Title
What a Clause Does: Raising Its Question and Answering It Too

Permalink
https://escholarship.org/uc/item/46d6f0b5

Author
Fan, Da

Publication Date
2021

Peer reviewed|Thesis/dissertation

What a Clause Does:
Raising Its Question and Answering It Too

By

DA FAN
DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Philosophy

in the

OFFICE OF GRADUATE STUDIES

of the

UNIVERSITY OF CALIFORNIA

DAVIS

Approved:

_____

Adam Sennet, Chair

_____

Cody Gilmore

_____

Rohan French

Committee in Charge

2021

i

Contents

What A Clause Does:

raising its question and answering it too

Abstract

The objective of this dissertation is to offer a novel pragmatic theory involving certain appropriateness conditions associated with embedded clauses. In short, the view is that an embedded clause raises a question and answers it too, in its local context. With respect to a local context, there are some conditions governing whether a question can be appropriately raised and whether an answer is an appropriate way of answering the question. The main task of this dissertation is to explore such pragmatic conditions on the relationship between embedded clauses, as question-answer contributors, and their local contexts. As the pragmatic theory is about what counts as appropriate questions and answers in local contexts, call it the *Good Question-Answer* view, or GQA for short.

   With this pragmatic theory, I am to apply it to various interesting phenomena invoked by the interaction between embedded clauses and their local contexts. In particular, I argue that GQA can provide an account of why linguistic devices such as focus, *it*-clefts, *rather-than* constructions, etc. are capable of affecting the truth-conditional interpretations of certain sentences, e.g. knowledge ascription sentences and counterfactuals. According to the account, such linguistic devices determine what question a clause raises. Consequently, they may bring about some accommodation effect on local contexts of embedded clauses. Whenever the truth-conditional interpretation of a sentence depends on such a local context, these question-regulating devices may have an effect on the interpretation.

   By proposing such a pragmatic account of the phenomena, we can resist the attempt to drastically change the standing semantic theories in response to the phenomena. For example, Schaffer's epistemic contrastivism, which takes "to know" as denoting a ternary relation, is partly motivated by the truth-conditional effect of question-regulating devices on knowledge ascription sentences. But if GQA is successful, such semantic revolutions

will turn out to be too hasty and unnecessary. Thus, my approach has a conservative bend: Given that a pragmatic account is available, some semantic revolutions will appear less attractive.

I will also apply GQA to some problems which have been thought of as semantic puzzles. It has been disputed whether the standing semantic theories of counterfactuals are correct in their commitment to the validity of some inference patterns. I argue that for some puzzles of this kind, it is hasty to take for granted that they are semantic puzzles and consequently to propose new semantic solutions. On the contrary, GQA is able to provide a pragmatic account of these puzzles, provided we take question-regulating devices, which are often ignored in the literature, into account.

Acknowledgments

Throughout the years of my Ph.D. study, I have received a lot of help and support from many people, without which I couldn't have completed this dissertation.

First, I would like to thank my advisor, Adam Sennet. In my first quarter at UC Davis, I took his seminars on the philosophy of language, which inspired me to start studying the topic of my dissertation. Since then, he has been guiding my research with his expertise, patience, and care. Repeatedly, it was his advice and encouragement that pulled me out of confusion and self-doubts. I am very grateful for having such a great mentor and nice friend.

I would also like to thank Nate Charlow, David Copp, Rohan French, and Cody Gilmore for their service on my dissertation committee. Their invaluable comments helped me improve the dissertation and also illuminate my future research.

I am indebted to other faculty members of the department as well, especially Hanti Lin, G.J. Mattey, and Robert May. They helped me learn not only how good philosophy should be done but also what good philosophers are like.

The completion of this dissertation would not be possible without the support of my fellow graduate students: Jordan Bell, Michael Bissonnette, I-Sen Chen, Tyrus Fisher, Natasha Haddal, Chanwoo Lee, Patrick Skeels, and S. Kaan Tabakci. It has been my privilege to be part of such a community of talented, supportive, and kind friends. I couldn't imagine surviving the program without our philosophical discussions in the department and fun times at bars.

Finally, words cannot express my gratitude to my wife, Yue San, and my mom, Xiang Ruan. I wouldn't have the courage to go through all these intense academic years without their unconditional support and love.

CHAPTER 1

Introduction

The objective of this dissertation is to propose and defend a view on the pragmatic effect of sentences and clauses. According to it, sentences and clauses, when used, do not only serve to contribute information. Rather, they play the role of first raising a certain question and then answering it. In a slogan, a sentence/clause raises its question and answers it too.

What question a sentence or clause raises is determined by some question triggering devices. A paradigm example is *focus*, which is realized by intonational stress. When a speaker utters a sentence, it comes with a certain intonation, a prominent aspect of which is that some words are stressed while others are not. For example, (1) and (2) are otherwise identical except that "marry" is stressed in the first but "Bertha" is stressed in the second. (Capital letters indicate stress.)

(1) Clyde MARRIED Bertha.

(2) Clyde married BERTHA.

Focus is a syntactic marking on certain components in a sentence. For now, we may assume that focused constituent in each of (1) and (2) is simply the word that is stressed.[1] The difference in stress indicates that the two sentences have different components being focused: the word *marry* is focused in (1) and *Bertha* in (2).

As the above examples involve different focused constituents, they raise different questions. In the literature on Question-Under-Discussion, focus has been regarded as a

---

[1] The relation between intonational stress (the highest pitch accent in a sentence or clause) and focus is complex, and in general, focus cannot be identified with the constituent that bears the stress. See Chapter 3 for discussion.

device that regulates what question a sentence is congruent with. In other words, the focus in a sentence signals which question can be appropriately answered with the sentence. Intuitively, as *marry* bears the focus in the sentence (1), the sentence can be used to appropriately answer the question *What did Clyde do to Bertha?*. Thus, when the sentence is uttered, it raises the question *What did Clyde do to Bertha?* and answers it with the proposition that Clyde married Bertha. In contrast, the question that can be appropriately answered by (2) is a different one, i.e. *Who did Clyde marry?*. Correspondingly, uttering (2) raises the question *Who did Clyde marry?* and answers it with the proposition that Clyde married Bertha.

Let us call the question that a sentence can appropriately answer the *Congruent Question* of the sentence. Thus, the question that a sentence/clause raises is its Congruent Question, and focus is a device that governs what the Congruent Question of a sentence/clause is.

Besides focus, there are other devices that play the same role of triggering Congruent Questions. For example, *rather-than* constructions and *it*-clefts both seem to regulate what question a sentence can appropriately answer. For example, the *rather-than* construction in (3) and the *it*-cleft in (4) both signal that the Congruent Question is *Who did Clyde marry?*.

(3) Clyde married Bertha rather than Beth.

(4) It was Bertha that Clyde married.

By taking Congruence Questions into account, the current theory treats an utterance of a sentence as bringing about its pragmatic effect in two steps. First, by raising its Congruent Question, the uttered sentence specifies a set of alternative ways of updating the context. For example, uttering (2) raises the question *Who did Clyde marry?*. The question can be answered by the following propositions: that Clyde married Alice, that Clyde married Bertha, etc.. Thus, raising the question in effect specifies a set of alternative ways that the context is to be updated: the proposition that Clyde married Alice can be added to it; or alternatively, the proposition that Clyde married Bertha can be added, etc..

2

In a nutshell, an utterance of a sentence first brings about the first pragmatic effect of signaling a set of alternative updates, which is determined by the Congruent Question.

Moreover, an utterance does not only raise a question; it answers the question as well. In other words, with a set of alternative updates specified, the utterance picks out a particular one and executes it so that the context is updated in this particular way. Continuing the example of (2), an utterance of it in effect chooses a particular answer to its Congruent Question, i.e. that Clyde married Bertha. In other words, among the alternative ways of updating the context, the utterance is to update the context by adding to the context the proposition that Clyde married Bertha.

In sum, this is what I mean by the slogan, that an uttered sentence raises a question and answers it too: an utterance of a sentence brings about the pragmatic effect of specifying certain alternative ways of updating the context and then picking out one of them to execute the update.

The question-raising role of utterances has long been noticed in the Question-Under-Discussion literature. According to it, the relation between the Congruent Question of an uttered sentence and the context of utterance is governed by some pragmatic principles. When such principles are not outright satisfied, an accommodation effect may ensue. For example, a pragmatic principle requires that when a sentence is uttered, the Congruent Question of it should be the Question-Under-Discussion in the current context—that is, the question that the relevant speakers are committed to answering in the discourse. Once the Congruent Question of the sentence and the context do not relate in this way, the Congruent Question may be accommodated into the context as the Question-Under-Discussion for the sake of felicity. For example, if the Question-Under-Discussion in a context is not *Who did Clyde marry?* but (2) is nonetheless uttered, within some constraints, the relevant speakers can treat the question as the Question-Under-Discussion in order to save the felicity of the utterance.

However, such a pragmatic theory so far only considers utterances of entire sentences as playing the question-answering role. The objective of the current research is to expand the scope of the view so that embedded clauses are treated in much the same way. In particular, I argue that an embedded clause brings about similar effects to its local context as what an entire sentence does to its global context. For an utterance of an entire sentence, it raises a question—its Congruent Question—against its context of utterance, specifying certain alternative ways of updating this global context. Correspondingly, I contend that an embedded clause, when used, raises its Congruent Question against its local context, specifying a set of alternative ways in which the local context can be updated.

Questions raised locally by embedded clauses are able to bring about certain accommodation effects as well. Here we take the local context of an embedded clause to be a body of information that is to be updated by the clause in the process of interpreting a whole utterance. For a complex sentence that embeds a clause, when it is uttered, the interpreter is to calculate how the embedded clause is to update its local context, generating a subsequent body of information based upon which an interpretation of the entire utterance can be determined. According to the current view, this local update involves two steps: the embedded clause first raises a question—its Congruent Question, specifying a set of alternative ways of updating the local context, and then picks out a particular one to execute. In other words, what an embedded clause contributes to its local context is a question-answer pair.

Much like the pragmatic principles governing how an utterance of an entire sentence should be related to its global context, there are pragmatic principles governing the relationship between embedded clauses and their local contexts. In a general form, the core pragmatic principle in the current view can be stated as follows: for an embedded clause, its Congruent Question should be a *good question* with respect to the local context, and the propositional content of it should be a *good answer* to the question with respect

4

to the local context. What count as good questions and good answers in contexts will be explained later, as there are various conditions on whether a question or answer is pragmatically appropriate relative to a context. Once such conditions are clear, we will see that the pragmatic principle enables some effect of accommodation. Roughly, when the Congruent Question of an embedded clause does not satisfy the good question conditions, or when the proposition expressed by the clause is not a good answer to the question in the local context, the interpreter may adjust the local context of the embedded clause, making the question-answer pair offered by the clause a good question and a good answer relative to the post-accommodation local context.

As what sits in the center of the current view is a pragmatic principle that requires embedded clauses to raise good questions and provide good answers in their local contexts, let us call it the Good Question-Answer view, or GQA for short.

In this dissertation, I will justify GQA from three perspectives. First, it is motivated from a theoretical perspective. In particular, I take GQA to be a natural step forward from two antecedently established frameworks in the studies of pragmatics. On the one hand, in studying the pragmatic effect of utterances of entire sentences on their global contexts, the Question-Under-Discussion approach, in essence, treats such utterances as playing a question-answering role. In other words, an utterance of an entire sentence raises a question and answers it too in its global context. On the other hand, the literature on presupposition projection has long been using local contexts as contexts of embedded clauses. For an utterance of a complex sentence, the overall pragmatic effect of the utterance is calculated based on how embedded clauses affect their respective local contexts, and an embedded clause is treated as updating its local context in much the same way as how a stand-alone utterance of an entire sentence updates its global context. Hence, if one accepts that an utterance of an entire sentence raises a question and answers it too in a global context, and if she also accepts that there is a plausible notion of local context that

captures what is subject to the update brought about by embedded clauses, it seems a very natural step forward to accept that an embedded clause raises a question and answers it too *in its local context*. If so, there is an issue of whether the question and the answer are pragmatically appropriate in the local context. GQA, then, provides some pragmatic principles governing this relation between question-answer pairs and local contexts. In Chapter 2, I will discuss in detail the parallel between utterances of entire sentences and their global contexts on the one hand, and embedded clauses and their local context on the other, with the goal of conceptually motivating GQA as a natural extension of the aforementioned, well-received frameworks.

Second, GQA is to be supported by its capability of explaining the truth-conditional effect of focus and other Congruent Question triggers. In examples of simple sentences, such as (1) and (2), focus does not affect truth-condition, although it determines what question such sentences can appropriately answer. However, in some intriguing complex sentences, focus involved in an embedded clause does seem to make a difference in what truth-conditional content a sentence expresses. For example, the truth-condition of a counterfactual may rely on how its constituents are focused.

[Marriage] Clyde and his girlfriend Bertha don't like close relationships. They see each other only twice a year, and they don't want to get married, etc.. However, Clyde found that if he were to get married soon, he would inherit a great amount of money. Then he married Bertha and got the money, while expecting that their 'loose' relationship would continue. (Adapted from Dretske, 1972.)

(5) If Clyde hadn't MARRIED Bertha, he wouldn't have inherited the money.

(6) If Clyde hadn't married BERTHA, he wouldn't have inherited the money.

When uttered in the given context, (5) sounds true while (6) false, and thus the pair of counterfactuals, differing from each other in focus only, have different truth-conditional contents. This looks mysterious. Focus does not seem to affect the propositional content

of the relevant embedded clauses (i.e. the antecedents in the above example), and the propositional content of a counterfactual should be a function of the contents of the constituents. Then, how can focus make a difference to the truth-conditional interpretation of a counterfactual?

Besides counterfactuals, the same phenomenon also appears in knowledge-*that* ascription sentences.

[Jewel Thief] Last night, Peter robbed the jewelry store. He smashed the window, forced open the locked safe, and stole the rubies inside. But Peter forgot to wear gloves. He also forgot about the security camera. Today, Mary the detective has been called to the scene to investigate. So far she has the following evidence. She has been told that there was a theft, she has found and identified Peter's fingerprints on the safe, and she has seen and recognized Peter on the security video, filmed in the act of forcing open the safe. She has no further information. (Schaffer and Knobe, 2012)

(7) Mary knows that PETER stole the rubies.

(8) Mary knows that Peter stole THE RUBIES.

Again, this minimal pair illustrates the truth-conditional effect: in the given context, (7) sounds true while (8) sounds false.

I will argue that GQA is able to explain this truth-conditional effect of focus. In short, the central contention of GQA is that an embedded clause has to relate to its local context in such a way that relative to the local context, the Congruent Question is a good question and the propositional content of the clause is a good answer to the question. If two embedded clauses are otherwise identical except that they involve different focus patterns, the Congruent Questions of them are different. Thus, their respective local contexts may have to satisfy different conditions in order to make question-answer pairs to be good questions and answers with respect to the local contexts. In other words, when two embedded clauses have different Congruent Questions, the pragmatic principle of

7

GQA enables different resolutions of what their local contexts are like. With more details explained later, this account is able to explain why (5) and (6), as well as (7) and (8), seem to have different truth-values. In such examples, the different Congruent Questions triggered by the relevant embedded clauses invoke different resolutions of local contexts. As local contexts matter for the truth-conditional interpretations of counterfactuals and knowledge ascription sentences, a truth-conditional effect is expected.

In Chapter 3, I will formulate the GQA theory in detail, explaining what counts as good questions and good answers in local contexts. Hence, the theory is utilized in Chapters 4 and 5 to account for the truth-conditional effect. In particular, Chapter 4 concerns the truth-conditional effect on knowledge ascription sentences. I argue that the pragmatic theory of GQA provides a satisfactory account of why truth-conditional interpretations of knowledge ascription sentences seem to vary according to focus and other Congruent Question triggers. Thereafter, in Chapter 5, the GQA theory is applied to counterfactuals; we will see how the theory is able to explain the focus-invoked difference between pairs of counterfactuals such as (5) and (6).

The third point of support for GQA is that it does not only explain apparent examples of the truth-conditional effect of focus and other Congruent Question triggers, but is also able to solve some puzzles that have been thought of as semantic conundrums. For example, in the literature on the semantics of counterfactuals, it has been disputed whether Conjunction Conditionalization (CC: $\varphi \wedge \psi \vdash \varphi > \psi$, where $>$ stands for counterfactual conditional) is a valid inference.

The problem is that the validity of CC entails that any counterfactual with a true antecedent and a true consequent is invariably true, while there are examples showing that ordinary speakers tend to judge some true-true counterfactuals to be false. Thus, the opponents of CC appeal to such counterexamples to argue that some received semantic theories for counterfactuals, which are committed to the validity of CC, have to be revised.

8

However, as I will argue in Chapter 6, the problem should not be taken for granted as a semantic puzzle. On the contrary, the GQA theory is able to explain away the counterexamples in terms of the pragmatic effect of Congruent Questions. If it is correct, the purported counterexample to CC is not a genuine challenge to the received semantic theories that validate CC. In particular, I will argue that some purported counterexamples of this sort crucially relies on certain focus pattern involved in the antecedents, which is oftentimes omitted in the written form of counterfactuals. Once we start to take focus into account, I argue, the pragmatic theory of GQA can provide a viable solution to the problem of CC, showing that the puzzle disputed in the semantic literature may deserve a pragmatic solution.

In a similar vein, I also apply GQA to another seemingly semantic puzzle. On the one hand, many semantic theories do not validate the inference of Simplification of Disjunctive Antecedent (SDA: $\varphi \vee \psi > \chi \vdash (\varphi > \chi) \wedge (\psi > \chi)$), since it is in conflict with some other very plausible semantic principles. On the other hand, in many examples, SDA seems to be an intuitively plausible inference, resulting in some disputes on whether SDA should be validated at the cost of other semantic principles. Regarding this issue, my view is that Congruent Questions involved in such examples have been ignored in the previous discussions so that the puzzle is taken for granted as a semantic one. However, once we identify the relevant Congruent Questions and apply the pragmatic theory of GQA to these examples, a pragmatic account of why SDA seems appealing will emerge.

Arguments of this sort are intended to show that GQA is not merely a pragmatic theory that is designed for explaining the cases where the focus (or any other Congruent Question trigger) is made explicit. Rather, it is a pragmatic theory that is powerful enough to apply to a wider scope, solving problems that have been perceived as semantic puzzles. Therefore, I take this as a point of support for GQA which demonstrates the explanatory power of the theory.

CHAPTER 2

Questions in Local Contexts

A declarative sentence is first and foremost a vehicle of information. That is, when asserted, its primary use is to contribute a certain piece of information to a discourse. However, as the literature on *Questions-under-discussion* suggests, a declarative sentence is associated with a question it is congruent with, such that asserting the sentence is not merely an act of contributing information, but rather an act of contributing certain information *as an answer to the question*. Although this question aspect of a sentence does not directly take part in determining what information the sentence is to deliver, it nonetheless plays a role in whether the sentence is appropriately used in a context. In a nutshell, for a sentence to be appropriately used, the question it is congruent with has to be the question-under-discussion in the context, or a question that can be accommodated, under some pragmatic constraints, as the question-under-discussion. In a sense, we can take the pragmatic effect of uttering a sentence as i) raising a question, and ii) contributing information by answering the question raised if the appropriateness conditions are met. In a slogan: a sentence raises a question and answers it too.

The current study amounts to a step forward from this point: a core thesis I am advocating in this dissertation is that a sentence contributes information by answering a question even when it is embedded as a *sub*-sentential clause. In particular, when a sentence is uttered, its appropriateness of use partially relies on whether each embedded clause involved, if any, is appropriately related to its *local context*, which in turn relies on whether the question congruent with the clause is suitable to raise in the local context. In

parallel with the question aspect of entire sentences that is noticed by Question-Under-Discussion researches, I propose that a sub-sentential clause also serves to i) raise a question against the local context of the clause, and ii) contribute information to the local context by answering the question, provided all the appropriateness conditions are met. In a slogan: a clause (locally) raises a question and answers it too.

The current view, then, is a pragmatic theory that works at the sub-sentential level, concerning how the question aspect of clauses affects the appropriateness of use in local contexts. This might sound alarming, in that it is not clear in what sense a clause is "used" in a local context. For the very notion of local context, according to some, is not a context of use in the regular sense, but rather is a step in calculating the semantic value of a complex sentence. It follows that at the very best what I mean by "appropriateness" of embedded clauses is a semantic notion, which breaks the parallel between sentences and embedded clauses regarding the effect of their question aspect. Therefore, before explaining the details about the pragmatics of clauses and the questions they locally raise, we need some conceptual clarifications first on the very notion of local context and how a pragmatic theory of locally raised questions is possible.

For that purpose, in this chapter, I will explain some key notions by reviewing and examining three theoretical foundations that my own view is built upon, in order to motivate my seemingly bold thesis from a theoretical perspective. The first foundation, briefly reviewed in §2.1, is Stalnaker's model of discourse dynamics, which is the underlying framework of representing context, appropriateness of use, and informational contribution made by sentences. The next three sections include discussions on what can be taken as two extensions of Stalnaker's model. On the one hand, studies on presupposition projection apply Stalnaker's model to the sub-sentential level: local context is represented in the same way as global context is on Stalnaker's model, and the interaction between clauses and local contexts are mostly the same as that between sentences and

11

global contexts. Although there are multiple different conceptions of local context at work in the presupposition literature, I will show that there is a plausible one that can serve my purpose of giving a pragmatic theory of local questions. On the other hand, Question-Under-Discussion theories extend Stalnaker's model in another direction, in the sense that a sentence, when uttered, is not a mere information contributor but rather contributes information by answering a question. I will briefly review the main spirit of this view; for, as mentioned above, it constitutes the starting point of my own theory. Local context is introduced in §2.2; and in §2.3, I will clarify the pragmatic concept of local context I adopt. Thereafter, the QUD approach is reviewed in §2.4. If both extensions of Stalnaker's model are acceptable, I argue, it will be natural to make a further move that combines them. §2.5 introduces and clarifies the way I combine the two approaches.

2.1. Stalnaker's Model

The pragmatic effect of utterances of sentences concerns the interaction between utterances and sentences. When a sentence is uttered in a context, it has both upstream and downstream effects. The downstream effect is about how the utterance makes changes to the context, and the upstream effect is governed by some appropriateness conditions: when such conditions are met, a sentence is appropriate to utter in a given context and thus brings about the downstream effect; otherwise, either it is judged as infelicitous, or accommodation takes place so that the context of utterance is changed in order to satisfy the appropriateness conditions. Stalnaker offers a model of both upstream and downstream effects. According to this model, contexts are understood as bodies of information and sentences have informational content, and hence uttering a sentence in a context is to contribute the piece of information encoded in the sentence to the body of information encoded in the context. First, let us take a look at the context side of the model.

The core notion in Stalnaker's model of context is *common ground*. For a group G of speakers, the common ground is a set of propositions, such that all the propositions in

this common ground are accepted by all the speakers in G, everyone in G believes that everyone in G accepts those propositions, and everyone in G believes that everyone in G believes so, etc..[1]

A key insight provided by this Stalnakerian view is that context is to be represented as a body of information. In this sense, as the common ground is a body of information a certain group of speakers share, the body of information encoded in a context is just the information in the common ground.[2] Given this informational nature of context, a formal representation of context emerges from a now routine possible worlds representation of information: informational objects are represented as sets of possible worlds, following Stalnaker (1978) and a tradition that dates back to the mid-20th century (e.g. Hintikka, 1962).

> **Context.** A context is represented as a context set c, which is a set of possible worlds such that for any possible world w, $w \in c$ $w \in p$ for every proposition p in the common ground of the context.

---

[1]The notion of acceptance should not be taken as belief. According to Stalnaker, "[a]cceptance, as I have used the term is a category of propositional attitudes and methodological stances toward a proposition, a category that includes belief, but also some attitudes (presumption, assumption, acceptance for the purposes of an argument or an inquiry) that contrast with belief, and with each other. To accept a proposition is to treat it as true for some reason. One ignores, at least temporarily, and perhaps in a limited context, the possibility that it is false."(Stalnaker, 2002, 716) A speaker could accept a proposition without believing it and the proposition could still be part of the common ground. By using this broader notion of acceptance, for example, an atheist can participate in a conversation about God that requires the proposition that God exists to be part of the common ground, insofar as she temporarily takes that proposition for granted for the purposes of the conversation.

[2] Strictly speaking, Stalnaker himself does not equate a context with common ground. In Stalnaker (1978), what underlies a discourse is "speaker's presupposition", i.e. what is *believed* by a speaker to be the common ground. But here I take context to be simply modeled by the actual common ground of a group of speakers, for two reasons. First, as we are interested in how communication *works* and only cases of successful communication matter in that regard, we can ignore differences between speakers' beliefs about their actual common ground. In other words, we can focus on what Stalnaker calls "*nondefective*" contexts, contexts in which speakers share the same beliefs of the common ground (Stalnaker, 1978, 332). Second, as Gauker (1998) and Simons (2003) observe, many theorists do treat context as actual common ground. Although this deviation from Stalnaker's own view makes a difference in how presupposition accommodation works, it matters less for my purposes of justifying the non-informational aspect of local context.

In other words, a context set is an information state that rules in every possible way of making all the propositions in the common ground collectively true while ruling out every possible way of making at least one of them false. Following Stalnaker, we may call each possible world in a context set a "live option" for the discourse participants—it is a way that the actual world could be given the participants' shared information.

With this informational model of context in hands, utterances, which serve to change the context, are taken accordingly as making informational contributions to the common ground: as sentences are uttered, new information is added into the context, and thus the body of information grows. Characteristically, then, the downstream effect of an utterance of a sentence is to perform an informational update on the context.

> **Update.** For a context c and a sentence that expresses proposition p, when the sentence is uttered in the context and accepted by the relevant speakers, then in the subsequent context c', the context set c contains all possible worlds in c where p is true. Formally, $c' = c \cap p$.

In other words, adding a piece of information amounts to ruling out some live options of what the world is like—adding p rules out all the possible ways that p is false.

