How similar is similar enough?

Michela Ippolito

University of Toronto

(Non-final draft)

1 Measuring similarity accross 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 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) $\phi \mathrel{\Box\!\!\rightarrow} \psi$ 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 $\varphi \rightarrow \psi$ (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 accross 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 Tichy (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 Tichy'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 Tichy'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), repeated in (4), true.

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

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 (4) should be false. One might elaborate on Lewis's story by saying that when selecting the worlds most similar to the actual world, we want to maximize similarity to the actual world even after the divergence: we keep as many facts true in the actual worlds as possible unless they make the actual laws vacuous.[1] Consider (5), from Arregui (2009).

(5) 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. According to the suggestion above, what might explain our judgment in this case is that we cannot keep the fact that heads came up because, since the outcome of Susan's coin tossing is up to the probabilistic laws that regulate this type of physical event, already assuming the outcome of such coin tossing event makes these laws vacuous.

Applied to the Jones examples above, this idea would explain why we cannot retain the fact that Jones is wearing his hat in Tichy's example but we do retain the fact that heads came up in Veltman's example: in worlds in which the weather is fine, and Jones tosses a coin, assuming the outcome of the coin tossing (heads) would make the probabilistic laws at work in those worlds vacuous, just like in the coin tossing case above. Hence, since we cannot assume that heads came up, it does not follow from the supposition that the weather is fine together with the laws of nature that Jones would be wearing his hat.

[1]This possible line of explanation has been suggested in Arregui (2009)

While interesting, this story does not go far. To see this consider (6), a variant of the coin tossing example given in (5), from Ippolito (2013).

(6) 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's 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 tosses the coin 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 it had been your turn to press the button, you would have lost $10.

Unlike (5), we judge (6) true. In (6), unlike what we saw in (5), we are allowed to keep the fact that the coin came up heads, 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 ((7)).

(7) 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.

(8) *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 (8) 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 (9). S is the set of possible situations and $\wp(S)$ is the power set of S, i.e. the set of propositions.

(9) 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.

(10) 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 (10-a), but not (10-b), 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.

(11)

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 (10-b), repeated below, would incorrectly be predicted to be true.

(12) 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 (69-b) 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 (13).

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

Now consider a variant of the original example.

(14) 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 (69-b) is no longer judged clearly false. Our judgments have "shifted" 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 (13) with a logically equivalent one in (15).

(15) 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 (9), 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 (14) 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 (10-a) and (10-b) 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 (16),

(16) *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.

(17) 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 statememt: 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 the cards that showed an odd number. The latter move, i.e. selecting the cards showing an odd number, would have allowed the subject to falsify the statement. Both of the subject's moves enabled the subject only to *confirm* the statement. Kratzer's suggestion 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.

(18) 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.

(19) 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 (14), or because of extra-linguistic requirements such as that propositions in the Base Set be "cognitive viable"),[2] in assembling a Base

[2]See Kratzer (2012), p. 132, for a discussion of the concept of cognitive viability.

Set that includes this non-accidental generalization, propositions confirming the non-accidental generalization that has been selected are privileged.

Schematic truth-conditions for *would* and *might* conditionals in Kratzer's premise semantics are given in (20) and (21).

(20) *Would-counterfactuals*

Given a world w and an admissible Base Set $\mathbf{F}_w$, a *would*-counterfactual with antecedent p and consequent q is true in w iff for every set in $\mathbf{F}_{w,p}$ there is a superset in $\mathbf{F}_{w,p}$ that logically implies q.

(21) *Might-counterfactuals*

Given a world w and an admissible Base Set $\mathbf{F}_w$, a *might*-counterfactual with antecedent p and consequent q is true in w iff there is a set in $\mathbf{F}_{w,p}$ such that q is compatible with all its supersets in $\mathbf{F}_{w,p}$.

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

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. Hence, the CPC will have truth-conditional effects.

Before turning to the King Ludwig of Bavaria example, let me illustrate how Kratzer's proposal account for Tichy'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 (22).

(22) 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, Non-Redundancy defined as below is going to rule out that the 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.

(23) *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 CPC and Redundancy 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 (18) 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 (10-a) is predicted to be true.

In the variant we have been considering, though, the non-accidental generalization is the one given in (19) and repeated below.

(24) 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 (25) will be judged true.

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

This prediction is not quite correct, though, since, as Kratzer points out, our judgments are uncertain in the case of (25). 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 (25) true. Kratzer's suggestion is that, behind the uncertainty in judging (25), lies the fact that the non-accidental generalization in (58) 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), p. 146. I will go back to this point and the discrepancy between the prediction made by the CPC and the uncertain judgments we get in cases like (25) 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 this time Peter bets for us. Peter bets on tails, presses a button in a random coin-tossing device, heads comes up and we lose. You don't like Peter and get really upset with him. Now suppose that Susan had pressed the button.

