Semantics & Pragmatics Volume 9, Article 6: 1–60, 2016
http://dx.doi.org/10.3765/sp.9.6

How similar is similar enough?*

Michela Ippolito
University of Toronto

Submitted 2014-07-02 / First decision 2014-11-11 / Revision received 2015-04-28 / Accepted 2015-04-29 / Early access published 2016-05-10 / Published 2017-03-12

Abstract I investigate the issue of the context-dependence of counterfactual conditionals and how the context constrains similarity in selecting the right set of worlds necessary to arrive at the correct truth-conditions. I propose that similarity is constrained by what I call Consistency and Non-Triviality. Assuming a model of the discourse along the lines proposed by Roberts (2012) and Büring (2003), according to which conversational moves are answers to often implicit questions under discussion, the idea behind Non-Triviality is that a counterfactual statement answers a conditional question under discussion and, therefore, is required to make a non-trivial assertion. I show that non-accidental generalizations which have often been taken to play an important role in the interpretation of counterfactuals, are crucial in selecting which conditional question is under discussion, and I propose a formal mechanism to identify the relevant question under discussion.

Keywords: counterfactuals, similarity, possible worlds, question under discussion, context-dependence, discourse tree, triviality

1 Measuring similarity across worlds

Lewis 1973 argues for a possible worlds semantics of counterfactual conditionals according to which a conditional of the form *if it had been the case*

* I would like to thank the reviewers for this journal for their very helpful reviews, as well as the editor Josh Dever for his insightful comments throughout the process. Many thanks as well to the audiences at some of the venues where this work was presented: the syntax-semantics project at the University of Toronto, Toronto-Ottawa-Montreal Semantics Workshop, the University of Milano-Bicocca, the Linguistics Department at UCSC. I would also like to thank the *Semantics and Linguistic Theory 23* reviewers for their comments on a much earlier version of this work. All mistakes are mine.

that ϕ, it would have been the case that ψ is true in a possible world w just in case the consequent ψ is true in all those ϕ-worlds that are as similar to the evaluation world w as allowed by the counterfactuality of ϕ. The formal truth-conditions for *if it had been the case that ϕ, it would have been the case that ψ* (or *if ϕ, would ψ* in short) are given in (1).

(1) $\phi \Box\!\rightarrow \psi$ is true at a world w (according to a given comparative similarity system) if and only if either (a) no ϕ-world belongs to S_w (the set of worlds accessible from w), or (b) there is a ϕ-world w' in S_w such that, for any world w'', if $w'' \leq_w w'$ then $\phi \rightarrow \psi$ (material implication) holds at w''.

What determines the set of ϕ-worlds in which the consequent ψ is required to be true is the relation of comparative similarity $\leq_w$, whose definition is given in (2).

(2) $w' \leq_w w''$ means the world w' is at least as similar to the world w as the world w''.

For any theory of counterfactual conditionals that employs the notion of comparative similarity in the sense of Lewis 1973 or some other mechanism where similarity across worlds is a key ingredient in the selection of the relevant set of antecedent worlds, it is crucial to say how we measure similarity across worlds. However, spelling out exactly which worlds are most similar to the evaluation world turns out to be a very difficult task. In the remaining of this section I will introduce some well-known counterfactual cases to illustrate the complexity of this task. The goal of this survey of cases is not to provide a review of the literature but to introduce the set of facts that my proposal aims to account.

In most cases, we have very clear intuitions about the truth or falsity of these conditionals, yet spelling out exactly the measure of similarity that is needed to account for our intuitions has turned out to be one of the most difficult problems for both philosophers and linguists. To appreciate the puzzle, consider Jones and the rain example from Tichý 1976. Jones always wears his hat if the weather is bad. If the weather is good, Jones wears his hat at random. Today the weather is bad and Jones is wearing his hat. In this context, suppose someone were to utter the counterfactual in (3).

(3) If the weather had been fine, Jones would be wearing his hat.

We judge (3) false, which means that in selecting counterfactual worlds in which the weather is fine that are otherwise maximally similar to the actual world, we disregard the fact that Jones is wearing his hat. Lewis 1979's list of priorities according to which similarity of laws trumps similarity of particular facts seems equipped to account for the judgment in (3): assuming determinism, the worlds we select are those worlds that shared the same history as the actual world up to the divergence time, i.e., the time when (thanks to a "miracle") the deterministic chain of events broke and these worlds took different paths following the actual laws. That is, when selecting the most similar worlds, we need to select those worlds that are just like the actual world up until they diverge from the actual world, but that follow their own course afterwards. Applied to Tichý's example, these worlds are going to be worlds that are just like the actual world up to the time when some miracle breaks the deterministic chain of events and the weather turns out to be fine but which, after that, follow undisturbed the actual laws. We should not try to make these worlds converge again just for the sake of maximizing the number of particular facts in common with the actual world since this would involve more "miracles" or inexplicable violations of the actual laws. Thus, since the actual laws say that if the weather is fine Jones might or might not wear his hat, the conditional in (3) comes out false.

However, things are not so simple. Consider a variant of Tichý's example from Veltman 2005. Every morning Jones tosses a coin. If heads comes up and the weather is fine, then he wears his hat. If the weather is bad, Jones always wears his hat (regardless of the outcome of the coin-tossing). Today the weather is bad, heads came up, and Jones is wearing his hat. In this context we judge (3) true. Lewis's system of priorities requires that we select those worlds that after the divergence proceed according to the actual laws. Since the outcome of a coin tossing is probabilistic, there are going to be worlds where heads comes up and worlds where tails comes up. Hence, the conditional in (3) should be false.

This tension between keeping and removing facts also emerges with respect to the same fact, as shown by the following two examples. Let us begin with (4), from Arregui 2009.

(4) Peter presses the button in a completely random coin-tossing device and the coin comes up heads. If Susan had pressed the button, the coin would have come up heads.

This counterfactual is judged false. One might suggest that the counterfactual is judged false because, since the outcome of Susan's coin tossing is up to the probabilistic laws that regulate this type of physical event, in keeping the fact that heads came up one would already assuming the outcome of such coin tossing event and therefore make these laws vacuous. Now consider (5), a variant of the coin tossing example given in (4), adapted from Ippolito 2013. The diacritic F on *you* indicates the presence of focus on the pronoun.

(5) Peter and Susan are taking turns at pressing a button on a completely random coin-tossing device. They both bet each time one presses the button, but (as part of their game) only the one actually pressing the button pays $10 if he or she loses. It is Peter's turn to press the button. Peter bets that the coin will come up heads, Susan bets that it will come up tails. Peter presses the button and heads comes up. Peter wins. Susan had bet on tails but since she wasn't the one pressing the button she does not have to pay $10. Now I say: Susan, you're lucky! If you$_F$ had pressed the button, you would have lost $10.

Unlike (4), we judge (5) true. In (5), unlike what we saw in (4), we do keep the fact that the coin came up heads, even despite the fact that this does make the probabilistic laws of nature we are assuming vacuous. The problem of explaining how we measure similarity across worlds remains unsolved.

One influential way of analyzing counterfactuals goes under the name of Premise Semantics and goes back to Ramsey 1929 and Goodman 1947, among others. There are different varieties of Premise Semantics, but the basic idea which is shared by all of them is that a counterfactual *if* ϕ, *would* ψ is true in w just in case ψ follows from ϕ together with a "suitable" set of premises. This is shown schematically in (6).

(6) ϕ
$\chi_1, \ldots, \chi_n$
$\therefore\ \psi$

The premises are propositions true in the actual world. Therefore, the question of selecting the suitable premises is the question of selecting which actual facts we keep and which ones we let go.

Kratzer 1989 proposes a version of Premise Semantics based on the notion of *lumping*. Here is a familiar illustration of the lumping relation.

(7) *Dialogue with a lunatic*

Lunatic: What did you do yesterday evening?
Paula: The only thing I did yesterday evening was paint this still life over there.
Lunatic: That is not true. You also painted these apples and you also painted these bananas. Hence painting this still life was not the only thing you did yesterday evening.

The lunatic's response in (7) is bizarre. Intuitively, this is because the proposition that Paula painted a still life and the proposition that Paula painted apples and bananas do not refer to separate facts. The proposition that Paula painted a still life *lumps* the proposition that Paula painted apples and bananas. The technical definition of lumping is given in (8). S is the set of possible situations and $\wp(S)$ is the power set of S, i.e., the set of propositions.

(8) For all propositions p and $q \in \wp(S)$ and all $w \in W$: p lumps q in w iff the following conditions hold:
(i) $w \in p$;
(ii) for all $s \in S$, if $s \leq w$ and $s \in p$, then $s \in q$.

When lumping is applied to the analysis of counterfactuals, the set of propositions relevant for the truth-conditions of a counterfactual *if* ϕ*, would* ψ is required to be (i) consistent, (ii) must include ϕ; (iii) must be closed under lumping; and (iv) must be closed under logical consequence.

Lumping is designed to account for our intuitions in a number of counterfactual cases, in particular the King Ludwig of Bavaria example. Here are the details of the case. King Ludwig of Bavaria likes to spend his weekends at Leoni Castle. Whenever the Royal Bavarian flag is up and the lights are on, the king is in the castle. At the moment, the lights are on, the flag is down, and the king is away. Suppose counterfactually that the flag were up.

(9) a. If the flag were up, the king would be in the castle.
b. If the flag were up, the lights would be off.

We judge (9a), but not (9b), true. Let us see how the lumping machine accounts for the King Ludwig's example. The propositions involved are: (a) whenever the flag is up and the lights are on, the King is in the castle; (b) the flag is down; (c) the lights are on; (d) the king is away; (e) the flag is up. The proposition in (e) is the counterfactual antecedent and it must be included in the set of propositions relevant for the truth-conditions of the counterfactual.

The question is: which propositions among the ones listed above must be removed in order to accommodate (e)? As we saw above, there are in principle two possibilities.

(10)

Possibility (i):	**Possibility (ii):**
(e) the flag is up	(e) the flag is up
(a) whenever the flag is up and the lights are on, the king is in the castle	(a) whenever the flag is up and the lights are on, the king is in the castle
(c) the lights are on	(d) the king is away

According to Possibility (ii), we should be able to remove the propositions that the lights are on and keep the proposition that the king is away. If this possibility were available, the counterfactual in (9b), repeated below, would incorrectly be predicted to be true.

(11) If the flag were up, the lights would be off.

Fortunately, Kratzer's lumping can rule out Possibility (ii), as follows. Proposition (a) and proposition (d) jointly imply that either the flag is down or the lights are off. Since the premise set is closed under logical consequence, we have to add this disjunctive proposition to the set. The problem is that this disjunctive proposition lumps (b), which is inconsistent with the counterfactual assumption (e) already in our set. However, since the set of propositions we select is closed under lumping, (b) must be included. The conclusion is that Possibility (ii) is ruled out by lumping and the counterfactual in (93b) is out.

The proposition that whenever the lights are on and the flag is up, the king is in the castle has a special role in arriving at the correct truth-conditions for counterfactuals because it is a non-accidental generalization and as such must be included in the set of propositions that are selected. Now, Kratzer 2012 observed that logically equivalent non-accidental generalizations can trigger different truth-value judgments if they have different forms. Consider again the King Ludwig of Bavaria example and the non-accidental generalization we have been using in (12).

(12) Whenever the lights are on and the flag is up, the king is in the castle.

Now consider a variant of the original example.

(13) Whenever the king is away, the lights are out or the flag is down. Right now, the king is away, the flag is down and the lights are on. What if the flag were up?

Kratzer's observation is that in this context, the sentence in (11) is no longer judged false. Our judgments have changed and we no longer have the clear judgments we had before. This is very surprising for we haven't changed any of the facts about this world, and we have only replaced our old statement in (12) with a logically equivalent one in (14).

(14) Whenever the king is away, the lights are out or the flag is down.

In the remaining of this section I will summarize the proposal that is sketched in Kratzer 2012 because, unlike Kratzer 1989 and other proposals (e.g., Veltman 2005) within premise semantics, Kratzer's more recent work has an important similarity with the main idea that I am going to defend in this paper. This idea is that what determines which propositions true in the actual world are going to be members of the premise set to which the counterfactual antecedent is added is neither merely determined by a logical relation between propositions nor is it merely determined by the relations between the facts or situations that these propositions are about, as in the case of the lumping relation in (8), nor solely by a combination of the two. Other constraints are at work in selecting the premise set. We learned from Lewis's and Stalnaker's work that counterfactuals are vague and the question, as Kratzer puts it, is which kind of explanation predicts the vagueness of counterfactuals best.

Here is where the contrast in truth-value judgment between the original King Ludwig of Bavaria example and the variant in (13) becomes crucial. Recall that in these two cases, the facts about the world are the same (the king is away, the flag is down, the lights are on) and the non-accidental generalizations are logically equivalent. Yet, our disposition to assent to the truth of (9a) and (9b) changes. Kratzer proposes to capture this fact by means of what she calls the *Confirming Proposition Constraint (CPC) for Base Sets*, where a *Base Set* is a set of propositions describing the facts of the world of evaluation. The CPC for Base Sets is given in (15),

(15) *CPC for Base Sets*
When constructing a Base Set, privilege confirming propositions for non-accidental generalizations.

The notion of a "confirming proposition" is crucial here. A proposition p confirms a proposition q iff p lumps q in every world where both p and q are true. Kratzer's example, Wason's Selection Task, helps to illustrate this concept.

(16) The subjects (students) were presented with an array of cards and told that every card had a letter on one side and a number on the other side, and that either would be face upwards. They were then instructed to decide which cards they would need to turn over in order to determine whether the experimenter was lying in uttering the following statement: If a card has a vowel on one side then it has an even number of the other side. (Wason 1966)

The vast majority of subjects selected cards that showed a vowel or an even number, even though the correct response should have been to select the cards that showed a vowel or an odd number. Selecting the cards showing an odd number, would have allowed the subject to falsify the statement. Kratzer's interpretation of this fact is that we are biased towards confirming propositions. Now, going back to the King Ludwig of Bavaria's example, the two logically equivalent non-accidental generalizations repeated below have different confirming propositions.

(17) Whenever the lights are on and the flag is up, the king is in the castle.
Example of confirming proposition: Right now, the lights are on, the flag is up, the king is in the castle.

(18) Whenever the king is away, the lights are out or the flag is down.
Example of confirming proposition: Right now, the king is away, the lights are out, the flag is up.