Besides, utterances of sentences can have *upstream* effects, too. That is, when a sentence is uttered, the sentence may carry some *appropriateness conditions* that have to be satisfied by its context of utterance—the context has to be such-and-so in order for the sentence to be felicitously uttered in it. If a sentence is uttered in a context that does not satisfy its appropriateness conditions, it either results in an infelicitous utterance and thus a case of communication failure, or the context is adjusted to the effect that the appropriateness conditions are satisfied. In the latter case, it has an upstream effect: the context in which the sentence is uttered is changed, for the sake of saving appropriateness.

The paradigm case of this upstream effect is presupposition accommodation. For a presupposition-carrying sentence to be felicitously uttered, the common ground has to

14

entail the presupposition of the sentence. For example, if it has not been established in the common ground that I have a dog, (9) by default cannot be appropriately uttered.

(9) I have to take my dog to a vet.

But when it is part of the common ground that I have a dog, uttering (9) sounds fine. When this occurs, we say that the presupposition is satisfied by the context. Thus, presupposition satisfaction is an appropriateness condition of presupposition-carrying sentences.

> **Presupposition Satisfaction.** For a context set c and a sentence that presupposes proposition p, the sentence can be felicitously uttered in the context only if c entails p.[3] Formally, $c \subseteq p$.

Given this condition, an upstream effect of *presupposition accommodation* is what occurs when the context is adjusted in order to to make the condition met.

> **Presupposition Accommodation.** For a the context set c and a sentence that presupposes proposition p where c does not entail p, when the sentence is uttered in the context, p may be accommodated as part of the context of utterance so that p is entailed by it. I.e., the post-accommodation context set $c' = c \cap p$.

Like the downstream effect, this upstream effect is also informational. It amounts to adding a piece of information, i.e. the presupposition in question, into the context, in order to save appropriateness.

At this point, let us take this model as a satisfactory representation of how the informational aspect of discourse dynamics works. That is, let us suppose that possible world semantics correctly captures informational contents of sentences and information states of discourse, that the satisfaction of sentence presuppositions is indeed determined by speakers' information states, and that adding a piece of information is a matter of ruling

---

[3]This formulation appears to take the notion of sentence presupposition as primitive, which deviates from Stalnaker's own view (1974) that sentence presupposition is derivative and can be defined based on speakers' presupposition. I take no stance on whether sentence presupposition should be primitive or be explained in pragmatics. This felicity condition is meant to capture the Stalnakerian understanding of what in the context is supposed to satisfy presuppositions, regardless of how presuppositions are triggered.

out some extant live options of what the world is like.[4] However, this model is limited in two ways. First, as will be clear in §2.4, interactions between utterances and contexts, both upstream and downstream, involves a non-informational aspect: utterances have to address the question-under-discussion in the context, and they may add or remove question-under-discussions. Second, for a complex sentence in which a presupposition-carrying clause is embedded, this model does not predict what the context has to be like in order to satisfy the presupposition. To address this last problem, we need to broaden the notion of context to include so-called local contexts. We now turn to this point.

2.2. Introducing Local Contexts

Previously in Stalnaker's informational representation of context, contexts are always *global*, in the sense that they are contexts *of entire utterances*: they are responsible for satisfying appropriateness conditions of the entire sentence that is uttered, and they are to be updated by the utterance of the entire sentence. The extension we are now considering is the introduction of *local contexts*, which are contexts *for embedded clauses*. This modification is by and large motivated by the problem of presupposition projection.

This *projection problem of presupposition*, dubbed by Langendoen and Savin (1971), is about how the presupposition of a complex sentence is determined by the presuppositions of its parts. Or in other words, how a complex sentence "inherits" the presuppositions of the presupposition-carrying clauses it embeds. As observed by Karttunen (1973), clause-embedding connectives or predicates display certain systematic patterns of presupposition inheritance behaviors. For example, internal negation directly inherits the presupposition of its embedded clause: (11) has the same presupposition p as (10).

(10) The king of France is bald.

---

[4]Of course, some tinkering might be worth having. For example, there appear to be more appropriateness conditions other than Presupposition Satisfaction. By uttering a sentence, the proposition it expresses must not contradict, or be entailed by, the common ground, because otherwise it would result in an absurd informational state or adds no new information to the current information state, each of which suffices for the infelicity of the utterance. See Stalnaker (1978).

presupposes: *p* France has a king.

(11) The king of France isn't bald.

Such operators are called presupposition *holes* by Karttunen (1973). In contrast, some other operators block the presupposition of embedded clauses from projecting, and are called presupposition *plugs*. As illustrated below, the indirect speech report operator is a plug, as (12) does not presuppose *p*.

(12) Luis said that the king of France was bald.

Besides these two categories, some other operators are more intriguing. For example, a conjunction directly inherits the presupposition of its first conjunct, as shown in (13), so it seems correct to say that the first argument place of conjunction is a presupposition hole. However, looking at (14) and (15), whether the presupposition of the second conjunct projects seems to depend on what the first conjunct is—it projects in (14) but not (15).

(13) The king of France is bald and today is cold.

(14) Today is cold and the king of France is bald.

(15) France has a king and the king of France is bald.

As Karttunen observes, the presupposition of the second conjunct projects unless it is entailed by the first conjunct. Such operators (or more accurately, such argument places of operators) are called presupposition *filters*, in the sense that the first conjunct "filters out" the presupposition of the second conjunct by entailment. A similar inheritance behavior is also found in conditionals: (16) inherits the presupposition of the antecedent, but when the clause is in the consequent place, the presupposition projects unless it is entailed by the antecedent.

(16) If the king of France is a bald, today is cold.

(17) If today is cold, the king of France is bald.

(18) If France has a king, the king of France is bald.

From the above examples, we obtain an arguably robust generalization of how presupposition projects in conjunctions and conditionals.

> For *A and B* or *If A then B*, it inherits
>
> > i) all the presuppositions of *A*, and
> >
> > ii) all the presuppositions of *B* except those which are contextually entailed by *A*.[5]

So far we have a fairly acceptable *description* of the projection behavior concerning some operators, but now we have to face the question of *why* these operators display such inheritance behaviors. This explanatory challenge can be met by the Stalnaker-Karttunen-Heim approach (Heim, 1983; Karttunen, 1974; Stalnaker, 1974), which crucially relies on two assumptions. First, as briefly mentioned in §2.1, presupposition is understood as a requirement on the context. Second, a presupposition-carrying clause embedded under an operator has to be satisfied by its *local context*. Now we turn to these two assumptions in turn.

In §2.1 we have seen that a key idea in the Stalnakerian pragmatics is that the information encoded in a context is what is taken for granted by speakers. For Stalnaker, presupposition is primarily a relation between speakers and propositions, and what a speaker presupposes is just what she believes to be the common ground—what she believes to be accepted by all the relevant speakers, to be believed by every speaker to be what is accepted by all, etc.. In an ideal circumstance (i.e. a non-defective context; see fn.2) where participants make no mistakes in what the common ground is, speakers' presupposition is just the actual common ground. Here we can see that presupposition in this

---

[5]The relation has to admit cases of *contextual* entailment in order to accommodate examples such as

> Geraldine is a Mormon and she has given up wearing her holy underwear.

The second conjunct presupposes that Geraldine has worn holy underwear, and this presupposition is not logically entailed by the first conjunct. But in a context where the speakers take for granted that all Mormons have worn holy underwear, it seems that the first conjunct does filter out the presupposition of the second. Hence the claim that in a conjunction, a presupposition of the second conjunct is filtered out if it is entailed by the first conjunct *together with the information encoded in the context*.

sense is just a kind of proposition attitude that speakers hold towards certain information, rather than something that a sentence carries. Indeed, Stalnaker (1973) treats the so-called *sentence presupposition* as derivative:

> ...we may say that a sentence has a presupposition in a derivative sense just in case the use of that sentence would for some reason normally be inappropriate unless the speaker presupposed a particular proposition. In such a case, I will say that a sentence requires a presupposition. This notion of presupposition *requirement* will be the explication of the linguists' notion of presupposition. (Stalnaker, 1973, 451)

For example, (10) carries its presupposition p in the sense that normally it would be inappropriate to utter in a context if p is not presupposed, in the primary sense, by the relevant speakers. In other words, an appropriate utterance of (10) *requires* a context where p is entailed by the common ground, and we say that a context *satisfies* the presupposition of a sentence when the requirement is met. As we have seen in §2.1, presupposition, understood in this way, is an upstream appropriateness condition imposed on the context of utterance.

Given that a presupposition-carrying sentence requires its context of utterance to entail the presupposition, we now ask what happens when the sentence is embedded. In other words, when embedded, on which context the clause imposes its presupposition requirement? Karttunen (1974) provides an answer to this question, by which the projection behaviors of the aforementioned operators are explained. On his view, when a complex sentence is uttered in a context, the relevant clause-embedding operator creates a *local context* for its embedded clause, which stands in a certain relation to the global context. In this sense, an embedded clause has a context of its own, which might not be identical to the *global* context in which the whole sentence is uttered. Moreover, the presupposition requirement of an embedded clause is imposed on its local context, and these requirements on local contexts, in turn, determine what the whole sentence requires

of its context of utterance: the global context of an utterance of a complex sentence must be such that every embedded clause has its presupposition satisfied in its local context.

Karttunen also specifies how local contexts are determined in the cases of logical connectives. Suppose a complex sentence is uttered in a context c, then

- In *Not A*, the local context of A is c;

- In *A and B* or *If A then B*, the local context of A is c, and that of B is $c \cap a$, where a is the proposition expressed by A.

- In *A or B*, the local context of A is c, and that of B is $c \cap \neg a$, where a is the proposition expressed by A.[6]

Given this way of determining local context, the presupposition of a complex sentence—that is, what it requires of its global context—can be calculated based on the embedded clauses' requirements on their respective local contexts. Consider *A and B*. The presupposition of A projects directly because the local context of A is identical to the global context. The presupposition of B is filtered out when entailed by A, because the global context itself does not have to entail the presupposition in such cases. The inheritance properties of other connectives are explained similarly.

As we can see, the crux of this account is that embedded clauses now have their own local contexts, on which the clauses' appropriateness conditions are imposed. Insofar as we can specify, for each embedding operator, how local contexts are derived from global contexts, we are able to infer how presuppositions project in complex sentences. To summarize, evidence for positing local contexts for embedded clauses that are distinct from global contexts of utterance comes from the local effect of presupposition satisfaction, or, say, the upstream effect of embedded clauses: the presupposition-invoked appropriateness conditions of a clause are to be satisfied locally by the clause's own local context.

---

[6]For simplicity, I just leave the case of disjunction unexplained. But it is easy to see that this local context rule predicts that in a disjunction, the first disjunct filters out a presupposition of the second disjunct just in case *its negation* contextually entails the presupposition, which fits the projection behaviors we observe with disjunctions. See Karttunen (1974).

Similar to the presupposition accommodation effect discussed before, when a local context does not satisfy the presupposition of a clause, adjustments can be made in order to save appropriateness. Note that the local context for an embedded clause is generated based on i) the global context of the whole utterance and ii) the embedding environment of the clause, including an embedding operator and sometimes other arguments of the operator. For example, in *If A then B* uttered in c, the local context of B, i.e. $a \cap c$, is determined by the global context c, as well as the fact that B is the consequent of a conditional where the first argument expresses proposition a. Then, an embedding environment determines a specific mapping from global contexts to local contexts of the clause embedded in it. The mapping enables a way of adjusting the global context when a clause's presupposition is not locally satisfied: when such a local presupposition failure occurs, we can conservatively add a piece of information into the global context, to the effect that its corresponding local context satisfies the presupposition of the clause in question. To illustrate, note that the consequent (call it B and the antecedent A) presupposes p:

(19) *If A then B*: If Geraldine is a Mormon, she has given up wearing her holy under-wear.

   p: Geraldine used to wear holy underwear.

Let a be the propositions expressed by A. When the global context is c, the local context of B is $a \cap c$. Now suppose $a \cap c$ does not entail p; that is, this local context does not satisfy the presupposition of clause B. Accommodation can happen by adding a piece of information q to the global context c, where $\rightarrow$ stands for material conditional:

$a \rightarrow p$: Geraldine is a Mormon $\rightarrow$ she used to wear holy underwear

Call this new, post-accommodation context c'. Hence, we can see that the local context corresponding to c' now satisfies the presupposition of B, as $a \cap c'$ entails p.

As illustrated in this case, an embedded clause can have an upstream pragmatic effect on the global context through the requirement it imposes on its own local context.

> **Presupposition Accommodation*.** Let A be a clause embedded in S, where the embedding environment of A maps a global context c to a local context $L(c)$. Suppose S is uttered in c, such that $L(c)$ does not satisfy the presupposition of A. The global context can be changed into $c \cap p$ where p is such that:
>
> i) $L(c \cap p)$ satisfies the presupposition of A, and
>
> ii) no proposition weaker than p satisfies condition i).

On the flip side, embedded clauses have their downstream effect as well. Consider a conjunction *A and B* uttered in a context c. Given that the local context of A is identical to the global context c, the downstream effect of it is to update c by adding to it a piece of information, i.e. the propositional content of A, and the local context for B results. Thus, with our commitment to local contexts, we are committed to both upstream and downstream effects of embedded clauses: presupposition-based appropriate conditions of a clause have to be satisfied by the local context, and the local context is to be informationally updated by the clause, the result of which sometimes becomes the local context for another clause. The downstream effect of clauses then is not much different from Update in the Stalnakerian model, except now we also accept informational updates on local contexts.

> **Update*.** For a clause that expresses proposition p and its local context c, the effect of updating c with the clause is $c \cap p$, which, depending on the embedding environment, may serve to be the local context of another clause.

Upstream and downstream sides combined, the sense in which this approach is an extension of the Stalnakerian model is clear. On the Stalnakerian model, what a sentence does is to add a piece of information to its context of utterance, provided that the context satisfies all the appropriateness conditions. The current approach just generalizes the idea

22

to include embedded clauses and local contexts: what a clause does is to add a piece of information to its local context, provided that the context satisfies all the appropiateness conditions. The effects of embedded clauses and entire sentences are the same, in the sense that they have the same effect on *their own contexts*, global or local.

2.3. A Pragmatic Notion of Local Context

As mentioned at the beginning of this chapter, whether the upstream effect of clauses on their local contexts can be called a pragmatic effect, as opposed to part of the semantic calculation, depends on what concept of local context is at work. The objective of this section is to motivate a pragmatic notion of local context, in terms of which the effect of clauses on their local contexts can be properly called a pragmatic effect.

In the preceding section, I treat local contexts in parallel with Stalnaker's notion of global context, as if a local context is just like its global counterpart except that it is the context of an embedded clause while a global context is a context in which an entire sentence is uttered. But the two are dissimilar in an obvious way: the notion of local context cannot have a straightforward reductive account, as what global context has in Stalnaker's theory. The notion of global context, on Stalnaker's view, is grounded by common ground, which is in turn grounded by speakers' propositional attitude of acceptance. But this does not apply to local contexts. For example, for a conditional $p \rightarrow q$ uttered in c, it is hard to say what propositional attitude speakers hold towards the local context of q, namely, towards $c \cap p$. It cannot be acceptance, because speakers do not have to accept p, even if they accept the whole conditional.

Thus, the challenge is to explain what a local context really is, if it cannot be grounded in the same way as the Stalnakerian global context is. I will first briefly review two approaches to the nature of local contexts. According to the dynamic semantics approach, local context is part of the compositional semantics of complex sentences, and for a given global context and an embedding environment, the semantics completely determines

23

how the local context in the environment is derived from the global context. On the other hand, a pragmatic approach takes a local context to be generated based on pragmatic reasoning. A crucial difference between the two is that the former approach relies on a rather unconventional semantic theory, while the latter is compatible with traditional truth-conditional semantics. As the goal of the current study is to explore a pragmatic account of various focus-related phenomena and resist semantic revolutions, I will avoid basing the notion of local context on dynamic semantics, but instead endorse a pragmatic understanding of local contexts.

The core assumption in dynamic semantics is that the semantic value of a sentence is not a truth-condition, which maps possible worlds to truth-values, but rather a Context Change Potential (CCP), which is a function that maps contexts to contexts. A convenient way of describing the theory is to say that dynamic semantics builds in the semantic value of a sentence what was traditionally thought of as the sentence's pragmatic effect.[7] For example, according to the Stalnakerian framework, a sentence that expresses proposition p, when asserted, has the pragmatic effect of changing the context c into $c \cap p$, and this effect is based on not only the semantic value of the sentence but also the pragmatics of assertion. However, in dynamic semantics, this effect is rather part of the sentence's meaning: the meaning of the sentence is a CCP that maps any context c to $c \cap p$. Let us use '[.]' to stand for the interpretation function that assigns each sentence its CCP. Hence '$[S]$' denotes the CCP of sentence S. Since a CCP is a function from contexts to contexts, let $c[S]$ be the context resulting from updating the context c with sentence S.

Just as the downstream effect of a sentence becomes part of the sentence's semantic value, the upstream effect can also be captured by CCPs. The crux is that, for a presupposition-carrying sentence, its CCP is a partial function: it is defined for a context just in case that context satisfies the presupposition of the sentence. For example, if a

---

[7]What is called "dynamic semantics" is the version given by Heim (1983). As it is designed for solving the problem of presupposition projection, the notion of local context used in the theory is what we are interested here. In this approach, the semantic analysis is at the propositional level.

context c does not entail that there is a unique king of France, for the sentence "The king of France is bald" (call it S), $c[S]$ is undefined. In general, the presupposition of a sentence, in this framework, is specified by the domain of the sentence's CCP.

When only atomic sentences are considered, the dynamic semantics is not much different from Stalnaker's pragmatics, except what is thought to be pragmatic is now defined within semantics by sentences' CCPs. But when it comes to complex sentences with embedded clauses and local contexts, the dynamic semantic framework can make nice predictions on presupposition projection. Such predictions are based on a semantics of sentential connectives, as follows.

$c[\neg p] = c - c[p]$

$c[p \wedge q] = c[p][q]$

$c[p \vee q] = c[p] \cup (c - c[p])[q]$

$c[p \rightarrow q] = c$ if $c[p][q] = c[p]$; $c[p \rightarrow q] = \emptyset$ if $c[p][q] \neq c[p]$

The presupposition of a complex sentence is determined by way of determining the domain of its CCP according to the above rules. For example, according to the first clause, the domain of the CCP of $\neg p$ is the same as that of p, which predicts that negation is a presupposition hole. For a conjunction, its CCP is defined for a context c just in case i) the CCP of the first conjunct is defined for c, and ii) the CCP of the second conjunct is defined for $c[p]$. This predicts that the first conjunct position is a presupposition hole while the second is a filter.

According to this approach, the notion of local context is a semantic one, in the sense that a clause's local context is derived from a global context in a way that is completely determined by the semantics of the relevant constituents. For example, when a conjunction $p \wedge q$ is uttered in a context c, the local context for the clause q is determined by the semantics of p as well as that of conjunction, and it does not rely on any pragmatic inferences.

Accordingly, the "appropriateness conditions" I mentioned, such as Presupposition Satisfaction, must also a semantic notion: in order to be semantically interpretable, a clause's presupposition has to be satisfied by what semantics determines to be its local context.

Although the dynamic semantics generates nice results about presupposition projection, they come with a cost: We have to accept this non-truth-conditional semantics in order to get the projection properties right, while the semantics is accused by many for including what should be pragmatic into the semantics.[8] A major task of the current study is to show that some phenomena that seemingly threaten traditional semantic theories can be explained away by the pragmatics I am offering. Therefore, it will better serve my purposes if the notion of local context does not rely on a revolutionary semantic theory such as dynamic semantics.

An alternative to dynamic semantics, regarding the notion of local context, can be found in Schlenker (2009, 2010). His approach assumes a traditional truth-conditional semantics for the above connectives, while using some pragmatic principles to derive local contexts, generating almost the same results as what dynamic semantics does.

According to Schlenker, the pragmatic derivation of a local context is based on some assumptions about how ordinary speakers calculate the effect of an utterance. First, taking the meaning of expressions as functions of various types (e.g. that of a clause is a function from possible worlds to truth-values), it is assumed that when processing such a function, it is easier if part of the domain of the function is disregarded. Second, it is assumed that before processing an embedded clause, the interpreter tries to make the task easier based on what she has already known about the meaning of the sentence. Third, it is assumed that processing is incremental: an utterance of a sentence is processed from left to right. Given these assumptions, "it follows that the interpreter will try to decide in advance of interpreting any expression E what is the smallest domain that he needs to consider

---

[8]Of course, such results are desirable, but it is debatable whether we really need a dynamic semantics or just a dynamic pragmatics to achieve the good results. See Lewis (2017) for a discussion on this issue.

when he assesses the meaning of E; this 'smallest domain' is our notion of a local context."

(Schlenker, 2009, 9-10)

To illustrate, suppose $p \wedge q$ is uttered in a context c, and right now an interpreter has processed $p\wedge$ and is about to process q. At this point, she already knows that only at the p-worlds in c, the truth-value of q is relevant to calculating the effect of the utterance; the truth-value of p at all other possible worlds plays no role, as they will not be included in the resulted context anyway, no matter what q is. Thus, she can safely disregard these not-p worlds (and worlds that are not in c at all), while not affecting the result of calculating what the whole utterance contributes to the context. Thus, when processing q, she disregards all the not-p worlds in c, and the domain of the meaning function of q can be restricted to the possible worlds in $c \cap p$—this set of possible worlds is the local context of q.

The contrast is sharp between Schlenker's approach and dynamic semantics. Local contexts in Schlenker's view are not completely determined by semantics; rather, it is a minimized domain of the function that the interpreter needs to process.

I will adopt Schlenker's notion for the most part, with the assumption of incremental processing loosened. The concern comes from the interpretation of counterfactual conditionals. For example, suppose one is at the point of processing the antecedent of the following counterfactual:

(20) If kangaroos had no tails, they would topple over.

If we stick to Schlenker's assumptions, it follows that when processing the antecedent, the whole domain of it, i.e. the set of all possible worlds, needs to be considered, as it is not a safe move to disregard any possible world before processing the antecedent. To see that, suppose she is right about to process the antecedent—that is, only "if" is processed so far. Suppose further that she disregards a possible world w. There could always be a continuation of the sentence that makes this move a mistake. Let p be a proposition such that w is one of the closest p-world to the actual world. Let proposition p' be $p - \{w\}$.

Suppose p and p' are expressable by two sentences, ϕ and ψ, respectively. Then, the counterfactual conditional *If ϕ, would ψ* shows that dropping w could be a mistake. For it is false at the actual world, but processing the antecedent without w will render it true—all the closest p-worlds to the actual worlds, *disregarding w*, are indeed p' worlds.[9] Since *If ϕ, would ψ* is a way that (20) could end up (as only "if" has been processed so far), dropping any world w will not be safe.

But this is intuitively not plausible: we do not need to consider all the possibilities when processing the antecedent in (20). Rather, it seems that there is a relatively small set of possible worlds that we consider: a set that includes worlds that are not very different from the actual world but includes some far-enough worlds where kangaroos do not have tails. In Lewis's terms, it is the smallest "sphere" of worlds that admits the truth of the antecedent. In processing the antecedent, we only need to know its truth-value at these worlds, without considering whether it is true or not at worlds that are very far from the actual world.

The problem, then, is that Schlenker's pragmatic theory predicts that the local context for the antecedent of a counterfactual needs to include all worlds, but it seems that the local context is limited. I propose to treat the assumption of incremental processing as secondary to the principles of domain-minimizing.[10] It is a plausible pragmatic principle that a better strategy in processing a clause is to minimize the domain of the relevant function before the interpreter comes to know what the rest of the sentence is. In such cases, as the interpreter in effect drops some worlds *blindly*, no matter how the sentence ends, as it makes the function easier to process and at the same time saves effort processing the next expression. However, when the this strategy of minimizing the domain is impossible,

---

[9]This argument assumes Lewis's semantics for counterfactuals, but there can easily be examples that work for different semantic theories.

[10]Schlenker explores local contexts derived without the assumption of incremental processing. According to him, sometimes the projection behaviors of presupposition are better explained without incrementality. See Schlenker (2009, 2010) and Emmanuel Chemla (2012).

the interpreter may plausibly take a peek at the next expression and decide what worlds are to be dropped. Especially, if the domain is a very large, cutting it down will become so desirable that the motivation to save effort by not looking at the next expression is outweighed. If this is correct, then incremental processing only describes an ideal case of minimizing the domain. However, (20) is not such an ideal case. For the interpreter, a better strategy seems to be taking the antecedent into account so that she does not have to deal with the set of all possible worlds. In other words, incremental processing may be sacrificed, for the sake of minimizing the domain.

Then, how exactly can the interpreter derive a smaller set of possible worlds as the local context for the antecedent in (20)? First, let us assume the Lewisian semantics for counterfactuals. Then, in order to determine if a world w in the context set is to be ruled out, the interpreter has to process the antecedent to the effect of generating the set all the closest antecedent-true worlds to w. There are two options for doing this. The first is to follow Schlenker's rules so that the interpreter must start with the set of all worlds, and use the semantic content of the antecedent to generate the result in one step. This seems to be a very demanding task. The second option, though, is to first cut down the size of the domain, generating the smallest sphere that admits the truth of the antecedent, and then add the proposition of the antecedent to this minimal sphere of possible worlds. In other words, the interpreter first makes use of the antecedent to select a limited set of possible worlds, i.e. the smallest set of possible worlds that contain some worlds where the antecedent is true. Subsequently, the interpreter processes the antecedent with regard to this set of possible worlds by perform a Stalnakerian update to it, ruling out the worlds in it where the antecedent is false. It appears to me that the second option is a better strategy: it first prepares the ground for a Stalnakerian update by minimizing the domain, and then perform the update on it.