(26) Poor Peter! I don't think you're being fair. If Susan had bet on tails and pressed the button, you would not have said a word to her.

(27) If Peter had bet on tails and 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 (26) true. Crucially, though, in the same context the same people do not judge (27) true. What explains these judgments? The contexts 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 like that). In (27), it must be the case that the proposition that heads came up is not privileged so that in the Crucial Set we'll have sets with the proposition that heads came up and sets with the propositions that tails came up (and the *would* counterfactual is false). But in (26), 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 your relation with Peter is relevant for (26), whereas the generalization about the random nature of coin-tossing is relevant for (27). 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 (26) and (27) 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 premise set. 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 will defend in this paper is better equipped to predict the shifts 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 not been clearly spelled out so far. 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 indentifying 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 ((26) and (27)) 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 that I will defend is that a crucial element driving the selection of the relevant premises (and, ultimately, possible worlds) is the need to avoid trivial moves in the discourse, where a move is to be understood in the sense of Roberts (1996b). A counterfactual conditional *if* ϕ*, would* ψ is understood to be an answer to a conditional question of the form *if* ϕ*, Q?*, where Q is a question raised relative to a temporary context in which ϕ is true. I will call this question the *conditional question under discussion* (CQUD). In other words, the counterfactual *if* ϕ*, would* ψ is an answer to the CQUD *if* ϕ*, Q?*, where Q and ψ are understood as a question and an answer raised in the same temporary ϕ-context.

We said that ψ is supposed to answer the modally subordinated question Q. Just like any other question, we require that Q be a non-trivial question, i.e. a question whose answer is not already entailed by the temporary context (the relevant ϕ-worlds). If any of the propositions true in these ϕ-worlds 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.

In what follows I will briefly introduce two notions that will be important for the discussion to follow: the notion of a question under discussion (QUD) and the notion of a conditional question (CQ). Let's start with the latter.

2.1 Conditional questions

I loosely follow the analysis of conditional questions in Isaacs and Rawlins (2008). For reasons of space, the following discussion will be informal and brief and will only be concerned with polar questions like (28).

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

According to Isaacs and 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 question means to update the context inquisitively, that is, to partition the context set into only two cells. 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 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

[3] See Roberts (1989) and Roberts (1996a) for an introduction to the concept of modal subordination.

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 (1996b), 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 answwearable). 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 (1996b), we can define Relevance a bit more formally as shown in (29).

(29) 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 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 (1996b) 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. In what follows, I will argue 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.

3 Back to counterfactuals

In the following sections, I will argue that 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, i.e. 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 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.

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 $\leq_A$** (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 $\leq_{A_{w_c}}$:

1. Consistency: for all propositions χ, $\chi \in A_{w_c}$, if $\phi \cap \chi = \emptyset$, then remove χ from A_{w_c}.

2. Non-Triviality: for all propositions χ, $\chi \in A_{w_c}$, if $\exists p \in$ Q-set s.t. $\chi \subseteq p$, then remove χ 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 that are inconsistent with the counterfactual antecedent ϕ and (ii) those propositions 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.

We can put the constrains on similarity given above in their final and more formal form in (34).

(34) **Constraining the similarity ordering $\leq_{A_{w_c}}$:**

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 $\phi \in \wp(W)$: $con_\phi(X) = \{p \in X : \phi \cap p \neq \emptyset\}$

(ii) for every $X \subseteq \wp(W)$ and question Q: $ntr_Q(X) = \{p \in X : \forall q \in Q\text{-alt}, p \nsubseteq q\}$

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 (34), the function *con* captures the Consistency constraint: $con_\phi(A_{w_c})$ is going to deliver the set of propositions in A_{w_c} that are consistent with the antecedent. The *ntr* function 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 in it entailing a possible answer to the question under discussion.

Finally, unlike Kratzer (2012) I will not assume that the premise set is subject to a non-redundancy constraint: in the present proposal if two propositions p and q are different yet p entails q, they can both be in A'. We will see in the next sections why this is important.

In what follows I will return to all 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 the coin-tossing examples, then move on to Jones and the weather's examples, and finally to the King Ludwig of Bavaria's example.

4 Jones and the weather

Recall Tichy's example in (3) and Veltman's variant in (4), repeated in (35) and (36).

(35) *Tichy'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.

(36) *Veltman's version of Tichy'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 (35) false and (36) 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?