Once a non-accidental generalization is selected (this could be either because the context has explicitly introduced it and therefore has made it salient, as in (13), or because of extra-linguistic requirements such as that propositions in the Base Set be "cognitive viable"),[1] propositions confirming the non-accidental generalization that has been selected are privileged in the process of assembling a Base Set that includes this non-accidental generalization.

Schematic truth-conditions for *would* and *might* conditionals in Kratzer's premise semantics are given in (19) and (20).

1 See Kratzer 2012: 132 for a discussion of the concept of cognitive viability.

(19) *Would-counterfactuals*
Given a world w and an admissible Base Set F_w, a *would*-counterfactual with antecedent p and consequent q is true in w iff for every set in $F_{w,p}$ there is a superset in $F_{w,p}$ that logically implies q.

(20) *Might-counterfactuals*
Given a world w and an admissible Base Set F_w, a *might*-counterfactual with antecedent p and consequent q is true in w iff there is a set in $F_{w,p}$ such that q is compatible with all its supersets in $F_{w,p}$.

The set $F_{w,p}$ is what Kratzer calls the "Crucial Set": it is the set of all subsets A of $F_w \cup \{p\}$ such that (i) A is consistent; (ii) $p \in$ A; (iii) A is closed under lumping in the evaluation world w (that is, for all $q \in$ A and $r \in F_w$: if q lumps r in w, then $r \in$ A).

The CPC can now be defined for the Crucial Set: when assembling the Crucial Set, privilege those premise sets that logically imply confirming propositions for the non-accidental generalizations they contain. The CPC will thus have truth-conditional effects.

Before turning to the King Ludwig of Bavaria example, let me illustrate how Kratzer's proposal account for Tichý's original example. Jones always wears his hat if the weather is bad. If the weather is good, Jones wears his hat at random. Today the weather is bad and Jones is wearing his hat. In this context, suppose someone were to utter the counterfactual in (21).

(21) If the weather had been fine, Jones would be wearing his hat.

This counterfactual is judged false. The confirming proposition for the non-accidental generalization that whenever the weather is bad Jones wears his hat, is that the weather is bad right now *and* Jones is wearing his hat. The CPC for Base Sets is going to privilege this proposition and, crucially, the Base Set is required to be non-redundant, where redundancy is defined as shown in (22).

(22) *Redundancy*
A set of propositions is redundant if it contains propositions p and q such that $p \neq q$ and $p \cap W \subseteq q \cap W$.

A non-redundant set of propositions will not be allowed to contain a proposition and its (proper) logical consequences. Hence, this requirement will rule out base sets can contain the two distinct propositions that the weather is bad right now and that Jones is wearing his hat, in addition to their conjunction. Since this conjunction cannot be added as it is inconsistent with the counterfactual antecedent, neither that the weather is bad nor that Jones is wearing his hat will be in the Base Set. The CPC and the non-redundancy requirement together make sure that if we remove the bad weather, we are no longer committed to Jones's hat.

Going back to the King Ludwig of Bavaria counterfactuals, how does the CPC explain the different judgments we have in the original example from Kratzer 1989 and in the variant discussed in Kratzer 2012? In the original example, the salient non-accidental generalization is (17) and the CPC requires that we choose Possibility (i) because it logically implies the confirming proposition that the flag is up, the lights are on and the king is in the castle. Because of the CPC, all privileged subsets in the Crucial Set can be expanded to a superset logically implying that the king is in the castle and, consequently, the *would* conditional in (9a) is predicted to be true.

In the variant we have been considering, though, the non-accidental generalization is the one given in (18) and repeated below.

(23) Whenever the king is away, the lights are out or the flag is down.

This time it is the proposition that the king is away and the lights are out that is the confirming proposition for the salient non-accidental generalization. Hence, since this proposition is logically implied by Possibility (ii), and since the CPC requires that we privilege premise sets that logically imply the relevant confirming propositions, the CPC predicts that Possibility (ii) will be chosen and that the counterfactual in (24) will be judged true.

(24) If the flag were up, the lights would be out.

This prediction is not quite correct, though, since, as Kratzer herself points out, our judgments are uncertain in the case of (24). In other words, while the CPC does account for the very unexpected contrast in truth-value judgments when we have formally different but logically equivalent propositions, we still need to say something about why we don't judge (24) true. Kratzer's suggestion is that, behind the uncertainty in judging (24), lies the fact that the non-accidental generalization in (18) is not as natural (that is, it does not describe the regularity in this case as naturally) as the basic generalization that whenever the lights are on and the flag is up, the king is in the castle. I refer the reader to the discussion of this particular point in Kratzer 2012: 146.

I will go back to the discrepancy between the prediction made by the CPC and the uncertain judgments we get in cases like (24) when discussing my proposal in section 6.

To sum up, Kratzer 2012 suggests that the shift in judgments in the King Ludwig of Bavaria example happens because (i) equivalent yet formally different non-accidental generalizations are salient in the context of utterance, and (ii) which propositions (of those true in the actual world) must be privileged when constructing a Base Set and when assembling the Crucial Set depends on the salient non-accidental generalization.

In what follows I want to push Kratzer's observation even further by showing that our truth-value judgments reveal to us that which propositions we select (or privilege, in Kratzer's terminology) changes even in contexts where the *same* non-accidental generalizations are salient. To see this, consider the following context: Peter, Susan, you and me are in the same team and we are playing a betting game. You like Susan but do not like Peter and you do not miss any opportunity to be mean to him. It is our team's turn to bet and we bet on tails. Peter presses a button in a random coin-tossing device for the team, heads comes up and we lose. You don't like Peter and get really upset with him. Now suppose counterfactually that Susan had pressed the button.

(25) Poor Peter! I don't think you're being fair. If Susan had pressed the button, you would not have said a single bad word to her.

(26) If Susan had pressed the button, the coin would have come up heads.

In the given context, where it is known that you have a tendency to be mean to Peter, people tend to judge the counterfactual in (25) true. Crucially, though, in the same context the same people do not judge (26) true. What explains these judgments? The context has made two non-accidental generalizations salient: (i) that coin-tossing is random and (ii) that whenever you have an opportunity, you are mean to Peter (or something along these lines). In (26), it must be the case that the proposition that heads came up is not privileged so that in the Crucial Set we'll have subsets with the proposition that heads came up and subsets with the propositions that tails came up (and the *would* counterfactual is false). But in (25), the proposition that heads came up must be privileged so as to end up in every member of the Crucial Set (and the *would* counterfactual is true). The point is that these two conditionals are uttered in exactly the same context with exactly the same salient non-accidental generalizations. However, different generalizations seem to be "relevant" for the two counterfactuals: the generalization about the hearer's relation with Peter is relevant for (25), whereas the generalization about the random nature of coin-tossing is relevant for (26). This difference is crucial to get the right truth-conditional judgments.

In other words, Kratzer's King Ludwig of Bavaria examples showed us that which propositions go in the premise set does not depend solely on the facts of the world of evaluation but also on some formal properties of the salient non-accidental generalizations. Now, the examples in (25) and (26) show that which propositions go in the premise set does not solely depend on (i) the facts of the world of evaluations (the facts are the same) and (ii) the semantic and formal properties of the non-accidental generalizations made salient in the context (both counterfactuals are judged as uttered in the same context), but it must depend on some other factor which, together with the previous ones, determines the selection of the relevant premises. In the next section, I will argue that understanding the context-dependence of counterfactuals is the key to figuring out what this other factor is in the mechanism selecting the relevant premises. I will argue that the account I am proposing is better equipped to predict the shift in our truth-value judgments.

2 The context-dependence of counterfactuals

Counterfactual conditionals are known to be context-dependent but the mechanism by which the context helps to assigns the correct truth-conditions to a counterfactual has remained opaque. In Lewis 1973 the connection between the selection of the relevant set of antecedent worlds (those antecedent worlds most similar to the actual world), and the context is explicitly made but Lewis does not provide a recipe for consistently identifying these worlds in all cases. The strategy in Kratzer 1989, based on lumping, is a clear recipe for selecting the right premises (and consequently the right possible worlds) but the context seems to play a more marginal role in accounting for our truth-value judgments. On the other hand, in Kratzer 2012, which is based on lumping supplemented with the notion of a confirming proposition (defined itself in terms of lumping), the role of the context becomes a bit more central in that the claim is that our truth-value judgments change when different (in either content or form) non-accidental generalizations are made salient in the context (since, as you will recall, different non-accidental generalizations will be confirmed by different propositions and confirming propositions are privileged by the CPC). However, the two examples we considered at the end of the last section ((25) and (26)) illustrated the need to say something more specific about the mechanism by which some non-accidental generalization get to play a role in the selections of the premises while others don't.

The general idea in the background is that a crucial element driving the selection of the relevant premises (and, ultimately, the set of relevant possible worlds) is the need to avoid trivial moves in the discourse, where a move is to be understood in the sense of Roberts 2012. Here are the main points of my proposal.

- An utterance of *if ϕ, would ψ* is required to be a relevant answer to what I will call a *conditional question under discussion* (CQUD) of the form *if ϕ, Q?*, where ψ is a possible answer to Q. ψ is a possible answer to Q if ψ is a member of the Q-alternative set (Q-alt) of the question Q, that is, the set of all the possible answers to Q. Furthermore, Q and ψ form a question-answer pair if they are interpreted relative to the same context, i.e., the relevant ϕ-worlds.

- This requirement is met if the conditional question (CQ) to which the utterance of *if ϕ, would ψ* is a relevant answer, is indeed under discussion. A CQ qualifies as being under discussion if CQ stands in a certain relation with the discourse: it is a subquestion in what I will call a $\mathcal{Q}$-tree, i.e., a question-tree similar to what Büring 2003 calls a *discourse tree*. This will be the core of our *$\mathcal{Q}$-Tree Constraint on counterfactuals* ($\mathcal{Q}$TC).

- The question Q must be a non-trivial question, and by this we mean a question whose answer is not already entailed by the "temporary context", the set of relevant ϕ-worlds to which ϕ is added and in which Q is raised. If any of the propositions true in this temporary context entail any of the answers to the modally subordinated question Q, then the question is trivial. If the question is trivial and ψ is a relevant answer to the question, then the counterfactual *if ϕ, would ψ* will be either vacuously true or vacuously false. Either way, it will be uninformative. (If ψ is not a relevant answer, then it will be vacuously false.) So, when constructing the correct premise set, a non-triviality constraint will ensure that the CQUD and the counterfactual answering it are not trivial by removing all those propositions true in the actual worlds (premises) that entail any answer to the conditional question. This is what I will call the *Non-Triviality constraint*.

Going back to our discussion of Kratzer's (2012) proposal, one key element of the account developed below is that what makes $\mathcal{Q}$-trees salient in the context are *non-accidental generalizations*, and this is because a non-accidental generalization presuppose those $\mathcal{Q}$-trees which it partially answers.

In sections 2.1 and 2.2 I will say a few words on the notions of a question under discussion (QUD) and a $\mathcal{Q}$-tree, and the notion of a conditional question (CQ), starting with the latter. I will go back to $\mathcal{Q}$-trees and non-accidental generalizations in section 3 when I will lay down the details of my proposal.

2.1 Conditional questions

I loosely follow the analysis of conditional questions in Isaacs & Rawlins 2008. For reasons of space, the following discussion will be informal and brief and will only be concerned with polar questions like (27).

(27) If Alfonso comes to the party, will Joanna leave?

According to Isaacs & Rawlins 2008, the question *will Joanna leave* is modally subordinated to the supposition expressed by the *if*-clause. The framework that these authors adopt is a variant of context change semantics, whereby the meaning of a sentence is its context change potential (cf. Heim 1992). The interpretation of a conditional question is done in two steps. The first step consists of creating a temporary copy of the context and update it with the antecedent. The second step consists of interpreting the question relative to this temporary context: this is meant to capture the intuition that the issue of whether Joanna will leave is only raised relative to the temporary context updated with the antecedent proposition that Alfonso will come to the party. Isaacs and Rawlins assume the analysis in Groenendijk 1999: to interpret a polar question means to update the context inquisitively, that is, to partition the context set into only two cells.[2] This operation creates an *inquisitive context*, that is, a context with more that one cell.

What is important for our purposes is the following. First, the question part in a conditional question is modally subordinated, that is, it is interpreted not relative to the main context but relative to a temporary context which consists of the main context updated with the proposition expressed by

2 For *wh*-questions the partitions will have to be more complex.

the antecedent *if*-clause.[3] Second, the answer to the conditional question is an answer to the modally subordinated question. Hence, the answer will be interpreted as eliminating one of the two cells in which the temporary context was partitioned by the question. In other words, both the question part in a conditional question and its answer are modally subordinated to the same temporary context. Applied to counterfactual conditionals, the temporary context (that is, the set of relevant ϕ-worlds) cannot be just the main context merely updated with ϕ. This temporary context will have to undergo some revisions at least so as to accommodate the counterfactual antecedent (see Heim 1992, Ippolito 2006 for discussion of this issue within the context change semantics framework). More generally, this is is the problem of similarity, that is, the problem of selecting the set of antecedent worlds maximally similar to the actual world. The goal of this paper is precisely to propose a mechanism for selecting the relevant set of antecedent worlds. In what follows, we will see that, in interpreting the counterfactual *if* ϕ*, would* ψ, it is the relevant set of antecedent worlds that will constitute the temporary context with respect to which a salient question is raised, a question to which ψ is understood to be the answer.

2.2 A question under discussion

Following Roberts 2012, let us assume that a discourse is a structure of questions and answers. The QUD-stack is the set of QUDs at a given point in the conversation. At each point in discourse, the question at the top of the stack is the (immediate) QUD. Once a question is raised and accepted, then the participants in the conversation are committed to answering it (if it is answerable). A discourse is structured coherently if it obeys the principle of Relevance, which informally requires that a given assertion select from the Q-alternative set of the QUD, where the Q-alternative set of a question is the set of possible answers denoted by the question, as in Hamblin 1973. Following Roberts 2012, we can define Relevance a bit more formally as shown in (28).

(28) A move m is relevant to the QUD q iff m either introduces a (partial) answer to q (m is an assertion) or is part of a strategy to answer q (m

3 I use the term *modal subordination* in the sense of Roberts 1989, Roberts 1996, and subsequent work.

is a question), where a strategy to answer *q* consists of answering all those subquestions whose answers constitute partial answers to *q*.

It follows that in a felicitous discourse, each move will be relevant to the current QUD.