In sum, I take minimizing the domain to be a stronger pragmatic motivation than saving effort of looking at the next expression. If the two reasons do not conflict with each other, I will follow Schlenker's derivation of local contexts; but when they do, an embedded clause can be part of what minimizes its own domain, given that the benefit of domain minimizing is great and the cost of considering the clause is small. In general, I take the local context of a clause to be a minimized set of possible worlds that is to be updated by the clause. In the rest of the dissertation, I will rely on this view on local contexts, as my own view, GQA, is by and large a theory of the pragmatic relationships between clauses and their local contexts.

2.4. Introducing Questions

The second extension of Stalnaker's model comes from the Question-under-Discussion (QUD) approach, which takes a question aspect of both context and sentences into account. By this extension, the QUD model of context is able to capture some pragmatic effects that cannot be properly captured by Stalnaker's purely informational model.

An immediate reason why we should not be satisfied with the informational model is that not all utterances are meant to provide information. For example, when an interrogative sentence is uttered, normally it is not meant to add information to the common ground. Granted, uttering an interrogative sentence indeed alters the common ground—the common ground is to include the proposition that the speaker in question was speaking, that she asked a question, etc.. But here Stalnaker's distinction between two ways that utterances change the context comes in handy. In one way, as Stalnaker points out, an assertion changes the context in virtue of the *performance* of the speech act. This effect is nothing different from non-linguistic events happening in the surrounding environment, the beliefs of which presumably will get into the common ground.

The fact that a speaker is speaking, saying the words he is saying in the way he is saying them, is a fact that is usually accessible to everyone present. Such observed facts

can be expected to change the presumed common ground knowledge of the speaker and his audience in the same way as any obviously observable change in the physical surroundings of the conversation will change the presumed common knowledge. (1978, 323)

This, however, is not what he calls the *essential effect* of assertion. The essential effect, in contrast, does not result from the very performance of the speech act, but rather from the *content* of the speech act, which is, in the case of assertion, a proposition, i.e. a piece of information.

Applying this distinction to utterances of interrogatives, we can see that they are informative only in the first sense but not the second. In other words, they do change the common ground by adding propositions such as that someone is saying something, etc.. However, their essential effect—the effect brought about by virtue of their content—is not informational: by asking a question, a speaker is not providing any piece of information and proposing adding it to the common ground; rather, she is signaling that a certain kind of information is wanting and requesting other speakers to provide such a piece of information.[11] This is to say that the Stalnakerian model, which takes the informational common ground to be what is changed by utterances, fails to capture the essential context change effect of interrogatives.

To be fair, this point by itself is not really a problem for Stalnaker. His model is only meant to account for how *assertion* interacts with context, and other kinds of speech acts, such as questions, are simply not in its scope. Moreover, given the sheer amount of kinds of speech acts, it is unfair to require a theoretic model to accommodate them all. So,

---

[11]Granted, an interrogative that carries some presupposition can achieve some informative effect. For example, saying "Who can pick up my sister from the airport?" can inform the hearer that the speaker has a sister, provided that they did not have such a belief before receiving the utterance. But it is not clear if this informative effect is part of the essential effect of utterances of interrogatives. If one prefers to count it as part of the essential effect, I would like to weaken my claim here: the essential effect of utterances of interrogatives is not *only* informative but rather involves an effect of requesting information. This weakened claims still serves the main purpose here, as it shows that a theory of utterances' pragmatic effect is not complete if it only captures the informative effects of utterances.

the defense goes, we can still take the Stalnakerian picture as a satisfactory account of assertions, while acknowledging the desirability of a broader model for the sake of other speech acts.

But the problem is more serious than that. With a model that takes contexts as mere informational states and assertions as mere information contributors, it cannot even adequately capture the pragmatic effect of assertion. This point is illustrated by the following example, where a question is asked, and this imposes some non-informational constraints on whether a sentence can be asserted thereafter. (Capital letters indicate which words are stressed.)

(21) Q: Who ate the cake?

A: JIM ate the cake.

A': Jim ate THE CAKE.

Although the two responses to (21Q) express the same proposition and thus informationally equivalent, it would be appropriate to utter (21A) but not (21A'). Moreover, the difference in felicity is intuitively due to the question asked in the context, as (21A) seems to be suitable for answering the question while (21A') does not.

This shows that there is some appropriateness condition based on what question is raised in the context and whether a sentence is suitable for answering it. What is important here is that this condition cannot be formulated in terms of information. For both sentences (21A) and (21A'), the information encoded in the context is the same, and the informational contents expressed by them are also the same. Therefore, we need to consider a non-informational aspect of contexts and sentences in order to capture this appropriateness condition.[12]

---

[12]This challenge to Stalnaker's model is still unfair if what the model is intended to achieve is merely to capture the informational aspect of the effect of assertions. But it at least shows that there are some non-informational aspect in the pragmatics of assertions. In this sense, the current point is not a challenge to Stalnaker's model as a model of the informational effect of assertions, but rather a challenge to the model as a model of the pragmatic effect of assertions, broadly speaking.

In the following, I will present a simplified version of the QUD model, which in effect enriches the context representation by adding a question component to it and takes the question that a sentence is suitable for answering into account in determining whether the sentence is appropriate to utter.

On the context representation side, this model includes a *question-under-discussion* (QUD) alongside a body of information in the representation of context. In a context, its QUD is the question that speakers are committed to answering. In the above example, (21Q) raises a question. Suppose the speakers decide to accept this question as what they should strive to answer in the following conversation, it becomes the QUD of the context. Hence, the context of utterance for (21A) and (21A) now has two components, an information component that represents the common ground in Stalnaker's sense, and a question component that marks what is asked by (21Q) as the QUD.

Following Roberts (2012), a question, as the semantic value of an interrogative such as (21Q), is defined as a set of propositions (cf. Hamblin, 1973).[13] Roughly, the question is defined by its possible answers; with each answer being a proposition, the question is the set of all these propositions.[14] For example, (21Q) has the possible answers that Alice ate the cake, that Bertha ate the cake, that Jim ate the cake, etc., so that all and only propositions like such are members of the question expressed by (21Q). Thus, the effect of uttering (21Q) is to set the QUD in the context to be this question, i.e. this set of propositions. With this new question component, we have the following representation of context.

---

[13]From now on, I will always use the term "question" to refer to the semantic entity, i.e. the set of propositions. In contrast, the term "interrogative" refers to the syntactic entity that expresses a question.

[14]I equate propositions that are members of a question to "possible answers" just for the sake of brevity. To be accurate, there is a compositional derivation of what question is denoted by an interrogative, while it may not include all those which we ordinarily feel to be "possible answers" to the interrogative. For example, the proposition that no one ate the cake is not a member proposition of the question that (21Q) denotes, but we may feel that a sentence that expresses this proposition counts as a possible answer to (21Q).

Context (QUD). A context is an ordered pair $\langle c, QUD \rangle$, where c is a context set and QUD is a question.[15]

Besides this modification on how contexts are represented, on the QUD approach, a declarative sentence is no longer understood as merely a vehicle of information. As illustrated in the above example, a declarative sentence has a certain question it is suitable for answering, and what this question is for a sentence may be determined by which constituent is focused (i.e. intonationally stressed, roughly). Specifically, (21A), with a focus on "Jim", is suitable for answering the question *Who ate the cake?*. In contrast, "the cake" is focused in (21A'), so it is suitable for answering a distinct question, *What did Jim eat?*. Although the derivation of such a question for a focus-bearing sentence is only defined later in 3, the idea is quite intuitive: two sentences having the same propositional content may nevertheless differ in terms of what question they can appropriately respond to. This difference is crucial for distinguishing (21A) and (21A) in terms of their felicity in the context. (21A) is appropriate to utter because is suitable for answering the question that is the QUD in the context, but (21A') is not because what it is suitable for answering is distinct from the QUD. Let's call such a question that a sentence is suitable for answering, the *congruent question* (CQ) of the sentence (cf. Simons et al., 2017).

Formally, the CQ of a sentence is a set of propositions obtained abstracting away the focused constituent from the sentence and then filling the position with entities of an appropriate type (cf. Rooth, 1985). Being a set of propositions, a CQ is a question in the semantic sense. For example, the CQ of (21A), as "Jim" is focused, is the set of propositions of the form *x ate the cake*, i.e. propositions that Alice ate the cake, that Bertha ate the cake,

---

[15]According to Roberts (2012), the QUD component in a context is not a single question but rather a "stack" of questions that bear certain relations to one another. But now the model is simplified, since only the most immediate question in the stack (the top one) directly bears on whether utterances are felicitous in a context. Besides, what Roberts (2012) provides is a diachronic model of context, which involves a track record of all the moves made by speakers in a conversation, and the context model specifies what the information state and QUD stack are at each point when a move is made. However, I adopt this simplified, synchronic model that only represents what the context is like at a certain time, because the current study is not meant to explore how a context evolves over time in reaction to a series of utterances.

that Jim ate the cake, etc.. In contrast, as the focused constituent in (21A) is "the cake", the CQ of it is a different set of propositions, which includes propositions of the form *Jim ate x*.

Given the new representation of context as well as the notion of CQ, we are able to formulate the non-informational appropriateness condition that the above example alludes to. On the one hand, a sentence, when uttered, is supposed to answer the QUD in the context. On the other hand, the question the sentence is suitable for answering is its CQ. Thus, the idea is, a sentence can be appropriately uttered only if its CQ *is* the QUD in the context.

> **Congruence.** For a context $\langle c, QUD \rangle$, if all the possible answers to QUD are $q_1, q_2, ...$, then a sentence S can be appropriately asserted in this context only if the CQ of the sentence is $\{q_1, q_2, ...\}$, namely, only if $CQ(S) = QUD$.[16]

Given the condition, the difference between (21A) and (21A') straightforwardly follows: the CQ of (21A) is the QUD, i.e. the very question raised by (21Q), while the CQ of (21A') is a distinct one.

As we have seen, an appropriateness condition does not only distinguish felicitous and infelicitous utterances but also enables an upstream effect of utterances. In particular, Congruence sometimes can be satisfied by accommodation. When a sentence is uttered but is not congruent with the QUD, within some pragmatic constraints, speakers can adjust the context to set the QUD to be the CQ of the sentence.[17]

---

[16]This congruence constraint is roughly the version proposed by Roberts (2012). In Rooth (1985, 1992), the constraint is weaker, in the sense that it only requires every possible answer to a question to be a proposition in the CQ, while the converse is not required.

[17]The pragmatic constraints on accommodating a CQ play a very significant role in QUD theories. According to Roberts (2012), a CQ can be accommodated only if it is *Relevant* to the original QUD in the context, in the sense that every complete answer to the accommodated CQ has to be able to settle the truth-value of at least one answer to the original QUD. Furthermore, this constraint gives the QUD model the power of explaining why the domain of quantification of a quantifier, such as "only", can vary according to focus. However, as the current study mainly concerns the effect of clauses on local contexts, constraints of this sort are not important, because in local contexts there do not exist questions "under discussion". See §2.5 for details.

Question Accommodation For a context $\langle c, QUD \rangle$, if a declarative sentence S is uttered in it but the congruent question of it is not QUD, then the congruent question may be accommodated as the QUD of the context. I.e., the post-accommodation context is $\langle c, QUD' \rangle$ where $QUD' = CQ(S)$.

Besides the new appropriateness condition and upstream effect, taking the question aspect of context and sentences also brings up some new downstream effects of assertions.[18] Roughly, when an assertion completely answers the QUD in its context or decides that it is unanswerable, it will remove the QUD from the context.[19] But I will not take this non-informational downstream effect into account, because it will only affect subsequent utterances whereas the current study only focuses on the interaction between single utterances and their context. Hence, I will pretend that the downstream effect of utterances is still a purely informational effect on the context set.

Taking stock, we can see that the QUD approach extends the Stalnakerian model on both the context side and the sentence side. For it adds a question component to the representation of context, and also associates congruent questions with sentences. However, so far this approach only has to do with global contexts and entire sentences. For an embedded clause, does it also have such a congruent question, a question that it is supposed to answer? Does a local context also have a question component which is part of what determines appropriateness? We now turn to these issues.

2.5. Local Congruent Questions

Regarding the pragmatic effect of sentences, the takeaway point is that the question aspect of a sentence, i.e. its CQ, is part of what determines appropriateness. In particular, the

---

[18]Since we have been focused on the effects of assertions, utterances of interrogative sentences are ignored. But they surely have downstream effects: uttering an interrogative, when appropriateness conditions are satisfied, can set the question it expresses as the new QUD in the context.

[19]This is part of the reason why Roberts needs a stack of questions in the context: after the removal of the top QUD, the next one becomes the top one and subsequent utterances are thereby constrained by this new one.

CQ of a sentence serves to specify what question the sentence is supposed to address. As the QUD in a context is the question that ought to be addressed at the moment, it follows that a sentence with a CQ that is distinct from the QUD violates this requirement, and is hence inappropriate to utter. In other words, the QUD specifies a set of acceptable ways of how the information state in the context can be updated, and the CQ of a sentence stands for a set of alternative updates that the sentence is picking from; then, the Congruence condition is the requirement that the set of alternative updates a sentence is picking from has to be what is determined by the context to be acceptable.

What I advocate here is that a sentence's CQ is part of what determines appropriateness, *even when the sentence is embedded as a sub-sentential clause in a complex construction.* In other words, the question aspect of a sentence, when the sentence is embedded, does not suddenly lose its power of regulating the relation between the sentence and the context. The CQ of an embedded clause has to be related to the *local* context of the clause in some way for the sake of appropriateness. But a problem rises immediately. So far, local contexts are understood as sets of possible worlds, i.e. information states. Namely, there is no QUD, or any question component, in local contexts. Then, how can there be appropriateness constraints regarding the relation between CQs and local contexts?

To make my claim clear, I do not intend to say that local contexts involve questions in the way that global contexts do. In a global context, the notion of QUD is well-grounded by speakers' intentions. Specifically, the QUD of a global context is the question that the speakers accept as what they ought to answer. However, there is no comparable ground for a question component in a local context. Therefore, I will keep a purely informational notion of local contexts, so that they are merely information states. Then the problem is how the CQ of a clause can be part of what determines appropriateness, given that there is no local QUD that requires a clause to address it.

The reason why a clause's CQ matters is the following. According to the spirit of the QUD approach, what a sentence does to its context is not just adding a piece of information to it; rather, it *raises a question* and then picks out an answer to it. In other words, a sentence first brings up a set of alternative ways of updating the information state in its context, and then picks out and executes one particular update. The absence of a QUD component in local contexts means that a local context itself does not specify a particular set of acceptable updates; in other words, it is not specified what question is raised in a local context. However, that does not mean that all questions can be raised in a local context with equal appropriateness. For example, when you walk into the room and hear the following utterance, you may reasonably infer that the information state in the context entails that someone ate the cake.

(22) JIM ate the cake.

This pragmatic inference is based on the appropriateness of a CQ relative to an informational state. In particular, as the CQ of (22) is *Who ate the cake?*,[20] you may consider in what information state this CQ can be raised with the best appropriateness. If the current information state does not rule out the possibility that no one ate the cake, it seems that the CQ is not a very good question to raise, because that possibility in the information state is simply ignored by all the alternative updates the CQ selects. Then, the reasoning goes, the information state should have established that someone ate the cake. More generally, if an information state is compatible with a possibility w, a good selection of updates, i.e. a good question to raise, should not exclude w by all the alternative updates.[21]

This is an illustration of why CQ is still part of what determines appropriateness even without a QUD in the context. The crux is that there are pragmatic constraints on the relation between a sentence's CQ and the information state that the sentence is updating. The local context of an embedded clause, as the minimal domain of possible worlds on

---

[20]How a sentence or clause maps to its CQ is explained in Chapter 3.
[21]This requirement is defeasible. See §3.4.2 for more details.

which the interpretation of the clause is based, is what the clause is directly updating. It thus follows that such pragmatic constraints apply to embedded clauses and local contexts, even though local contexts are informational states without QUDs. In Chapter 3, I will specify what such pragmatic constraints are, and then the rest of the dissertation aims at showing some desirable results of applying the pragmatic constraints to concrete examples. What is important here is that the CQ aspect of a clause is still operative when the clause is embedded: the appropriateness of using the clause in its local context has to do with how the CQ fares with the information state of the local context.

Moreover, the appropriateness conditions, applying to local contexts, will make some special upstream effect possible, which can ultimately affect the truth-conditional interpretation of some kinds of utterances. As we have seen, appropriateness conditions, in general, can bring about accommodation effects, which in effect changes local or global contexts that are to be updated. When question-based appropriateness is considered, an embedded clause's CQ can affect the local context of the clause by accommodation. Local contexts, in many cases, can affect truth-conditional interpretations, so it follows that accommodation of the question-based appropriateness has some truth-conditional effects. This pragmatic process will play a significant role when we consider examples such as knowledge ascription sentences and counterfactual conditionals in Chapters 4 through 6. Without concerning much detail, the general theme there is to exploit the above pragmatic process in order to explain truth-conditional effects of focus and other CQ-regulating devices. By exploring such a pragmatic account, I hope to showcase that many puzzles that seem to deserve semantic solutions can actually be plausibly explained in terms of pragmatics, and thus overhauls of semantic theories may be unnecessary. But before that, let us turn to a detailed theory of CQs and the appropriateness conditions they give rise to in the next chapter.

CHAPTER 3

Congruent Questions: Semantics and Pragmatics

According to the framework established in the preceding chapter, the pragmatic effect of a clause is to contribute a piece of information to its local context by answering a question. That is to say, a clause does not merely update the information encoded in its context but also carries a Congruent Question (CQ), which has to do with its appropriateness. Given this framework, I claim that it can be explained pragmatically why a mere difference in focus can result in different truth-conditional interpretations. As focus triggers CQs, it is able to affect local contexts based on some pragmatic constraints governing the relation between local contexts and CQs of embedded clauses. This is an upstream effect of embedded clauses that is brought about by their CQs.

As will be shown next, this framework is able to explain the phenomenon that some question-regulating features of embedded clauses, including focus, can have a truth-conditional effect, which is puzzling at first sight. By compositionality, the meaning of a whole sentence is determined by the meanings of its parts and the way they are combined. However, there are examples (repeat from Chapter 1) where compositionality is apparently violated due to the effect of focus. (Capital words are focused—stress them when you read).

[MARRIAGE] Clyde and his girlfriend Bertha don't like close relationships. They see each other only twice a year and they don't want to get married, etc.. However, Clyde found that if he were to get married soon, he would inherit a great amount of money. Then he married Bertha and got the money while expecting that their 'loose' relationship would continue. (Adapted from Dretske, 1972.)

(23) If Clyde hadn't MARRIED Bertha, he wouldn't have inherited the money.

(24) If Clyde hadn't married BERTHA, he wouldn't have inherited the money.

The same effect also appears in knowledge-that ascription sentences.

[JEWEL THIEF] Last night, Peter robbed the jewelry store. He smashed the window, forced open the locked safe, and stole the rubies inside. But Peter forgot to wear gloves. He also forgot about the security camera. Today, Mary the detective has been called to the scene to investigate. So far she has the following evidence. She has been told that there was a theft, she has found and identified Peter's fingerprints on the safe, and she has seen and recognized Peter on the security video, filmed in the act of forcing open the safe. She has no further information. (Schaffer and Knobe, 2012)[1]

(25) Mary knows that PETER stole the rubies.

(26) Mary knows that Peter stole THE RUBIES.[2]

Given this phenomenon, a major objective of this dissertation is to provide a pragmatic theory, which is based on the upstream effect of CQs, to explain this focus effect.[3]

Regarding the above examples, the extant semantic theories are unable to predict, or even inconsistent with, language users' judgments. For instance, semantic theories

---

[1]This is a verbatim copy of what Schaffer and Knobe (2012) use as a vignette in their experiment to test how native speakers judge the truth values of knowledge ascription sentences given this contextual background. The results show that the participants' truth-condition judgments are sensitive to various linguistic devices which presumably trigger 'contrasts', such as *rather-than*, *wh*-clauses, and questions-under-discussion. Although the experiment, unfortunately, does not test on people's judgments on knowledge ascription sentences involving focus such as (25) and (26), Schaffer indicates that focus is one of the linguistic devices that are able to trigger such a contrastive effect in his other papers, such as Schaffer (2004, 2005, 2008).

[2] This effect of focus on truth-conditional meanings has been observed in a variety of other kinds of expression as well. To mention some: adverbs of quantification (*always*, *usually*,...), auxiliary modal verbs (*must*,...), attitude verbs (*believe*,...), determiner quantifiers (*many*, *most*,...), exclusives (*only*,...), and additives (*too*, *even*,...).

[3]Given the sheer diversity of expressions displaying the focus effect (see f.n. 2), it could be the case that there is no uniform way of explaining how focus contributes to truth-valuable content. For example, Beaver and Clark (2008) propose a hybrid model, according to which the focus effect displayed in various expressions is realized by three distinct semantic/pragmatic mechanisms. However, the scope of this dissertation is restricted to only counterfactuals and knowledge-that ascription sentences, the focus effect on which, according to the account I shall propose, deserves a uniform explanation.

for counterfactuals takes what the antecedent of a counterfactual contributes semantically to be the proposition the antecedent expresses, and consequently if two otherwise identical counterfactuals' antecedents have different focus patterns but are nonetheless truth-conditionally equivalent, the two entire counterfactuals should have the same truth-value. This, however, is not what we observed in the pair of (23) and (24).

Facing this puzzle, there could be two reactions. Semantic progressivists may posit new (and oftentimes more sophisticated) semantic theories. In contrast, their conservativist opponents are skeptical of the necessity of such progressive moves, as the data might be explained by independently motivated pragmatic theories, without semantic revolutions. Such tugs-of-war are common at the semantic-pragmatic interface, and they are mostly local wars: disputes are based on particular sets of data, and the victory of either camp at one place does not automatically translate to its victory in other areas. In dealing with the phenomenon illustrated above, I shall join the conservativist camp in the dispute around the issue of how focus and other question-triggers can have a truth-conditional effect on certain sentences, arguing that this phenomenon should be explained by appeal to the pragmatic framework introduced in Chapter 2, according to which an embedded clause is understood as making its informational contribution *by answering a question*. In particular, the account I advocate involves two main theses. First, focus, as well as some other features, determines the CQ of a clause. Second, some pragmatic principles mandate that the CQ be related to the local context in a certain way, so that speakers, when interpreting an utterance, may infer what the local context is like from the clause's CQ, and thus the truth-conditional interpretation is sensitive to focus if it is sensitive to what the local context is.

Accordingly, the account involves two components. First, it is to give a theory of how clauses' question-triggers, such as focus, determine CQs. This calls for a theory of the focus-induced interpretation. I will adopt the Alternative Semantics developed in

Rooth (1985), which essentially takes a focus-triggered question to be a set of alternative propositions. The second component is a pragmatic theory of questions; or more precisely, a specification of pragmatic principles constraining relations between clauses' CQs and the local contexts they are to update. Very roughly, for an embedded clause, the pragmatic principle requires the question raised by the clause, i.e. the CQ, to be a *good question* to be asked against the information state of the local context; also they require the answer provided by the clause to be a *good answer* to the CQ, with respect to the information encoded in the local context. In a slogan, a clause should raise a good question and provide a good answer to it. What is crucial to this view is that whether a question or an answer is good is relative to informational states: the information established in a context determines what questions and answers are good. Therefore, based on what question a clause raises and what answer it provides, speakers can assume that both the question and the answer are good in the local context of the clause, and infer what the local context must be like, resulting in an upstream effect. Then, if the truth-condition of a sentence is sensitive to such a local context, question-triggers in a clause can influence the truth-conditional interpretation in an indirect way, bridged by pragmatic constraints on what questions and answers count as good. As is obvious, my account crucially relies on a specification of what counts as "good", and then it can explain the phenomena by appeal to pragmatic inferences which assume that, by default, an embedded clause raises a good question provides a good answer to it. Call this view, the Good Question-Answer view, or GQA for short.

Towards a full exposition of GQA, this chapter is to accomplish the following four tasks. First, as we have seen, what question a clause congruently answers is regulated by some features of the clause, such as focus. But focus is not the only feature that plays the question-regulating role. In §3.1, I provide a brief survey on some of these features. Then, §3.2 includes the semantic component of GQA, which is to offer a semantic representation

43

and derivation of CQs of clauses. I will largely follow the Alternative Semantics, adopting a multi-dimensional semantic framework that treats CQs as a dimension of the clause's semantic interpretation, alongside its truth-conditional dimension. With the semantics in hand, §3.4 turns to the core of GQA, i.e. the pragmatic principles governing what questions and answers count as good, relative to contexts. Lastly, in §3.5, I will briefly discuss what this pragmatic theory is to achieve.

3.1. Focus and Other Question Triggers

What question a clause or sentence can appropriately address, i.e. the CQ, is regulated by some features, among which focus is a paradigm. In this section, I will lay out some basics about focus and the question-regulating role it plays, and then turn to other constructions that are functionally similar.

Focus is a syntactic feature that marks some constituents in a sentence. The theoretical role it plays is to explain some regularities between prosodic patterns of sentences and the pragmatic role sentences can play in contexts. For example, the following two sentences bear different prosodies, where capital letters specify where the highest pitch accent in a sentence is.

(27) **a.** Alice introduced BETH to Carl.

b. Alice introduced Beth to CARL.

The difference in prosody is correlated with whether these sentences can be felicitously used in certain contexts. In particular, (27a) can be felicitously uttered as a response to each of the following utterances, while (27b) cannot.

(28) **a.** Whom did Alice introduce to Carl?

b. Alice introduced Barry to Carl.

44

To account for such regularities, it is proposed that the syntactic difference between the two expressions in (27) is that different constituents in them are focused. Using the notation [.]$_F$ to pick out the focused constituent, those in (27) are represented as follows:

(29) **a.** Alice introduced [Beth]$_F$ to Carl.

     b. Alice introduced Beth to [Carl]$_F$.