Counterfactuals are interpreted as being answers to the conditional question under discussion at the time of utterance. In the case of both (35) and (36), the CQUD is a question about what would be the case with respect to a certain issue, if the weather had been fine. What issue? The context has made salient non-accidental generalizations involving the weather and Jones's wearing his hat. It is the connection between these two facts that is under discussion in the context of utterance. Hence the CQUD here is (37).

(37) If the weather had been fine, <u>would Jones be wearing his hat?</u>

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$\}$. Note that, in giving the Q-set of the modally subordinated question here and in the following examples, I ignore the nonindicative morphology on the modal auxiliary *would*. I assume that it is there for independent reasons and that its contribution is not relevant to the questions we aim to answer in this paper.

For both (35) and (36), the schematic truth-conditions will be the following.

(38) When evaluating the counterfactual conditional **If the weather had been fine, Jones would be wearing his hat** in a context c where *if ϕ, would Q?* 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'

However, since the facts in the two scenarios we have been considering are different, these schematic truth-conditions are going to deliver different truth-conditional judgments in the two examples. Let's start with (35). The relevant propositions true in the actual world are given in (39).

(39) 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 do not keep (39-a) because it is inconsistent with the antecedent proposition (that the weather is fine). Non-Triviality, on the other hand, rules out the proposition in (39-b) since it entails one member of the Q-alternative set. Nothing rules out (39-c), so we keep it. Since we do not retain (39-b), the counterfactual in (35) is false. This is because, as a result of removing (39-b), 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 (36). The relevant propositions this time are shown in (40).

(40) 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

Since the CQUD for (36) is also (37), 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, (40-a) is ruled out by Consistency, and (40-b) is ruled out by Non-Triviality. Crucially, though, the proposition in (40-c) 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 (36).

4.1 Redundancy

In order to account for our truth-conditional judgment in (36) I proposed above that the premise set we start with includes the propositions listed in (40). Now, suppose that instead of these propositions we have the following.

(41) 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 (41) 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 (41-a) and not as an indepedent member of the set. This set seems to characterize the actual world as well as (40). However, if (41) is the premise set we choose, Tichy'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 (41-a) and, with it, we would be losing the proposition that head came up.

My claim is that what prevents (41) from being selected as the relevant set is that, *contra* Kratzer (2012)m there is a preferance for redundancy in constructing the right set of premises. Here is the definition we introduced before from Kratzer (2012).

(42) *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).

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

It follow that in Veltman's version of Tichy's example, we should either select (40) where the only occurence of the proposition that the weather is bad is independent of the proposition that heads came up in that set or the set in (44).

(44) 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 (44-a) as well as (44-c) will be removed when adding the counterfactual proposition that the weather is fine will not cause the counterfactual in (36) to be false because the proposition that heads came up in (44-d) will not be removed.

Here is another illustration of a set of premises that initially seems to challenge our proposal. Suppose that we are considering the following set when evaluating Tichy's original example.

(45) 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 is that, once we add the counterfactual proposition that the weather is fine to this set, the combination of this proposition and (45-a) entails that Jones is wearing is hat, which would then incorrectly predic that we should judge TIchy's counterfactual true.

This seems a problem but notice what happens once we require Redundancy, as shown in (46).

(46) 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 (46-b) is excluded by consistency and the proposition that Jones is wearing his hat in (46-c) is excluded by Non-Triviality. Hence, since consistency forced us to remove both (46-b) and (46-c), it will force us to remove the disjunction in (46-a) as well.

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.

(47) 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.

(48) 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 (47). The relevant propositions are given in (49).

(49) 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 (47) is uttered. Assuming that (47) is an answer to the conditional question under discussion at the time of utterance means that both the consequent in (47) and the question part of the CQUD are modally subordinated to the same temporary context. Hence, the question is going to be of the form *If I had bet on heads, ... ?* World-knowledge tells us that there is a non-accidental correlation between betting on something and either winning or losing. Hence, the CQUD will be (50).

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

Consistency will rule out (49-a). 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 (49-d). Both (49-b) and (49-c) 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 (48)? What issue does the (counterfactual) supposition that Susan pressed the button raise? Given our world-knowledge (that each coin tossing event has a random outcome) and given the context of utterance, the CQUD is (51).

(51) If Susan had pressed the button, what would have come up?

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 (52).

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

This time Consistency rules out (52-a); Non-Triviality rules out (52-b). It follows that $sim_{\leq_{ntr_Q(con_\phi(A_{w_c}))}}(\lambda w.$ Susan pressed the button in $w) \not\subseteq \{w'$: 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 be under discussion in a given context. To see this point clearly, suppose we construct a given context and then stipulate that a certain CQ is under discussion in the context of utterance. For example, imagine Peter pushed the button in a random coin-tossing device and it came up heads. Suppose also that I bet on tails and so I lost \$10. Now, let's *stipulate* that the CQUD is *If Susan had pressed the button, would I have lost \$10?* Our intuition is that in this context, the counterfactual in (53) is false.