Roberts 2012 applies these ideas about information structure to the phenomenon of association with focus, and in particular to the analysis of the focus-sensitive particle *only*. She agrees with Rooth 1985 that association with focus is the result of how prosodic focus affects the restriction of the domain of the relevant operator but argues that this can be explained in terms of information structure and the notions introduced above. The proposal in this paper can be viewed as arguing that something very similar to what Roberts has proposed for focus operators can be used to account for our intuitions about the truth and falsehood of counterfactual conditionals: in order to successfully restrict the domain of the modal operator we need to identify the QUD.

Another way of representing the structure of a discourse comes from Büring 2003 where a discourse is represented as a tree, i.e., a hierarchical structure of questions and answers, as shown in (29). I will call diagrams like (29) *Q-trees* and I will refer to the 'discourse' node at the top as the *root node*.

(29)

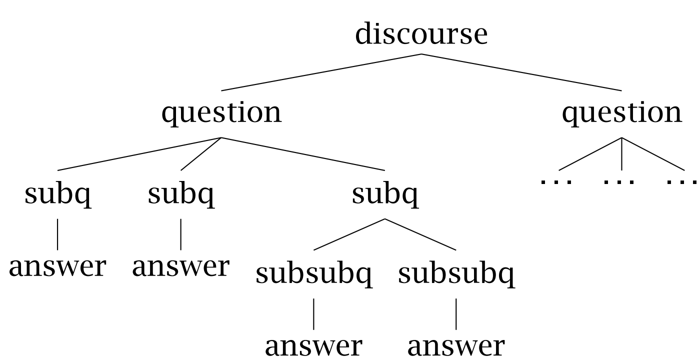

What will be important in our account is the idea that questions can be arranged hierarchically and that these hierarchies are part of the discourse.

3 Back to counterfactuals

A counterfactual of the form *if* ϕ*, would* ψ uttered in a given context is understood as a *conditional answer* to what I will call the *conditional question under discussion* (CQUD). The discourse in which an utterance of *if* ϕ*, would*

ψ is made is felicitous if it obeys Relevance, as defined above. Now, this is the case if the counterfactual is a relevant conditional answer to the CQUD. This means that the answer (ψ) must eliminate one of the cells in which the temporary context was partitioned by the modally subordinated question. For this to be possible, two requirements must be satisfied. First, it must be the case that both the modally subordinated answer ψ and the modally subordinated question are interpreted relative to the same set of worlds, that is, a temporary context revised and updated with ϕ. This is the case if the CQUD to which *if* ϕ*, would* ψ is a relevant answer is of the form *if* ϕ*, would* Q*?* Second, it must also be the case that the modally subordinated answer ψ selects from the Q-alternative set of the question [Q?].

Before proceeding with the arguments, let me spell out those assumptions about the semantics for counterfactuals that will be relevant in constructing my proposal. For ease of exposition, I will combine elements from Lewis 1973, Kratzer 1991, and von Fintel 2001.

A counterfactual of the form *if* ϕ*, would* ψ is true in the actual world w_c just in case ψ is true in all ϕ-worlds most similar to the actual world w_c. This is schematically given in (30).

(30) $[\![\text{if } \phi,\ \text{would } \psi]\!]^c = 1$ in w_c iff $\forall w' \in sim_{\leq_{A_{w_c}}}([\![\phi]\!]^c) : [\![\psi]\!]^c(w') = 1$

The crucial notion here is that of comparative similarity between worlds, $w' \leq_A w''$, which is defined as follows.

(31) For all $w, w' \in W$, for any $A \subseteq \wp(W)$:
$w \leq_A w'$ iff $\{p : p \in A \text{ and } w' \in p\} \subseteq \{p : p \in A \text{ and } w \in p\}$

Comparative similarity is defined relative to a set of propositions A: for any two worlds w and w', w is ranked as high as w' just in case the number of propositions in A true in w is at least as high as the number of propositions in A true in w'. Since in counterfactuals the ordering is given by a relation of comparative similarity to the actual world, let A_{w_c} be a set of propositions which fully describe the actual world w_c. The similarity function $sim_{\leq_{A_{w_c}}}$ applies to the antecedent proposition and returns the set of antecedent worlds which are at least as similar to w_c as any other (accessible) antecedent-world.

(32) $sim_{\leq_{A_{w_c}}}(p) = \{w' : p(w') = 1 \;\&\; \forall w'' : p(w'') = 1 \rightarrow w' \leq_{A_{w_c}} w''\}$

We saw above that $\leq_{A_{w_c}}$ needs to be constrained. In what follows I will propose an algorithm which systematically constrains the ordering source by constraining A_{w_c}.

As mentioned above, the crucial idea is that an utterance of a counterfactual conditional *if* ϕ*, would* ψ in a context c is interpreted relative to the CQUD in c at the time the utterance is made. I propose that, when evaluating a counterfactual of the form *if* ϕ*, would* ψ, constraining the similarity ordering $\leq_{A_{w_c}}$ is not only guided by the need to avoid inconsistencies, but also by the need to avoid trivial moves (both question-moves and answer-moves) in the conversational context. The proposal is summarized in (33).

(33) **Constraining the similarity ordering** $\leq$ (informal):
When evaluating a counterfactual *if* ϕ*, would* ψ in a context c, where *if* ϕ*, would Q?* is the current CQUD, relative to the similarity ordering $\leq_{A_{w_c}}$:
Consistency: for all propositions p, $p \in A_{w_c}$, if either $p \cap \phi = \emptyset$ or $\neg\phi \subseteq p$, then remove p from A_{w_c}.
Non-Triviality: for all propositions p, $p \in A_{w_c}$, if $\exists r \in$ Q-set such that either $p \subseteq r$ or $r \subseteq p$, then remove p from A_{w_c}.

When selecting the set of ϕ-worlds maximally similar to w_c to be quantified over by the necessity modal, we will keep all propositions true in w_c except (i) those propositions p that are inconsistent with the counterfactual antecedent ϕ or are entailed by $\neg\phi$ and (ii) those propositions p that entail a member of the Q-alternative set of the question under discussion (that is, those propositions that entail an answer to the question under discussion), or are entailed by it. In this paper, we will focus primarily on Non-Triviality. I will say here a few words on the way I formulated Consistency in (33). Requiring that we remove only the propositions that are inconsistent with the antecedent is not enough. Suppose that someone must publish at least two books to be nominated for a prestigious award, and that Mary published three books and was nominated for the award. In this case, we judge (34) false.

(34) If Mary hadn't published three books, she would (still) have been nominated for the award.

If in these cases we only removed propositions inconsistent with the antecedent, we would only remove the proposition that Mary published three books, but not the proposition that she published two books, thus incorrectly predicting that (34) should be true. However, (33) requires that we remove not only the proposition that Mary published three books but also any proposition entailed by it, i.e., that Mary published two books and that Mary published one book. Once we remove the proposition that Mary published two books and the proposition that Mary published one book, then we correctly predict that (34) should be false.

We can put the constraints on similarity given above in their final and more formal form as shown in (35).

(35) **Constraining the similarity ordering $\leq$:**
When evaluating a counterfactual *if ϕ, would ψ* in a context c, where *if ϕ, would Q?* is the current CQUD, relative to the similarity ordering $\leq_{A_{w_c}}$:

Step 1: Revise A_{w_c}
$A' = ntr_Q(con_\phi(A_{w_c}))$
where:

(i) for every $X \subseteq \wp(W)$ and $\phi \in \wp(W)$:
$$con_\phi(X) = \{p \in X : p \cap \phi \neq \emptyset \text{ and } \neg\phi \not\subset p\}$$

(ii) for every $X \subseteq \wp(W)$ and question Q:
$$ntr_Q(X) = \{p \in X : \neg\exists r \in Q\text{-alt, either } p \subseteq r \text{ or } r \subseteq p\}$$

Step 2: Define $sim_\leq$ relative to the revised set A':
$$sim_{\leq_{A'}}(\phi) = \{w' : \phi(w') = 1 \ \&\ \forall w'' : \phi(w'') = 1 \rightarrow w' \leq_{A'} w''\}$$

In (35), the function *con* captures the Consistency constraint: $con_\phi(A_{w_c})$ is going to deliver a set of propositions in A_{w_c} consistent with the antecedent. The function *ntr* is designed to capture the basic idea behind the Non-Triviality constraint in (33): *ntr* will constrain $con_\phi(A_{w_c})$ by ruling out those propositions p in $con_\phi(A_{w_c})$ entailing, and being entailed by, a possible answer to the question under discussion. I will also that the the revised set A' is closed under logical consequence: for any proposition p and q such that p is in A' and $p \subseteq q$, q is in A' too.[4]

4 Take the counterfactual *If Mary were not in Toronto, Sue would be sad* and suppose that Mary and Sue are in Toronto. When removing the proposition that Mary is in Toronto, Consistency is also going to remove the proposition that someone is in Toronto. But because we have the proposition that Sue is in Toronto, closure under logical consequence will ensure that the proposition that someone is in Toronto is in the revised set as well.

Unlike Kratzer 2012 I will not assume that the premise set is subject to a non-redundancy constraint. Instead, I will assume Redundancy as defined in (22) above: if p is in A', A' will contain p's logical consequences as well. We will see in section 4.1 why this is important.

Now, for an utterance of *if* ϕ*, would* ψ to be felicitous, the conditional question with respect to which the counterfactual is evaluated must be under discussion, which we are assuming means that it must be a subquestion of a $\mathcal{Q}$-tree, a family of questions arranged hierarchically, salient in the context of utterance. We will formulate this requirement as shown in the $\mathcal{Q}$-Tree Constraint for counterfactuals below.

(36) *The* $\mathcal{Q}$*-Tree Constraint on counterfactuals* ($\mathcal{Q}$TC):
A conditional question CQ which an utterance of the counterfactual *if* ϕ*, would* ψ is a relevant answer to in a context c qualifies as being under discussion in c only if there is at least one $\mathcal{Q}$-tree $\mathcal{Q}'$ salient in c such that CQ is a subquestion of $\mathcal{Q}'$, where for every question Q in $\mathcal{Q}'$, Q is a subquestion of $\mathcal{Q}'$ just in case Q is a question in $\mathcal{Q}'$ and Q is not the root node.

According to (36), a conditional question is under discussion if it is a subquestion of at least one $\mathcal{Q}$-tree salient in the context of utterance. Any question in the tree, except the root node, counts as a subquestion. In some cases different $\mathcal{Q}$-trees can be constructed to represent the same discourse topic, as we will see in sections 4 and 5. For now, what is important to notice is that, even though there might be more than one $\mathcal{Q}$-tree salient in c, the $\mathcal{Q}$TC is satisfied as long as there is *at least* one Q-tree of which the conditional question that the counterfactual is a relevant answer to is a subquestion.

My proposal is that non-accidental generalizations of the *whenever/if* kind presuppose (at least) a $\mathcal{Q}$-tree which they partially answer: hence, whenever a non-accidental generalization is salient in c, a $\mathcal{Q}$-tree is going to be accommodated in c. The intuition behind the proposal that a generalization of the form *whenever p, q* presupposes a $\mathcal{Q}$-tree is that knowing that whenever p is true, q is true provides a partial answer to the larger question about the circumstances under which q is true (and those under which q isn't). For example, the non-accidental generalization that whenever the weather is bad, Jones wears his hat, partially answers the larger question about Jones's hat wearing habits. The salient $\mathcal{Q}$-tree is shown in (37).[5]

(37) when does Jones wear his hat?

- does Jones wears his hat when the weather is bad?
  - 'when the weather is bad Jones wears is hat'
- does Jones wear his hat when the weather is fine?
  - 'when the weather is fine Jones wears his hat at random'

Not just any $\mathcal{Q}$-tree can be presupposed by a given non-accidental generalization. In particular, a $\mathcal{Q}$-tree $\mathcal{Q}'$ qualifies as being presupposed by a non-accidental generalization g' just in case (i) the subquestion that g' answers is one of the terminal subquestions of $\mathcal{Q}'$ and (ii) every subquestion in $\mathcal{Q}$ corresponds to one of the restrictors of the *whenever* operator. The $\mathcal{Q}$-tree in (37) meets these requirements since the non-accidental generalization that whenever the weather is bad, Jones wears his hat answers one of the tree's terminal subquestions and, since there is only one restrictor (*that the weather is bad*), there is only one level of subquestions in the tree. The general schema then is that, given a non-accidental generalization g' of the form *whenever* p_1 *and* p_2, ... *and* p_n, q, a $\mathcal{Q}$-tree $\mathcal{Q}'$ is presupposed by g' just in case there are n levels of subquestions in the $\mathcal{Q}'$ (excluding the root node) and g' answers one of the terminal subquestions in $\mathcal{Q}'$.

When a non-accidental generalization is made salient in the discourse, the $\mathcal{Q}$-trees which the generalization partially answers are accommodated in the discourse. Recall that in Roberts's view of the discourse, the discourse is a structure of questions and answers and each move in the conversation is a step toward answering the QUD. Thus, the move of making a non-accidental generalization G salient in the discourse is understood as a move toward answering a QUD, which — if not already in the discourse — will have to be accommodated. This proposal is strongly reminiscent of Büring's theory of *Contrastive Topic* (CT). Consider the example in (38) (originally from Jackendoff 1972) where CT marks the contrastive topic and F marks the focused phrase.

(38) Q: What about Fred? What did he eat?

5 I am assuming here that *when* and *if* are interchangeable in the questions I will be considering.

A: FRED$_{\text{CT}}$ ate the BEANS$_{\text{F}}$.

The CT-value of A's utterance in (38) is a set of questions as shown in (39).

(39)

who ate what?

what did Fred eat | what did Mary eat?

FRED$_{\text{CT}}$ ate the BEANS$_{\text{F}}$ | MARY$_{\text{CT}}$ ate . . .

Büring's proposal is that for a question-answer sequence Q-A to be well-formed, there must be one question-tree containing Q-A as a sub-tree. Since this is the case for the Q-A sequence in (38), which corresponds to the left branch of (39), that sequence is well-formed. Büring's contrastive topic indicates that the sequence Q-A is part of a larger discourse and the implication that, whereas Fred ate the beans, other people ate different things, is a conversational implicature.

What I would like to suggest here is that non-accidental generalizations of the *whenever/if* type have a (default) CT-F structure, in particular one where the antecedent (or one of its parts) is the CT and the F is in the consequent clause, as shown in (40).

(40) Whenever the weather is [bad]$_{\text{CT}}$, [Jones wears his hat]$_{\text{F}}$.

Given what we said above, then, the CT-value of (40) will be the $\mathcal{Q}$-tree in (37), repeated below, and just like in Jackendoff's example, the implication that when the weather is fine, Jones does not wear his hat is a conversational implicature.