The explanation for the regularities between felicity judgments and prosodic patterns appeals to two sets of rules. On the one hand, there are rules governing how focus can be realized by pitch accent configurations. They are rules of *focus realization*. The details about these rules do not have to concern us here, since examples used here all involve very straightforward correlations between focus and what is called *nuclear pitch accent*, i.e. the strongest pitch accent in a short sentence or a short unit of a longer sentence (such as the antecedent or consequent clause in a conditional). In particular, in the simplest case where a single word is focused, that word bears the nuclear pitch accent—e.g. "Beth" and "Carl" in (29a) and (29b), respectively. When focus is on a constituent larger than a single word, the nuclear pitch accent must be within the scope of the focus. But exactly which word it falls on within the focus is a complex issue. As a rule of thumb, the nuclear pitch accent tends to be on the final, or close to final, word within the focus. For example:

(30) (Who introduced Beth to Carl?) [The girl walking over THERE]$_F$ introduced Beth to Carl.

This rule is highly defeasible, as the distribution of pitch accents is constrained by not only focus but other factors as well.[4] In what follows I will avoid using examples involving such larger scopes of focus, so that where the nuclear pitch accent is will be easy to decide

---

[4]For example, Schwarzschild (1999) suggests that GIVENNESS—a marking that is licensed by a salient discourse antecedent that bears some relevance to the GIVENNESS-marked constituent—is such a factor. He proposes the rule that if a constituent is GIVEN, it is deaccented even if it appears within the scope of focus. This explains why the nuclear pitch accent cannot fall on "him" but must move leftward in the following example.

     (What did John's mother do?) She [PRAISED him]$_F$. (Schwarzschild, 1999, 145)

In addition, default accenting rules also constrain where the nuclear pitch accent is within focus. See Büring (2016) for a survey on this issue.

based on the focus marking, despite that focus realization rules are not discussed here in detail.

The other set of rules govern the relations between focus and context. They are conditions for when focus is licensed by contextual factors. As (28) illustrates, the focus on "Beth", as opposed to "Carl", can be licensed by various items that occur in the discourse, such as certain questions and assertions. This follows from the rules governing Question-Answer Congruence and Contrast: focusing on "Beth" is licensed as it makes the sentence a congruent answer to (28a), or marks the constituent that contrasting (28b). In contrast, focusing on "Carl" does not relate the sentence to these discourse antecedents in the way that these rules permit. This explains why focusing on "Beth" or "Carl" affects felicity. Hence, such rules of focus licensing, together with the rules concerning focus realization, explain the regularities between prosodic patterns (esp. the distribution of pitch accents) and contexts of utterance.

A clarification before moving on. As illustrated in (30), the constituent that bears focus can be larger than the constituent that bears the nuclear pitch accent. But the very placement of nuclear pitch accent is compatible with smaller scopes of focus. In this sense, prosodic patterns underdetermine focus marking; or in other words, there can be a focus ambiguity even when the prosody of a sentence is fixed. For example, stressing the word "Sam" might give rise to multiple possible foci.

(31) **a.** (Whom did Kim call?) Kim called [SAM]$_F$.

   b. (What did Kim do?) Kim [called SAM]$_F$.

   c. (What happened?) [Kim called SAM]$_F$.

Given this focus ambiguity, it is worth keeping in mind that focus identification is not merely a matter of prosody. However, prosody constrains what can be focused by the focus realization rules, according to which the nuclear pitch accent must fall within the

focused constituent. Thus, the size of focus can be as narrow as just one word, or as large as a whole sentence.

To summarize, the theoretical role focus plays involves two aspects: it is a syntactic feature that, on the one hand, is realized phonetically by certain distribution of pitch accents and, on the other, is related to the context in various ways, such as regulating Question-Answer Congruence and Contrast. The phonetic aspect of focus—how it is realized in prosody—is rather language-specific and the focus realization in a language might interact with some quite contingent features of the language. In the current approach, the emphasis is on the pragmatic consequences focus may give rise to, so the phonetic realization of it will be omitted. In particular, for the example sentences considered in the following, only focus marking is specified, and the prosody, or pitch accents distribution, will not be included—in most cases it will be very straightforward for the reader to determine, based on focus, where a nuclear pitch accent is.

On the pragmatic side of focus, it is clear that focus plays a discourse-regulating role—it indicates a sentence is answering a certain question or contrasting to an assertion made in a discourse. I propose that these different ways of relating focus to context be subsumed under a general role that focus plays: it determines what question a sentence is answering. For example, the focus on "Beth" in (29a) determines that the sentence congruently answers the question *Whom did Alice introduce to Carl?*. Given that, we can explain why (29a) is felicitous in responding to the utterances in (28). When the question (28a) is raised in the discourse, (29a) is felicitous because the question it is congruent with is identical to (28a). When the assertion (28b), that Alice introduced Barry to Carl, is made in the discourse, (29a) is felicitous because it congruently answers a question to which the previous assertion (28b) provides a different answer. (29a) contrasts with (28b) in the sense that it provides an alternative answer to the same question. On the other hand, as "Carl" is focused in (29b), it congruently answers a different question, i.e. *To whom Alice*

47

introduced Beth?. Then, (29b) is not felicitous as a response to (28a) or (28b): it cannot congruently answer the former, and it does not contrast with the latter.

In sum, we can take focus as a (congruent) question trigger, determining what question a sentence is congruently answering. Based on that, various ways in which contexts license focus can be explained. Understanding focus in such a way as a question trigger is in accordance with the Alternative Semantics tradition, originated in Rooth (1985, 1992), provided questions are understood in terms of the Hamblin semantics (Hamblin, 1973). On the one hand, Alternative Semantics treats focus as a trigger of alternative propositions. For instance, focusing on "Beth" in (29a) triggers a set of propositions of the form *Alice introduced x to Carl*, where *x* is any individual. On the other hand, such sets of alternative propositions are the semantic entities that are denoted by questions, on the Hamblin semantics. As mentioned in Chapter 2, a question is defined by its possible answers, which is a set of propositions. For example, the set of propositions just mentioned defines the question *Who did Alice introduce to Carl?*. Then, taking focus as a question trigger is essentially understanding it in terms of Alternative Semantics—it is a trigger of alternative propositions, which in turn define a question.

How exactly focus triggers questions will be specified formally in the next section. Before that, let us note that focus is not the only linguistic device that plays this question-triggering role. Notice that focus plays the role of regulating Question-Answer Congruence and Contrast. If an expression or construction plays a similar role, it is evidence that it can be understood as a question trigger as well. Such question triggers include *it*-cleft, *rather-than* construction, *and-not* construction, among others. In the following examples, the a-sentences can be felicitously uttered as a response to the question (28a) or a contrast statement of (28b), while the b-sentences cannot.

(32) **a.** It was Beth that was introduced by Alice to Carl.

   b. It was Carl that Alice introduced Beth to.

(33) **a.** Alice introduced Beth rather than anyone else to Carl.

    b. Alice introduced Beth to Carl rather than anyone else.

(34) **a.** Alice introduced Beth to Carl and didn't introduce Barry to Carl.

    b. Alice introduced Beth to Carl and didn't introduce Beth to Ceth.

This similarity between focus and the above linguistic devices suggests they should be understood similarly: they are question triggers. [5] Because of this similarity, the following semantic theory of question triggers is expected to apply to focus and other question triggers alike.

3.2. A Semantics of Congruent Questions

In this section, I shall present a common ground theory of focus that is intended to answer two questions. First, what semantic interpretation does focus induce? Second, in what way is the truth-conditional effect of focus possible, i.e. does focus affect truth-conditional content by semantics, or pragmatics? In this common ground theory, the first question is answered by Alternative Semantics, the gist of which is that focus induces a set of alternatives, which is a semantic value of sentences alongside the traditional, truth-conditional semantic value. Regarding the second question, the common ground theory holds that the truth-conditional effect of focus is realized by pragmatic reasoning.

---

[5] However, I am not claiming that triggering a question is *all* these devices contribute. For example, *it-cleft* is commonly thought of as a semantic trigger of an existential presupposition, or a so-called "hard" presupposition trigger, such that the presupposition it triggers presents in every context and cannot be cancelled: (31a) semantically presupposes that Alice introduced someone to Carl. The other two devices, in contrast, are arguably only "soft" triggers, since they can be felicitously uttered even if the relevant existential presupposition is not satisfied in a context:

> (I don't know if Alice introduced anyone to Carl.) But if Alice introduced Beth to Carl and didn't introduce Barry to Carl, Barry would be upset.

Since presupposition projects from the antecedent position in a conditional, the felicity of the above utterance in its context shows that the *and-not* construction does not semantically trigger an existential presupposition. Because of this difference, Abusch (2009) uses *it-cleft* as a comparison to soft triggers such as the *and-not* construction and explains that they invoke certain alternatives which can pragmatically trigger existential presuppositions in some contexts. The approach is similar to the current view that these linguistic devices are question triggers, since questions, in our sense, are just sets of alternative propositions.

As stated above, it is assumed that, at the syntactic level, focus is represented as a marking on some constituents in a structure. Call it *F-marking*.[6] Beyond the syntactic level, I follow the Alternative Semantics approach developed by Rooth (1985) to formulate the semantic value focus induces.[7]

Formally, for any expression α, it has an *Ordinary Semantic Value* (OSV henceforth), noted as $[\![\alpha]\!]^O$, which is whatever semantic value α has in the traditional truth-conditional semantics. On the other hand, another interpretation function gives α a separate *Focus Semantic Value* (FSV henceforth), noted as $[\![\alpha]\!]^F$, which is a set of alternatives determined by how the constituents in α are F-marked. For a simple constituent α (a terminal node in a parse tree), if it is not F-marked, then $[\![\alpha]\!]^F=\{[\![\alpha]\!]^O\}$; otherwise, $[\![[\alpha]_F]\!]^F$ is a set of all the entities that belong to the same semantic type as $[\![[\alpha]_F]\!]^O$. To illustrate, taking **a**, **b**, **c**, and **introduce** as terms in the metalanguage which respectively denote Alice, Beth, Carl, and the ternary relation of introduction, the interpretations of the simple constituents in (35) are as follows.

(35) Alice introduced Beth to [Carl]$_F$.

$[\![\text{Alice}]\!]^O=$**a** $[\![\text{Alice}]\!]^F=\{$**a**$\}$

$[\![\text{Beth}]\!]^O=$**b** $[\![\text{Beth}]\!]^F=\{$**b**$\}$

$[\![[\text{Carl}]_F]\!]^O=$**c** $[\![[\text{Carl}]_F]\!]^F=\{$**a, b, c,**...$\}$

$[\![\text{introduce}]\!]^O=$**introduce**[8] $[\![\text{introduce}]\!]^F=\{$**introduce**$\}$

---

[6]As explained in the preceding section, focus does not always coincide with the strongest pitch accent. Because what concerns us here is the semantic/pragmatic properties of focus, I will take F-markings in syntactic structures as given, ignoring the relation between F-markings and their phonetic realizations.

[7]By taking the Alternative Semantics as a part of the common ground theory, I do not mean that everyone agrees that it is 'the correct' semantics of focus-induced interpretation. In particular, other theories of focus-induced interpretation include the Structured Meanings approach developed by von Stechow (1981, 1991), Cresswell and Stechow (1982), Krifka (1992), and others, as well as the event-based semantics by Bonomi and Casalegno (1993) and Herburger (2000). See also Beaver and Clark (2008) for a comparison between them and the Alternative Semantics.

[8]Assume that the lexical meaning of *introduce* is **introduce**= $[\lambda x.\lambda y.\lambda z.z$ introduced x to $y]$.

Given that each simple constituent is assigned both an OSV and an FSV, the two dimensions of meaning are determined compositionally for complex expressions. In particular, for any α whose daughters are β and γ, if $[\![\beta]\!]^O$ is a function that takes $[\![\gamma]\!]^O$ as an argument such that $[\![\alpha]\!]^O=[\![\beta]\!]^O([\![\gamma]\!]^O)$, then $[\![\alpha]\!]^F=\{x(y)\colon x\in[\![\beta]\!]^F$ and $y\in[\![\gamma]\!]^F\}$.[9]

$[\![$introduced Beth$]\!]^O=$**introduce(b)**

$[\![$introduced Beth to [Carl]$_F]\!]^O=$**introduce(b)(c)**

$[\![(35)]\!]^O=$**introduce(b)(c)(a)**

$[\![$introduced Beth$]\!]^F=\{$**introduce(b)**$\}$

$[\![$introduced Beth to [Carl]$_F]\!]^F=\{$**introduce(b)(a),introduce(b)(b),introduce(b)(c),**...$\}$

$[\![(35)]\!]^F=\{$**introduce(b)(a)(a),introduce(b)(b)(a),introduce(b)(c)(a),**...$\}$

In words, for example, the FSV of (35) is the set collecting all the propositions expressible by *Alice introduced Beth to x*, where *x* denotes any individual or plurality of individuals.

According to this way of interpreting focus, the OSV dimension of an expression is to be calculated independently of focus. In contrast, the FSV of it crucially depends on which part is focused. For example, if (35) has "Beth" focused instead, as in (36), its OSV remains to be the proposition, that Alice introduced Beth to Carl, but the FSV is a different set of alternative propositions.

(36) Alice introduced [Beth]$_F$ to Carl.

$[\![(36)]\!]^O=$**introduce(b)(c)(a)**

$[\![(36)]\!]^F=\{$**introduce(a)(c)(a),introduce(b)(c)(a),introduce(c)(c)(a),**...$\}$

In words, the FSV of (36) is the set of all the propositions expressible by *Alice introduced x to Carl*, where *x* denotes any individual or plurality of individuals.

---

[9]This only covers complex constituents whose OSVs are determined by function application in the truth-functional semantics. But it can be easily generalized to other cases. For example, if $[\![\beta]\!]^O$ and $[\![\gamma]\!]^O$ are of the same type and thus determine $[\![\alpha]\!]^O$ by predicate modification such that $[\![\alpha]\!]^O=[\![\beta]\!]^O\cap[\![\gamma]\!]^O$, then the members of $[\![\alpha]\!]^F$ are generated by applying predicate modification to any pair formed by one member in $[\![\beta]\!]^F$ and one in $[\![\gamma]\!]^F$: $[\![\alpha]\!]^F=\{x\cap y\colon x\in[\![\beta]\!]^F$ and $y\in[\![\gamma]\!]^F\}$. In general, for a complex expression whose OSV is determined by a certain rule of semantic composition, its FSV is the result of the pair-wise application of that rule to the FSVs of the components.

In sum, on Alternative Semantics, focus is understood as a trigger of alternatives. At the sentential level, focusing on a constituent triggers a set of alternative propositions. This fits in with the view advocated in the preceding section: focus is a question trigger. In the current approach, I follow Hamblin (1973) to identify questions with sets of alternative propositions, and a proposition is an answer to a question just in case it is a member of the set.

Question&Answer. A question Q is a set of propositions. A proposition p is an answer to Q just in case $p \in Q$.

The idea here is that a question can be defined by its possible answers. For example, the question *Whom did Alice introduce to Carl?* can be directly answered by propositions of the form *Alice introduced x to Carl*, so the question is understood as the set of all such propositions. In contrast, the question *Who did Alice introduce Beth to?* is a different set that includes propositions of the form *Alice introduced Beth to x* as its members.[10]

Now it is evident that focus, understood in Alternative Semantics, is closely related to questions, understood in the Hamblin semantics. When interpreted at the sentential level, focus contributes a set of alternative propositions. Hence, as a set of propositions defines a question, it is to say that focus is a question trigger. For example, the focus on "Carl" in (35) triggers the question *Who did Alice introduce Beth to?*, and that on "Beth" in (36) triggers the question *Who did Alice introduce to Carl?*, as these questions are the FSVs of the two sentences, respectively. I will call the question triggered by focus a *Congruent Question* (CQ), which is formally identical to the FSV of a sentence. This notion will play a fundamental role in the pragmatics of focus. As will see in §3.3, the CQ of an embedded

---

[10]From now on, the term "question" to refer to semantic objects, i.e. sets of propositions, rather than interrogative sentences in the syntactic sense, though it is assumed, following Hamblin, that the semantic value of an interrogative sentence is a question, i.e. a set of propositions. Accordingly, an answer to a question is always a proposition rather than a declarative sentence, though a declarative sentence can express a proposition that answers a question.

clause indicates a set of acceptable informational update on the local context of the clause, which enables some pragmatic effects of accommodation.

The key feature of Alternative Semantics is that every expression is treated as having two semantic values, in which sense it is a multi-dimensional semantics. As such, the semantics on its own does not imply any interaction between the two dimensions: no mechanism is provided to predict how the focus-induced question (FSV) can affect the truth-conditional meaning (OSV). It is not a problem for expressions (35) and (36), in which focus does not affect truth-conditional meanings. However, as the examples at the beginning of this chapter show, focus can have a truth-conditional effect. In other words, the question triggered by a focus must be related to truth-conditional meaning in some way, so that two sentences that differ only in focus can have different truth-conditions. There are some options regarding how the two dimensions are related so that the truth-conditional effect can be accounted for, to which I turn in the next section.

3.3. The Focus Effect: semantics or pragmatics?

The semantics of focus introduced above is not committed to any connection between truth-conditions and focus-induced questions, but focus does have a truth-conditional effect. For instance, the following pair have different truth-conditions: (37) is true just in case Alice introduced Beth and no one else to Carl, but (38) is true just in case Alice introduced Beth to Carl and no one other else.

(37) Alice only introduced [Beth]$_F$ to Carl.

(38) Alice only introduced Beth to [Carl]$_F$.

Such a truth-conditional effect of focus can be explained by either semantics or pragmatics. On a semantic theory, the truth-conditional interpretation is to take the focus-triggered question as an argument, so that shifting the focus will result in a different question being triggered, and thus a different truth-condition follows. This kind of account will straightforwardly assume that whenever there is a case displaying the truth-conditional

53

effect of focus, some expressions that are semantically sensitive to focus must be involved: such expressions are able to take FSVs in their scopes as arguments to determine the OSV of a whole structure. For example, in order to explain the above pair, the semantic approach would assume that in the semantics of (preverbal) *only* there is an argument place for the FSV triggered in its scope:

Semantics of *only*. *NP only VP* is true iff

(i) $[\![VP]\!]^O([\![NP]\!]^O)=$**true**,

(ii) $[\![VP]\!]^F \neq \{[\![VP]\!]^O\}$, and

(iii) for all $f \in [\![VP]\!]^F$, if $f([\![NP]\!]^O)=$**true**, then $f = [\![VP]\!]^O$.

In words, a sentence of the form *x only F'ed* is true just in case (i) x indeed F'ed, (ii) there is a focus involved in the VP *F'ed*, and (iii) among all the focus-triggered alternatives, *F'ed* is the only predicate that is true of *x*. This semantics for *only* correctly captures the truth-conditions of (37) and (38): (37) is true just in case Alice introduced Beth but no one else to Carl, but (38) is true just in case Alice introduced Beth to Carl but no one else.

Although this semantic theory only explains why focus affects the truth-conditions of "*only* sentences", the general methodology is clear: whenever we find that focus is able to affect the truth-conditional interpretation of a sentence, there must be an expression involved in the sentence, such as *only*, that is responsible for the phenomenon, in the way that it is semantically sensitive to focus-triggered FSVs. This kind of approach belongs to what Rooth (1992) calls *Weak Theories*, which are committed to the following hypothesis.

Weak Hypothesis. The focus effect on truth-condition is due to the fact that the OSVs of some expressions are semantically sensitive to FSVs in their scopes, so that FSVs participate in the composition of truth-conditional meaning.

However, I will not pursue this option, due to three concerns. First, such theories are not general, in the sense that we have to pack some semantic clauses referring to FSVs into the lexical semantics of a term once we find it displaying the focus effect. Second, weak

theories are not theoretically conservative, in the sense that it amounts to a drastic change of the truth-conditional semantics if we make it refer to the non-truth-conditional dimension. Third, what is more serious is that weak theories fail to appreciate the *optionality* of the focus effect, and thus wrong predictions ensue.

(39) Alice always takes [Beth]$_F$ to movies.

Alice always takes Beth to [movies]$_F$.

(40) Mary always remembers to go to [church]$_F$. (Beaver and Clark, 2003)

(39) shows that *always*, somehow similar to *only*, displays the focus effect. But if we follow the Weak Hypothesis to build focus sensitivity into the semantics of *always* in a similar way as the above semantics for *only*,[11] (40) is predicted to mean that whenever Mary remembers to go somewhere, it is a time when she remembers to go to church, which is, however, not what (40) intuitively means.[12] Examples like (40) show the optionality of the focus effect: the focus effect predicted by weak theories does not always appear. Consequently, weak theories are unattractive, as they would make the terms like *always* and *only* be obligatorily sensitive to focus.[13]

Now we need a *Strong Theory*, with the commitment to the following hypothesis.

Strong Hypothesis. For any expression or construction in natural languages, its truth-conditional semantics does not refer to focus. (cf. Rooth, 1992)[14]

---

[11]For example, *NP always VP* is true iff, if at any time there is an $f \in \llbracket VP \rrbracket^F$ such that $f(\llbracket NP \rrbracket^O)=$**true**, then it is a time when $\llbracket VP \rrbracket^O(\llbracket NP \rrbracket^O)$.

[12]It seems that, in (40), the focus effect on the domain of the quantifier *always*, predicted by weak theories, is trumped by the effect brought about by the presupposition triggered by *remember*.

[13]Although Beaver and Clark (2003) use (40) and other examples to argue that the weak hypothesis does not apply to *always*, they nevertheless show that parallel arguments are *not* applicable to *only*: the focus sensitivity of *only* seems indeed non-optional. However, this does not jeopardize my rejection of the weak hypothesis, so long as terms behaving like *only* (including, arguably, *even* and *also*; cf. Beaver and Clark, 2008) do not concern us here—my project is mainly on knowledge ascription sentences and counterfactuals.

[14]This formulation of the Strong Hypothesis might be excessively strong regarding my purposes, because whether some expressions, such as *only* and *even*, are semantically sensitive to focus does not matter to my current goal of explaining only the focus effect displayed by counterfactuals and knowledge ascription sentences.

Although the Strong Hypothesis is a negative claim, it implies that the focus effect on truth-conditions has to be explained by appeal to some pragmatic mechanisms beyond the realm of truth-conditional semantics. A general assumption accepted by many theories in the literature is that the truth-conditional effect of focus should be explained by i) the context-sensitivity of certain expressions and ii) a pragmatics of how focus is able to influence the resolution of some contextual parameters. Thus, focus is only to influence truth-conditions indirectly through certain pragmatic principles governing the relation between focus and context, as well as the ability of context to affect truth conditions. Thus, a major task of a strong theory is to postulate a certain pragmatic relation, call it *F-relation*, between context and focus.

At this point, it is worth pointing out that an F-marking in a sentence in fact induces multiple FSVs at various levels. For example, although there is only one F-marking on *Carl* in (35), multiple FSVs are generated respectively by the noun *[Carl]$_F$*, the VP *introduced Beth to [Carl]$_F$*, and the complete sentence *Alice introduced Beth to [Carl]$_F$*. Due to this fact, let an *F-domain* be an expression (a word, a phrase, a clause, etc.) at which level an FSV is generated.[15] The F-relation we are interested in, then, is defined as a relation between a context and an F-domain. The basic rationale in the common ground theory is that when a sentence is uttered, contexts can be resolved based on the F-markings involved, given the pragmatic constraint that some F-domain must be F-related to its context.[16]

For now, I am to leave the F-relation unspecified because the current goal is just to lay out a common ground framework without cashing out the details. Roughly speaking, particular strong theories are different ways to specify the F-relation, whereas the common framework they share can be captured by the following two general theses.

---

[15]The terms *F-relation* and *F-domain* are borrowed from Büring (2016).

[16]Different strong theories endorse different assumptions about what F-domains are required to be F-related to contexts: for some (e.g. Rooth, 1985, 1992), an F-marking in a sentence is pragmatically licensed once there exists an F-domain containing the F-marked constituent, be it a word, a phrase, or a complete clause/sentence, stands in the F-relation to its context of utterance; for some others (e.g. Roberts, 2012), the F-domains that are required to be F-related to contexts are always complete sentences.

Context Sensitivity. The truth-conditional interpretation of a focus-sensitive expression depends on certain contextual parameters whose values can be pragmatically resolved.

Pragmatic Constraint. An utterance of a sentence is felicitous only if a certain F-relation holds between an F-domain involved in the sentence and its context.

The two theses combined imply that focus could have a truth-conditional effect in the following way: given a focus-involving utterance, a hearer, in order to preserve its felicity, may accommodate, *à la* Lewis (1979), into the context whatever it takes for the F-relation to hold (or equivalently, for the above pragmatic constraint to be satisfied), and thus the truth-conditional content of the utterance is affected due to the context sensitivity of the uttered sentence. Theories built on the common ground framework can be found in Rooth (1992), von Fintel (1994), Kadmon (2001), Roberts (2012), and (partially) Beaver and Clark (2003, 2008). Lastly, it is worth noting that such strong theories allow the optionality of the focus effect because they appeal to the pragmatic process of accommodation, which is in principle optional and might be defeated by various pragmatic reasons.

3.4. Pragmatics of Focus

Given the above framework, my own theory amounts to a specification of the F-relation. Before formulating the view in detail, it is worth pointing out up front that my theory is only concerned with F-domains that are *complete clauses*. Consequently, the FSVs I am going to consider are only sets of alternative *propositions*, rather than sets of properties, individuals, or others. This restriction is due to the limited goal of this project concerning only the focus effect on counterfactuals and knowledge ascription, where relevant F-markings are contained in clauses embedded under conditional antecedent or the verb *know*, and relevant FSVs are supposedly obtained at the embedded clause level.[17]

---

[17]Consequently, this view actually has nothing to say about examples like (37) and (38) above, unless *only* is not interpreted *in situ* but raised up to take complete clauses as its scope.