(53) If Susan had pressed the button, I would 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 any members of the Q-alternative set of the CQ *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 (53) should come out true. How can our proposal explain our judgment here? The crucial point here is to understand what it means for a question to be "under discussion". In the scenario we constructed above we stipulated that the CQUD is *If Susan had pressed the button, would I have lost \$10?* but now we need to ask what it means for this question to be under discussion in that context. Since no particular non-accidental generalization has been made salient in the context, we can assume that the non-accidental generalizations at work in this case will be the ones which express our general world knowledge. One such generalization will be, for example, that whenever there is a coin tossing the outcome is random. Another one is that whenever the outcome of the coin tossing does not match your bet, you lose. However, that whenever someone (for example, Susan) presses the button, you lose \$10 is not a non-accidental generalization made available by our world knowledge. In the given context and given our world knoweldge, entertaining the counterfactual supposition that Susan pressed the button will raise the issue of whether I would have lost \$10 only in so far as it raises the issue of what the outcome of that coin tossing would have been. In other words, in this context, if the CQ *If Susan had pressed the button, would I have lost \$10?* is under discussion, that is because the CQ *If Susan had pressed the button, what would have come up?* is under discussion. This tells us that the non-accidental generalizations salient in the context of utterance are a crucial factor in making a CQ under discussion. Since in our context the money question can only be under discussion if the outcome question is too, Non-Triviality will require that we remove from the premise set not only the proposition that I lost \$10 but also the proposition that heads came up. Thus, in the given context and given our world knowledge, the counterfactual in (53) is no longer predicted to be true.

This analysis predicts that (53) will be true if evaluated in a context that encourages us to separate the issue of whether someone would lose money from the issue of the outcome of the coin tossing. If we are right about the centrality of non-accidental generalizations in the evaluation of counterfactuals, such a context would have to be one where other non-accidental generalizations are salient. The example we introduced earlier in (6), repeated below in (54), shows that this is precisely what happens.

(54) 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's 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 tosses the coin 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 it had been your turn to press the button, you would have lost $10.

In this context, we judge the counterfactual that if it had been your turn to press 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 if the outcome of the coin tossing does not match your bet, you lose money only if you were the one tossing the coin. In this context, it is possible to entertain the counterfactual possibility that Susan pressed the button, and raise the issue of whether I would have lost $10 *without* also raising the issue of what the outcome of the coin tossing would have been: this is because there is a direct relation between who the agent of the tossing is and whether someone will lose $10. Hence, in this context, because raising the issue of whether I would have lost $10 does not (necessarily) raise the issue of what the outcome of that coin tossing would have been, the latter issue will not be raised and the proposition that heads came up will not be required by Non-Triviality to be removed.

In other words, the problem with (53) was that we stipulated what the CQUD is without spelling out which non-accidental generalizations would have to be salient in the context of utterance for that conditional question to be under discussion. Once we spelled out those non-accidental generalizations, we saw that in that context raising a certain issue could not be done without raising other issues as well, thus requiring Non-Triviality to remove more propositions that we initially thought. To conclude, the general point that pairs like (53) and (54) 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 ageneralizations. In Kratzer's story, non-accidental generalizations play a key role in accounting for the different ways in which we construct premise sets: what I will argue in the rest of this paper is that there is a "formal link" between a non-accidental generalization and the CQUD and that the hypothesis that the counterfactual is a relevant answer to the CQUD helps us explain why, in a context where multiple non-accidental generalizations are salient, we select some but not others. Because, as we observed above, there is a formal connection between the CQUD and the counterfactual (i.e. the consequent is understood to be an answer to the question and, consequently, both are interpreted as being modally subordinated to the relevant antecedent worlds), it follows that there must be a formal connection between a non-accidental generalization and the counterfactual. In the remaining of this section I will introduce an example that will motivate what I will call the *Matching Condition*, which we will discuss at length in section 7.

Consider the following scenario. Peter, Susan, you and me are on 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 are evaluating the following two counterfactuals.

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

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

We judge (55) true, but (56) 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 (56) but not in evaluating (55). 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 (57) or (58) (repeated below) to be the salient generalization (we will go back to Kratzer's case in section 6).

(57) 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.

(58) 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, generalizations as (57) and (58) have different confirming propositions and therefore will affect which propositions will be kept and which ones will be removed.