(41)

when does Jones wear his hat?

does Jones wears his hat when the weather is bad? | does Jones wear his hat when the weather is fine?

when the weather is [bad]$_{\text{CT}}$ [Jones wears is hat]$_{\text{F}}$ | when the weather is fine$_{\text{CT}}$ [Jones wears his hat at random]$_{\text{F}}$

To sum up, a salient non-accidental generalization of the *whenever/if* kind presupposes a $\mathcal{Q}$-tree such that the salient non-accidental generalization answers one of this $\mathcal{Q}$-tree's subquestions. This sheds some light on the observation made by many that non-accidental generalizations are crucial in the interpretation of counterfactual conditionals (cf. Pollock 1976, Lewis 1979, Kratzer 1989, Kratzer 2012, just to cite a few). As we will see in detail in the next sections, our proposal is that non-accidental generalizations are central because they presuppose a $\mathcal{Q}$-tree which must contain as one of its subquestions the conditional question that the counterfactual is an answer to.

In what follows, I will return to the examples that we considered above, and I will show how the present proposal provides a natural way of accounting for these and other cases that we will introduce later. I will start with Jones and the weather's examples, then move on to the coin-tossing examples, and finally to the King Ludwig of Bavaria's example.

4 Jones and the weather

Recall Tichý's and Veltman's examples discussed in section 1, repeated in (42) and (43).

(42) *Veltman's example*:
Suppose that Jones always wears his hat if the weather is bad. If the weather is fine, he wears his hat at random. Today the weather is bad and Jones is wearing his hat. Suppose counterfactually that the weather had been fine.
If the weather had been fine, Jones would be wearing his hat.

(43) *Veltman's version of Veltman's example*:
Suppose that Jones tosses a coin every morning before he checks the weather. If heads comes up and the weather is fine, Jones wears his hat. Jones always wears his hat if the weather is bad. Today heads came up, the weather is bad and Jones is wearing his hat. Suppose counterfactually that the weather had been fine.
If the weather had been fine, Jones would be wearing his hat.

The observation is that we judge (42) false and (43) true. This means that, in selecting the set of worlds most similar to w_c in which the weather is fine, we do not keep the proposition that Jones is wearing his hat, but we keep the proposition that heads came up and thus require all antecedent worlds to be worlds in which heads came up. In other words, we treat the proposition that Jones is wearing his hat and the proposition that heads came up differently, despite the fact that they are both consistent with the counterfactual antecedent. Why?

Let us start with the counterfactual in (42). Assuming relevance, the conditional question to which (42) is an answer should be (44).

(44) If the weather had been fine, would Jones be wearing his hat?

Now, according to the $\mathcal{Q}$TC, repeated in (45), (44) is only under discussion if it is part of a $\mathcal{Q}$-tree in the context of utterance.

(45) *The $\mathcal{Q}$-Tree Constraint on counterfactuals* ($\mathcal{Q}$TC):
A conditional question CQ which an utterance of the counterfactual *if* ϕ, *would* ψ is a relevant answer to in a context c qualifies as being under discussion in c only if there is at least one $\mathcal{Q}$-tree $\mathcal{Q}'$ salient in c such that CQ is a subquestion of $\mathcal{Q}'$, where for every question Q in $\mathcal{Q}'$, Q is a subquestion of $\mathcal{Q}'$ just in case Q is a question in $\mathcal{Q}'$ and Q is not the root node.

Non-accidental generalizations presuppose $\mathcal{Q}$-trees and, therefore, whenever a non-accidental generalization is salient in c, a $\mathcal{Q}$-tree is accommodated in c. The $\mathcal{Q}$-tree for the non-accidental generalization in (42) that whenever the weather is bad, Jones wears his hat, is shown in (46).

(46) when does Jones wear his hat?

does Jones wears his hat when the weather is bad? | **does Jones wear his hat when the weather is fine?**

Now, $\mathcal{Q}$TC is satisfied since (44) is one of the subquestions in the $\mathcal{Q}$-tree presupposed by the non-accidental generalization about Jones's hat wearing habits. Note that here and in the following discussion, I will ignore the non-indicative morphology on the tense in the antecedent and on the modal *would* in the consequent in (44), since I assume that this morphology is there for independent reasons and that its contribution is not relevant to the questions we aim to answer in this paper.

The Q-alternative set for the modally subordinated question *would Jones be wearing his hat?* is {that Jones is wearing his hat; that Jones is not wearing his hat}. The schematic truth-conditions for (42) will be the following.

(47) When evaluating the counterfactual conditional **If the weather had been fine, Jones would be wearing his hat** in a context c where *would Jones be wearing his hat if the weather had been fine?* is the current CQUD, relative to the similarity ordering $\leq$ and a set of propositions A_{w_c} describing the facts of the actual world:
$[\![$**[If the weather had been fine]**$_\phi$, **[Jones would be wearing his hat]**$_\psi]\!]^c = 1$ in w_c iff $\forall w' \in sim_{\leq_{ntr_Q(con_\phi(A_{w_C}))}}(\lambda w.$ the weather is fine in $w)$, Jones is wearing his hat in w'.

In the scenario for the counterfactual in (42), the relevant propositions true in the actual world are given in (48).

(48)
a. that the weather is bad
b. that Jones is wearing his hat
c. that whenever the weather is bad, Jones wears his hat

Now, Consistency requires that we remove (48a) because it is inconsistent with the antecedent proposition (that the weather is fine). Non-Triviality, on the other hand, rules out the proposition in (48b) since it entails one member of the Q-alternative set. Nothing rules out (48c), so we keep it. Since we remove (48b), the counterfactual in (42) is false. This is because, as a result of removing (48b), the set of antecedent-worlds maximally similar to the actual world (i.e., $sim_{\leq_{ntr_Q(con_\phi(A_{w_C}))}}(\lambda w.$ the weather is fine in $w)$, in the truth-conditions above) is going to include both worlds where Jones is wearing his hat and worlds where Jones is not wearing his hat.

Let us turn to (43). Just like in the previous case, the conditional question to which (43) is a relevant answer must also be (44). Does this question satisfy $\mathcal{Q}$TC in this case? The non-accidental generalization we are assuming here is more complex: if the weather is fine and heads comes up, Jones wears his hat. There are two restrictors in the *whenever*-clause (*that the weather is fine*; *that heads comes up*) and, therefore, two levels of subquestions are required to occur in the presupposed $\mathcal{Q}$-tree, one introducing each of the restrictors. Because there are two levels of subquestions, we can construct more than one $\mathcal{Q}$-tree corresponding to the non-accidental generalization, as shown in (49) and (50).

(49)

- when does J wear his hat?
  - **does J wear his hat when the weather is fine?**
    - does J wear his hat when the weather is fine and heads comes up?
    - does J wear his hat when the weather is fine and tails comes up?
  - does J wear his hat when the weather is bad?
    - …
    - …

(50)

- when does J wear his hat?
  - does J wear his hat when heads comes up?
    - does J wear his hat when heads comes up and the weather is fine?
    - does J wear his hat when tails comes up and the weather is bad?
  - does J wear his hat when tails comes up?
    - does J wear his hat when tails comes up and the weather is fine?
    - does J wear his hat when tails comes up and the weather is bad?

Even though there are two possible $\mathcal{Q}$-trees that could represent the content of the root question *when does Jones wear his hat?* which the non-accidental generalization presupposes, and even though (44) (i.e., the question to which the counterfactual we need to evaluate is assumed to be a relevant answer) is not a subquestion of (50), the $\mathcal{Q}$TC is satisfied because there is at least on $\mathcal{Q}$-tree, i.e., (49), such that (44) is one of its subquestions. Similarity can now be constrained according to (35).[6] The relevant propositions this time are shown in (51).

(51)
a. that the weather is bad
b. that Jones is wearing his hat
c. that heads came up
d. that whenever the weather is bad, Jones wears his hat
e. that whenever the weather is fine and heads comes up, Jones wears his hat

6 If the counterfactual we are evaluating were *if tails had come up, Jones would be wearing his hat*, the CQ *if tails had come up, would Jones be wearing his hat?* (which the counterfactual is a relevant answer to) would qualify as being under discussing because of the $\mathcal{Q}$-tree in (50). This counterfactual is correctly predicted to be true, because the proposition that the weather is bad is removed by neither Consistency nor Non-Triviality.

Since the CQUD for (43) is also (44), the Q-alternative set is {that Jones is wearing his hat; that Jones is not wearing his hat} in this case too. Now, of the propositions listed above, (51a) is ruled out by Consistency, and (51b) is ruled out by Non-Triviality. Crucially, though, the proposition in (51c) is neither ruled by Consistency nor is it ruled out by Non-Triviality. Therefore, we keep it. As a result, all worlds in which the weather is fine and that are otherwise maximally similar to the actual world (again, $sim_{\leq_{ntr_Q(con_\phi(A_{w_C}))}}(\lambda w.$ the weather is fine in w)), are going to be worlds where heads came up. Hence, the truth of the counterfactual in (43).

4.1 Redundancy and Strength

In this section I will argue that the selection of the relevant premise set is subject to two additional constraint: the *Redundancy Constraint* and the *Strength Constraint*. Strength requires that, among the propositions true in the actual world, we select those that are stronger where, given two propositions p and q, p is stronger than q just in case p entails q. Redundancy, on the other hand, will ensure that weaker propositions are going to be selected as well.

In order to account for our truth-conditional judgment in (43) I proposed above that the premise set we start with includes the propositions listed in (51). Now, suppose that instead of those propositions we have the following set.

(52) a. that the weather is bad and heads came up
b. that Jones is wearing his hat
c. that whenever the weather is bad, Jones wears his hat
d. that whenever the weather is fine and heads comes up, Jones wears his hat

The non-accidental generalizations in (52) are the same as before but what is different is that in this set the proposition that heads came up only appears as part of the conjunction in (52a) and not as an independent member of the set. This set seems to characterize the actual world as well as (51). However, if (52) is the premise set we choose, Tichý's counterfactual is incorrectly predicted to be judged true. This is because in order to remove the proposition that the weather is bad (inconsistent with the counterfactual antecedent that the weather is fine), we would have to remove (52a) and, with it, we would be losing the proposition that head came up.

What prevents (52) from being selected as the relevant set is that, *contra* Kratzer 2012, there is a preference for redundancy in constructing the right set of premises. Here is the definition we introduced before from Kratzer 2012.

(53) *Redundancy*
A set of propositions is redundant if it contains propositions p and q such that $p \neq q$ and $p \cap W \subseteq q \cap W$.

The point is that, when selecting a premise set, we should privilege those sets where propositions true in the world of evaluation are individual members of such sets (as well as maybe being lumped together).

(54) *Redundancy Constraint*
When describing a world, privilege those premise sets A such that for any proposition p and q, such that $p \in A$ and $p \cap W \subseteq q \cap W$, $q \in A$.

It follows that in Veltman's version of Tichý's example, we should either select (51) where the only occurrence of the proposition that the weather is bad is independent of the proposition that heads came up in that set or the set in (55).

(55) a. that the weather is bad and heads came up
b. that Jones is wearing his hat
c. that the weather is bad
d. that heads came up
e. that whenever the weather is bad, Jones wears his hat
f. that whenever the weather is fine and heads comes up, Jones wears his hat

Now, the fact that (55a) as well as (55c) will be removed when adding the counterfactual proposition that the weather is fine will not cause the counterfactual in (43) to be false because the proposition that heads came up in (55d) will not be removed.

The selection of the right set of premises is also constrained by *Strength* which requires that we privilege those premise sets that contain stronger propositions, where strength is based on the notion of entailment.

(56) *Strength Constraint*
When describing a world w, for any proposition p and q such that

$p \in w$ and $q \in w$ and $p \subseteq q$, privilege those premise sets A such that $p \in A$.

Here is an illustration of the role played by *Strength*. Suppose that we are considering the following set when evaluating Tichý's original example.

(57) a. that the weather is bad or Jones is wearing his hat
b. that whenever the weather is bad, Jones wears his hat

The problem with this set is that, once we add the counterfactual proposition that the weather is fine to this set, the combination of this proposition and (57a) entails that Jones is wearing is hat, which would then incorrectly predict that we should judge Tichý's counterfactual true. However, *Strength* rules out (57) and forces us to have the set in (58): given that (58a) and (58b) are both true and (58b) entails (58a), then *Strength* requires that (58b) be also in the premise set.

(58) a. that the weather is bad or Jones is wearing his hat
b. that the weather is bad
c. that Jones is wearing his hat
d. that whenever the weather is bad, Jones wears his hat

The proposition that the weather is bad in (58b) is excluded by Consistency, and the proposition that Jones is wearing his hat in (58c) is excluded by Non-Triviality. As for (58a), it is excluded by both Consistency and Non-Triviality since it is both entailed by (58b) and (58c).

Successfully ruling out the disjunctive proposition in (58a) depended on being able to rule out not only all those propositions inconsistent with the antecedent or entailing an answer to CQUD ((58b) and (58c)) but, crucially, all the propositions entailed by them as well. According to our proposal, when you remove a proposition (either because it's inconsistent with the antecedent or because it already answers the question under discussion), you remove everything that it entails as well. The need to do that is not just limited to cases like (58) but is more general. Consider Tichý's original example again and the following set of propositions.

(59) a. that the weather is bad
b. that Jones is wearing his hat
c. that Jones is wearing something
d. that whenever the weather is bad, Jones wears his hat.

Unlike the previous sets of propositions we have considered in the Tichý cases, (59) explicitly lists the proposition that Jones is wearing something, which is true in the actual world since it's entailed by the proposition that Jones is wearing his hat, also true in the actual world. The CQUD is *would Jones be wearing his hat if the weather had been fine?* Consistency requires that we remove the proposition that the weather is bad. What about Non-Triviality? If Non-Triviality only required that we remove propositions entailing any member of the Q-alternative set, then we would remove the proposition that Jones is wearing his hat, but not the proposition that Jones is wearing something. Hence, even though we would correctly predict that the counterfactual *If the weather had been fine, Jones would be wearing his hat* is false, we would also predict that the counterfactual *If the weather had been fine, Jones would be wearing something*: but if this counterfactual is true, it is certainly not true because Jones is wearing his hat in the actual world. However, Non-Triviality requires that we remove (59c) too since this proposition is entailed by (59b) which in turns entails a member of the Q-alternative set (that Jones is wearing his hat).

Having introduced all the components that go into the selection of the right premise set and, consequently, the right set of possible worlds, we can summarize the selection process as follows.

