards the back of the movie theater. Hans comes in a minute later, and sits down beside her. The movie starts, and Babette realizes that Hans is having trouble seeing the screen. Babette had forgotten he has eye problems, and thinks to herself:

(51) I should have sat closer to the front.

Babette doesn't have eye problems. Worlds in which Babette sits closer to the front and Hans stays where he is are not better than worlds in which they both sit at the back. But we judge (51) true because we assume that they will sit together, and worlds in which Babette is sitting closer to the front will also be worlds in which Hans is sitting closer to the front and sees the screen properly. In judging (51) true, we give priority to the fact that they sit together over the fact that Hans is actually sitting towards the back. We can account for our intuitions in a Veltman-style proposal (as in (33)) if we assume that there is some law-like relation between the fact that Babette and Hans come to the cinema together and that they sit together (*if you go with someone to the cinema, you sit where they sit*). This isn't a logical law or a natural law. But when we judge (51) true, we assume that if Hans goes with Babette to the cinema, where he sits is not actually an independent fact. Where he sits depends on where she sits. In Veltman's terms, this means that the fact that he sits in the back of the cinema will not be part of the basis for the evaluation world in (51). The revision set for (51) will include worlds in which Babette sits near the front and Hans sits with her and worlds in which Babette sits near the back and Hans sits with her. In the best worlds in the revision set, Babette sits near the front, Hans is with her, and he sees the screen properly. This predicts that (51) is true.

What about the *puzzle of facts that are grouped together*? Let us consider another example. Suppose that Regina and I have gone for a walk and have come across a hanging bridge. I go first, and Regina waits. I am halfway across, when I realize that a puma has crept up behind Regina without her noticing and she is in danger. But it is also dangerous for Regina to start crossing the bridge while I am still on it. I think to myself:

(52) I should not still be on the bridge.

I am in no danger from the puma, and I am in no danger on the bridge. We judge (52) true only if we assume that if I had crossed more quickly, Regina might not still be waiting and so might have been out of reach of the puma. We can capture this in our Veltman-style analysis of *should* if we recognize that in the actual world, Regina waiting is a dependent fact. If I were not on the bridge, she might not be waiting. The relevant generalization will not be a law of logic, nor a law of nature, but something weaker, along the lines of *if there is somebody crossing a hanging bridge, you wait*. Given this generalization (or something like it), the fact that Regina is waiting to cross the bridge will not part of the basis for the evaluation world in (52). In the revision set for the *should*-statement in (52) we will find worlds in which I am still on the bridge and Regina is waiting, worlds in which I have crossed the bridge and Regina is waiting, and worlds in which I have crossed the bridge and Regina is not waiting. In the best ones, I will have crossed the bridge and Regina will not be waiting. The prediction is that (52) is true.

Let us turn now to the last puzzle illustrated by Kratzer for counterfactuals, the *puzzle of the irrelevant facts*. Suppose that there is a wedding dinner with guest spread over several tables. The bride, my best friend, has asked me to sit with her three brothers in law: John, Jack and Jim. They are very shy and she wants to make sure somebody friendly has dinner with them. As it turns out, two of the brothers (John and Jack) sit at one table, and Jim sits at another. I sit with John and Jack. It might not be reasonable for my friend to come up to me and say (53) (there is a sense in which we would judge (53) false):

(53) You should be sitting with Jim.

However, she could indicate that she was unhappy with the situation by uttering (54):

(54) (This is such a pity!) You should be sitting with John, Jack and Jim.