The puzzle that (55) and (56) 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 judgements, we must keep the proposition that heads came up in the premise set in (55) (it is because we keep this proposition that we judge (55) true) but not in (56) (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 involved in something wrong, you don't get upset – one might deny that they are both salient for the two counterfactuals we are considering. In particular, one is tempted to say that the generalization about the nature of coin-tossing is relevant for (56), whereas the generalization about the speaker's reaction to Susan's behavior is relevant for (55). This is not so obvious: after all, the generalization about the nature of coin-tossing should also be relevant when evaluating (55) 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. What is still 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), that is, we assume 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 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 (i.e. the ϕ-worlds). 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?]. 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 (59) true but (60) false.

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

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

Take (59): since the discourse is subject to Relevance, the CQUD will be understood to be *If Susan had bet on tails and then pressed the button, would you have been so upset?* 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. In (60), Relevance requires that the CQUD be *If Susan had bet on tails and then pressed the button, what would have* *come up?* Non-Triviality will remove the proposition that heads came up (and not the proposition that you got so upset), and the counterfactual correctly comes out false.

Viewing the utterance of a counterfactual as an answer to a CQUD plays a crucial role in our account of counterfactuals like (59) and (60) 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 *choose* the CQUD between two conditional questions which could equally well be under discussion in the context of utterance. There are two (related) things to note here. First, it could be that the conditional question to which the counterfactual we are evaluating is a relevant answer could not be under discussion given some things we assume to be true. In this case, we predict that the counterfactual will come out as either trivially true or trivially false. We will look at cases that fall into this category in section 7. Second, there is a strong connection between CQUD and non-accidental generalizations, reminiscent of Kratzer's connection between non-accidental generalizations and confirming propositions. In section 7 I will propose to capture this connection formally by means of what I will call the *Matching Conditions on CQUDs*.

6 King Ludwig of Bavaria

Recall Kratzer's King Ludwig of Bavaria example given above in (10-a). 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.

(61) 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 (61-a) but not (61-b). Let us begin with (61-a). Intuitively, we would like to say that, since the context has established that the position of the flag is one factor in determining the location of the king, changing the status of the flag raises the issue of where the king is (either in the castle or away). Let us assume that the CQUD in the context at the time of utterance is (62).[4]

[4]We will go back to the mechanism responsible for selecting the appropriate CQUD later in this section.

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

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 (63).

(63) 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 (63-b) is ruled out by Consistency and, crucially, proposition (63-c) is ruled out by Non-Triviality since it entails a member of the Q-alternative set of the question under discussion. Both (63-a) and (63-d) 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 nonaccidental generalization holds.

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

Compare (61-a) to (61-b). Because the QUD is determined by the context of utterance, and since (61-b) is uttered in the same context as (61-a), the CQUD is going to be the same as before, i.e. (62). Recall that according to (33), how we constrain similarity is determined by the counterfactual antecedent and the CQUD. This means that, just like in the case of (61-a), Consistency will rule out (63-b) and Non-Triviality will rule out (63-c). As a result, all worlds selected by the similarity function will be worlds where the lights are on, making the counterfactual in (61-b) false, as shown below.

(65) $sim_{\leq_{ntr_Q(con_\phi(A_{w_c}))}}(\lambda w.$ the flag is up in $w) \cap \{w :$ the lights are off in $w'\} = \emptyset$

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 (61-a) false and (61-b) true. 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.

There is another way to manipulate 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 our of the flag is down, we do not get the same truth-condiitonal judgments.

(66) 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 (61-a) is no longer judged true. In the next section, 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 address the important issue of what makes a conditional question under discussion and we will propose that whether a given conditional question can be under discussion depends on the non-accidental generalizations available in the context of utterance.

7 Non-accidental generalizations and CQUD

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.

(67) *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 (66-a) and (66-b) 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 (66-b), the counterfactual in (61-b) should come out true. This, however, is not what we observe: even in this context, (61-b) 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 –(66-b)– is a "less natural way of describing the regularity in the Kinf Ludwig case" (Kratzer (2012), p.146) and that there is a cognitive bias against it.

Our proposal offers a promising way of accounting for Kratzer's data and in particular the "insecure" judgments we have when interpreting (61-b) in the context of the generalization (66-b). Let us start with our intuition that (61-a) would be true when uttered in the context of the non-accidental generalization in (66-a). I am going to propose here that the way in which non-accidental generalizations play a crucial role in the selection of the appropriate premises is that the form of the non-accidental generalization determines the form of the CQUD, as indicated in (68). I will call this the *Matching Condition on the CQUD*.

(68) *Matching Condition on CQUD* (MC)

For a conditional question *if* ϕ*, Q?* to be under discussion in a given context c, there must be a non-accidental generalization of the form *whenever/if p, q* salient in c such that: (i) p entails ϕ, and (ii) q is a possible answer to Q.