- There are two basic ingredients of this mechanism: (i) an initial set of propositions A_{w_c} describing the actual world and (ii) a salient $\mathcal{Q}$-tree.
  - A_{w_c} must satisfy the *Redundancy* and *Strength* constraints.
  - A $\mathcal{Q}$-tree is salient in a context c if it presupposed by a non-accidental generalization salient in c (where a non-accidental generalization presupposes a $\mathcal{Q}$-tree just in case it answers one of its terminal subquestions and the $\mathcal{Q}$-tree has as many levels of subquestions as there are restrictors in the non-accidental generalization). A non-accidental generalization presupposes all those $\mathcal{Q}$-trees which it partially answers (and which satisfy the constraints above).
- With the concept of a $\mathcal{Q}$-tree, we can define a CQUD. The conditional question that the counterfactual *if* ϕ, *would* ψ is a relevant answer to is *under discussion* just in case this question satisfies the $\mathcal{Q}$-Tree Constraint. According to the $\mathcal{Q}$TC, a CQ is a CQUD just in case it is a subquestion of a salient $\mathcal{Q}$-tree.

- With a set of propositions A_{w_c} (satisfying both Redundancy and Strength) and a CQUD, we can now revise A_{w_c} by applying *Consistency* and *Non-Triviality*. Let's call this revision A'.
  - Consistency removes from A_{w_c} any proposition p that is inconsistent with the antecedent ϕ of the counterfactual, and any propositions q that is entailed by p. Informally, Consistency removes any proposition p that is inconsistent with the antecedent, and all that is true in virtue of p.
  - Non-Triviality removes from A_{w_c} any proposition p that entails a possible answer to the CQUD, and any propositions q that is entailed by p. Informally, Non-Triviality removes any proposition p that entails a possible answer to the CQUD, and all that is true in virtue of p.

  The revised set A' is closed under logical consequence.
- Finally, we define $sim_{\leq_{A'}}(\phi)$ as the set of ϕ-worlds closest to the ideal, which is the revised set A'. The counterfactual *if ϕ, would ψ* is true iff $sim_{\leq_{A'}}(\phi) \subseteq \psi$; false, otherwise.

5 Coin tossing

Suppose that this morning, I bet on tails, I tossed a coin, heads came up, and I lost \$10. Now, suppose counterfactually that I had bet on heads.

(60) If I had bet on heads, I wouldn't have lost \$10.

We judge this counterfactual true. Now, consider Arregui's variant from Arregui 2009. Assume that Peter presses a button in a random coin-tossing device and heads comes up. Now, suppose counterfactually that Susan had pressed the button.

(61) If Susan had pressed the button, the coin would have come up heads.

We judge this counterfactual false. In the remaining part of this section I will show how our proposal about constraining similarity on the basis of the CQUD explains our judgments in these cases.

Let us start with (60). The relevant propositions are given in (62).

(62) a. that I bet on tails

b. that I tossed a coin
c. that heads came up
d. that I lost \$10

The crucial step is to identify the CQUD at the time (60) is uttered. Assume that (60) is an answer to the conditional question under discussion at the time of utterance. The conditional question to which (60) is a relevant answer is (63).

(63) If I had bet on heads, would I have lost \$10?

The question is whether (63) satisfies $\mathcal{Q}$TC. Our world-knowledge entails that there is a non-accidental correlation between betting on something and either winning or losing: that is, you win if the outcome of the coin tossing matches your bet. As we said above, this non-accidental generalization presupposes a topic under discussion, which can be represented as a $\mathcal{Q}$-tree and which the non-accidental generalization partially answers. We can construct more than one $\mathcal{Q}$-tree representing the topic under discussion, but this is not a problem since the $\mathcal{Q}$TC requires that for a conditional question to be under discussion there must be at least *a* $\mathcal{Q}$-tree such that the conditional question is one of its subquestions. One of the $\mathcal{Q}$-trees representing the topic under discussion and presupposed by our non-accidental generalization (i.e., the question- tree that this non-accidental generalization at least partially answers) is shown in (64).[7]

(64)

- when does someone lose in a random coin-tossing?
  - **does someone lose if they bet on heads?**
    - does someone lose if they bet on heads and heads comes up?
    - does someone lose if they bet on heads and tails comes up?
  - does someone lose if they bet on tails?
    - does someone lose if they bet on tails and tails comes up?
    - does someone lose if they bet on tails and heads comes up?

Technically, (63) is not part of this $\mathcal{Q}$-tree: however, (63) is a subquestion of the general question *does someone lose if they bet on heads?*, which is a subquestion of the $\mathcal{Q}$-tree in (64).

7 I take random coin-tossing and pressing a button in a random coin-tossing device to be equivalent here.

(65) *question*: does someone lose if they bet on heads?

a. *sub-question*: do I lose if I bet on heads?
b. *sub-question*: does Susan lose if she bets on heads?
c. *sub-question*: does Peter lose if he bets on heads? . . .

Hence, (63) satisfies $\mathcal{Q}$TC, and can be under discussion and become part of the algorithm to constrain similarity. Consistency will rule out (62a). Because the Q-alternative set for the modally subordinated question will be {that I lost \$10, that I didn't lose \$10}, Non-Triviality will then rule out (62d). Both (62b) and (62c) will stay. Hence, all worlds in $sim_{\leq_{ntr_Q(con_\phi(A_{w_C}))}}(\lambda w.$ I bet on heads in $w)$ will be worlds where I tossed a coin and heads came up and in none of those worlds I lost \$10.

What is different in (61)? The conditional question that the counterfactual is a relevant answer to is given in (66).

(66) If Susan had pressed the button, would the coin have come up heads?

Given our world-knowledge, one salient non-accidental generalization is about the random nature of coin-tossing. Hence, the $\mathcal{Q}$-tree representing the presupposed discourse topic is shown in (67).

(67) what's the outcome in a random coin-tossing?

- **when tossed, does the coin come up heads?**
- when tossed, does the coin come up tails?

The conditional question in (66) counts as a subquestion of this $\mathcal{Q}$-tree, and hence it counts as being under discussion. Since we now have a CQUD, we can run Non-Triviality. The Q-alternative set of the modally subordinated question *what would have come up?* is {that it came up heads; that it came up tails}. The relevant propositions are given in (68).

(68) a. that Peter pressed the button
b. that heads came up

This time Consistency rules out (68a) and Non-Triviality rules out (68b). As a result, it follows that $sim_{\leq_{ntr_Q(con_\phi(A_{w_C}))}}(\{w$: Susan pressed the button in $w\}) \notin \{w'$: the coin came up heads in $w'\}$ and the counterfactual is correctly predicted to be false.

In the present proposal, the CQUD is essential in choosing which propositions must be removed or kept in evaluating a counterfactual. In turn, properties of the context in which the counterfactual is evaluated are crucial in making a CQ under discussion, so that not every CQ will satisfy the $\mathcal{Q}$TC. There are two cases I will consider: the first case is one where the CQ we are checking is "related" to the discourse topic (the root of a $\mathcal{Q}$-tree) but the CQ is not a *subquestion* in any salient $\mathcal{Q}$-tree; the second case is one in which the CQ is simply unrelated to any $\mathcal{Q}$-tree presupposed by a salient non-accidental generalization. The latter case has to do with "odd" counterfactuals such as Kratzer's *If Paula were buying a pound of apples, the Atlantic Ocean might be drying up*, which I will discuss in section 6.1. In this section, we will look at the first kind of cases. Let us assume that Peter pushed the button in a random coin-tossing device and it came up heads. Assume also that I bet on tails and so I lost \$10. Our intuition is that in this context, the counterfactual in (69) is false.

(69) If Susan had pressed the button, I would have lost \$10.

Now, let's *stipulate* that the CQUD is *If Susan had pressed the button, would I have lost \$10?* This seems to be a problem at first for the proposal I am defending because the proposition that heads came up does not entail (and is not entailed by) any members of the Q-alternative set of the question *would I have lost \$10?* (i.e., {I lost \$10; I did not lose \$10}). Hence, that proposition should be kept (together with the proposition that I bet on tails) and the counterfactual in (69) should come out true. How can our proposal explain our judgment here? I will argue that the conditional question that (69) is a relevant answer to, i.e., *If Susan had pressed the button, would I have lost \$10?*, does not satisfy the $\mathcal{Q}$TC and therefore is not under discussion. The non-accidental generalization salient in this context is one about losing and it involves two facts: in coin-tossing, you lose whenever (i) the outcome of the tossing is x and (ii) your bet is y. In other words, you lose whenever the outcome of the coin-tossing does not match your bet. The discourse topic presupposed by this non-accidental generalization is about the circumstances under which you lose in a coin-tossing. There are two possible $\mathcal{Q}$-trees that can be constructed, (70) and (71).

(70) **when does someone lose in a random coin-tossing**?

- does someone lose if they bet on heads?
  - does someone lose if they bet on heads and heads comes up?
  - does someone lose if they bet on heads and tails comes up?
- does someone lose if they bet on tails?
  - does someone lose if they bet on tails and tails comes up?
  - does someone lose if they bet on tails and heads comes up?

(71) **when does someone lose in a random coin-tossing**?

- does someone lose if heads comes up?
  - does someone lose if heads comes up and they bet on heads?
  - does someone lose if heads comes up and they bet on tails?
- does someone lose if tails comes up?
  - does someone lose if tails comes up and they bet on tails?
  - does someone lose if tails comes up and they bet on heads?

The observation is that the conditional question *If Susan had pressed the button, would I have lost \$10?* is not a subquestion of either tree. This brings up again an important point about the relation between non-accidental generalizations and $\mathcal{Q}$-trees: non-accidental generalizations are answers to a terminal question in the $\mathcal{Q}$-tree(s) that they presuppose and the number of levels of subquestions in the $\mathcal{Q}$-tree(s) must be the same as the number of restrictors in the non-accidental generalization. In the current example, the generalization tells us that losing in a coin-tossing happens whenever *someone bets on either heads or tails* and *the outcome does not match the bet*. Facts about the identity of the agent of the coin-tossing are not part of this non-accidental generalization and, for this reason, the question *If Susan had pressed the button, would someone lose?* cannot be in the $\mathcal{Q}$-tree presupposed by such a non-accidental generalization. We will see below (example (74)) that if we manipulate the non-accidental generalization in such a way that the identity of the coin-tosser matters, then this fact will find its way into the corresponding $\mathcal{Q}$-tree and will therefore change our judgments.

Now, because the agent of the coin-tossing is irrelevant to the non-accidental generalization in question, the question *If Susan had pressed the button, would I have lost \$10?* can only relate to the root question, i.e., the discourse topic, *when does someone lose in a coin-tossing?* However, because root questions do not count as subquestions of a $\mathcal{Q}$-tree, the $\mathcal{Q}$TC is violated.

Let us make the assumption that whenever the conditional question that the counterfactual is a relevant answer to is a root question (of possibly multiple *Q*-trees), every question in the *Q*-tree stemming from that root is under discussion. In our particular example, this means that all the four terminal questions are under discussion.

(72) CQUDs:

a. does someone lose when heads comes up and they bet on heads?
b. does someone lose when heads comes up and they bet on tails?
c. does someone lose when tails comes up and they bet on heads?
d. does someone lose when tails comes up and they bet on tails?

Since all these questions are under discussion, the counterfactual in (69) is going to be a relevant answer only if we *accommodate* some material in its restriction, that is, only if we understand the conditional as follows.

(73) If Susan had tossed the coin **<and [heads had come up and I had bet on heads] or [heads had come up and I had bet on tails] or [tails had come up and I had bet on heads] or [tails had come up and I had bet on tails]>**, I would have lost $10.

Given the accommodated disjunctive material, the counterfactual is rescued and correctly predicted to be false.[8]

In other words, in (69), the relevant non-accidental generalization is that you lose whenever the outcome of a coin tossing does not match your bet, presupposes a discourse-topic (losing in coin-tossing) that can be represented by the two *Q*-trees in (70) and (71), which are distinct but stem from the same root question *when does someone lose in a random coin tossing?* The conditional question *if Susan had pressed the button, would I have lost $10?* is closely related to the root question. Thus, the *Q*TC is not satisfied. This would rule out the counterfactual in (69) as infelicitous since it violates Relevance. However, the counterfactual can be rescued if we make it a relevant answer to the questions under discussion, and we do that by accommodating in the antecedent all that is under discussion, as shown in (73). Accommodating this material, makes the counterfactual relevant but false.

If we are right about the central role played by non-accidental generalizations in the evaluation of counterfactuals, a context where different

8 For the interpretation of counterfactuals with disjunctive antecedents, see for example Alonso-Ovalle 2006.

non-accidental generalizations are salient might make a counterfactual like (69) true. The example we introduced earlier in (5), repeated below in (74), shows that this is precisely what happens.

(74) Peter and Susan are taking turns at pressing a button on a completely random coin-tossing device. They both bet each time one presses the button, but (as part of their game) only the one actually pressing the button pays $10 if he or she loses. It is Peter's turn to press the button. Peter bets that the coin will come up heads, Susan bets that it will come up tails. Peter presses the button and heads comes up. Peter wins. Susan had bet on tails but since she wasn't the one pressing the button she does not have to pay $10. Now I say: Susan, you're lucky! If you$_F$ had pressed the button, you would have lost $10.

In this context, we judge the counterfactual that if Susan had pressed the button, you would have lost $10, true. In addition to our usual non-accidental generalization that whenever you toss a coin the outcome is random, one more generalization is salient in this context, i.e., that you lose money only if you are the one pressing the button, as shown in (75).

(75) Whenever your bet does not match the outcome of the coin-tossing *and* you are the one tossing the coin, you lose.

This non-accidental generalizations is based on three facts: (i) your bet; (ii) the outcome of the coin-tossing; (iii) whether you are the agent of the coin-tossing. The $\mathcal{Q}$-tree presupposed by (75) must be such that the non-accidental generalization answers one of its terminal subquestions and there there are three levels of subquestions, one for each restrictor of the *whenever* operator.

The following $\mathcal{Q}$-tree is one of the $\mathcal{Q}$-trees compatible with the discourse topic.

(76) when does someone lose in a random coin-tossing?

- **does someone lose if they press the button?**
  - does someone lose if they press the button and heads comes up?
  - does someone lose if they press the button and tails comes up?
- does someone lose if someone else presses the button?
  - does someone lose if someone else presses the button and tails comes up?
  - does someone lose if someone else presses the button and heads comes up?

Since the CQ *if you had pressed the button, would you have lost $10?* is a subquestion of this $\mathcal{Q}$-tree, the $\mathcal{Q}$TC is satisfied. As a result, Non-Triviality will not require that the proposition that heads came up be removed.