Let us consider (53). Currently, I am not sitting with Jim. There are various ways in which things could be different so that I would be sitting with Jim: I could have sat with Jim at his table (and not with John and Jack), or Jim could have sat with John, Jack and me at our table. If we allow for all such options, it is true that in the best worlds I would be sitting with Jim (he would have sat with us at our table). The fact that we can judge (53) false indicates that this is not the only way to resolve the *should*-statement. We can judge (53) false because we can take it to claim that I should have sat with Jim instead of sitting with John and Jack. To make sense of this judgment, we need to claim that in evaluating (53) we take for granted that if I sit with John and Jack (given the way things were regarding the distribution of people at the table), I don't sit with Jim (*if you sit with John and Jack, you don't sit with Jim*). If we take these generalizations for granted, the fact that I don't sit with Jim is dependent on my sitting with John and Jack. To evaluate (53), it is necessary to exclude from the basis of the evaluation world the fact that I sit with John and Jack (since it would bring with it the fact that I don't sit with Jim). The revision set for (53) will include worlds in which I sit with John and Jack and worlds in which I sit with Jim. Assuming that there is nothing better about sitting with either group of brothers, it will be false that in the best worlds in the revision set, I sit with Jim. The prediction is that (53) is false.

The case of (54) is different. To evaluate (54) we have to exclude all facts incompatible with the proposition that I sat with John, Jack and Jim. This means we have to exclude the actual facts regarding the distribution of people around tables (notice that in judging (54) true, we do not take for granted that Jim would have sat at our table or that we would have sat at his, it is equally possible that we all sat at some third table). The facts regarding where people were actually sitting that night are not relevant, we are allowed to consider possibilities with all kinds of sitting arrangements. In the best ones, I sit with the three brothers and (54) is true.

6. CONCLUSION

In this paper I have argued for a parallelism between deontic statements and counterfactuals. I have shown that the strategies we pursue in identifying the domain of quantification of counterfactual conditionals are also operational in the way we identify the domain of quantification of deontic statements. In identifying the domain of quantification of deontic statements, we pay attention to facts and to the dependencies between facts in a manner that is similar to what we do with counterfactual conditionals. In both cases, we ignore facts in order to allow law-like patterns to hold.

We have examined two types of empirical evidence in favor of the ‘counterfactual-style’ analysis for *should*. We began our presentation with a discussion of our intuitions regarding primary and secondary duties. We observed that we could make sense of those intuitions with a semantics for *should* that paid attention to dependencies between facts. After presenting such an analysis, we then expanded the empirical base to include examples of deontic statements modeled after counterfactual puzzles discussed by Kratzer (Kratzer 1981b). We observed that the account we had proposed for *should* could also shed light on our intuitions regarding these examples.

The idea that the semantics for counterfactuals needs to pay attention to dependencies between facts has been noted in the literature. We have followed Veltman (2005) in providing a solution that recognizes the role of law-like dependencies between facts (where those laws need not be natural laws or laws of logic). The flexibility of the notion of ‘law’ in Veltman’s analysis has proven crucial in the success of the proposal, as it has allowed us to invoke laws to recognize a variety of relations between facts. As in Veltman’s proposal, the notion of what is a ‘law’ itself is kept rather intuitive and is not analyzed. There are alternative accounts of the role of the dependencies between facts in the semantics of counterfactuals (e.g. Kratzer 1989, 1990, 2002, 2005, Schulz 2007). A discussion of such alternatives in the domain of deontics remains outside the scope of this paper. My objective has been to highlight the similarity between the two types of modality in terms of how we pay attention to facts. Dependencies between facts are taken into account in the interpretation of both counterfactuals and deontics. I have captured this with an analysis that places a similar notion of revision at the heart of both types of modalities.

REFERENCES

Åqvist, Lennart (2002). Deontic Logic. In Gabbay, D. & Guenther, F. (eds.) *Handbook of Philosophical Logic*, Volume 8, Kluwer, The Netherlands, 147-264.

Arregui, Ana (2008). Chisholm's Paradox: On detaching obligations from deontic conditionals. In Friedman, T. & Ito, S., (eds.), *Proceedings of Semantics and Linguistic Theory 18*, 55-72. Ithaca, N.Y.: CLS, Cornell University.

Arregui, Ana (2009). On similarity in counterfactuals. *Linguistics and Philosophy*, 32(3), 245-278.

Arregui, Ana (2010). Detaching *if*-clauses from *should*. *Natural Language Semantics*, 18(3), 241-293.

van Benthem, Johan, Patrick Girard & Olivier Roy (2009). Everything else being equal: A modal logic for ceteris paribus preferences. *Journal of Philosophical Logic* 38, 83–125.