What MC requires is that there is a formal relation between the CQUD and a non-accidental generalization salient in the context. Only CQs that satisfy the MC will be under discussion.

Going back to our example, when evaluating the counterfactual in (61-a), repeated below in (69-a), in the context of (66-a), the MC will be satisfied if the CQUD is *if the flag were up, where would the king be?* since the restrictor of the adverb *whenever* entails the antecedent ϕ and its nuclear scope is a possible answer to the subordinated question *where would the king be?*

(69) 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 version of our proposal, which now has the MC as one of its parts, (69-a) comes out true, whereas (69-b) comes out trivially false, as explained above.

Now, suppose that the non-accidental generalization salient in the context is (66-b). MC requires that the restrictor of *whenever* entails the *if-clause* of the CQUD. However, we are also assuming that the counterfactual in (69-b) is a relevant answer to a CQUD and for this to be the case it must be the case that both the consequent and the modally subordinated question are interpreted relative to the same set of antecedent worlds. This means that, for the counterfactual to be relevant to the CQUD, the antecedent of the CQUD would have to be *if the flag were up*. The problem is that there is no logical relation between the proposition that the flag is up and the proposition that the king is away. Hence, the MC is violated. Note that if the MC is satisfied, the conditional is out because it violates Relevance. Similarly, note that if we interpreted (69-b) relative to (66-a), the counterfactual would be ruled out because either the MC or Relevance would be violated.

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 (66-b), without needing to get into the murky classification of non-accidental generalizations as "natural" or "less natural".

Before moving to the next section, let's quickly review the examples we have considered so far in light of the MC. The non-accidental generalizations in the weather examples were as shown below. (70-a) and (70-b) are relevant for Tichy's original example; (70-c) is relevant for Veltman's version.

(70) a. Whenever the weather is bad, Jones wears his hat.

b. If the weather is fine, Jones wears his hat at random.

c. Whenever the weather is fine and heads comes up, Jones wears his hat.

Given these generalizations, the conditional question *if the weather had been fine, would Jones wear his hat?* does satisfy the MC in both cases since in both (70-b) and (70-c) the restrictor of the generalization entails the proposition in the antecedent clause of the conditional question, and the nuclear scope is an answer to the subordinated question.

In most of the coin-tossing examples we considered above, the relevant non-accidental generalizations were (i) that all coin-tossing events are random and (ii) whenever you bet on something, you either win or lose. We judged (71) true, but (72) false (evaluated in a context where Peter presses a button in a random coin-tossing device and heads comes up).

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

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

Given the two generalizations above, the conditional questions that satisfy MC are (i) *If Susan had pressed the button, what would have come up?* and (ii) *If I had bet on heads, would I have lost?*. The former is the CQUD to which (72) is a relevant answer; the latter is the CQUD to which (71) is a relevant answer. If you now consider the variant in (73), the non-accidental generalization in this case is that whenever it is your turn to press the button and you bet on something, you either lose or win.

(73) 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's 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 tosses the coin 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 it had been your turn to press the button, you would have lost $10.

The conditional question that we identified above as the relevant CQUD, i.e. *If you had pressed the button, would you have lost?*, satisfies the MC and is the question to which the counterfactual in (73) is a relevant answer to.

Before concluding, consider the counterfactual in (74): Paula is buying a pound of apples and the Atlantic Ocean is not drying up. The example is repeated in (74).

(74) 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 (74) to be a felicitous (relevant) answer, the CQUD at the time of utterance would have to be something like (75).

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

The problem with (75) is that, given our world-knowledge, this conditional question cannot plausibly be under discussion. Supposing counterfactually that Paula is not buying a pound of apples does not raise the issue of whether the Atlantic Ocean will dry up. In other words, there is no non-accidental generalization such that it and (75) together satisfy MC. Of course, it might if our assumptions about the world were to change, in which case our judgment of (74) might change as well.

Similarly for (76), 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).

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

Just like (74), (76) is not only false but strange. For (76) to be a relevant answer, the CQUD at the time of utterance would have to be (77).

(77) 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, entertaining the counterfactual supposition that Clara is in this room does not raise the issue of what she is professionally. For (77) to be the CQUD, there would need to be a non-accidental generalization like "if someone is in this room, she is a painter" in the context of utterance, but clearly a geenralization of this kind is neither part of the context nor part of our world knowledge.

Examples like (74) 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-knoweldge there is no non-accidental generalization connecting these facts.