In other words, the problem with (69) in the context we described above is that it violates the $\mathcal{Q}$TC and the example is rescued by accommodating implicit premises in the antecedent. The general point that pairs like (69) and (74) raise is that by changing the relevant non-accidental generalizations we can change the CQUD, which may results in different truth-conditional judgments.

In what follows I will show that assuming that a counterfactual is a relevant answer to a CQUD provides a systematic way of explaining why, in evaluating counterfactuals, we impose different requirements on the premise sets, even when these counterfactuals are uttered in the same context with the same salient non-accidental generalizations. Consider the following scenario. Peter, Susan, you and me are on the same team. You like Susan but not Peter, and you regrettably are incapable of concealing this fact. It is our team's turn to bet and we bet on tails. Peter presses a button in a random coin-tossing device; heads comes up and we lose. You get very upset with Peter. In this context, we are evaluating the following two counterfactuals.

(77) (What a scene!) If Susan had bet on tails and then pressed the button, you would not have been so upset.

(78) If Susan had bet on tails and then pressed the button, the coin would have come up heads.

We judge (77) true, but (78) false. We haven't changed the facts or the non-accidental generalizations, yet the judgments reveal that the proposition that heads came up is removed in evaluating (78) but not in evaluating (77). Recall Kratzer's proposal that, when constructing a Base Set, we privilege confirming propositions for the relevant and salient non-accidental generalizations. She accounts for the change in intuitions in the King Ludwig of Bavaria example by suggesting that whether we judge the counterfactual true or false depends on whether we take (79) or (23) (repeated below as (80)) to be the salient generalization (we will go back to Kratzer's case in section 6).

(79) Whenever the lights are on and the flag is up, the king is in the castle.
Example of confirming proposition: Right now, the lights are on, the flag is up, the king is away.

(80) Whenever the king is away, the lights are out or the flag is down.
Example of confirming proposition: Right now, the king is away, the lights are out, the flag is up.

Truth-conditionally equivalent, yet formally different, the generalizations (79) and (80) have different confirming propositions and this affects which propositions will be kept and which ones will be removed.

The puzzle that (77) and (78) raise is that, while we have not changed the facts or the generalizations salient in the context, the selection of the premises seems to proceed very differently: in order to account for our truth-conditional judgments, we must keep the proposition that heads came up in the premise set in (77) (it is because we keep this proposition that we judge (77) true) but not in (78) (since, if we did, we would judge the counterfactual true). We are not changing the facts or the generalizations; yet, we construct different base sets, to use Kratzer's term. Even though the context has made two non-accidental generalizations salient — (i) whenever a coin is tossed, the outcome is random and (ii) whenever Susan is responsible for her team's losing, you don't get upset — one might try to deny that they are both salient for the two counterfactuals we are considering. In particular, one might be tempted to say that the generalization about the nature of coin-tossing is relevant for (78), whereas the generalization about the speaker's reaction to Susan's behavior is relevant for (77). This is dubious, though: after all, the generalization about the nature of coin-tossing should also be relevant when evaluating (77) since the counterfactual worlds in which Susan bets and then presses the button are still understood to be worlds in which the device Susan is pressing is a random coin-tossing device. Hence, what is missing is a systematic account of how we privilege some of the information available in the context in the evaluation of counterfactuals.

Here is where thinking of a counterfactual conditional as an answer to a CQUD is helpful. When judging an utterance of a counterfactual *if* ϕ*, would* ψ in a context c, we assume that the discourse of which this utterance is part obeys Relevance (as defined above), and therefore require that an utterance of *if* ϕ*, would* ψ is a *relevant* answer to a CQUD in c. As explained in section 3, for this to be possible it must be the case that the conditional question *if* ϕ*, Q?* to which the counterfactual *if* ϕ*, would* ψ is an answer is under discussion. This question-answer pair is felicitous if the modally subordinated answer ψ selects from the Q-alternative set of the question Q? In other words, when we evaluate *if* ϕ*, would* ψ we take ψ to be the relevant answer to the question of whether ψ or any of its alternatives (in the meaning of the question) would be true if it were the case that ϕ. Now, let us go back to our case. Peter, Susan, you and me are in the same team. You like Susan but not Peter, and you regrettably do not normally conceal this fact. It is our team's turn to bet and this time Peter bets for us. Peter bets on tails; he presses a button in a random coin-tossing device; heads comes up and we lose. You get very upset with Peter. In this context, we judge (81) true but (82) false.

(81) (What a scene!) If Susan had pressed the button, you would not have been so upset.

(82) If Susan had pressed the button, the coin would have come up heads.

Take (81): since the discourse is subject to Relevance, (81) will be understood as an answer to the question *If Susan had pressed the button, would you have been so upset?* Does this question satisfy $\mathcal{Q}$TC? The relevant non-accidental generalization is that you get very upset whenever Peter, but not Susan, does something bad. The $\mathcal{Q}$-tree in (83) is one way of representing the family of questions presupposed by this non-accidental generalization. Since the conditional question *if Susan had bet on tails and then pressed the button, would you have been so upset?* is one of its subquestions, it satisfies the $\mathcal{Q}$TC and it counts as being under discussion.

(83) when do you get upset?

- … when S presses the button?
  - … when S presses the button and bets on heads?
    - … and heads comes up?
    - … and tails comes up
  - … when S presses the button and bets on tails?
    - … and heads comes up?
    - … and tails comes up?
- … when someone else presses the button?
  - **… when someone else presses the button and bets on tails?**
    - … and heads comes up?
    - … and tails comes up?
  - … when someone else presses the button and bets on heads?
    - … and heads comes up?
    - … and tails comes up?

Once we have the CQUD, then Non-Triviality does the rest: the proposition that you got so upset will be removed, the proposition that heads came up will not, and the counterfactual comes out true. Notice that the non-accidental generalization that the outcome of a coin tossing is random is also in principle available in the context and the $\mathcal{Q}$-tree it presupposes is given in

(84). However, uttering the counterfactual itself requires that it be an answer to the CQUD. This means that the CQ the counterfactual answers must be a subquestion of a Q-tree made salient by a non-accidental generalization. In other words, it is an utterance of (81) itself that selects (83), rather than (84).

In (82), on the other hand, Relevance requires that the CQ be *If Susan had bet on tails and then pressed the button, what would have come up?* and since there is a Q-tree presupposed in c (cf. (84)) of which this CQ is a subtree, this CQ qualifies as being under discussion.

(84) what's the outcome in a random coin-tossing?

- if tossed, does the coin come up heads?
- if tossed, does the coin come up tails?

Non-Triviality will remove the proposition that heads came up (and not the proposition that you got so upset), and the counterfactual is correctly predicted to be false.

Viewing the utterance of a counterfactual as an answer to a CQUD plays a crucial role in our account of counterfactuals like (81) and (82) because discourse structure is required to be subject to Relevance which, in our case, means that an utterance of a counterfactual is understood to be relevant to a CQUD. The role of Relevance in our case, then, is to select a CQUD between two conditional questions which could equally well be under discussion in the context of utterance. Note that the conditional question to which the counterfactual we are evaluating is a relevant answer might not be under discussion given some things we assume to be true. In this case, we predict that the counterfactual will come out as odd (and either trivially true or trivially false). We will look at cases that fall into this category in section 7.

6 King Ludwig of Bavaria

Recall Kratzer's King Ludwig of Bavaria example given above in (9a). Whenever the lights are on and the flag is up, the king is in the castle. Now, the lights are on, the flag is down, and the king is away. Now, suppose counterfactually that the flag were up. The relevant examples are repeated below.

(85) a. If the flag were up, the king would be in the castle.
b. If the flag were up, the lights would be off.

In this context, we accept (85a) but not (85b). Let us begin with (85a). Assuming that (85a) is a relevant answer to a CQ, the CQ would have to be (86).

(86) If the flag were up, where would the king be?

Does (86) satisfy $\mathcal{Q}$TC, which requires that it be a subquestion of a $\mathcal{Q}$-tree presupposed in the context? The non-accidental generalization that whenever the lights are on and the flag is up the king is in the castle presupposes the $\mathcal{Q}$-tree in (87).

(87)
- when is the king in the castle?
  - is the king in the castle if the lights are on?
    - is the king in the castle if the lights are on and the flag is up?
    - is the king in the castle if the lights are on and the flag is down?
  - **is the king in the castle if the flag is up?**
    - is the king in the castle if the flag is up and the lights are on?
    - is the king in the castle if the flag is up and the lights are off?

Since the $\mathcal{Q}$-tree in (87) is a possible hierarchy of questions presupposed by the non-accidental generalization that whenever the lights are on and the flag is up the king is in the castle, and since the CQ above is indeed part of (87), the $\mathcal{Q}$TC is satisfied and this question qualifies as being a conditional question under discussion.

The Q-alternative set for the modally subordinated question is {that the king is in the castle, that the king is away}. The relevant propositions are given in (88).

(88) a. that the lights are on
b. that the flag is down
c. that the king is away
d. that whenever the flag is up and the lights are on, the king is in the castle

Proposition (88b) is ruled out by Consistency and, crucially, proposition (88c) is ruled out by Non-Triviality since it entails a member of the Q-alternative set of the question under discussion. Both (88a) and (88d) are ruled out neither by Consistency nor by Non-Triviality. We keep them both and as a result all worlds in which the flag is up and that are otherwise maximally similar to the actual world are worlds where the lights are on (as shown below) and where the non-accidental generalization holds.

(89) $sim_{\leq_{ntr_Q(con_\phi(A_{w_c}))}}(\{w : \text{the flag is up in } w\}) \subseteq \{w' : \text{the king is in the castle in } w'\}$

Compare (85a) to (85b). The conditional question that (85b) is a relevant answer to would be *if the flag were up, what would the lights be?* The problem is that, since (85b) is uttered in the same context as (85a) with the same salient non-accidental generalization, that question is not part of any of the Q-trees presupposed in the context of utterance (in particular, the one we drew above). Hence, the question *if the flag were up, what would the lights be?* does not satisfy the QTC and does not qualify as the CQUD. The Q-tree tells us that a conditional question is under discussion in the given context if it is a question about the circumstances in which the king is in the castle. It follows that an utterance of (85b) in the given context violates Relevance.

This accounts captures the intuition that (85b) is an odd thing to say in the given context. (85b) violates Relevance and, as such, it is felt as the wrong move to make in the conversation.

As we expect, if we manipulate the facts of this world, our intuitions about the truth of these conditionals will change too. This is because by changing the facts we change what is under discussion in the context and therefore what the CQUD is. Suppose that the king of Bavaria no longer goes to Leoni Castle as he prefers to spend his days at his other residences. However, someone still takes care of the castle and they play with the flag and the lights always making sure, though, that the lights are never on when the flag is up and that the flag is never up when the lights are on (since the king is always away). In this context, I think our judgments would be reversed: we would judge (85b) true but but not (85a). This is because, given the current assumptions, counterfactually assuming that the flag is up does not raise the issue of where the king is (we assume he is always away), but does raise the issue of the status of the lights. In other words, what is under discussion are the circumstances under which the flag would be up. Notice that in this case, since we also have the non-accidental generalization that the king is never in Leoni Castle, we correctly predict that the counterfactual in (85a) is false.

There is another way of manipulating our intuitions. In section 1, we discussed Kratzer 2012's observation that if we replace the non-accidental generalization (i) that whenever the lights are on and the flag is up, the king is in the castle with the logically equivalent generalization (ii) that whenever the king is away, either the lights are out or the flag is down, we do not get the same truth-conditional judgments.

(90) a. Whenever the lights are on and the flag is up, the king is in the castle.
b. Whenever the king is away, either the lights are off or the flag is down.

In particular, when evaluated in a context that has made salient the latter generalization, the counterfactual (85a) is no longer judged true. In section 6.1, we will review Kratzer's account and we will offer an alternative explanation based on the CQUD proposal developed above. In doing so, we will provide further evidence that whether a conditional question is under discussion or not depends on the $\mathcal{Q}$-tree made salient by the non-accidental generalizations available in the context of utterance.

6.1 Shifty intuitions

Kratzer's account of the judgments we have in the King Ludwig of Bavaria examples was that, when constructing a Base Set, we privilege propositions which confirm our non-accidental generalizations, as repeated below.

(91) *CPC for Base Sets*
When constructing a Base Set, privilege confirming propositions for non-accidental generalizations.

Despite being logically equivalent, the two non-accidental generalizations in (90a) and (90b) above are confirmed by different propositions. Hence, the CPC for Base Sets will instruct us to construct different Base Sets in the two cases. We already noticed above that Kratzer's proposal alone predicts that, uttered in a context where the salient non-accidental generalization is (90b), the counterfactual in (85b) should come out true. This, however, is not what we observe: even in this context, (85b) remains strange and our judgments "insecure". Kratzer's suggestion is that the problem with this variant of the example is that the relevant non-accidental generalization (90b) is a "less natural way of describing the regularity in the King Ludwig case" (Kratzer 2012: 146) and that there is a cognitive bias against it.

Our proposal offers a different way of accounting for Kratzer's data and in particular the "insecure" judgments we have when interpreting (85b) in the context of the generalization (90b). Let us start with our intuition that (85a) would be true when uttered in the context of the non-accidental generalization in (90a). Recall the $\mathcal{Q}$TC.

(92) *The $\mathcal{Q}$-Tree Constraint on counterfactuals* ($\mathcal{Q}$TC):
A conditional question CQ which an utterance of the counterfactual *if ϕ, would ψ* is a relevant answer to in a context c qualifies as being under discussion in c only if there is at least one $\mathcal{Q}$-tree $\mathcal{Q}'$ salient in c such that CQ is a subquestion of $\mathcal{Q}'$, where for every question Q in $\mathcal{Q}'$, Q is a subquestion of $\mathcal{Q}'$ just in case Q is a question in $\mathcal{Q}'$ and Q is not the root node.

As we saw above, a $\mathcal{Q}$-tree is a family of questions arranged in a hierarchical structure representing what is under discussion. At the top of the structure we have the root node and lower nodes are occupied by its subquestions. Exploiting the similarity with Büring's contrastive topic, I suggested that a non-accidental generalization has a CT-F structure and that the $\mathcal{Q}$-tree it presupposes has the question formed by replacing the CT with a *wh*-word as the root node and subquestions replacing the CT with some alternative as its daughters. What the $\mathcal{Q}$TC requires is that for a conditional question to be under discussion, it must be a subquestion of a salient $\mathcal{Q}$-tree. Therefore, since an utterance of a counterfactual is relevant only if it is an answer to a conditional question under discussion, the counterfactual will be felicitous only if it answers a subquestion of a salient $\mathcal{Q}$-tree.