Bhatt, Rajesh & Roumyana Pancheva (2006). Conditionals. In *The Blackwell Companion to Syntax*, v. I, Blackwell, 638-687.

Bienvenu, Meghyn, Jérôme Lang & Nic Wilson (2010). From Preference Logics to Preference Languages, and Back. In Lin, F., Sattler, U. & Truszczyński, M. (eds.), *Principles of Knowledge Representation and Reasoning: Proceedings of the Twelfth International Conference (KR2010)*, Association for the Advancement of Artificial Intelligence Publications, 414-424.

Carmó, José & Andrew Jones (2002). Deontic Logic and Contrary-to-Duties. In Gabbay, D. & 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.

Copley, Bridget (2006). What Should *Should* mean? (ms.)

Feldman, Fred (1986). *Doing the best we can*. Dordrecht: Reidel.

von Stechow, Kai (1999). "The Presupposition of Subjunctive Conditionals". In *The Interpretive Tract*, edited by Sauerland, U. & Percus, O., MIT Working Papers in Linguistics 25, pp. 29–44. Cambridge, MA: MITWPL.

von Stechow, Kai (2001). Counterfactuals in a dynamic context. In Kenstowicz, M. (ed.) *Ken Hale: A Life in Language*, MIT Press, 123-152.

von Stechow, Kai & Sabine Iatridou (2008). How to Say *Ought* in Foreign: The Composition of Weak Necessity Modals. In Guéron, J. & Lecarme, J. (eds.) *Time and Modality*, Springer, 115-141.

Kratzer, Angelika (1981a). The Notional Category of Modality. In Eikmeyer, H.J. & Rieser, H. (eds.) *Words, Worlds, and Contexts*, de Gruyter, Berlin-New York, 38-74. (Reprinted in Portner, P. & Partee, B. (eds.) (2002). *Formal Semantics. The Essential Readings*. Blackwell, Oxford.)

Kratzer, Angelika (1981b). Partition and revision: the semantics of counterfactuals, *Journal of Philosophical Logic* 10(2), 201-216.

Kratzer, Angelika (1989). An investigations of the lumps of thought. *Linguistics and Philosophy* 12(5), 607-653.

Kratzer, Angelika (1990). How specific is a fact? short conference version in the *Proceedings of the 1990 Conference on Theories of Partial Information*. Center for Cognitive Science and College of Liberal Arts at the University of Texas at Austin.

Kratzer, Angelika (1991). Modality. In von Stechow, A. & Wunderlich, D. (eds.) *Handbuch Semantik/Handbook Semantics*. de Gruyter, Berlin and New York, 639-650

Kratzer, Angelika (2002). Facts: Particulars of Information units? *Linguistics and Philosophy*, 25(5-6), 655-670.

Kratzer, Angelika (2005). Constraining Premise Sets for Counterfactuals. *Journal of Semantics* 22(2), 153-158.

Pearl, Judea (2000). *Causality: Models, Reasoning, and Inference*. Cambridge, UK: Cambridge University Press.

Prakken, Henry & Mark Sergott (1996). Contrary-to-Duty Obligations. *Studia Logica* 57(1), 91-115.

Schulz, Katrin (2007). Minimal models in semantics and pragmatics: free choice, exhaustivity and conditionals. PhD thesis, University of Amsterdam, Institute for Logic, Language and Computation, Amsterdam. ULR. <http://www.ilc.uva.nl/publications/dissertations/DS-2007-04.text.pdf>.

Tichy, Pavel (1976). A counterexample to the Stalnaker-Lewis analysis of counterfactuals. *Philosophical Studies* 9(4), 271-273.

Veltman, Frank (2005). Making counterfactual assumptions. *Journal of Semantics* 22(2), 159-180.

von Wright, Georg H. (1963). *The logic of preference*. Edinburgh: Edinburgh University Press.

Zvolenszky, Zsófia (2002). 'Is a Possible-Worlds Semantics of Modality Possible? A Problem for Kratzer's Semantics'. In Jackson, B. (ed.), *Proceedings of SALT 12*, 339-358. Ithaca, N.Y.: CLC Publications.