To conclude, according to the MC non-accidental generalizations play a central role in the interpretation of counterfactual conditionals, in line with what has been observed by many (Pollock (1976), Lewis (1979), Kratzer (1989), Kratzer (2012), just to cite a few).

8 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 (78).

(78) 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 (78) 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 constrains on similarity given in (34) are repeated in (79).

(79) **Constraining the similarity ordering $\leq_{A_{w_c}}$:**

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 $\phi \in \wp(W)$: $con_\phi(X) = \{p \in X : \phi \cap p \neq \emptyset\}$

(ii) for every $X \subseteq \wp(W)$ and question Q: $ntr_Q(X) = \{p \in X : \forall q \in Q\text{-alt}, p \not\subseteq q\}$

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 constrain, whereas the function *ntr* captures the Non-Triviality constrain. What is wrong with counterfactuals like (78) 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 (74) in section 7. In (74) the relevant question was not under discussion because no non-accidental generalization connected buying apples to oceans drying up; in (78), the relevant question is not under discussion because the antecedent already entails the answer to the question. In (78), 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 vacuously true. As for (74), *ntr* applied vacuously 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 (74) on the one hand, and (78) on the other, 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. (74)) or because the connection between the two is trivial (cf. (78)). 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 (79) 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 (80) and (81) (construed with a metaphysical modal base) are not truth conditionally equivalent.

(80) John has the flu (now).

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

Condoravdi proposes that the modal operator in (81) 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 (80) 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 (81) is ruled out and the only available reading is the epistemic one. If applied to examples like (78), 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, nothing is stipulated: 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 (81)) 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.[5] This point is important for the semantics of counteractuals. 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. Let's then say that an assertion is trivially true relative to a given context if either (i) it is a tautology; or (ii) it is entailed by a proposition already assumed in the context, whose parts are not themselves already entailed by the context.

Hence, the difference between informative countefactuals 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, i.e. the antecedent. Being trivially true because one implicit premise entailed by itself the consequent is precisely the

[5]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 I contend that this does not make the proposition that I take the umbrella trivially true because neither conjunct entails the proposition in question by itself.

problem that Non-Triviality, and in particular the function *ntr*, is designed to avoid.

9 "Symmetrical" cases

Before moving to the next section, I am going to look at some "symmetrical" counterfactuals and show how our proposal handles these cases. The two cases I am going to look at are given in (82) and (83) below.

(82) a. If New York had been in Georgia, New York would have been in the South.

b. If New York had been in Georgia, Georgia would have been in the North.

(83) a. If Verdi and Bizet had been compatriots, Verdi would have been French.

b. If Verdi and Bizet had been compatriots, Bizet would have been Italian.

In (82), we judge (82-a) true but (82-b) false. There are two ways of constructing the set of premises.

(84)

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

(85) a. New York 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 is in New York State. At this point, Non-Triviality requires that we remove any propositions entailing a possible answer to the CQUD. Given the non-accidental generalizations salient above and the MC, the CQUD is *If New York had been in Georgia, where would New York have been?*. Since the Q-alt of the subordinated question is {New York is in the North, New York is in the South}, Non-Triviality requires that we remove the proposition that New York is in the North. However, we do keep the proposition that New York State is in the North and the proposition that Georgia is in the South. Hence, we correctly predict that (82-a) is true but (82-b) is false.

Let's look at (83) now. If the CQUD is *If Verdi and Bizet were compatriots, what would Verdi and Bizet be?*, then the Q-alt of the subordinated question can be argued to be {Verdi is Italian; Verdi is French; Bizet is Italian; Bizet is French; Verdi and Bizet are Italian; Verdi and Bizet are French}. Non-Triviality will then remove the prepositions that Verdi is Italian and that Bizet is French.

(86)

Possibility (i):	**Possibility (ii)**:
(e) Verdi and Bizet are compatriots	(e) Verdi and Bizet are compatriots
(a) Verdi is Italian	(a) Bizet is French

As a result, both possibilities above will be ruled out and both counterfactuals in (83) will be false.

10 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 (87).

(87) 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 (87), a variant of an example in Bennett (2003), shows that the very same conditional in (87) is ambiguous between a causal and a non-causal interpretation.

(88) 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 (88) 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 (87). 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 asnwers 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 (88). Its causal interpretation is straightforward. Assuming that the discourse of which (88) 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 (88) 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 (89).

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

If (89) 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 (89) 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 (90) but less explicit about the goal of the conversation, which is to enable you to eat (in case you are hungry).

(90) 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 (91).