Going back to our example, when evaluating the counterfactual in (85a), repeated below in (93a), in the context of (90a), $\mathcal{Q}$TC is satisfied if the conditional question *if the flag were up, where would the king be?* is part of the $\mathcal{Q}$-tree in (87). Since it is, $\mathcal{Q}$TC is satisfied, the counterfactual satisfies Relevance and, once similarity is constrained as proposed in (35), the counterfactual in (93a) correctly comes out true. As for (93b), the CQ it answers does not satisfy the $\mathcal{Q}$TC; hence, (93b) is correctly predicted to be odd.

(93) a. If the flag were up, the king would be in the castle.
b. If the flag were up, the lights would be off.

Now, suppose that the non-accidental generalization salient in the context is (90b): whenever the king is away, either the lights are off or the flag is down. The $\mathcal{Q}$-tree for this new generalization is given below.

(94) when are the lights out or the flag down?

- are the lights out or is the flag down when the king is away?
- are the lights out or is the flag down when the king is in the castle?

The question that the counterfactual in (93b) answers is *if the flag were up, what about the lights?* The problem is that this question does not satisfy the $\mathcal{Q}$TC since it is not in the $\mathcal{Q}$-tree associated with the non-accidental generalization in (90b). With respect to the $\mathcal{Q}$-tree in (94), (93b) is ruled out as violating Relevance. Similarly, note that if we interpreted (93b) relative to (90a), the counterfactual would also be ruled out as violating Relevance.

The current proposal neatly accounts for the shifty intuitions in the King Ludwig of Bavaria example and, more importantly, for the "uncertain" judgments we get when considering the non-accidental generalization in (90b), without needing to get into the murky classification of non-accidental generalizations as "natural" or "less natural".

Before concluding, consider the counterfactual in (95): Paula is buying a pound of apples and the Atlantic Ocean is not drying up. The example is repeated in (95).

(95) If Paula were not buying a pound of apples, the Atlantic Ocean might be drying up.

This conditional is false. More importantly, it is strange. Why? For (95) to be a felicitous (relevant) answer, the CQUD at the time of utterance would have to be something like (96).

(96) If Paula were not buying a pound of apples, would the Atlantic Ocean be drying up?

The problem with (96) is that there would need to be a non-accidental generalization salient in the context such that the $\mathcal{Q}$-tree associated with it includes the conditional question in (96). That is, in the context in which (95) is uttered, the conditional question in (96) would have to be under discussion.

Since this is not true, $\mathcal{Q}$TC is not satisfied and the conditional is strange. Of course, (44) might be under discussion if our assumptions about the world were to change, in which case our judgment of (95) would change as well.

Similarly for (97), also from Kratzer 1989. Paula and Otto are the only people in this room. They are both painters. Clara is not in the room and she is a sculptor (not a painter).

(97) If Clara were also in this room, she might be a painter.

Just like (95), (97) is not only false but strange. For (97) to be a relevant answer, the CQUD at the time of utterance would have to be (98).

(98) If Clara were also in this room, what would she be?

Just like in the previous example, what is wrong with this is that, given our world-knowledge, there isn't a $\mathcal{Q}$-tree salient in the context of which (98) is part of. For (98) to be the CQUD, there would need to be a non-accidental generalization like "if someone is in this room, that person is a painter" in the context of utterance, but a generalization of this kind is neither part of the context nor part of our world knowledge. In other words, entertaining the counterfactual supposition that Clara is in this room does not raise the issue of what she is professionally.

Examples like (95) show that identifying the correct CQUD is crucial not only in selecting the right set of worlds (by selecting the correct premise set) but also in explaining the independent observation that counterfactual conditionals where the truth of the consequent has no connection with the antecedent are odd: counterfactually supposing the truth of the antecedent fails to raise the issue that the consequent is supposed to be an answer to. In other words, in these cases asserting *if ϕ, would ψ* is an infelicitous move since the conditional question it is supposed to answer, i.e., *if ϕ, would ψ?*, was not under discussion in the context of utterance.

What is under discussion (relative to the counterfactual assumption that we are making) is determined by contextual assumptions together with our world-knowledge. Why do the suppositions that Paula is not buying a pound of apples and that Clara is in this room not raise the issues of whether the Atlantic Ocean is drying up and whether Clara is a painter respectively? The answer is that in our world-knowledge there is no non-accidental generalization connecting these facts and raising the relevant issues.

The conclusion of the last few sections is that, through the $\mathcal{Q}$TC, we have proposed a mechanism through which non-accidental generalizations play a central role in the interpretation of counterfactual conditionals as often observed by those working on the semantics of counterfactuals.

7 Triviality and the Diversity Condition

In this section, we will look at cases where there is a logical connection between the antecedent and the consequent clauses, in that the former entails the latter. Consider (99).

(99) If Paula had eaten candy and cake, she would have eaten cake.

The proposition that Paula ate cake is clearly connected to the proposition that Paula ate candy and cake since the latter entails the former. The counterfactual in (99) is true, but trivially so.

Recall the intuition behind the Non-Triviality constraint on similarity proposed in (33): if a proposition entails an answer to the modally subordinated question, then it will be removed. The constraints on similarity given in (35) are repeated in (100).

(100) **Constraining the similarity ordering $\leq$:**
When evaluating a counterfactual *if ϕ, would ψ* in a context c, where *if ϕ, would Q?* is the current CQUD, relative to the similarity ordering $\leq_{A_{w_c}}$:

Step 1: Revise A_{w_c}
$A' = ntr_Q(con_\phi(A_{w_c}))$
where:

(i) for every $X \subseteq \wp(W)$ and $\phi \in \wp(W)$:
$$con_\phi(X) = \{p \in X : p \cap \phi \neq \emptyset \text{ and } \neg\phi \not\subset p\}$$

(ii) for every $X \subseteq \wp(W)$ and question Q:
$$ntr_Q(X) = \{p \in X : \neg\exists r \in Q\text{-alt, either } p \subseteq r \text{ or } r \subseteq p\}$$

Step 2: Define $sim_\leq$ relative to the revised set A':
$$sim_{\leq_{A'}}(\phi) = \{w' : \phi(w') = 1 \;\&\; \forall w'' : \phi(w'') = 1 \rightarrow w' \leq_{A'} w''\}$$

As explained above, the function *con* captures the Consistency constraint, whereas the function *ntr* captures the Non-Triviality constraint. What is wrong with counterfactuals like (99) is that the conditional question to which they are supposed to be an answer, i.e., *If Paula had eaten candy and cake, would she have eaten cake?*, is not under discussion, just like in the Atlantic Ocean example in (95) in section 7. In the latter the relevant question was not under discussion because no non-accidental generalization connected buying apples to oceans drying up; in (99), the relevant question is not under discussion because the antecedent already entails the answer to the question. In (99), the fact that there is no CQUD that the counterfactual could be a relevant answer to, means that *ntr* applies vacuously (since Q-alt in this case is $\emptyset$). These two facts combined cause the conditional to come out trivially true. As for (95), *ntr* applied trivially in that case as well, but the conditional came out false (and odd).

To recap, in the previous sections we saw that identifying the correct CQUD is essential in being able to rule out the right facts and to select the correct set of antecedent-worlds for the truth-conditions of the counterfactual conditional. What odd examples like (95) and (99) show is that identifying the correct CQUD is crucial in explaining the observation that counterfactual conditionals where either (i) the truth of the consequent has no connection with the antecedent or (ii) the truth of the consequent is entailed by the antecedent alone are odd: counterfactually supposing the truth of the antecedent fails to raise the issue that the consequent is supposed to be an answer to either because there is no connection between the two propositions in the context of utterance (cf. (95)) or because the connection between the two is trivial (cf. (99)). In all these cases, asserting *if* ϕ*, would* ψ is an infelicitous move since the conditional question that the conditional assertion is supposed to answer, i.e., *if* ϕ*, would* ψ*?*, was not under discussion in the context of utterance.

One last point before moving on: the constraint in (100) captures Condoravdi 2002's Diversity Condition without actually stipulating it. Ignoring some details that are not immediately relevant to the present discussion, Condoravdi's Diversity Condition stipulates that, given a modal sentence of the form [Modal p], the modal base cannot entail p. This condition is intended to explain why the two sentences in (101) and (102) (construed with a metaphysical modal base) are not truth conditionally equivalent.

(101) John has the flu (now).

(102) John might have the flu (now).

Condoravdi proposes that the modal operator in (102) quantifies over worlds metaphysically accessible at the utterance time, that is, worlds that share the same history as the actual world up to and including the utterance time. Since the proposition embedded under the modal is about the utterance time, if it is true that John has the flu now, then the two sentences should be equivalent and (101) should be able to be used to mean that John has the flu, which it is not. By stipulating that p can never be "settled" relative to the modal base of the sentence, the metaphysical reading for (102) is ruled out and the only available reading is the epistemic one. If applied to examples like (99), the Diversity Condition would require that the relevant set of antecedent-worlds not entail the proposition in the consequent, a requirement clearly violated in this example since the consequent is entailed by the antecedent itself. Note also that every time *any* of the implicit premises entails the consequent (which is the case if any premise entails an answer to the CQUD and the counterfactual is a relevant answer), the Diversity Condition is violated and the counterfactual is correctly ruled out, just like in our proposal.

The problem with Condoravdi's condition is that it is stipulative. In the current proposal, on the other hand, the selection of the relevant antecedent worlds is constrained by Non-Triviality, according to which, to be felicitous, a modal assertion of the form *if* ϕ, ψ (where the restriction of the modal can be overt as in a conditional or covert as in (102)) should not make a trivial contribution relative to the conditional question (*if* ϕ, *would Q?*) that is under discussion in the context of utterance.

A note about entailment and triviality. We want to make sure that just being entailed by the context does not rule out an assertion as being trivially true. For example, suppose that the context entails that if it rains I take the umbrella and that it is raining. The proposition that I take the umbrella is entailed by the context but asserting it is not a vacuous move. Compare this case with asserting that it is raining in a context already entailing that it is raining. While they are both entailed by the overall context of utterance, only asserting the latter feels like a vacuous conversational move. I suggest that the reason why the assertion that I take the umbrella in the first context is not as vacuous as asserting that it is raining in the second context, is that the proposition that I take the umbrella is entailed by the conjunction of two propositions true in the context of utterance (first, that if it rains I take the umbrella, and second, that it is raining) but by neither of them individually.[9]

9 The proposition formed by the conjunction of these two, that is, *If it rains I take the umbrella and it is raining*, does entail the proposition that I take the umbrella but this does not make

This point is important for the semantics of counterfactuals. Simplifying for a moment the Lewisian/Kratzerian semantics for counterfactuals, a conditional of the form *if ϕ, would ψ* is true just in case all the ϕ-worlds are ψ worlds. In light of what I have just suggested, if the consequent follows from the antecedent together with a suitable set of premises but is entailed by neither the antecedent or any of the premises by themselves, then the *would*-conditional is not trivially true.

Hence, the difference between informative counterfactuals like *If it had rained, I would have taken the umbrella* and trivially true counterfactuals like *If it had rained, it would have rained* or *If Paula had eaten candy, she would have eaten candy* is that, while in all cases the counterfactual is true just in case the consequent is true in all the relevant antecedent worlds, it is only in the trivially true cases that the consequent is entailed by a single premise by itself (the antecedent in this case). Being trivially true because one implicit premise entailed by itself the consequent is precisely the problem that Non-Triviality, and in particular the function *ntr*, is designed to avoid.

8 Some additional cases

Firstly, consider cases of subjunctive conditionals that are not counterfactual (cf. Anderson 1951).

(103) If Smith had taken arsenic, he would show the symptoms which he does in fact show.

The question to which (103) is a relevant answer is *if Smith had taken arsenic, which symptoms would he show?* We assume that if someone takes arsenic, they will show a certain set of symptoms (which are the ones which Smith is showing in the actual world). Hence, this generalization presupposes the question *which symptoms do someone show?* with all its subquestions, as indicated below.

(104) which symptoms does someone show?

. . . if they take arsenic? . . . if they take . . . ?

the proposition that I take the umbrella trivially true because neither conjunct entails the proposition in question by itself.

Since the question to which the counterfactual is a relevant answer is a subquestion in this Q-tree presupposed by the non-accidental generalization we are assuming, this question satisfies the QTC and qualifies as the CQUD. Hence, *ntr* will remove the proposition that Smith shows the symptoms he shown, but since we are retaining the non-accidental generalization above (that if you take arsenic you show these symptoms), the counterfactual will come out true.

Before moving to the next section, I am going to look at one "symmetrical" counterfactual and show that our proposal can handle these cases as well. The case I am going to look at is given in (105).

(105) a. If New York City had been in Georgia, New York City would have been in the South.
b. If New York City had been in Georgia, Georgia would have been in the North.

In (105), we judge (105a) true but (105b) false. There are two ways of constructing the set of premises.

(106)

Possibility (i):	**Possibility (ii):**
(e) New York City is in Georgia	(e) New York City is in Georgia
(a) Georgia is in the South	(a) New York City is in the North

Our truth-value judgments tell us that we must choose Possibility (i). Here is a way of ruling out Possibility (ii). Let us start with some relevant propositions true in the actual world.

(107) a. New York City is in New York State.
b. New York State is in the North.
c. Georgia is in the South.
d. If a city is in New York State, it is in the North.
e. If a city is in Georgia, it is in the South.

In order to accommodate the counterfactual antecedent, Consistency requires that we remove the proposition that New York City is in New York State. At this point, Non-Triviality requires that we remove any propositions entailing a possible answer to the CQUD. Take (105a): the question that this counterfactual is a relevant answer to is *if New York City had been in Georgia, would New York City have been in the South?* Let's check whether it satisfies the $\mathcal{Q}$TC. The relevant non-accidental generalizations are given in (107d) and (107e) and one $\mathcal{Q}$-tree that is presupposed is given below.

(108) when is a city in the South?

...if it's in NY state? **...if it's in Georgia?** ...if it's in ...?

The question *if New York City had been in Georgia, would New York City have been in the South?* counts as a subquestion of (108) and therefore qualifies as being under discussion. In constraining similarity, *ntr* will then remove the proposition that New York City is in the North but not the proposition that Georgia is in the South. Hence, (105a) will correctly come out true. The problem with (105b) is that it is about the location of states whereas the non-accidental generalizations we are assuming are about cities and their locations. In our proposal, what is wrong with (105b) is that it cannot be an answer to any of the subquestions presupposed by these non-accidental generalization. Hence, the feeling that the counterfactual is incongruent with what we assume, and in the end false (since we do keep the proposition that New York State is in the North and the proposition that Georgia is in the South).