My view on the F-relation is called *the Good Question-Answer View* (GQA). It is based on the idea presented in Chapter 2, that sentences and clauses are answers to their CQs. Such a CQ specifies multiple acceptable ways that a context is to be updated, and an answer to the question is to pick up one specific way out of the many. An understanding of the F-relation between a sentence or clause and its context is implicit in this characterization of the question-answering role sentences and clauses play. First of all, given a sentence or a clause that is an F-domain, the CQ of it is its FSV, and the answer it provides is its OSV, i.e. its propositional content. Considering the question-answering function of a clause in a context, the CQ of the clause specifies multiple ways of updating the context, and the clause's propositional content picks out one among them. Then, GQA is the view that such a CQ must be a pragmatically appropriate question to raise with respect to the context of the sentence or clause—it has to be a *good question*. In other words, for a context, not every set of propositions is appropriate for the purpose of specifying acceptable ways of updating the context. So, the requirement is that a sentence or clause's CQ should be appropriate in a certain way, with regard to the context of it. On the other hand, as a sentence or clause does not only raise a question but also answers it, there are pragmatic appropriateness concerns regarding the answer, i.e. the propositional content. For among many ways of updating a context as specified by a question, some of them are better than others. Then, the requirement on the answer's side is that the propositional content of a sentence or clause has to be a *good answer* to its FSV, with respect to the context. To summarize, GQA is the thesis that a sentence or clause, when used, is pragmatically required to have an FSV (i.e. CQ) that is a good question with respect to its context, and to have an OSV (i.e. propositional content) that is a good answer to the question with respect to its context.

F-Relation. For context *c* and clause *S*, the F-relation holds between them just in case

i) $[\![S]\!]^F$ is a *Good Question* with respect to c, and

ii) $[\![S]\!]^O$ is a *Good Answer* to $[\![S]\!]^F$ with respect to c.

Of course, what counts as good questions and answers in context is yet to specify and argue for, to which I will turn shortly. Before that, it is worth pausing to note that the catch-all word "good" is used here because I intend GQA to be an open-ended program that subsumes multiple pragmatic principles capturing various ways how questions and answers can be appropriate or inappropriate in contexts. What I will do here is to specify several principles of this kind, while leaving the list open for further additions or modifications. What is most important to this project is the proposal that when sentences and clauses are used, the role they play is not only about their informational contribution but also their question-answering properties. On that basis, together with some pragmatic constraints on how questions and answers have to be related to contexts, we can give a pragmatic account for certain intriguing phenomena. Now we turn to some of such pragmatic constraints.

3.4.1. Good Question: Non-Triviality. To see what questions and answers are pragmatically good, the core methodological point is that we can obtain the conditions that good questions and answers normally have to satisfy by looking at the goals they typically accomplish. Starting with questions, the idea is that a question serves to specify certain possible ways that contexts could be changed. Consider a context c, understood in the Stalnakerian way: it is a set including all possible worlds that are live options regarding what the actual world is like, given what a group of speakers commonly take for granted. Suppose (41) is the question that the speakers intend to answer in the context.

(41) Who did Alice introduce to Carl?

According to the Hamblin semantics, the question is the set of propositions of the form *Alice introduced x to Carl*, meaning that each proposition of this form is an answer to the question. Thus, the question typically serves to regulate the discourse to the effect that

59

the next move in the discourse is expected to pick out an answer to (41), and thus to add the corresponding piece of information to c. In other words, (41) carves up the possible ways the context c could evolve upon the next update of information: the next update is acceptable just in case it picks out an answer to (41), so that $c \cap p$ is a possible subsequent context of c only if p is an answer to (41). Given that the goal of a question is to specify possible context changes, whether a question is good with respect to a context should be assessed in terms of how appropriate it is as a good specification of such possible context changes.

There might be various aspects concerning this appropriateness, but a question can be inappropriate in an obvious way: if the current context entails the truth or falsity of every possible answer to a question, then the question is trivial and should not be raised in the context. In other words, the question does not provide a good selection of possible context changes, as no answer to the question will bring about a substantive change to the context without crashing it. Hence, a good question should not be such a trivial question.

> **Triviality.** A question Q is trivial in a context c if and only if c entails the truth or falsity of every answer to Q. Formally, Q is trivial in c if and only if for every $p \in Q$, either $c \subseteq p$ or $c \subseteq \neg p$.

Since a good question cannot be trivial, we have a necessary condition for good questions.

3.4.2. Good Question: Cover. Besides being trivial, there is another way in which a question is inappropriate. Suppose a given context involves a live option that is incompatible with every answer to the given question. That is, the context does not rule out the possibility that all the answers to the question are false. Formally, suppose we have a context c and a question Q such that for some $w \in c$, $w \notin p$ for every $p \in Q$. My claim is that Q is *not* a good question with respect to c in such a situation. To see that, first note that for each $w \in c$, it is a live option that is compatible with what the relevant speakers commonly take for granted. If $w \notin p$ for every $p \in Q$, then every answer to the question

60

Q is false at w. Therefore, updating c with any answer to Q will result in a new context that excludes w—that is, every answer will rule out w. However, if w is to be ruled out by every possible update given by the question Q, how could it be included as a live option in c in the first place? In this sense, the question Q cannot be an appropriate specification of all the possible updates of the context c, and thus Q fails to properly play the role of question in the context—it is not a good question.[18] Given this consideration, a good question should not preclude any live option in a given context:

Cover. For a set Q of propositions and a context c, Q Covers c only if $c \subseteq \bigcup Q$.

Pictorially speaking, the condition is satisfied when all the answers to a question jointly cover a given context set. I take Cover as a second condition for good question, alongside the condition of non-triviality.

By taking Cover to be a necessary condition for good questions, and by requiring a clause's CQ to be a good question for felicity, it means that for a clause of which the CQ is Q, the context of the clause has to entail the disjunction of all the propositions in Q, i.e. the disjunction of all the possible answers to Q. As the question is derived by substituting a focused constituent, the disjunction, note as $\bigvee Q$, is an existential proposition. For example, (42) triggers the question Q, which is the set of propositions of the form *Alice introduced x to Carl*. Cover requires the context of (42) to entail that at least one of the propositions is true, which is tantamount to the requirement that the context entails $\bigvee Q$, that Alice introduced someone to Carl.

(42) Alice introduced [Beth]$_F$ to Carl.

      Q: Whom did Alice introduce to Carl?

      $\bigvee Q$: Alice introduced someone to Carl.

In general, taking Cover as a necessary condition for good questions then is to claim that there is a focus-induced existential presupposition: a clause presupposes the proposition

---

[18]A similar argument can be found in Ciardelli et al. (2019, 21).

that is derived by substituting a variable for its focused constituent and then existentially binding that variable.

In the literature, it has been noticed by many that focus has a strong tendency to trigger such an existential presupposition, but the presupposition is not always present. Consider an example from (Rooth, 1999, 241).

In my department, a football pool is held each week, where people bet on the outcomes of games. It is set up so that at most one person can win; if nobody wins, the prize money is carried over to the next week.

(43) (A: Did anyone win the football pool this week?) B: Probably not, because it's unlikely that [Mary]$_F$ won it, and she's the only person who ever wins.

Given that "it is unlikely that" is a presupposition hole, the felicity of B's response shows that the existential presupposition in question is not present in this case. Rooth strengthens this point by comparing focus with *it*-cleft, which is presumably a paradigm device that triggers an existential presupposition by grammar.

(44) (A: Did anyone win the football pool this week?) B: # Probably not, because it's unlikely that it's Mary who won it, and she's the only person who ever wins.

B's response is infelicitous, and together with (43), the minimal pair show that focus and *it*-cleft should not be taken as triggering an existential presupposition in the same way. In particular, the existential presupposition of focus does not appear when it would result in incoherence in the context. In this sense, focus is a soft trigger of existential presupposition. On the other hand, *it*-cleft is a hard trigger, because its presupposition is uncancellable, as (44) illustrates.

Given that, focus tends to but does not necessarily invoke an existential presupposition. Consequently, Cover should not be a necessary condition for good question, but just a default rule.[19] By allowing exceptions, however, it is not to jeopardize the motivation for

---

[19]Perhaps one can insist that focus is an existential presupposition trigger and try to explain away the above example. For instance, we could say that the focus in (43) should be interpreted with respect to the

Cover. The condition of Cover is meant to capture what the *typical* function of questions requires of how a question is related to a context. As said above, this typical function is to specify possible ways of context change upon the next move in a discourse. When a question serves this function, it is reasonable to require that no live options in an original context be left out by the possible ways of context change as specified by the answers to the question. However, in (43), the question triggered by the sentence does not serve this typical function. It is clear that B claims that it is improbable that anyone won, and then she moves on to justify this point by saying "it's unlikely that [Mary]$_F$ won it, and she's the only person who ever wins." Although the question *Who won the football pool?* is triggered, it does not play the role of specifying possible context changes that comes next. What B intends to achieve by addressing that question is not to change the context; rather, she is merely entertaining this *who*-question and trying to undermine the plausibility of all its answers. In such a circumstance, the failure of Cover as an appropriateness condition is not a surprise, and this exception does not pull away the motivation for Cover in cases where a sentence or clause is supposed to answer a question that serves the typical function of question.[20]

_____

F-domain that includes not only the *that*-clause but a larger one: "it's unlikely that [Mary]$_F$ won it." If so, the existential presupposition triggered in this case will be that there is a person such that it's unlikely that she won the pool. This would explain why (43) does not sound like a case of presupposition failure, as this existential presupposition is probably satisfied in the context. If the existential presupposition view is viable, it means that the pragmatic condition of Cover is a necessary condition for appropriateness, and consequently it does not need to be qualified. This is a welcome for my purposes. But for those who do not agree with the existential presupposition view, I am happy to restrict the condition as on that works by default, as opposed to a necessary condition.

[20]As to why focus tends to but not necessarily triggers an existential presupposition, my account is in spirit the same as many others in the literature, e.g. Rooth (1992), Roberts (2012), Kadmon and Sevi (2011), and Abusch (2009). The common idea is that focus does not semantically contribute an existential presupposition, but merely a question, i.e. a set of alternative propositions. The existential presupposition is a secondary pragmatic effect of the question. In particular, such accounts roughly agree on the point the focus-triggered question tends to indicate what the current QUD is, and the existential presupposition is pragmatically derived due to the reasoning that when such a QUD is raised, then, *by default*, the context should be such that not all answers to the question are false. For example, as the FSV of (42) indicates that the topical question is *Who did Alice introduced to Carl?*, we are pragmatically motivated to infer that the current context is one that entails the existence of someone who was introduced by Alice to Carl—otherwise, the topical question should better be *Did Alice introduce anyone to Carl?*. This account is equivalent to mine if QUD equals "possible ways of context change". But the two notions do not coincide if local contexts are

3.4.3. Good Answer: Logical Strongness. Now we turn to good answers. The function that an answer typically serves is to update the context to a certain effect. We have seen that a question sets up possible ways that context can be changed, and an answer is just to pick out one and execute that context change. From this perspective, as the typical function of an answer is to result in a context change, it seems that all the answers to a question are equally good in the sense that each is able to update the context, in a way that is allowed by the question. However, there are cases in which some answers are better than others. This happens when multiple distinct answers to a question bring about the same update effect to a context. In such cases, such contextually equivalent answers may not be equally good. To illustrate, consider the following example.

> [It's common knowledge that Clyde goes to a party whenever Bertha goes. Suppose the person who answers the following question wants to express that both Bertha and Clyde will go to the party.]

(45) Q: Who will go to the party?

    A1: # Bertha will go.

    A2: Bertha and Clyde will go.

Given that (45Q) denotes the set of all the propositions of the form *x will go to the party*, both (45A1) and (45A2) are answers to it. Moreover, since the context entails that Clyde will go to the party if Bertha goes, (45A1) and (45A1) are contextually equivalent: if (45A1) is true at some w in the context, so is (45A2), and the reverse also holds because (45A2) logically entails (45A1). Given that, both answers have the same effect of updating the given context, as they both restrict the current context set to the worlds where both Bertha and Clyde will go to the party. However, the two sentences intuitively do not equally

considered. At least on Roberts's theory, QUD is only attributed to global contexts, while it makes less sense to say that an embedded clause is responding to a question "under discussion". In contrast, the notion of "possible ways of context change" is broader and is intended to apply to local contexts as well. On the GQA view, every clause changes its context based on a presumed specification of ways in which the context can be changed, which is regulated by focus and other question-triggers involved in the clause.

well serve the purpose of achieving that very update effect. In particular, (45A1) does not sound good if the speaker wants to restrict the context set to only the both-go worlds, as the interlocutor may have to ask a further question—*Will Clyde go as well?*—in order to make sure if the common knowledge in the prior context is indeed part of the information state that is shared by all the speakers, or otherwise it needs to be nullified or corrected. Given that (45A2) is logically stronger than but contextually equivalent to (45A1), we may think that (45A2) is providing more information than needed. But the bit of "redundant" information nonetheless serves to confirm the prior context, which explains why dropping it is not pragmatically adequate.

The above example shows how contextually equivalent but logically inequivalent answers may have a difference in appropriateness. In particular, a good answer should be a logically stronger answer than all its contextual equivalent answers, with the aim of providing more information that may confirm the prior context:

> **Logical Strongness.** For question Q and $p \in Q$, p is a Logically Strong answer to Q with respect to a context c only if for all $q \in Q$, if $p \cap c = q \cap c$, then $q \not\subset p$.

In words, an answer satisfies this condition only if it is one of the logically strongest among its contextually equivalent answers to the same question. Based on the above consideration, Logical Strongness is taken to be a necessary condition for good answers, as we pragmatically prefer, among contextual equivalents, those which carry redundant information that can confirm the prior context as much as possible.

Like Cover, Logical Strongness also needs to be qualified. In its current form, Logical Strongness cannot be a hard constraint on good answers, because in some cases a sentence or clause can fail to provide a Logically Strong answer, without invoking infelicity.

> [It's common knowledge that Clyde goes to a party whenever Bertha goes. Suppose the person who answers the following question wants to express that both Bertha and Clyde will go to the party.]

65

(46) Q: Who will go to the party? I know that Clyde will go if Bertha does.

A1: Bertha will go.

A2: ?? Bertha and Clyde will go.

In this discourse, the first speaker explicitly affirms that Clyde will go if Bertha does, so that it is clear now that the context unequivocally contains this piece of information. Then, (46A1) sounds a felicitous response, while (46A2), however, does not sound as good. But as in the previous example, (46A1) is not a Logically Strong answer, which shows that Logical Strongness should not be a hard constraint on good answers.

We can see in both (45) and (46), there is a key piece of information in the context that makes the two answers contextually equivalent: the information that Clyde goes to a party whenever Bertha does. The difference between the two cases is how entrenched this key piece of information is in the context. For (45), the piece of information is involved in the context, but it is presumably a highly defeasible part of the whole context: to the relevant speakers, it seems quite plausible for someone else to nullify this piece of information as part of what has been established in the context. When one answers the question (45Q) with (A1), it raises the concern that the speaker might have stopped accepting the key piece of information as part of the context. For if she still accepts it, there would be a contextually equivalent answer, (A2), which would not raise this concern. Given that (A1) brings about the confusion of whether the speaker still accepts that key piece of information, and since (A2) is contextually equivalent to it without invoking such a confusion, (A1) is not as good as (A2).

On the other hand, in (46), the crucial step in the above reasoning is missing. That is, in (46), that key piece of information is made explicit by the first speaker, so that it is no longer plausible to suspect that one may have stopped accepting it. Then, (A1) sounds felicitous in this case, while its competitor, (A2), has no advantage over it, and may be a worse choice because of its verbosity, compared to (A1).

66

Based on the contrast between the two cases, Logical Strongness should be a good answer condition that only applies to the comparison between answers whose contextual equivalence is established by defeasible information in the context. In other words, a good answer has to be Logically Strong, while the comparison ignores all other answers that are indefeasibly equivalent to it.

This proviso is informal, as the current formal framework only captures the informational aspect of contexts. Whether a piece of information in a context is defeasible is not a matter of what information a context encodes. Rather, it has to do with how firmly a group of speakers hold the piece of information as an unchallengeable part of the context. This might depend on various other aspects of a context that I will not explore here. In Chapter 4, this principle will be applied to counterfactuals, where the relevant local contexts we consider will be highly malleable and virtually no information should be indefeasible. Hence for my purposes, the proviso here is mainly to restrict the condition of Logical Strongness to a plausible scope so that it can be used later without the threat of obvious counterexamples.

3.4.4. Good Answer: Completeness. Besides Logical Strongness, another good answer condition is Completeness. In the literature, it is common practice to distinguish *Complete* answers from merely *Partial* answers. Intuitively, a complete answer is an answer which contextually entails the truth or falsity of every other answer, while a partial answer only has to contextually settle the truth-values of some. Given that each answer to a question entails at least itself, any answer to a question is a partial answer. However, not every answer is complete.[21]

> **Completeness.** For a question Q and $p \in Q$, p is a Complete answer to Q with respect to a context c iff for every $q \in Q$, $p \cap c \subseteq q$ or $p \cap c \subseteq \neg q$.

---

[21]This definition of partial and complete answers is, however, slightly different from the common practice. For example, Roberts (2012) defines these notions in the way that even non-answers (according to Question&Answer) can be partial or complete answers. But here I restrict partial and complete answers to only answers as defined in Question&Answer.

Intuitively, Complete answers are more informative than merely partial answers, and this informativeness, in general, makes a Complete answer pragmatically preferable to partial answers. Thus, a good answer by default should be a Complete answer. Then, if we take Completeness to be a condition for good answers, it amounts to a requirement on the informativeness of answers. Given the pragmatic consideration that more informative answers are generally preferred, I assume that being a Complete answer is also a condition for good answers.

Like the two previous conditions, Completeness needs a proviso. Sometimes a question may not be intended to solicit as much information as a Complete answer proffers. In a situation such as the following, the question raised, a so-called "mention-some" question, does not seem to require a complete answer.

(47) Q: Who has a pencil?

A: I have one.

The answer (A) is only a partial answer but it is felicitous. It is obvious that in this case, it is easy to identify the purpose of asking the question is to borrow a pencil from anyone around the speaker, rather than to start an inquiry that aims at a maximally informative answer that settles the question of which of the relevant people possess a pencil. Importantly, there is a practical purpose of raising the question, the satisfaction of which by a partial answer is considered to be the whole point of raising the question. Thus, the following sounds a fair qualification for when Completeness can be a good answer condition: in a context where the purpose of raising a question is to solicit information, a good answer to the question should be Complete.

3.4.5. Summary. What we have so far includes some conditions for good questions and good answers with certain qualifications. Combining all those, we have a formulation of the notions of good question and good answer.

Good Question. A question Q is a good question with respect to a context c only if

(i) Q is not Trivial in c, and

(ii) Q Covers c —unless the question does not serve the function of specifying possible ways c is to be changed.

Good Answer. For $p \in Q$, p is a good answer to Q with respect to a context c only if

(i) p is a *Logically Strong*—ignoring other answers to Q that are contextually equivalent to p where the equivalence is indefeasible; and

(ii) p is a *Complete* answer to Q with respect to c—unless the purpose of raising Q is not to solicit information that updates c.

Although this formulation looks somewhat unwieldy with all those provisos, the provisos nonetheless do not matter much for my purposes. As mentioned, the main concern of the GQA view is the interaction between an embedded clause's question-answer pair and its local context. Especially, in what follows we will examine some results of the GQA view when knowledge ascription sentences and counterfactuals are assessed. The provisos above, however, will not be relevant there.

So far, combining the above constraints, we have a somewhat detailed characterization of how focus, or the question-answer pair associated with it, should be related to context. Again, this should not be taken as a complete formulation, since there might be more conditions other than Non-Triviality, Cover, Logical Strongness, and Completeness. Based on what we have so far, the GQA view specifies an F-relation, although the notions of good question and good answer involved in it are programmatic and open to further augmentations or revisions:

F-Relation. For a clause S and a context c, the F-relation holds between them only if $[\![S]\!]^F$ is a good question and $[\![S]\!]^O$ is a good answer to $[\![S]\!]^F$ with respect to c.

In other words, the F-relation holds between a clausal F-domain S and its context c only if S raises a Non-Trivial question that Covers c and provides a Logically Strong and Complete answer to the question with respect to c, given that the qualifications are satisfied.

3.5. GQA: a conservativist take on semantics and more

The GQA view is a pragmatic theory regarding how contexts are supposed to interact with question-answer pairs contributed by clauses. Importantly, I claim that the pragmatic constraints proposed above apply to embedded clauses and their local contexts. The view enables pragmatic inferences about what a certain local context is like, based on the question-answer pair triggered by an embedded clause, which is determined by focus or other question-triggers. Thus, the result of a pragmatic resolution of what a local context is like is dependent upon its embedded clause's focus pattern, which consequently leads to an account of cases where focus affects the truth-condition of a sentence.

To repeat the example mentioned at beginning of this chapter, given an appropriate context such as [Marriage], the following counterfactuals appear to have different truth-values.

(48) If Clyde hadn't [married]$_F$ Bertha, he wouldn't have inherited the money.

(49) If Clyde hadn't married [Bertha]$_F$, he wouldn't have inherited the money.

To assess a counterfactual, one needs to consider a set of counterfactual situations, which is then subject to an update by the antecedent of the counterfactual. In other words, such a set of counterfactual situations constitute the local context of the antecedent. The antecedent is to update it to the effect that those where the antecedent is true are ruled in while others are ruled out, and the consequent is then assessed based on the resulted context. According to GQA, the antecedent and its local context have to be related in a certain way: the local context has to be a context in which the question and answer triggered by the antecedent (the FSV and the OSV) are a good question-answer pair. Given such a relationship, speakers are thus allowed to infer what the local context must be like, based on the given antecedent.

This suggests an account for the contrast between the above pair of counterfactuals. As their antecedents are focused in different ways, the questions corresponding to the two

antecedents are thus not identical, which in turn may affect how speakers infer what the local context is like, based on the Good Question and Good Answer conditions formulated above. The detail in such inferences will be explained in Chapter 5; however, given that the local contexts are resolved differently for the two antecedents, it is not a surprise that the truth-conditions of the above counterfactuals differ, as the truth-condition of a counterfactual, no matter what semantic theory one adopts, crucially depends on a selection of relevant counterfactual situations.

Therefore, GQA is primarily an account of the truth-conditional effect of focus on certain kinds of sentences, such as counterfactuals. Moreover, being a *pragmatic* account of the effect, it makes a focus-sensitive semantics of such sentences unnecessary. According to traditional semantic theories of counterfactuals and knowledge ascription sentences, focus is not part of the calculation of the semantic values of these sentences. In the subsequent chapters, we will see exactly how GQA pragmatically explains the truth-conditional effect of focus, without the need of changing such traditional semantic theories. In this regard, GQA serves to resist potential semantic revolutions that might be taken as necessary reactions to the truth-conditional effect of focus.

Beyond the realm of theoretic pragmatics and semantics, GQA also has implications in some relevant philosophical areas. As one may have noticed, sentences that display the focus effect are philosophically interesting. A debate in contemporary epistemology concerns the competition among different semantic models for knowledge ascription sentences. One of the proposals, known as contrastivism, is the view that knowledge is a three-place relation and its semantics for knowledge ascription is correspondingly committed to an extra argument place of the semantics of the verb "to know", in addition to an argument place for the knower and another for the content. This proposal is partially motivated by the focus effect displayed by knowledge ascription sentences. Based on the phenomenon that the truth-condition of a knowledge ascription can be changed by how

it is focused, the contrastivist claims that the three-place model is more plausible than the traditional two-place models. This contrastivist attempt is an instance of the semantic revolutions I mentioned, as it deviates from the traditional semantic theories of knowledge ascription in order to account for the focus effect. However, GQA resists such revolutions: the focus effect is to be explained pragmatically by GQA, and thus contrastivism, as an unnecessary drastic change of the semantics, loses its appeal. Therefore, GQA serves to fend off the contrastivist revolution by offering a way to defend a conservativist take on the semantics of knowledge ascription, according to which knowledge is simply a binary relation.

Besides, counterfactual is another philosophically interesting sentence type that displays the focus effect. The semantics of counterfactuals matters for many philosophical issues, such as causation, explanation, moral responsibility, etc., are closely related to counterfactuals, if not directly analyzable in terms of them. The fact that focus has a truth-conditional effect on counterfactuals might lead to a semantic revolution, too. For the traditional semantic theories, as in Stalnaker (1968), Lewis (1973), and Kratzer (1981a), by themselves do not predict such an effect. As briefly explained above, the GQA account implies that the local context of the antecedent of a counterfactual is subject to pragmatic resolution, which is based on the question-answer pair triggered by the antecedent. Hence, as the truth-condition of the counterfactual is dependent on what the antecedent's local context is, the truth-conditional effect of focus results. Therefore, GQA explains the effect in terms of pragmatic inferences based on the Good Question/Good Answer constraints, and thus undermines the necessity of a drastic change of the semantics for counterfactuals.

Furthermore, the application of GQA is not limited to pairs of counterfactuals that differ in truth-value due to focus. An important aspect of GQA is that the local context of the antecedent of a counterfactual is resolved based on the question-answer pair triggered

by the antecedent. This, I believe, also sheds light on some conundrums that have been taken by previous approaches as semantic puzzles of counterfactuals. For example, as will be discussed in §6.1, counterfactuals with a true antecedent and a true consequent (true-true counterfactuals) are invariably true, according to the standard Stalnaker-Lewis semantics. However, there are alleged counterexamples against it, resulting in attempts to revise the semantics to accommodate the falsity of some true-true counterfactuals. However, as will be clear, this semantic revolution would lead to an overhaul of the Stalnaker-Lewis semantics, of which a lot of desirable properties would be lost. I hence argue that, quipped with GQA, we do not have to treat this issue as a semantic puzzle. Rather, a pragmatic solution emerges from GQA, which accounts for the seeming falsity of true-true counterfactuals by appeal to the pragmatic constraints on how antecedents should be related to their local contexts.

Another issue, which has been thought of as a semantic puzzle as well, has to do with two inference patterns involving counterfactuals. Both SEA and SDA have some intuitive appeal, but they are jointly unacceptable because they together entail Antecedent Strengthening, an inference pattern that is rejected by most semanticists.