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

11 Conclusion

It is commonly held that counterfactuals are context-dependence in 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 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, to which the counterfactual conditional is supposed to be an answer. In other words, any proposition that entails a member of the Q-alternative set of the subordinated question under discussion cannot be part of the set of propositions that determines the similarity ordering. Assuming a model of discourse structure similar to the one proposed by Roberts (1996a), 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 supposed to answer a conditional question under discussion and therefore, for the discourse to be felicitous, it must be 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 nonaccidental 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. In particular, I have proposed what I have called a Matching Condition constraining the form of the CQUD: for a CQ *if* ϕ, *Q?* to be under discussion in a given context c, there must be a non-accidental generalization *whenever/if* p, q salient in c such that $p$entails ϕ and q is a possible answer to the subordinated question Q. This formal link between non-accidental generalizations and conditional questions constrains which conditional questions can be under discussion. This captures our intuition that non-accidental generalizations are crucial in identifying which issues are raised by entertaining a counterfactual assumption.

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.

One issue that this proposal raises is about the role of the CQUD in the interpretation of embedded counterfactuals as in (92) below.

(92) Mary thinks that if the whether had been nice, Jones would be wearing his hat.

In what context are we going to look for the relevant CQ and in what sense is this question under discussion?

I will conclude by stressing that the context-dependence of similarity 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 match 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 (1996b).

References

Arregui, Ana: 2009, 'On similarity in counterfactuals', *Linguistics and Philosophy* 245–278.

Bennett, Jonathan: 2003, *A Philosophical Guide to Conditionals*. Oxfrod University Press.

Condoravdi, Cleo: 2002, 'Temporal Interpretation of Modals', in D. Beaver, S. Kaufmann, and L. Casillas (eds.), *Stanford Papers on Semantics*. CSLI Publishers.

von Fintel, Kai: 2001, 'Counterfactuals in a Dynamic Context', in M. Kenstowicz (ed.), *Ken Hale: A Life in Language*, 123–152. MIT Press, Cambridge, MA.

Goodman, Nelson: 1947, 'The Problem of Counterfactual Conditionals', *Journal of Philosophy* **44**, 113–128.

Groenendijk, Jeroen: 1999, 'The logic of interrogation', in T. Matthews and D. L. Strolovitch (eds.), *Proceedings of SALT IX*, 109–126. CLC Publications.

Hamblin, C. L.: 1973, 'Questions in Montague English', *Foundations of Language* **10**, 41–53.

Heim, Irene: 1992, 'Presupposition Projection and the Semantics of Attitude Verbs', *Journal of Semantics* **9**, 183–221.

Ippolito, Michela: 2006, 'Semantic Composition and Presupposition Projection in Subjunctive Conditionals', *Linguistics and Philosophy* **29**, 631–672.

Ippolito, Michela: 2013, *Subjunctive Conditionals: a Linguistic Analysis.*, Linguistic Inquiry Monograph. MIT Press.

Isaacs, James and Kyle Rawlins: 2008, 'Conditional Questions', *Journal of Semantics* **25**, 269–319.

Kratzer, Angelika: 1989, 'An Investigation of the Lumps of Thought', *Linguistics and Philosophy* **12**, 607–653.

Kratzer, Angelika: 1991, 'Modality', in A. v. Stechow and D. Wunderlich (eds.), *Semantics. An International Handbook of Contemporary Research*, 639–650. De Gruyter, Berlin.

Kratzer, Angelika: 2012, *Modals and Conditionals*. Oxford University Press.

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

Lewis, David: 1979, 'Counterfactual Dependence and Time's Arrow', *Noûs* **13**, 455–476.

Pollock, J.: 1976, *Subjunctive Reasoning*. Reidel.

Ramsey, F. P.: 1929, 'General propositions and causality', in R. B. Braithwaite (ed.), *The foundations of Mathematics and other Logical Essays*, 237–255. Routledge & Kegan Paul, London.

Roberts, Craige: 1989, 'Modal subordination and pronominal anaphora in discourse', *Linguistics and Philosophy* **12**, 683–721.

Roberts, Craige: 1996a, 'Anaphora in Intensional Contexts', in S. Lappin (ed.), *The handbook of contemporary semantic theory*, 215–246. Blackwell, Oxford, UK.

Roberts, Craige: 1996b, 'Information Structure in Discourse: Towards an Integrated Formal Theory of Pragmatics'. The Ohio State University.

Rooth, Mats: 1985, *Association with focus*. GLSA Publications.

Tichy, Pavel: 1976, 'A counterexample to the Stalnaker-Lewis analysis of counterfactuals', *Philosophical Studies* **29**, 271–273.

Veltman, Frank: 2005, 'Making Counterfactual Assumptions', *Journal of Semantics* **22**, 159–180.

Wason, Peter: 1966, 'Reasoning', in H. Penguin (ed.), *New Horizons in Psychology 1*, 135–51.