9 Non-causal/non-interference conditionals

In this section I will explore the consequences of our proposal for those conditionals where there is no causal connection between the antecedent and the consequent propositions. An example is given in (109).

(109) A and B took a car trip and they strapped their baby in her car seat. Everything went fine. When they returned home, they heard of a car accident where a child was injured because she was not strapped in her car seat.

A: Thank god we didn't have an accident and our baby wasn't hurt!

B: Sure but remember that if we had had an accident, our baby would have been strapped in her car seat.

Bennett 2003 has called these conditionals *non-interference conditionals*. The distinctive property of these conditionals is that there is no making-true relation between ϕ and ψ. The example in (109), a variant of an example in

Bennett 2003, shows that the very same conditional in (109) is ambiguous between a causal and a non-causal interpretation.

(110) If the dam had broken, the water would have been low.

According to the causal interpretation, the water would have been low as a result of the breaking of the dam. To appreciate the non-causal interpretation, let us imagine that the dam did not break, and now you are worrying about all the damage to the fields that would have occurred had the dam broken. By uttering (110) I mean to reassure you that since the water was low, had the dam broken, the damage would have been modest. The contrast between these readings is real, but how can a theory of counterfactuals explain it?

The proposal that I would like to make is that the difference between the causal and non-causal counterfactuals lies in their relation to the CQUD. A causal counterfactual answers the CQUD *directly*, whereas a non-causal counterfactual answers the CQUD *indirectly* by spelling out a premise assuming which the CQUD is then answered.

Let's start with the non-causal conditional in (109). A's utterance tells us that the conditional question that A takes to be under discussion at the time of utterance is about whether if there had been an accident, the baby would have been hurt. B's utterance answers this question but indirectly. B utters a conditional whose consequent is actually the premise assuming which the conditional question to which A's utterance is relevant is answered. B answers A's CQ indirectly by providing a premise which, together with the proposition that there was an accident, will entail the answer to the CQUD: since the baby was strapped in, if there had been an accident, the baby would not have been hurt. The reason why B chose to answer the CQUD indirectly is that by choosing this strategy, not only did B effectively answer the CQUD but B also explicitly provided the reason for that answer.

Let's look at (110). Its causal interpretation is straightforward. Assuming that the discourse of which (110) is part obeys relevance, the conditional question under discussion at the time of utterance would have to be about how tall the water would have been as a result of the breaking of the dam. As for its non-causal interpretation, just like in the previous example, the purpose of the counterfactual in the context of utterance is to spell out one of the premises assuming which the conditional question salient in the context is answered. But answering this question is done indirectly. In this example, the CQUD is about whether there would have been damage if the dam had broken. The speaker does not answer this question directly; instead, he asserts (110) whose consequent spells out the premise assuming which the issue of whether there would have been damage as a result of the dam breaking is resolved, i.e., the CQUD is answered.

I believe that this mechanism is similar to the one responsible for the interpretation of so-called relevance conditionals. Take (111).

(111) If you're hungry, there is food in the fridge.

If (111) is part of a coherent discourse, the conditional question under discussion at the time of utterance is about whether there is anything to eat in case you are hungry. The speaker of (111) answers the CQUD indirectly by making explicit a premise assuming which the conditional question is answered. The result is an answer that is as informative as (112) but less explicit about the goal of the conversation, which is to enable you to eat (in case you are hungry).

(112) If you're hungry, you can eat what's in the fridge.

The idea that we can answer a CQUD indirectly is also useful in providing an account of cases of disagreement such as the one in (113).

(113) Eva was invited at Kai's house yesterday for his birthday party. Kai has two cats and Eva is allergic to cats. Eva, though, couldn't go to the party because of a prior engagement.

A: If Eva had gone to Kai's birthday party, she would have had an allergic reaction.

B: No, If Eva had gone to Kai's birthday party, she would have taken allergy medication.

The CQUD here is whether Eva would have had an allergic reaction had she gone to Kai's birthday party. The goal of the exchange between A and B is to establish the answer to this question. B is not directly disagreeing with A as if she would have done had she uttered the conditional *if Eva had gone to Kai's birthday party, she would not have had an allergic reaction*. B is disagreeing *indirectly* by asserting a conditional whose consequence is a proposition such that, when added to the relevant set of premises, will force A to conclude that Eva would not have had an allergic reaction, had she gone to Kai's birthday party. If A accepts B's assertion (because she recognizes the truth of the non-accidental generalization on which B's claim is based, i.e., that whenever

Eva goes where there are cats, she takes allergy medication), then B will have been successful in doing that.

10 Concluding remarks

It is commonly held that counterfactuals are context-dependent and that which possible worlds (or which premises) are selected in order to arrive at the correct truth-conditions is, implicitly or explicitly, taken to depend on the particular assumptions that are made in the context of utterance. The goal of this paper was to advance our understanding of the role of the context in figuring out the truth-conditions of counterfactual conditionals.

Combining elements from Lewis 1973, Kratzer 1991, and von Fintel 2001, I proposed a possible worlds semantics for counterfactuals conditionals where the relevant antecedent worlds are selected based on how similar they are to the actual world. The relation of similarity is constrained by what I called Consistency and Non-Triviality. Consistency is what every theory of counterfactuals must include: when selecting the relevant set of worlds according to how similar they are to the actual world, we must rule out all those propositions that are true in the actual world that are inconsistent with the counterfactual antecedent. The contribution of this paper revolves around Non-Triviality. According to Non-Triviality, we must rule out all those propositions that are true in the actual world but that entail a possible answer to the conditional question under discussion. In other words, any proposition that entails a member of the Q-alternative set of the subordinated question under discussion (i.e., any proposition that entails an answer to its Q-alternative set) cannot be part of the set of propositions determining the similarity ordering. Assuming a model of discourse structure similar to the one proposed in Roberts 1996 and Büring 2003 and related work, according to which all conversational moves (questions and assertions) are answers to (often implicit) questions under discussion, the idea behind Non-Triviality is that a counterfactual statement is relevant if it answers a conditional question under discussion and that, for the discourse to be felicitous, the conditional must make a non-trivial assertion.

I have proposed that identifying the conditional question under discussion depends on features of the context and world-knowledge. In particular, I showed that non-accidental generalizations which have often been taken to play an important role in the interpretation of counterfactuals, are crucial in selecting which conditional question is under discussion. I have proposed that a non-accidental generalization salient in a context c presupposes a $\mathcal{Q}$-tree in c, that is, a family of questions structured hierarchically representing what is under discussion in c. According to the $\mathcal{Q}$TC I have proposed, the CQ that an utterance of the counterfactual *if* ϕ, *would* ψ is a relevant answer to qualifies as being under discussion if the CQ is a subquestion of a salient $\mathcal{Q}$-tree.

This proposal proposes a formal link between non-accidental generalizations and counterfactual conditionals, capturing our intuition that non-accidental generalizations are crucial in identifying which issues are raised by entertaining a counterfactual assumption and are therefore central in evaluating the truth of these counterfactuals.

I also showed that Non-Triviality (i) rules out all those odd counterfactuals where the antecedent is not relevant to the truth of the consequent and (ii) captures Condoravdi's Diversity Condition without having to stipulate it.

This proposal has consequences for the interpretation of embedded counterfactuals like (114) in particular with respect to the role that conditional questions play in the computation of similarity in the theory I have proposed in this paper.

(114) Mary believes that if I had bet on heads, I wouldn't have lost \$10.

Let us suppose that Mary's doxastic alternatives (Mary's belief worlds) in (114) are all worlds where the same non-accidental generalization we assumed for (60) in section 5 is true: that you win if your bet matches the outcome of the coin-tossing. The relevant $\mathcal{Q}$-tree corresponding to this non-accidental generalization is just like the one we constructed for (60), repeated in (115).

(115) when does someone lose in a random coin-tossing?

- **does someone lose if they bet on heads**
  - does someone lose if they bet on heads and heads comes up?
  - does someone lose if they bet on heads and tails comes up?
- does someone lose if they bet on tails
  - does someone lose if they bet on tails and tails comes up?
  - does someone lose if they bet on tails and heads comes up?

One possibility is that this $\mathcal{Q}$-tree above is somehow represented in the embedded context with respect to which the embedded counterfactual is evaluated. In other words, this $\mathcal{Q}$-tree needs to be somehow represented in

Mary's doxastic state so as to capture the intuition that (114) says that Mary believes that with respect to the issue of when I would lose money in coin-tossing, if I had bet on heads, I wouldn't have lost. Again abstracting away for the counterfactual morphology for now, the conditional question that the counterfactual is supposed to answer (if I had bet on heads, would I have lost?) is a subquestion of (115) and, therefore, the computation of similarity can be done with respect to this question. In other words, the suggestion here is that the conditional question we use to compute similarity is made salient by the non-accidental generalization true in Mary's doxastic worlds. I leave the exploration and formalization of this idea for the future.

I will conclude by stressing that the context-dependence of similarity in this proposal is reduced to the context-dependence of the CQUD. I argued that in searching for the appropriate CQUD two criteria are crucial: (i) that the CQUD be a subquestion of a salient Q-tree presupposed by a non-accidental generalization contextually salient or part of our world-knowledge and (ii) that the counterfactual we are evaluating be a relevant answer to such CQUD, where "relevant" is to be understood in the sense of Roberts 2012.

References

Alonso-Ovalle, Luis. 2006. *Disjunction in alternative semantics*. University of Massachusetts Amherst dissertation. http://semanticsarchive.net/Archive/TVkY2ZlM/alonso-ovalle2006.pdf.

Anderson, Alan Ross. 1951. A note on subjunctive and counterfactual conditionals. *Analysis* 11. 35–38.

Arregui, Ana. 2009. On similarity in counterfactuals. *Linguistics and Philosophy*. 245–278. http://dx.doi.org/10.1007/s10988-009-9060-7.

Bennett, Jonathan. 2003. *A philosophical guide to conditionals*. Oxford University Press. http://dx.doi.org/10.1093/0199258872.001.0001.

Büring, Daniel. 2003. On D-trees, beans, and B-accents. *Linguistics and Philosophy* 26. 511–545. http://dx.doi.org/10.1023/A:1025887707652.

Condoravdi, Cleo. 2002. Temporal interpretation of modals: Modals for the present and for the past. In David I. Beaver, Luis D. Casillas Martínez, Brady Z. Clark & Stefan Kaufmann (eds.), *The construction of meaning*, 59–88. CSLI Publications.

von Fintel, Kai. 2001. Counterfactuals in a dynamic context. In Michael Kenstowicz (ed.), *Ken Hale: A life in language*, 123–152. Cambridge, MA: MIT Press.

Goodman, Nelson. 1947. The problem of counterfactual conditionals. *Journal of Philosophy* 44. 113–128. http://dx.doi.org/10.2307/2019988.

Groenendijk, Jeroen. 1999. The logic of interrogation. In Tanya Matthews & Devon Strolovitch (eds.), *Semantics and Linguistic Theory (SALT) 9*, 109–126. Ithaca, NY: CLC Publications. http://dx.doi.org/10.3765/salt.v9i0.2835.

Hamblin, C. L. 1973. Questions in Montague English. *Foundations of Language* 10(1). 41–53. http://www.jstor.org/stable/25000703.

Heim, Irene. 1992. Presupposition projection and the semantics of attitude verbs. *Journal of Semantics* 9(3). 183–221. http://dx.doi.org/10.1093/jos/9.3.183.

Ippolito, Michela. 2006. Semantic composition and presupposition projection in subjunctive conditionals. *Linguistics and Philosophy* 29(6). 631–672. http://dx.doi.org/10.1007/s10988-006-9006-2.

Ippolito, Michela. 2013. *Subjunctive conditionals: A linguistic analysis* (Linguistic Inquiry Monographs). MIT Press.

Isaacs, James & Kyle Rawlins. 2008. Conditional questions. *Journal of Semantics* 25(3). 269–319. http://dx.doi.org/10.1093/jos/ffn003.

Jackendoff, Ray. 1972. *Semantic interpretation in generative grammar* (Current Studies in Linguistics). Cambridge, MA: MIT Press.

Kratzer, Angelika. 1989. An investigation of the lumps of thought. *Linguistics and Philosophy* 12(5). 607–653. http://dx.doi.org/10.1007/BF00627775.

Kratzer, Angelika. 1991. Modality. In Arnim von Stechow & Dieter Wunderlich (eds.), *Semantics: An international handbook of contemporary research*, 639–650. Berlin: De Gruyter.

Kratzer, Angelika. 2012. *Modals and conditionals*. Oxford University Press.

Lewis, David. 1973. *Counterfactuals*. Cambridge, MA: Harvard University Press.

Lewis, David. 1979. Counterfactual dependence and time's arrow. *Noûs* 13(4). 455–476. http://dx.doi.org/10.2307/2215339.

Pollock, John L. 1976. *Subjunctive reasoning*. Dordrecht: Reidel.

Ramsey, Frank Plumpton. 1929. General propositions and causality. In Richard Bevan Braithwaite (ed.), *The foundations of mathematics and other logical essays*, 237–255. London: Routledge & Kegan Paul.

Roberts, Craige. 1989. Modal subordination and pronominal anaphora in discourse. *Linguistics and Philosophy* 12(6). 683–721. http://dx.doi.org/10.1007/BF00632602.

Roberts, Craige. 1996. Anaphora in intensional contexts. In Shalom Lappin (ed.), *The handbook of contemporary semantic theory*, 215–246. Oxford, UK: Blackwell.

Roberts, Craige. 2012. Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics* 5(6). 1–69. http://dx.doi.org/10.3765/sp.5.6.

Rooth, Mats. 1985. *Association with focus*. University of Massachusetts Amherst dissertation. http://dx.doi.org/1813/28568.

Tichý, Pavel. 1976. A counterexample to the Stalnaker-Lewis analysis of counterfactuals. *Philosophical Studies* 29(4). 271–273. http://dx.doi.org/10.1007/BF00411887.

Veltman, Frank. 2005. Making counterfactual assumptions. *Journal of Semantics* 22(2). 159–180. http://dx.doi.org/10.1093/jos/ffh022.

Wason, Peter C. 1966. Reasoning. In Brian M. Foss (ed.), *New horizons in psychology*, 135–151. Penguin Books.

Michela Ippolito
Department of Linguistics
University of Toronto
Sidney Smith Hall
100 St. George Street
Toronto, ON M5S 3G3
Canada
michela.ippolito@utoronto.ca