Substitution of Equivalent Antecedents (SEA). $\varphi > \psi \vdash \varphi' > \psi$, if φ is truth-conditionally equivalent with φ'.

Simplification of Disjunctive Antecedent (SDA). $(\varphi \vee \psi) > \chi \vdash (\varphi > \chi) \wedge (\psi > \chi)$.

Antecedent Strengthening (AS). $\varphi > \psi \vdash (\varphi \wedge \chi) > \psi$.

As the standard Stalnaker-Lewis semantics validates SEA but not SDA, there are disputes regarding whether it is the right choice in semantics. In §6.2, I will show that the examples where SDA seems appealing can be explained away if we pay attention to question-answer pairs triggered by the antecedents in those cases, together with the aforementioned pragmatic process of local context resolution based on GQA. Once we take the question-answer pragmatics into account, those puzzling cases are no longer calling for a semantic

revolution. That is, cases in which SDA looks plausible, do not require a new semantic theory, as they are explained away by GQA, a pragmatic account. Therefore, GQA is not only a pragmatic account of the focus effect on counterfactuals; also, by explaining the focus effect, it can dissolve puzzles that have been thought of as only solvable by new, revolutionary semantics. In this sense, GQA lends force to a conservative take on the semantics of counterfactuals, resisting the temptation of creating new semantic theories in reaction to counterexamples.

In sum, GQA is intended to be a novel pragmatic theory of the question-answering role of clauses and its truth-conditional effect, and a conservative proposal against semantic revolutions. It also matters to philosophers, since many sentence types that display the truth-conditional effect are philosophical interesting: great ramifications in various philosophical areas would follow if the semantics were to change, and it would be a hasty move unless we carefully consider possible pragmatic accounts of the truth-conditional effect. What I am offering, i.e. GQA, is such an account.

CHAPTER 4

Knowledge Ascription and Congruent Questions

Although epistemologists rarely agree with each other, this is almost a common ground that appears unquestionable: Knowledge-*that* is a binary relation between an agent and a proposition that serves as the content of knowledge. Accordingly, in the search for an appropriate linguistic model for knowledge ascriptions, it is natural to take "to know" as a binary verb, which is to combine a subject and a content *that*-clause to form a complete knowledge ascription sentence. But this view is challenged by *epistemic contrastivism*, which claims that binary ascriptions, such as "Moore knows that he has hands," is incomplete in itself, in the very same way that "I prefer ice creams" is incomplete—"Prefer ice creams *to what*?" Whether I prefer ice creams depends on what alternatives are under consideration: I prefer ice creams to pretzels but not to cinnamon rolls. According to the contrastivist, knowledge ascription is similar in that respect: it is true that Moore knows that he has hands rather than stumps, but that doesn't entail that he knows he has hands rather than, say, brain-in-a-vat hand-images. Then, knowledge ascription appears to be sensitive to contrasts: One may know a proposition with respect to some contrasts but not to others, just as one may prefer ice creams to some contrasts but not to others. Hence, the contrastivist claims that knowledge is a *ternary relation* between an agent, a content proposition, and a contrast proposition, in the same way as preference is a ternary relation between an agent, an object that is preferred, and an object that is preferred to. It is natural, according to the contrastivist, to assume that the best linguistic model for knowledge ascription sentences should treat "to know" as a ternary verb, and the complete form of knowledge ascription is thus *Kspq*, which says "*s* knows that *p* rather than *q*."

75

As it appears, contrastivism constitutes a radical move. If it were right, then we would have to leave altogether the previous debates, which have mostly taken for granted that "to know" deserves a binary semantics. However, this move is a semantic revolution based on evidence that is better explained by pragmatics—specifically, GQA, a pragmatics of how clauses interact with their local contexts. Thus, in this chapter, I argue against epistemic contrastivism, defending the received wisdom that an appropriate linguistic model for knowledge ascription is binary, so that the debate should be restricted to binary theories.

There are two kinds of arguments allegedly supporting epistemic contrastivism. First, as illustrated above, our truth-value judgments about knowledge ascription sentences shift according to different contrasts provided, and this intuitive contrast-sensitivity is supported by empirical evidence (Schaffer and Knobe, 2012). As epistemic contrastivism predicts this empirically testified shiftiness, this lends some strength to the position on grounds of empirical adequacy.

Second, the contrastivist also argues that there are desirable theoretical consequences if contrastivism is true. In particular, it is promised that contrastivism can offer a plausible version of epistemic closure that does not necessitate the sacrifice of *epistemic modesty*—that one knows that she has hands but doesn't know that she is not a brain in a vat. Epistemic modesty is clearly desirable, insofar as we are not willing to accept either skepticism, according to which one doesn't know that she has hands, or dogmatism, according to which one knows that she isn't a brain in a vat. Some sort of epistemic closure is also necessary, as knowledge does extend by logical entailment, in the sense that if one knows p, she doesn't need to gather any new evidence in order to know some logical consequence q of p—if she knows logic well and forms new beliefs correctly based on logical inference, all the evidence that justifies p also justifies her belief in q. Thus, converting to contrastivism seems an immensely attractive option, if it can secure both epistemic modesty and closure.

These arguments come with different strengths and might be understood as supporting different variants of contrastivism. If the view of contrastivism under consideration is about the semantic interpretation of our ordinary knowledge ascription, then only the first kind of arguments are able to provide proper supports, as they are concerned with how actual knowledge ascription sentences behave. On the other hand, the second kind of arguments concern which concept of knowledge can achieve some theoretical goal, i.e. resisting skepticism and dogmatism while preserving epistemic closure. If arguments of this sort are effective, it only follows that anyone sharing the theoretical goal in question will wish that contrastivism were true, which may or may not be in line with the concept of knowledge that we actually use when saying someone knows something. But this distinction will be ignored in this paper, as my goal here is to object to both kinds of arguments for contrastivism. As will be clear later, contrastivism is neither empirically adequate as a descriptive linguistic theory of knowledge ascription sentences nor theoretically desirable as a way out from philosophical conundrums.

Although I argue that contrastivism is not sufficiently supported empirically, the data it alludes to are still illuminating: at least, they show that knowledge ascriptions appear to be sensitive to contrasts. Thus, with contrastivism rejected, we still need an alternative account of why knowledge ascription sentences behave as such. I will propose a binary treatment within the contextualist framework. The account aims at explaining the apparent contrast-sensitivity by appeal to GQA. First, it adopts a contextualist semantic framework for knowledge ascription sentences, according to which the content of such a sentence depends on a contextual parameter that can be resolved by some pragmatic processes. Second, it refers to the pragmatic rules in GQA to explain how contrast-triggering constructions (e.g. *...rather than...*) can affect the contextual parameter to which knowledge ascription sentences are sensitive. Hence the apparent contrast-sensitivity is explained

through a pragmatic detour: contrastive constructions are able to influence some components in the context, which in turn affect the truth-conditional content of knowledge ascriptions.

The organization of this chapter is as follows. In §4.1, I will briefly lay out the basic tenets of epistemic contrastivism. The desirability-based motivation for contrastivism is explained and argued against in §4.2; we will see that contrastivism does not reconcile the apparent tension between epistemic modesty and closure. §4.3 turns to its purported empirical evidence, and I argue that empirical data are in fact against the view. In §4.4, I provide a GQA-based account of the apparent contrast-sensitivity of knowledge ascription. Specifically, after summarizing some success conditions for a positive account in §4.4.1, the GQA-based account of the contrast-sensitivity is formulated and applied to the relevant data in §4.4.2.

4.1. Contrastivism

Since epistemic contrastivism is formulated and motivated in various ways in the literature, I will not try to survey and lay out every variant of it, but only focus on two prominent aspects of the view.[1] First, epistemic contrastivism is a view on *the concept of knowledge*, according to which knowledge is a ternary relation. Call this thesis TERNICITY.

> **Ternicity.** Knowledge is a ternary relation of the form *Kspq*, which holds between
> an agent *s*, a content proposition *p*, and a contrast proposition *q*.[2]

---

[1]Major advocates of epistemic contrastivism include Karjalainen and Morton (2003), Sinnott-Armstrong (2004, 2008), Schaffer (2004, 2005, 2008), and Morton (2012).

[2]Contrastivists may not have a consensus on what "knowledge is a ternary relation" means. For example, Schaffer (2004) defends the view that this ternary relation is the denotation of the verb "to know" in our ordinary language, according to which contrastivism amounts to a descriptive theory of our knowledge ascription language and the concept of knowledge that is denoted by it. On the other hand, Sinnott-Armstrong argues for a revisionist version of contrastivism, which indicates that the ternary relation is what we should make use of if the goal is to "describe a person's epistemic position as precisely as possible" (2008, 268). This paper is focused on the descriptive version of contrastivism that is put forth in a series of papers by Schaffer.

Although Ternicity is only a claim about the form of the knowledge relation and a complete contrastive theory of knowledge ultimately needs a plausible analysis of this ternary relation *K*, the evaluation of contrastivism in this paper does not depend much on such an analysis. But roughly speaking, Ternicity is inspired by the idea that knowledge is relative to a class of contrast propositions, or "relevant alternatives", in the sense that one's knowing something amounts to a certain class of alternative propositions being ruled out by her epistemic state (cf. Dretske, 1972). Assuming an internalist view, the epistemic state of an agent may include everything she has as evidence, such as her sensory experiences, memory, etc., and that epistemic state's ruling out a proposition means that in every possible world where the agent has exactly that epistemic state as she actually has, the proposition is false.

Contrastivism can be understood as an approach that treats relevant alternatives as an integral part of knowledge, so what one knows is not only some content proposition *simpliciter* but rather a content proposition *in contrast to* a certain class of alternatives. Hence, to know is always to know some *p* by ruling out some propositions that are incompatible with *p*, the disjunction of which is the contrast proposition *q*. One's ruling out one class of contrasts does not mean that she can rule out others; thus, when ascribing knowledge, we are ascribing a ternary relation *Kspq*, meaning that *s knows that p rather than q*; this, however, by no means implies that *s* also knows *p* rather than *q*'. For example, Moore knows that he has hands rather than stumps, but he doesn't know that he has hands rather than BIV hand-images, because his epistemic state does rule out the stump-possibilities but in no way rules out the BIV possibilities.

Ternicity on its own is not sufficient for distinguishing contrastivism from its closely related view, i.e. contextualism. Knowledge ascription sentences, in ordinary language use, are often binary in the surface form *s knows that p*, which does not appear to involve a contrast proposition. Contrastivists, such as Schaffer (2004), take the binary surface form

as incomplete and claim that the contrast proposition is implicitly fixed by the context in such cases. Then, the truth-value of an apparently binary knowledge ascription sentence will shift with contexts, as different contexts may supply different values for the contrast. For example, an ordinary context may provide the binary ascription in (50) with a contrast proposition that just involves some ordinary alternatives: that Moore has stumps or hook hands or scissor hands.

(50) Moore knows that he has hands.

Insofar as these alternatives are ruled out by Moore's epistemic state, (50) is true when uttered in an ordinary context. However, a different context may supply a different contrast proposition. For instance, if we are in a "philosophical" context, the contrast proposition will be a different one that involves some skeptic alternatives, such as that Moore has BIV hand-images. As he cannot rule out the skeptic alternatives, (50) is not true when uttered in this context.

In this sense, epistemic contrastivism is similar to contextualism, because they both allow the content of a (binary) knowledge ascription sentence to shift with contexts. Moreover, if the contrast can only be determined contextually, then contrastivism is nothing more than a particular version of contextualism. However, Schaffer (2004) distinguishes contrastivism from contextualism in terms of how contrasts are fixed.[3]

> [C]ontext-dependence is generated by the absence of an explicit setting for q, which is then implicitly saturated by different alternatives in different contexts. The model for this mechanism is 'prefers', which licenses reduced expression (such as 'Ann prefers chocolate') where the covert foil takes different values in different contexts. (Schaffer, 2004, 82–83)

---

[3]Sometimes the contrastivist takes the view as a variant of contextualism in the sense that they both allow the shiftiness of binary knowledge ascription: "This is because the contrastive view allows that one ascriber could truly say 's knows that p,' while a second ascriber in a second context (with a different range of relevant alternatives) could truly deny 's knows that p.'" (Schaffer and Knobe, 2012, 687, italics mine.) However, this is just a terminological difference. I will use the term "contextualism" in a stronger sense that excludes contrastivism, and the distinction is made by the thesis of SATURATION below.

Since Ternicity is only a claim about the *concept* of knowledge, what the contrastivist needs in addition is an account of what determines the contrasts for knowledge ascription *sentences*. According to the above excerpt from Schaffer, the linguistic model of the verb "to know" directly reflects the ternicity of the concept of knowledge, in the same way as "to prefer" directly reflects the ternary relation of preference. Based on this view, there are two ways that the contrast proposition is determined in knowledge ascription. First, if a knowledge ascription sentence is ternary in its surface form, the contrast proposition is directly determined by the contrast clause. For example, the *rather-than* clause in (51) determines the contrast proposition to be that Moore has stumps.

(51) Moore knows that he has hands rather than (he has) stumps.

(52) Moore prefers ice creams to pretzels.

Nothing is mysterious in this case: the contrast proposition is explicitly determined by the *rather-than* clause, just as the third argument in the preference relation expressed in (52) is determined by the explicit complement phrase.

However, for knowledge ascription sentences without *rather-than* or similar constructions, the contrast is determined by the context. For example, when asked if Moore can tell whether hands or stumps are attached to the end of his arms, we can leave out the contrast and just say (53) to expresses the same content as (51). This is comparable to apparently binary preference ascription sentences, such as (54). When the context is clear—e.g. when asked which of ice creams and pretzels Moore prefers, (54) is felicitous to utter and in effect expresses the same content as (52).

(53) Moore knows that he has hands.

(54) Moore prefers ice creams.

Thus, just as the ternicity of the preference relation is compatible with apparently binary preference ascription, the contrastivist treats binary knowledge ascription in the same way.

81

Summarizing the two ways in which the contrast is determined, the contrastivist is committed to the following linguistic claim:

> **Saturation.** The verb "to know" denotes the ternary relation *Kspq*. In a knowledge ascription sentence, if there is a contrast expression (e.g. *rather-than*), the value of *q* is saturated by it; otherwise, the value is determined by the context.

With Saturation, contrastivism can be separated from contextualism, even though both allow knowledge ascription to be context-dependent. In short, binary ascription is the paradigm for the contextualist, and the shiftiness of binary ascription sentences is explained by the context-dependency of the verb "to know". For a ternary ascription sentence, insofar as the contextualist allows the explicit contrast expression to affect the context, the truth-conditional effect of the contrast expression can be explained by the context-dependency of "to know". On the other hand, the contrastivist views ternary ascription as the paradigm and explains binary ascription, and the context-dependency thereof, in terms of the context playing the role of completing an incomplete sentence.

With the view of contrastivism laid out and separated from contextualism, the next step is to examine some arguments for it. In the literature, contrastivism is motivated and argued for from various perspectives. It is infeasible to review them all. I will focus on and object to arguments that come in two kinds—one arguing that contrastivism leads to desirable consequences, and the other appealing to empirical evidence of speakers' truth-value judgments.

4.2. Modesty and Closure

4.2.1. The tension and the contrastivist solution. Skepticism is false: we do have a lot of ordinary knowledge, e.g. Moore knows that he has hands. Dogmatism is false, too: we do not know whatever is not supported by evidence, e.g. Moore does not know that he is not a BIV. Hence EPISTEMIC MODESTY:

Epistemic Modesty (EM). Moore knows that he has hands. Moore does not know that he is not a BIV.[4]

On the other hand, knowledge is closed under entailment, in the sense that if someone knows that p, and if that p entails that q, then she is in the position to know that q based upon deductive inference. Hence EPISTEMIC CLOSURE:

Epistemic Closure (EC). If one knows that p and if that p entails that q, then, *with certain other conditions satisfied*, she knows that q.

The spirit of EC is that knowledge extends freely under entailment, but it is not completely free: we do know something without knowing all its logical consequences. To extend knowledge from p to its logical consequence q, the subject arguably has to know that p entails q, deduce q from p, and thereby come to believe that q. Such conditions, and perhaps more, are encoded in EC with the "other conditions" proviso. But all these conditions are about how logically competent the subject is and how she comes to believe that q solely based on her knowledge of p as well as logic. Thus, in principle, knowledge extends along entailment, in the sense that no new evidence is needed for such extension.

The above two theses have been extensively motivated either by intuition or theoretical considerations in epistemology, even though they are still disputable. I shall not defend them, but just assume that epistemologists, in general, would prefer them being true rather than being false.

Famously, there is a tension between the two theses. First, by the first clause in EM, Moore knows that he has hands. Suppose that Moore knows that having hands entails not being a BIV, he comes to believe that he is not a BIV by inference from that he has

---

[4]Of course, this statement is only a particular instance of the general thesis of Epistemic Modesty. I use Moore's hand-knowledge and BIV-ignorance just for concreteness, without bothering to give a general characterization of what counts as ordinary or skeptical.

hands, etc.. That is, Moore satisfies all the "other conditions", with respect to the hand-proposition and the not-BIV-proposition. Then by EC, it follows that Moore knows that he is not a BIV, contradicting the second clause in EM.

The most notable solution to this problem is from epistemic contextualism.[5] Roughly speaking, the contextualist claims that whether we can truthfully ascribe knowledge or ignorance to some person depends on some contextually given standard, and EC holds when the context is fixed. If the context sets a very high standard for knowledge, Moore doesn't know that he isn't a BIV and doesn't know that he has hands either. But when the standard is low, as in ordinary contexts, Moore knows that he has hands and also knows he is not a BIV.[6] In short, the contextualist attempts to reconcile the tension between EM and EC by i) restricting the application of EC to cases where the relevant parameters in the context are fixed and ii) rejecting that the two clauses in EM are true in the same context.

Schaffer (2004, 2007a) argues that contrastivism can also solve the problem. Given Ternicity and Saturation, the binary principle of closure EM needs an overhaul. In particular, Schaffer (2007a) advocates the following two principles as what capture the closure properties of knowledge.[7]

---

[5]As both contextualism and contrastivism are views on the *semantics* of knowledge ascription, their solutions are actually at the metalanguage level, explaining why, on the one hand, "Moore knows that he has hands" is true "Moore knows that he is not a BIV" is false, and, on the other, some version of Epistemic Closure holds. Strictly speaking, both Epistemic Modesty and Closure should be formulated at the metalanguage level, in order for them to be exactly what the two views aim to save. To confess, I will freely speak at the object language level for the sake of easiness of expression, while inviting the reader to recover the semantic ascent readings of relevant sentences when appropriate.

[6]This doesn't mean that we can truthfully say that Moore knows that he isn't a BIV even in an ordinary context, as according to some contextualists (e.g. Lewis, 1979), mentioning the BIV possibility would change the context into a high-standard one.

[7]There are two more principles that concern with multi-premise closure properties—what extension of knowledge is allowed given multiple pieces of knowledge that one has:

INTERSECT-p: If one knows that p_1 rather than q, and if she also knows that p_2 rather than q, then, *with certain other conditions satisfied*, she knows that $p_1 \wedge p_2$ rather than q.

UNION-q: If one knows that p rather than q_1, and if she also knows that p rather than q_2, then, *with certain other conditions satisfied*, she knows that p rather than $q_1 \vee q_2$.

But since only single-premise closure will be discussed here, these two principles are omitted.

Expand-*p*. If one knows that *p* rather than *q*, and if *p* entails *p*′ and *p*′ is incompatible with *q*, then, *with certain other conditions satisfied*, she knows that *p*′ rather than *q*.

Contract-*q*. If one knows that *p* rather than *q*, and if *q* is entailed by *q*′ and *q*′ is not necessarily false, then, *with certain other conditions satisfied*, she knows that *p* rather than *q*′.[8]

In the spirit of Epistemic Modesty, we ascribe ignorance to Moore by saying that he doesn't know that he isn't a BIV. But the contrastivist would point out that the complete form of this ascription is that Moore doesn't know that he isn't a BIV *rather than he is a BIV*. This ascription is true because Moore can't rule out the BIV possibilities. Furthermore, the contrastivist emphasizes that nothing forbids Moore from having ordinary knowledge. By the above contrastivist version of closure, esp. Expand-*p*, what follows from Moore's non-BIV ignorance is only that he doesn't know that he has hands *rather than he is a BIV*. Indeed, Moore *is* ignorant of this, because again, he can't rule out the BIV possibility. However, this is not what we ascribe to Moore as his ordinary knowledge. Rather, the contrastivist continues, in ascribing the hand-knowledge to Moore, we mean that Moore knows that he has hands *rather than stumps* (or other ordinary, non-skeptical contrasts). Then, according to contrastivism, this account successfully reconciles the tension between epistemic closure and epistemic modesty, while accounting for the intuition that backs them: knowledge extends by entailment, and Moore has the ordinary knowledge that he has hands rather than stumps, as well as the non-BIV ignorance that he is not a BIV rather than he is a BIV, at the same time.

---

[8]According to Expand-*p* and Contract-*q*, any knowledge ascription that is derived from one of the two principles should be such that the content and the contrast are incompatible and the contrast is not necessarily false. These limits come from the requirement that the contrast proposition has to be a genuine contrast to the content. On the one hand, if the contrast is compatible with the content, then it does not make sense to say that one knows the content by ruling out the contrast. On the other hand, if the contrast is necessarily false, then it is not a genuine contrast that one needs to rule out in order to know the content. See Schaffer (2007a, 242–244).

But given that having hands entails not being a BIV, does the contrastivist version of closure implausibly extends Moore's ordinary knowledge so that he knows that he isn't a BIV rather than he has stumps, according to Expand-p? The answer is no. Notice that Expand-p only applies when the entailed proposition is incompatible with the contrast. Since not being a BIV is compatible with having stumps, Expand-p doesn't lead to this implausible result. Indeed, the contrastivist requires every true knowledge ascription to have incompatible content and contrast propositions, because the latter is supposed to be a *contrast*, i.e. an alternative, to the content. Thus, Expand-p seems to extend knowledge within some plausible limits, without judging the subject as knowing what they actually do not know.

A new feature of this contrastivist version of closure is that knowledge is now closed under entailment along two dimensions. As knowledge ascription involves two argument places for propositions, it is not only closed under entailment at the content position but also at the contrast position, as stated in Contract-q. Intuitively, if one knows that p rather than q, then her epistemic state rules out every q-possibility. If q' entails q, every q'-possibility is *a fortiori* a q-possibility and thus already has been ruled out by the subject, so that she knows that p rather than q', too. For example, since Moore knows that he has hands rather than stumps, it follows naturally that he knows that he has hands rather than brown stumps.

4.2.2. Old problems replicated (twice over). Both Expand-p and Contract-q, however, are challenged in the literature: the contrastivist closure does not only leave the old tension between modesty and closure intact but also duplicates the same problem into the contrast argument place.

To see the problem, suppose one has some ordinary knowledge where the content is p and the contrast q, where p and q are incompatible. Accordingly, we can design a skeptical statement: one is deceived by an evil demon into believing that p while in fact it is the

case that q. For Moore's ordinary knowledge that he has hands rather than stumps, (55) is its corresponding skeptical proposition:

> (55) Moore is deceived by an evil demon into believing that he has hands while in fact he has stumps.

Let $S(p,q)$ be the skeptical proposition generated by p and q. As Hughes (2013) shows, we can make use of this proposition to form a counterexample to Expand-p. $S(p,q)$ is incompatible with p. Thus, $p \wedge \neg S(p,q)$ is equivalent to and thus entailed by p. Furthermore, $p \wedge \neg S(p,q)$ is incompatible with q. By Expand-p, we can infer (57) (brackets are just for readability) from Moore's ordinary knowledge (56).

> (56) $Kmpq$: Moore knows that he has hands rather than he has stumps.

> (57) $Km(p \wedge \neg S(p,q))q$: Moore knows that [he has hands and is not deceived by an evil demon into believing so while he in fact has stumps] rather than he has stumps.

This problem is not much different from the old problem of closure that we encountered. The old problem is that every ordinary knowledge content entails the negation of a skeptical proposition, which is, according to modesty, not known. Now, the contrastivist closure doesn't entail that Moore knows the negation of the skeptical proposition itself, i.e. $Km\neg S(p,q)q$, because $\neg S(p,q)$ is compatible with q. But it is not a substantial improvement, as the incompatibility condition is easily satisfied by the conjunction $p \wedge \neg S(p,q)$, and thus $Km(p \wedge \neg S(p,q))q$, i.e. (57), is entailed. However, this resulted knowledge ascription seems false, as our intuition of epistemic modesty doesn't approve this ascription.[9] Call this Problem-p.

Problem-p generalizes immediately. Even if we don't consider those skeptical possibilities, it is easy to see that Expand-p extends knowledge in a way that may be illicit.

---

[9]In Hughes (2013, 586–589), what functions as the skeptical proposition is that Moore (or any other subject) is a BIV. I think both counterexamples work well, but the skeptical proposition I choose here better reflects the role it plays: $S(p,q)$ is a proposition that the subject in question cannot differentiate by evidence from q while being incompatible with p.

Adapting an example from Dretske (1970), starting from the ordinary knowledge (58), we can infer (59) by using Expand-p.

(58) Moore knows that it is a zebra rather than a mule.

(59) Moore knows that it is a zebra and not a cleverly disguised mule rather than a mule.

(59) is intuitively false (given that Moore is no expert in zoology) and apparently not modest. Indeed, it requires more evidence possessed by Moore than (58) does, so it should not be a consequence of (58) that is guaranteed by a plausible version of Closure.

Problem-p seems to be enough to show that the old problem of closure reappears in the contrastive version. Even worse, the contrastivist closure is subject to another way of illicitly expanding knowledge, as it adds an extra argument place to knowledge ascription. This time, instead of adding the negation of the skeptical proposition into the content position, Kelp (2011) notices that we can add the skeptical proposition itself to the contrast position. By doing that, the contrast becomes logically stronger, and thus Contract-q entails that the resulted knowledge ascription is true.

(60) $Kmp(q \wedge S(p,q))$: Moore knows that he has hands rather than he is deceived by an evil demon into believing so while he in fact has stumps.

(60) is intuitively false and immodest because it implies that Moore can rule out a skeptical possibility. Call this Problem-q, since it shows that the contrastivist closure illicitly expands knowledge at the q-position. Again, this problem does not necessarily rely on skeptical possibilities in a strong sense.

(61) Moore knows that it is a zebra rather than a cleverly disguised mule.

Contract-q entails that (61) follows from (58) because being a cleverly disguised mule entails being a mule. But it's not right, since Moore isn't supposed to have evidence to distinguish a zebra from a cleverly disguised mule.

Granted, the presence of counterexamples is not the end of the world for a theory, as some modifications might rectify the problem while preserving the core spirit. But the above problems run deeper, which shows that the contrastivist is heading in the wrong direction in response to the original problem of closure.

Let us first focus on Problem-q. By adding a third argument place for "to know" and allowing it to be saturated by explicit contrast expressions, the contrastivist then allows the elimination of the shiftiness of knowledge ascription by spelling out explicitly what the contrast is. Hence, there seem to be some ordinary contrasts that a subject normally rules out, e.g. that Moore has stumps, as well as skeptical contrasts that people can't rule out, e.g. that Moore is a BIV. In pursuing modesty, the contrastivist deems that ordinary knowledge ascription is true because they have ordinary contrasts, and also this piece of ordinary knowledge doesn't entail immodest knowledge, since ruling out an ordinary contrast doesn't entail that any skeptical contrast is thereby ruled out. But this is not right. As (60) illustrates, even an apparently ordinary contrast, such as that Moore has stumps, is entailed by a skeptical contrast, that Moore is deceived by an evil demon into believing that he has hands while he in fact has stumps. Thus, it seems that the contrast proposition that Moore has stumps is not really an "ordinary" contrast after all, as some stump-possibilities are skeptical scenarios that people can't rule out.

The contrastivist might concede that the seemingly ordinary knowledge ascription (56) is actually not ordinary and thus not true, as Moore cannot rule out *every* stump-possibility. Then the question is, what is a genuine instance of ordinary knowledge ascription? There is hardly any. Considering how we constructed a skeptical proposition out of the seemingly benign ordinary contrast that Moore has stumps, the strategy obviously applies to most, if not all, other apparently ordinary contrast propositions as well.[10] Hence, the contrastivist

---

[10]The strategy does have a limitation. Assume for the sake of argument that there is a sentence, call it q, which uniquely describes a possible world. Then in $Kspq$, q is a genuinely ordinary contrast, provided the subject can rule out q. However, if we are left with only this kind of "ordinary" knowledge, contrastivism apparently loses its attraction to many of us who want to defend modesty.

would have to take the vast majority of knowledge ascriptions that we intuitively think are ordinary and true as false ascriptions. This is a typical skepticist result, directly violating our intuition of modesty.

On the other hand, Problem-p illustrates how our intuitive judgments on a knowledge ascription are sensitive to not only what the content proposition is but also the way in which it is expressed. Although p (that Moore has hands) entails $\neg S(p,q)$ (that it's not the case that he is only deceived into believing that he has hands while in fact he has stumps) so that p and $p \wedge \neg S(p,q)$ express the same proposition, the two express it in very different ways. In particular, while p plainly states that Moore has hands and thus picks out all the hand-possibilities, $p \wedge \neg S(p,q)$ picks out the same set of possibilities by additionally mentioning, and excluding, some skeptical possibilities, i.e. those where Moore is handless but deceived to believe that he has hands. As Problem-p shows, turning the knowledge relation into a ternary one does not help account for this phenomenon.

However, the contrastivist might want to treat the *and-not* construction in (57) as a syntactic device that signals contrast, in the same way that *rather-than* does. If so, the logical form of (57) can be treated along the following line:

(62) $Kmp(S(p,q) \vee q)$: Moore knows that he has hands rather than [(he is deceived by an evil demon into believing so while he in fact has stumps) or (he has stumps)].

If that is the case, then Expand-p does not allow the inference from the ordinary ascription (56) to (62). I have no qualm with this solution itself: If *rather-than* signals a contrast, why can't *and-not* also do? But this move seems to be a misdiagnosis of the problem. The problem is actually not due to the conjunct $\neg S(p,q)$ being added to the p-position, but rather because of a skeptical possibility's being brought up in some way into salience. Compare the ordinary ascription (56) and the following (63).

(63) Moore knows that he has hands, which BIVs can't have, rather than stumps.

90

Unlike the ordinary ascription (56), (63) sounds false, while the only difference between the two is the non-restrictive relative clause "which BIVs can't have" present in (63). The contrastivist might be tempted to treat the non-restrictive relative clause, again, as some contrast-signaling device like *rather-than*. But this move is not in line with the mainstream semantic theories of non-restrictive relative clauses. Such clauses are viewed as parenthetical, in the sense that they do not participate in the composition of the truth-conditional meaning of the whole sentence in which they are embedded (e.g. Potts, 2005). In other words, the constituents in the main clause of (63) are not supposed to take the relative clause to fill in any argument places. Then, how could the clause in (63) provide the value of the contrast for "to know"?

For ones who embrace the spirit of both epistemic closure and epistemic modesty, the above arguments show that contrastivism is not a desirable option. But if one rejects either closure or modesty entirely, these arguments don't make any recommendation as to whether contrastivism is a plausible theory. The next section is devoted to arguing that contrastivism is not an empirically adequate theory of knowledge ascription, so that we should reject the theory no matter whether we intend to save closure and modesty.

4.3. Empirical Adequacy

4.3.1. The empirical data. In a survey done by Schaffer and Knobe (2012), participants are asked to read the following [Jewel Thief] vignette and then judge to what degree they agree with various knowledge ascriptions. The results show that the truth-value judgments of speakers are sensitive to contrasts triggered in several different ways. In what follows, I use '#' to mark what the participants tend to disagree with, according to the data.

[Jewel Thief] Last night, Peter robbed the jewelry store. He smashed the window, forced open the locked safe, and stole the rubies inside. But Peter forgot to wear gloves. He also forgot about the security camera. Today, Mary the detective has been called to the

scene to investigate. So far she has the following evidence. She has been told that there was a theft, she has found and identified Peter's fingerprints on the safe, and she has seen and recognized Peter on the security video, filmed in the act of forcing open the safe. She has no further information.

(64a) Mary knows that Peter rather than anyone else stole the rubies.[11]

(64b) # Mary knows that Peter stole the rubies rather than anything else.

(65a) Mary knows who stole the rubies.

(65b) # Mary knows what Peter stole.

The above two pairs are cases where contrasts are signaled explicitly by the *rather-than-* construction and *wh*-clauses. Limiting the scope of this paper to knowledge-*that* ascription, I simply list (65a) and (65b) here without discussing them.[12]

In addition, (66a) and (66b) come as a pair of utterances where different contrasts are implicitly given by different QUDs provided in [C1] and [C2].

[C1] Everyone is now asking the big question: Who stole the rubies? The news reporter is about to write a story about Mary. He is wondering if Mary now knows who stole the rubies. He writes:

(66a) Mary knows that Peter stole the rubies.

---

[11]Originally in Schaffer and Knobe (2012), these sentences read 'Mary now knows...' But in the current formulations the word 'now' is dropped. Given that Gerken and Beebe (2016, 139–142) are able to replicate the experimental results, and given that the conditions they use are exactly like those used by Schaffer and Knobe except that 'now' is dropped, it seems fair to present Schaffer and Knobe's data with this minor change.

[12]Schaffer and Knobe (2012) do think that contrastivism successfully explains the contrastive effect shown by the pair in (65). One reason why I set them aside is that it seems not clear to me how the contrastivist theses, i.e. Ternicity and Saturation could be extended to cover knowledge-*wh* ascriptions, especially when the aim is a compositional semantics that preserves the uniformity of 'know'. For example, Schaffer (2009) proposes that a knowledge-*wh* ascription, with Q being the question that corresponds to its embedded *wh*-clause, is true iff there is a proposition p such that $KspQ$ and p is the true answer to Q. But I do not see how this existentially quantified truth condition could be unified with the version of contrastivism for knowledge-*that* ascriptions, so it seems better, for my purposes, to focus on only knowledge-*that* ascriptions at this point. For Schaffer's view on knowledge-*wh* ascriptions and relevant discussions, see Schaffer (2007b, 2009), Brogaard (2009), Kallestrup (2009), Aloni and Égré (2010), and Steglich-Petersen (2014).

[C2] Everyone is now asking the big question: What did Peter steal? The news reporter is about to write a story about Mary. He is wondering if Mary now knows what Peter stole. He writes:

(66b) # Mary knows that Peter stole the rubies.

In addition, Schaffer (2008) indicates that *it*-cleft and focus can also elicit similar contrastive effects. This claim, unfortunately, is not tested in the above survey, but I nevertheless assume its truth and thus include the following pairs, uttered in the context given by [Jewel Thief], as part of the data, on a par with (64)–(66) above.

(67a) Mary knows that it was Peter that stole the rubies.

(67b) # Mary knows that it was the rubies that Peter stole.

(68a) Mary knows that [Peter]$_F$ stole the rubies.

(68b) # Mary knows that Peter stole [the rubies]$_F$.

Call these results *the contrastive data*. Contrastivism, i.e. Ternicity and Saturation combined, provides a straightforward account of the difference between these a-sentences and b-sentences, assuming that *rather-than*, *it*-cleft, and focus serve to signal contrasts and that topical questions implicit in the context can determine contrast when a knowledge ascription is binary in its surface form. In particular, the assumption is that *it*-cleft and focus both trigger a set of alternative propositions obtained by substituting the constituents that are marked by these two features. For (67a) and (68a), the alternative propositions include that Peter stole the rubies, that Alice stole the rubies, that Bob stole the rubies, etc.. But for (67b) and (68b), the alternatives are that Peter stole the rubies, that Peter stole the sapphires, that Peter stole the diamonds, etc.. Then, the proposition expressed by the embedded *that*-clause fills the content argument place of the ternary knowledge relation, while the disjunction of all other alternative propositions fills the contrast argument. Therefore, the content proposition of all the above ascriptions are the same, i.e. the proposition that Peter stole the rubies, while the contrasts are different: all the

a-sentences have the contrast proposition that someone other than Peter stole the rubies, which is ruled out by Mary's epistemic state, but all the b-sentences have the contrast that Peter stole something other than the rubies, which isn't ruled out by Mary. This explains why all the a-sentences are true while the b-sentences are false: the ternary knowledge relation holds between Mary, the content proposition that Peter stole the rubies, and the contrast proposition that someone else stole the rubies, but it doesn't hold if the contrast proposition is that Peter stole something else.

If that is the case, then the intuitive difference between the a-sentences and b-sentences is to be explained by the difference in their semantic content. In other words, while all the a-sentences express a proposition *Kspq*, which is true, the b-sentences express a different one *Kspq′*, which is false. However, in the rest of this section, I will use two tests to show that the intuitive difference should nonetheless be accounted for in terms of pragmatics, rather than the semantic content, as the contrastivist claims.

4.3.2. Lying vs. misleading. The first test concerns the distinction between lying and misleading. The underlying idea is this: if one utters a sentence *S*, which is false, but doesn't believe in the semantic content that is expressed by *S*, then by saying *S*, she is lying. Otherwise, if she believes in the semantic content, but does not believe in some other false content that is pragmatically generated from the utterance of *S*, then she is not lying but merely being misleading. Borrowing an example from (Saul, 2012, 70), suppose an elderly woman is dying and she asks you if her son is well. Her son actually had a car accident and died after you saw him yesterday, and you know it. Consider the following two sentences that you might say to the old woman.

(69) He's happy and healthy.

(70) I saw him yesterday and he was happy and healthy.

Obviously, (69) is a lie, since you believe that its semantic content is false. But (70) is merely misleading, because you believe in its semantic content, or what it literally says,

94

but what is pragmatically implied by it is believed by you to be false, i.e. that the son is happy and healthy right now (or so to your best knowledge). Moreover, suppose further that there is someone else who knows that you actually believe that the son is not happy and healthy and pointed it out with undeniable evidence (say, it was her who told you that the son died). Then depending on which of (69) and (70) you uttered, the response you can give is different. If you said (69) previously, then you have to retract what you said. But if you said (70), you are able to stand by it, saying "Well, what I said was (what I believed to be) true."

As we have some fairly clear intuition about if some utterance is a lie or merely misleading, this distinction can be utilized to tell whether some content conveyed by an utterance belongs to the semantic content of the uttered sentence or is merely something generated by certain pragmatic mechanisms. According to the contrastivist, the a-sentences included in the contrastive data above all express the same true proposition *Kspq*. Suppose someone knows, and *a fortiori* believes in, all the facts involved in the scenario [Jewel Thief]. She uttered:

(71) Mary doesn't know that Peter rather than anyone else stole the rubies.

This is the negation of (64a). If the semantic content of (64a) is true, then it is believed by the utterer to be true because, by assumption, she has complete knowledge about the scenario. Then, it would be a lie if (71) is uttered. But is (71) a lie, or merely misleading? My intuition suggests that it is merely misleading. Just in case it is not so clear, think what if she had been challenged by some other person. Suppose someone asked, "But Mary knows for sure that Peter was the thief, how could she not know that Peter rather than anyone else stole the rubies?" She would be able to respond by saying "Mary doesn't know that it was the rubies that were stolen, so she doesn't know that Peter rather than anyone else stole the rubies" and thus stand by what she said, i.e. (71), without retracting it.

It becomes more obvious if we compare (71) to other examples that involve uncontroversially ternary verbs, such as "to introduce". Suppose I know that Alice did introduce Beth to Carl but I uttered (72).

(72) Alice didn't introduce Beth to Carl.

Of course I'm lying. The semantic content of (72) is known by me to be false. I cannot defend my utterance in any sensible way but only retract it once challenged by undeniable evidence that I knew that Alice did introduce Beth to Carl.

Consequently, the lying vs. misleading test indicates that the semantic content of (71), whatever that is, is *not* what the utterer believes to be false. Given that the utterer knows all the facts about [Jewel Thief], the semantic content of (71) is not false, and thus the semantic content of its opposite, (64a), is not true. As the contrastivist accounts for the speakers' acceptance of (64a) by saying that (64a) expresses the true proposition *Kspq*, it is wrongheaded in explaining the contrastive data. The same equally applies to other examples included in the contrastive data.

4.3.3. Insaturability. The above test indicates that the *perceived* content of the a-sentences, which is true, must be different from their semantic content, which should be false. In other words, *what is communicated* by them according to speakers' judgments, must not be explained solely based on the semantics of "to know"; rather, some pragmatic mechanism must be responsible. The point is to be further attested by more examples, which are intended to show that the perceived content of the a-sentences indeed varies with the context in which it is uttered.

First, note that speakers judged (64a) as true but (64b) false, according to the contrastive data. This happens when we judge them individually. However, if the negation of (64b) is uttered first and followed by (64a), the previous judgment no longer holds.

(73) Mary doesn't know that Peter stole the rubies rather than anything else. But she knows that Peter rather than anyone else stole the rubies.

Something strange is easily detected, and we are much less willing to say that (73) is true than we used to when saying (64b) is false and (64a) true. The contrastivist takes the logical forms of (64a) and (64b) as *Kspq* and *Kspq*′ respectively and deems the former to be true and the latter false, given the cover story [Jewel Thief]. It follows that (73) is just saying ¬*Kspq*′ ∧ *Kspq*, which should be true. This prediction, however, fails.

Compared, again, to the paradigm ternary predicate "to introduce", it is more clear that "to know" deserves a different treatment. Suppose Alice introduced Beth to Carl but didn't introduce Beth to Catherine. Then (74) is true and (75) is false.

(74) Alice introduced Beth to Carl.

(75) Alice introduced Beth to Catherine.

(76) Alice didn't introduce Beth to Catherine. But Alice introduced Beth to Carl.

Obviously, uttering the negation of (75) and then (74) in one breath, as in (76), comes out true. This result is well expected, given the ternicity of "to introduce": In both (74) and (75), all the argument places of the predicate are explicitly filled by appropriate objects, and consequently they are not affected by contexts, so when uttered in the sequence (76), the truth values do not change.

Turning back to (73), we see that its constituents, (64a) and (64b), do not stay the same as when they are uttered individually, unlike the case of "to introduce". Then, it seems, "to know" shouldn't be treated as a ternary predicate.

Before giving up, the contrastivist might want to rescue the view by appeal to context. Indeed, uttered individually, the difference between (64a) and (64b) is accounted for by the contrastivist in terms of the contrast argument of "to know" taking different values that are determined by the explicit *rather-than* clauses in these two sentences. But the contrastivist could postulate that even if there is an explicit expression such as *rather-than* in a knowledge ascription, the contrast may not be completely determined by it. On the contrary, in some contexts, the contrast is determined by *both* an explicit expression *and* the

97

context. In particular, the contrastivist might say that in the sequence (73), as the sentence (64a) comes after the negation (64b), the context of utterance for (64a) is influenced by (64b): the contrast expression in (64b) uttered earlier makes salient the proposition that Peter stole something other than the rubies, and thus this proposition becomes a part of the context that is fed into the contrast argument of the sentence (64a). Then, in this case, the contrast proposition of (64a) is the disjunction of two propositions, one determined by *rather-than* and the other by the context: that someone else stole the rubies or Peter stole something else. In this way, the falsity of this knowledge ascription is accounted for within the contrastivist framework: it is false because Mary can't rule out this contrast proposition.

There is nothing unreasonable to explain away a couple of counterexamples in this way, by appeal to contexts. However, an in-principle problem is that the contrastivist might not be able to find a case where a knowledge ascription sentence really has the contrast argument *saturated* by an explicit expression and is no longer sensitive to the context. We have seen in §4.2.2, as the example (63) shows, that the presence of a *rather-than* clause fails to fully determine the contrast, as a non-restrictive relative clause, which presumably cannot be a semantic determinant of contrast, affects what the contrast is. Then, it seems that explicit expressions can never fully determine the contrast and the context always plays a role. This point is furthered by (73) above. As it shows, some relevant possibilities mentioned in the context around a knowledge ascription do affect on the contrast, even if an explicit contrast expression is already there in the ascription sentence. Furthermore, we can generate counterexamples like (73) easily in a systematic way. Let *Kspq* be a knowledge ascription that the contrastivist deems as true. According to contrastivism, for *Kspq* to be true, the person *s* doesn't have to rule out all the ¬*p*-possibilities; rather, ruling out all the *q*-possibilities is enough. Then, here is a way to turn *Kspq* false by manipulating the context: add something to the context indicating that *s* can't rule out *q′*, where (i) *q′*

98

is inconsistent with *p* and (ii) doesn't entail *q*. By (ii), we avoid contradicting the truth of *Kspq*, as by the contrastivist's lights *Kspq* entails that *s* can rule out *q*, and thus everything that entails *q*. By (i), we push the speakers' intuition towards the rejection of *Kspq*: since *s* can't rule out *q'* and *q'* is inconsistent with the content *p* of the purported knowledge, then *p* can't be ascribed to *s* as knowledge.

To illustrate, suppose Goldman is seeing a (real) barn and he can rule out the possibility that the object is a house. Then (77) sounds true, in a normal non-philosophical context.

(77) Goldman knows that it's a barn rather than a house.

Applying the above method, we need to find a proposition that is inconsistent with that the object is a barn, and that doesn't entail that it is a house. Unsurprisingly: the proposition that the object is a barn-facade. Adding that to the context, we get

(78) Goldman doesn't know that it's a barn rather than a barn-facade. But he knows that it's a barn rather than a house.

This, like (73), doesn't sound good, or at least sounds much less acceptable than (77) uttered alone. Thus, this is another example where the contrastivist has to appeal to contexts: the facade-possibility is added into the contrast of the second sentence in (78) because it is mentioned previously.

This is not a happy result for the contrastivist. If for every knowledge ascription, the context matters for what its contrast is, then that means that the contrast argument place is *unsaturable* by explicit expressions. In other words, the contrastivist has to be committed to an argument place of "to know" that is *inherently context-sensitive*. It follows that "to know" has to be context-sensitive even if it is ternary. But this undermines the contrastivist approach. First, as seen in (74)–(76), paradigmatic ternary predicates don't seem to work this way; namely, their argument places are saturable and thus their shiftiness is eliminated whenever explicit contrast expressions are present. This suggests that "to know" shouldn't be treated as a ternary predicate. Second, if there is no way to

99

get rid of the context-sensitivity of knowledge ascription, then the good old contextualist view should be superior to contrastivism, because (i) a contextualist interpretation is *required* even if knowledge is ternary, and (ii) a contextualist interpretation is *sufficient* even if knowledge is binary. In particular, a contextualist approach can insist that "to know" is binary (*pace* Saturation) and semantically context-sensitive, while accounting for its sensitivity to explicit contrast expressions (e.g. *rather-than* clauses) by appeal to some pragmatic mechanisms by which the contrast expressions may affect context. In the next section, I develop such a contextualist account, added with the GQA pragmatics that explains how contrast expressions could affect context.

4.4. A GQA Pragmatics for Contextualism

To reiterate the gist of the above counterarguments, we have seen that contrastivism fails because, on the one hand, a genuinely ternary verb wouldn't be shifty once all argument places of it were saturated, and on the other, articulating contrasts doesn't eliminate the shiftiness of knowledge ascription. Then the question is, given that the shiftiness can't be attributed to a third argument place of "to know", where should it be located? There are, of course, multiple available options. One is to follow epistemic contextualism, attributing the shiftiness to the context-sensitivity of "to know" in a certain way. For example, one contextualist approach is to treat "to know" as an indexical, in that it expresses different semantic contents in different contexts. Or alternatively, a contextualist could postulate a hidden variable to the verb, of which the value is inherently determined by the context. In the rest of this chapter, I will pursue a hidden-variable semantics in which "to know" is treated as involving a hidden variable whose value is contextually determined, while providing a GQA-based pragmatic theory concerns how the value of the hidden variable is contextually determined. But before moving on to the account itself, let me make clear some success conditions for such a theory.

4.4.1. Desiderata. First, as the contrastivist approach of adding an argument place for contrasts is wrongheaded, the current account will stick to the traditional binary framework: "to know" is a binary verb, and thus a knowledge ascription of a binary surface form is taken as syntactically complete. Then, to explain the shiftiness of such binary knowledge ascription sentences is to give an account of what contextual parameter is responsible for affecting the truth-value of binary ascriptions.

> **Desideratum 1.** A binary account must make clear to what contextual parameter knowledge ascriptions are sensitive.

On the other hand, such an account should be friendly to the contrastive data that are alleged to favor the contrastivist view. As the contrastive data suggest, explicit contrast-signaling expressions, such as *rather-than* clauses, appear to be capable of affecting the truth-value of knowledge ascription. Since in a binary account there is no place for contrasts at the semantic level, the truth-value effect has to be explained in terms of some pragmatic mechanism, i.e. a certain pragmatic principle that enables contrast expressions to affect the contextual parameter mentioned in Desideratum 1.

> **Desideratum 2.** A binary account must explain in what way an explicit contrast expression can affect the value of the contextual parameter to which knowledge ascriptions are sensitive, so that its effect on the truth-value can be realized in a pragmatic way.

Once the basics of a binary account are settled, we are in the position to evaluate some of its consequences. First, as demonstrated in §4.3.2, a correct account should treat utterances like (71) as merely misleading rather than a lie.

> **Desideratum 3.** A binary account must not entail that (71), uttered in [Jewel Thief], constitutes a lie.

101

Second, many would prefer a theory that reconciles the alleged tension between Epistemic Modesty and Epistemic Closure. Although this might not be a necessary condition for the success of an account, it would nevertheless add some attraction.

> **Desideratum 4.** A binary contextualist account had better allow both Epistemic Modesty and Epistemic Closure to be true.

4.4.2. A contextualist-GQA solution. Let's start with a familiar idea: an epistemic agent has a body of evidence—including sensory experience, memory, etc.—that rules out some possibilities, i.e. the possibilities in which she doesn't have the body of evidence that she actually has. Let each possibility be represented by a possible world, and call the set of all the possible worlds that are not ruled out by an agent's evidence her *epistemic state*. Since the epistemic state is relative to epistemic agents as well as possible worlds, let $E_{\langle a,w \rangle}$ denote the epistemic state of the agent a at world w.

The core doctrine of contextualism is that the truth of ascribing p as knowledge to an agent does not require the agent's evidence to rule out *all* the not-p possibilities. Rather, what possibilities her evidence has to rule out in order for the truth of knowledge ascription depends on the context. For example, Lewis (1996) contends that a context determines a set of possibilities that are properly ignored, and saying that a knows that p in that context is to say that every not-p possibility is either ruled out by a's evidence *or properly ignored in the context*. As there are many different versions of contextualism, we don't have to be committed to Lewis's view that it is being properly ignored that makes some possibilities not relevant to the evaluation of knowledge ascriptions in a context. But a more general contextualist contention is this: in a context, an utterance of a knowledge ascription sentence is interpreted with respect to a contextually determined *range of relevant (in some sense) possibilities*, and a knowledge ascription of p is true just in case all the not-p possibilities that are *within* the range of relevant possibilities are ruled out by the evidence of the agent in question. Let a set of possible worlds R represent the set of *r*elevant

102

possibilities determined by a context. Let us also assume that the information encoded in a context (i.e. the Stalnakerian context set) entails what possibilities are relevant in the above sense. Then, in a context where the context set is c, the set of relevant possibilities is R_c. Hence the truth condition:

> **Contextualist Semantics.** In a context c, an utterance of "a knows that p" is true at w if and only if p is true at all the possible worlds in $R_c \cap E_{\langle a,w \rangle}$.

Gloss: An utterance of "a knows that p" is true just in case among all the relevant possibilities in its context, only those p-possibilities are compatible with a's evidence—in other words, a's evidence rules out all the relevant not-p possibilities.[13] Given this rudimentary contextualist semantics, I consider Desideratum 1 satisfied: the contextual parameter to which the knowledge ascriptions are sensitive is R_c, a set of relevant possibilities.[14]

Turing to Desideratum 2, my proposal is that the pairs of sentences in the contrastive data all involve a shift in CQ, which is realized by focus, *rather-than*, and *it*-cleft. This shift in CQ results in a difference in how R_c is resolved, hence a difference in the truth-conditional interpretation.

In particular, for a knowledge ascription, its embedded *that*-clause determines a CQ. I argue that R_c, the set of relevant possibilities, serves as the local context for the embedded *that*-clause. In Chapter 3, we have seen some Good Question conditions, which are pragmatic constraints governing how a clause's CQ should relate to its local context. These

---

[13]A serious contextualist semantics (e.g. Schaffer and Szabo, 2013) must be much more sophisticated than this. For example, the current semantics doesn't even entail the factivity of knowledge, which is a shortcoming. However, we can be satisfied with the current version of the contextualist semantics because revisions can be made to address such concerns.

[14]This is not a reductive approach that specifies what R_c really is, since I did not provide a theory of how the set of relevant possibilities should be understood. However, if such a theory is preferred, there are multiple options in the literature. To mention some, for Lewis (1996), R_c is governed by multiple rules, one of which is *Rule of Attention*: if relevant speakers in a conversation attend to some possibilities, then those possibilities must be included in R_c. For Cohen (1999), possibilities in R_c are those which are salient in the context c. And for Blome-Tillmann (2009), they are the possibilities that are compatible with the relevant speakers' presuppositions. But regardless of the nature of the parameter, what is important to my theory is that R_c, whatever it is, serves as the local context for the embedded *that*-clause in a knowledge ascription. If this is correct, the contrastive data are easily explained by GQA, within the binary, contextualist framework.

pragmatic constraints enable an upstream effect of accommodation. For a knowledge ascription sentence to be appropriate to utter in a context, the CQ of its embedded *that*-clause has to be a Good Question with respect to the local context. Thus, for the sake of appropriateness, the interpreter of the knowledge ascription sentence may adjust the local context R_c to the effect that the Good Question conditions are satisfied. This upstream effect will explain the contrastive data: in each pair of sentences there, since their embedded clauses have different CQs, the accommodation effect on their R_c will be different; as R_c is part of what determines the truth-value of knowledge ascription sentences, two sentences in the pair, therefore, have different truth-values. In what follows, I will explain the two crucial steps in this account: first, why the local context is R_c, and second, how this GQA-based upstream effect results in different resolutions of R_c for knowledge ascription sentences in the contrastive data.

As mentioned in §2.3, the local context for a clause is a minimized set of worlds, i.e. an information state, for the clause to perform its Stalnakerian update, and the information state is determined partly by pragmatic reasons concerning the easiness of processing the clause. To see why R_c is identified as the local context for knowledge ascriptions, we may first consider the goal of processing a knowledge ascription in a context. According to the Stalnakerian framework, uttering a knowledge ascription sentence is to perform an update on the global context. If the context is c, then uttering "a knows that p" is to keep all the possible worlds in c where the sentence is true, while ruling out others. Then, processing the utterance is to calculate the resulted context set. In order to achieve this goal, we need to determine, for each $w \in c$, whether p is true at every world in $R_c \cap E_{\langle a,w \rangle}$—if it is the case, w is kept in the subsequent global context; otherwise, it is ruled out. Then, what information state does p update? There are three options:

(i) First, the interpreter calculates $R_c \cap E_{\langle a,w \rangle}$ for each $w \in c$, and then for each of the resulted information states, updates it with p. If the second update does not make any change, the corresponding world w is kept.

According to this, p serves to perform a Stalnakerian update on multiple information states: for each $w \in c$, there is a corresponding $R_c \cap E_{\langle a,w \rangle}$ which p updates.

(ii) First, the interpreter uses p to update $E_{\langle a,w \rangle}$ for each $w \in c$, and then updates R_c with each of the resulted information states $p \cap E_{\langle a,w \rangle}$. If the second update does not make any change, the corresponding world w is kept.

In this way, p updates multiple information states as well: for each $w \in c$, there is a corresponding $E_{\langle a,w \rangle}$ which p updates.

(iii) First, the interpreter uses p to update R_c, and then updates $E_{\langle a,w \rangle}$, for each w, with the resulted information state $p \cap R_c$. If the second update does not make any change, the corresponding world w is kept.

In this way, p updates a unique information state, namely, R_c.

The three possible ways of calculating the final pragmatic effect of the given knowledge ascription all have the same result, while the processes are different. It appears to me that the third one is the best strategy. Notice that each of (i)–(iii) involves two steps. However, according to (i) and (ii), both steps have to do with an awful amount of information states because $E_{\langle a,w \rangle}$, a world-specific value, is involved in both steps. In contrast, the first step in (iii) is quite simple: it is just a Stalnakerian update on R_c with p, as $E_{\langle a,w \rangle}$ is not involved. Therefore, (iii) is a better strategy of processing the clause, compared to the other two options. According to our working concept of local context, I take R_c to be the local context of p, when "a knows that p" is uttered in c.

With the local context being R_c, according to GQA, the embedded clause in a knowledge ascription is able to have some upstream effect on the local context, which in turn affects the truth-value. I will use this pragmatic account to explain the contrastive data. In particular,

it is to use GQA to explain how the R_c parameter can be affected by the following factors: (i) contrast-signaling devices occurring in the embedded clause of a knowledge ascription, as illustrated in the contrastive data (64), (67), and (68), and (ii) questions raised in the discourse, as in (66).

First, uttered in [Jewel Thief], the pair of sentences in (64), repeated as (79a) and (79b), appear to have different truth-values.

(79a) Mary knows that Peter rather than anyone else stole the rubies.

(79b) # Mary knows that Peter stole the rubies rather than anything else.

The two embedded clauses have different *rather-than* structures. According to the semantics of CQ provided in Chapter 3, the CQ of (79a) is *Was it Peter or anyone else who stole the rubies?*. But for (79b), the CQ is *Did Peter stole the rubies or anything else?*. Hence, by Cover, the embedded clauses in the two knowledge ascriptions have different requirements on their local context R_c. Specifically, for (79a), R_c must not involve any possibilities where no one stole the rubies. In other words, in a context in which (79a) can be felicitously uttered, all the relevant possibilities in its R_c are someone-stole-rubies possibilities. On the other hand, the local context R_c in the context for (79b) isn't constrained in such a way: it can involve some possibilities where no one stole the rubies.

This explains why the pair appear to have different truth-values. Upon hearing (79a), the hearer tends to accommodate the appropriateness conditions by making the relevant pragmatic rules satisfied, including Cover. The result is that all the relevant possibilities in the context are someone-stole-rubies possibilities. Among those, Mary's epistemic state is only compatible with the Peter-stole-rubies possibilities and rules out all others. Therefore, according to Contextualist Semantics, (79a) comes out true. In contrast, the felicity of (79b) doesn't require all the relevant possibilities in the context to be someone-stole-rubies possibilities. Rather, there might well be possibilities where what was stolen

was something else. Mary's epistemic states cannot rule out these possibilities. Therefore, (79b) comes out false.

Now it is clear why *rather-than* construction can have a truth-conditional effect on knowledge ascriptions. Other contrast-signaling devices, e.g. focus and *it*-cleft, can be explained in the same way, as all of them are devices that determine the CQ of a clause. I thus leave the details aside. What's worth mentioning is that this pragmatic account explains the examples raised earlier in which skeptical possibilities are brought up in knowledge ascriptions in some peculiar ways. One such example is (57), repeated as (80):

(80) $Km(p \wedge \neg S(p,q))q$: Moore knows that [he has hands and is not deceived by an evil demon into believing so while he in fact has stumps] rather than he has stumps.

In this knowledge ascription, the clause that comes before *rather-than* involves an *and-not* construction, which according to Abusch (2010) triggers a set of two alternative propositions: that Moore has hands (call it p), and that he's deceived by an evil demon into believing so while he in fact has stumps (call it $S(p,q)$). Naturally, this *and-not* construction works in the same way as *rather-than*, in the sense that it congruently addresses the question *Is p the case, or is $S(p,q)$ the case?*. Thus, it is treated as a CQ-triggering device. In addition, the *rather-than* clause in (80) also triggers a CQ—*Does Moore have hands or stumps?*. Given that there are two CQ-triggers, I assume that the whole clause embedded under "to know" in this example has two CQs, $\{p, S(p,q)\}$ and $\{p, q\}$. Then according to Triviality proposed in Chapter 3, for the CQs not to be trivial questions in the local context R_c, there must be some skeptical possibilities in R_c where Moore is deceived by an evil demon. Therefore, with respect to such an R_c, since Moore can't rule out all the no-hands possibilities, (80) comes out false. Similar treatment also applies to the ascription (63), provided we assume that the non-restrictive relative clause there also triggers a CQ.[15]

---

[15]I am not inclined to say that non-restrictive relative clauses are CQ triggers. For (63), what happens there seems to have to do with the negation in the non-restrictive relative clause. In particular, as the main clause states that Moore has hands, and the relative clause states that BIVs can*not* have hands, it is similar to the *and-not* construction, and this is, I think, the reason why the relative clause in this example triggers a CQ.

Turning to the second effect to be explained, it is more straightforward. In the contexts [C1] and [C2], two questions are respectively what "everyone is now asking": *Who stole the rubies?* and *What did Peter steal?*. I submit that these questions influence how readers resolve the focus patterns. If the first question is asked, when reading the subsequent knowledge ascription, it is natural to put a focus on "Peter"; but if the second is asked, it is natural to focus "the rubies". Hence, Cover, in the same manner as presented above, explains how the R_c parameter has to be adjusted according to these two different focus-induced CQs, which in turn results in the difference in truth-value of "Mary knows that Peter stole the rubies" when uttered in the two contexts.[16]

So much for the first two desiderata. Desideratum 3 requires the current account not to entail that uttering (70) in [Jewel Thief], repeated as (81), constitutes a lie.

(81) Mary doesn't know that Peter rather than anyone else stole the rubies.

To satisfy this desideratum is to explain why the semantic content of (81) is not false. Given the contextualist semantics, the semantic content of it crucially depends on how the contextual parameter R_c is resolved. We have seen one resolution that makes it exclude all the possibilities where the stolen jewel was not the rubies. Based on that resolution of R_c, (81) is false and consequently a lie. But note that this specific resolution is derived from the interpreter's pragmatic inference. In other words, in order to accommodate the felicity, the speakers infer that the R_c should be restricted as such, so that (81) turns out false. Such an inference determines a semantic content of (81) based on the assumption that the utterance of (81) is felicitous, in compliance with the pragmatic rule of Cover. Thus, this content is just the result of our charitable interpretation, and other interpretations are

---

[16]Schaffer and Knobe (2012) take the two questions as Questions-Under-Discussion. But I disagree: the knowledge ascription "Mary knows that Peter stole the rubies" is not a congruent answer to either question, in the sense of Roberts (2012, 31–32). Although before the knowledge ascription is uttered, they seem to be the QUD, because of the incongruence of the ascription to it, some other questions must be accommodated as the immediate QUDs, say, *Who was the person that stole the rubies, according to what Mary knows?* and *What did Peter steal, according to what Mary knows?*. These newly accommodated QUDs are what regulates the appropriate focus pattern of the subsequent knowledge ascription. Thus, the difference between the two ascriptions, in this case, comes down to a difference in focus, as I claimed.

possible as well. This explains why the utterer of (81) can stand by what she said after being challenged: what she does by refusing to retract the utterance is to make salient that other semantic contents, derived from other resolutions of R_c, are possible, and some of these contents are not false. A consequence of this account is that after the utterer stands by and doesn't retract what she said, the felicity of the utterance is sacrificed. That does seem to be the case. Although the utterer can stand by what she said, the interlocutor would then get an uncomfortable feeling of the utterance: "Now I see what she meant, and that's indeed not false. But it's a really strange way of saying it!"

As to Desideratum 4, the account should be compatible with both Epistemic Closure and Epistemic Modesty. Regarding Closure, I do not have much to add to the contextualist solutions in the literature: it is entailed by the contextualist semantics, with just one standard contextualist proviso—it holds when the contextual parameter R_c is fixed. Regarding Epistemic Modesty, let's assume that in ordinary contexts the R_c parameter doesn't involve any skeptical possibilities. In the relevant knowledge ascriptions, as the focus is not explicit in the written form, I assume that the entire embedded clause is focused, so that the CQ is a polar question.[17] Thus, the ordinary knowledge ascription "Moore knows that he has hands" is true when uttered in an ordinary context: neither Cover and Triviality requires that R_c accommodate any skeptical possibility. On the other hand, when "Moore knows that he isn't a BIV" is uttered, Triviality requires R_c to include some BIV-possibilities: as the CQ is *Is Moore a BIV or not?*, R_c has to be compatible with both possible answers, which means that it must involve some skeptical and some ordinary possibilities. Therefore, this ascription comes out false even if the initial context is ordinary.

With the four desiderata satisfied, I conclude that the two motivations mentioned earlier, i.e. explaining the contrastive data and reconciling Closure and Modesty, are not

---

[17]When an entire sentence is focused, it has a neutral intonation. Thus, it seems very natural to assume this wide scope focus when it is not explicit.

strong enough for initiating a semantic revolution that adds a contrast argument to the verb "to know". On the one hand, the contrastivist semantics itself has various problems: it does not help reconcile the tension between closure, and it generates some empirically inadequate predictions. On the other hand, GQA does achieve both goals when combined with a more traditional contextualist semantics. Given such a situation, a contrastivist revolution, being a progressive proposal, is not genuine progress.

CHAPTER 5

Counterfactuals and Congruent Questions

We have seen how the effect of focus, as well as other CQ-triggers, on knowledge ascription sentences can be explained by the GQA-based pragmatics, together with a conservative, binary semantics. This chapter aims at a similar account of the effect of CQs-triggers on counterfactuals. A famous example that exhibits the effect of focus on the truth-condition of counterfactuals is from Dretske (1972).

> [Marriage] Clyde and his girlfriend Bertha don't like close relationships. They see each other only twice a year, and they don't want to get married, etc.. However, Clyde found that if he were to get married soon, he would inherit a great amount of money. Then he married Bertha and got the money, while expecting that their 'loose' relationship would continue. (Dretske, 1972, 432, slightly adapted.)

> (82) If Clyde hadn't [married]$_F$ Bertha, he wouldn't have inherited the money.

> (83) If Clyde hadn't married [Bertha]$_F$, he wouldn't have inherited the money.

This pair of counterfactuals, when uttered in the context [Marriage], seem to have different truth values. In particular, as "married" is focused in (82), the counterfactual seems to concern situations where Clyde didn't marry Bertha or anyone else. Consequently, given the cover story, he wouldn't be able to inherit the money. The counterfactual comes out true. In contrast, "Bertha" is focused in (83), so the relevant counterfactual situations seem to be cases where Clyde married someone other than Bertha. In that case, he would inherit the money. This counterfactual comes out false.

I argue that this effect is brought about by the local upstream effect of focus: as the two antecedents have different focus patterns, they trigger different CQs, which in

turn invokes different upstream effects on the local contexts that are made available by pragmatic constraints concerning Good Question and Good Answer. In other words, the CQs impose different requirements on their local contexts, leading to different resolutions of the local contexts. As will be clear, the local context of the antecedent plays an important role in determining the truth-condition of a counterfactual. Thus, different resolutions of the local contexts result in different truth-value judgments.

A prominent feature of this account is that it explains the effect by a pragmatic process of resolving the local context based on the CQ of an antecedent. Thus, it is expected to be compatible with a wide range of semantic theories for counterfactuals in the literature, insofar as these semantic theories are committed to a contextual parameter which can serve as the local context of antecedent and can be part of what determines the truth-condition of counterfactuals. Indeed, various semantic theories involve such a contextual parameter, although they disagree on what it exactly is. Without discriminating between these semantic theories, I submit that the GQA-based pragmatics protects them from radical semantic revolutions. By explaining the effect of CQs pragmatically, the GQA-based account cancels the necessity of potential changes to the extent semantic theories for the purpose of semantically explaining the effect.

One may doubt whether there is really a threat from such semantic revolutions, as no one, to my knowledge, has proposed a new semantics for counterfactuals that takes the CQ effect into account. Thus, by adopting the GQA-based pragmatics in explaining the CQ effect, it seems to be a merely preventive move. However, this is not the case. In the literature, some researchers question the mainstream semantic theories of counterfactuals, e.g. the Stalnaker-Lewis semantics, because they fail to predict some of our truth-value judgments. For example, some counterfactuals with a true antecedent and a true consequent appear to be false, while those semantic theories predict that they are invariably true. Puzzles as such are commonly taken to be semantic problems and thus motivate

new semantic theories in alignment with ordinary truth-value judgments. However, I argue that some of these puzzles do not sufficiently motivate semantic revolutions. In particular, I argue that focus plays an important role in our interpretation of the examples that are alleged to be problematic for the standing semantic theories. As the focus aspect is often ignored in the written form of sentences, the puzzles seem to deserve a semantic solution; but once we take focus into account, there is a pragmatic way of explaining those examples away, without deviating from the standing semantic theories. In that light, there are semantic revolutions incited by the effect of CQ-triggers such as focus, although the role played by CQ-triggers is often ignored. Therefore, the GQA-based pragmatic account does not only prevent potential semantic revolutions that explicitly address the effect of CQ-triggers, but also oppose ongoing revolutionary acts that implicitly rely on examples involving CQ-triggers.

It will be discussed in Chapter 6 how GQA fares with these puzzles involving implicit CQ-triggers. In the current chapter, let us first focus on examples where CQ-triggers are well-specified, such as (82) and (83), in order to make clear how GQA accounts for the effect of CQ-triggers on counterfactuals. The remainder of this chapter is organized as follows. In §5.1, I will first consider two pragmatic accounts for the effect of CQ-triggers. Both of them share the same general framework with GQA: focus is interpreted as a question-trigger, and its effect is explained in terms of some pragmatic constraints on the relation between the focus-triggered question and a contextual parameter that is capable of affecting the truth-condition of counterfactuals. However, these two accounts only work for a special kind of counterfactuals, while the CQ effect on other counterfactuals is left unexplained. This motivates a new pragmatic account, i.e. GQA, which, I believe, captures a more general pragmatic relation between antecedents of counterfactuals and their contexts. In §5.2, I will first introduce some semantic assumptions which are presumably compatible with various standing semantic theories of counterfactuals. Thereafter,

113

we can see how GQA accounts for the effect of CQ on a contextual parameter, which in turn results in an effect on the truth-value assessments on counterfactuals. Finally, in §5.3, I will argue that for the successful predictions on truth-values made by the standing semantic theories, taking the GQA pragmatics into account will not undermine the desirable results.

5.1. Accounts of Focused Counterfactuals

According to §3.1, there are multiple expressions and constructions that can trigger CQs. However, in this chapter, we will only discuss the effect of focus on counterfactuals, as focus is the most frequently discussed CQ-trigger in the literature regarding counterfactuals. That said, I see no obstacle for carrying over the discussion to other CQ-triggers. In what follows, I will discuss two accounts of the effect of focus on counterfactuals' truth-conditional interpretations. One of them is from Rooth (1999), which is based on the premise semantics (Kratzer, 1981a,b) and explains the pragmatic relation between focus and context in terms of an anaphoric relation. The other account is due to von Fintel (1994). He is committed to a strict-conditional semantics of counterfactuals and takes focus to be capable of pragmatically restricting the domain of quantification. I will introduce these two accounts in §5.1.1, and point out a common problem of them in §5.1.2.

5.1.1. Two accounts. The account provided by Rooth (1999) uses a pragmatic theory of focus anaphoricity to explain the focus effect. First, this account, as a pragmatic account, works in the same way as GQA in explaining the focus effect. Namely, it relies on a semantics which assumes that some contextual parameter is part of what determining the truth-condition of relevant expressions (counterfactuals for the current discussion), and uses some pragmatic principles governing the relation between this contextual parameter and focus in order to explain how focus can have its truth-conditional effect. The semantics

114

assumed by Rooth (1999) closely follows Kratzer's premise semantics (Kratzer, 1981a,b). Let > stand for the natural language counterfactual construction *If...would....*

> **Premise Semantics.** An ordering source, H, is a contextually determined set of propositions that jointly characterize the actual world (i.e. $\bigcap H = \{w\}$ where w is the actual world). A counterfactual $\varphi > \psi$ is true if and only if for every H' that is a maximal subset of H that is consistent with $[\![\varphi]\!]$, $\bigcap H' \cap [\![\varphi]\!] \subseteq [\![\psi]\!]$.

In words, a context determines an ordering source H, which is a set of propositions that are jointly true only at the actual world. The propositions in H can be understood as descriptions of all the facts in the actual world. Then, in order to assess a counterfactual $\varphi > \psi$ is true, we consider some counterfactual situations. Such counterfactual situations may be different from the actual world, as φ should be true in them. Thus, we need to abandon some propositions in H in order to describe such counterfactual situations. On the other hand, we do not want these counterfactual situations to deviate too much from the actual world. Thus, in describing such counterfactual situations, we choose the maximal subsets of H that are consistent with φ plus the proposition expressed by φ itself—these sets of propositions are taken to be what describe the relevant counterfactual situations. Thus, the counterfactual is true just in case ψ is true in all these situations.

What is important for our purposes is that this semantics specifies what contribution the context makes to the truth-condition of a counterfactual. In particular, the truth-value of a counterfactual is sensitive to the contextually supplied ordering source H. This allows upstream effects of focus on H to affect the truth-value of a counterfactual, by which the effect of focus on the truth-value can be explained. The account by Rooth (1999) does exactly this. According to the account, focus is understood as an anaphoric feature, which has to be licensed by some antecedent in the context. Specifically, Rooth suggests that, in a counterfactual, focus involved in the antecedent is anaphoric on the contextual parameter H, in the following way:

115

Focus Anaphoricity in Counterfactuals. For a counterfactual $\varphi > \psi$ that is uttered in a context, the focus involved in φ is licensed just in case the focus-triggered *focal closure* is a member of the ordering source H that is determined by the context, where the focal closure is the union of the focus-triggered FSV.

This pragmatic principle can explain why (82) and (83) have different truth-values, with the assumption that the relevant F-domains for the pair of counterfactuals are the clauses embedded under negation in the antecedents, i.e. (84) and (86), of which the FSVs are (85) and (87), respectively.

(84) Clyde had [married]$_F$ Bertha.

(85) $[\![(84)]\!]^F=\{p : p$ is a proposition expressible by a sentence of the form *Clyde x'ed Bertha*$\}$

(86) Clyde had married [Bertha]$_F$.

(87) $[\![(84)]\!]^F=\{p : p$ is a proposition expressible by a sentence of the form *Clyde married x*$\}$

Given the above FSVs, the focal closures of the two clauses are derived: for (84), it is the proposition that Clyde did something to Bertha, and for (86), it is the proposition that Clyde married someone. Then, by the anaphoricity requirement, the ordering source H for (83) has to include the proposition that Clyde married someone. Thus, in a resulted maximal subset H', the proposition is likely a member as well.[1] If that is correct, Clyde married someone in all the counterfactual situations that are relevant to the assessment of the counterfactual, which means that he inherited the money in all these situations as well. Therefore, (83) comes out false, as its consequent says that Clyde wouldn't have inherited the money. Similarly, for (82), the pragmatic principle also requires the ordering source to include the focal closure, but in this case, the focal closure is the proposition that Clyde did something to Bertha. Plausibly, adding this to H will have little effect: it does

---

[1]This depends on what else is in H. According to the premise semantics, it is not guaranteed that all the maximal subsets of H will preserve the focal closure. However, we may expect that in a normal context, the focal closure is kept in every maximal subset, as it should not be a threat to the consistency of such a subset with the antecedent of (83): if a subset of H that does not include the focal closure is consistent with the antecedent, then adding the focal closure should not break the consistency.

not make the counterfactual situations described by H' situations where Clyde married anyone at all. Thus, (82) is false.

von Fintel (1994), on the other hand, treats counterfactuals as having a tripartite structure: a quantifier *would*, an *if*-clause that restricts the domain of the quantifier, and a main clause that is the nuclear scope of the quantificational structure. According to him, the quantifier *would* have a covert argument, i.e. a *resource domain*, that is to be filled by the context. Let us take this resource domain to be a set R of possible worlds.[2] Then, the semantics of counterfactuals is as follows.

> **Quantificational Semantics** Uttered in a context in which the resource domain is R, a counterfactual $\varphi > \psi$ is true if and only if $R \cap [\![\varphi]\!] \subseteq [\![\psi]\!]$.

Thus, *if*-clause, say, φ, restricts R by intersection, resulting in a set of possible worlds $R \cap [\![\varphi]\!]$ that serves as the domain of quantification for *would*: if all the worlds in it are worlds where the consequent is true, the whole conditional is true; otherwise, the conditional is false.

The focus effect is explained by a pragmatic constraint on the relation between the focus and the contextually supplied resource domain R. Very roughly, von Fintel's claim is that the focal closure triggered in the antecedent of a counterfactual *restricts* the resource domain through some pragmatic mechanism.

> **Focus as Q-Domain Restriction.** For a context with a resource domain R and a counterfactual $\varphi > \psi$ uttered in it, R is a subset of the focal closure triggered in φ.

In words, the resource domain is restricted to the effect that the focal closure of the antecedent is true at every possible world in it. Given that, the focus effect is explained. On the one hand, for (83), as the focal closure is the proposition that Clyde married someone, it follows that Clyde married someone in all the possible worlds in the domain

---

[2]It should be noted that von Fintel (1994) uses the Situation Semantics framework, and thus *would* is a quantifier over situations rather than possible worlds. But this does not make a difference in explaining the effect of focus.

of quantification for *would*. Consequently, Clyde inherited the money in all such worlds and thus the counterfactual (83) comes out true. In contrast, the focal closure triggered in the antecedent of (82) is the proposition that Clyde did something to Bertha. Therefore, there may be worlds in the domain of quantification where Clyde didn't marry anyone and didn't inherit the money. This explains why (82) comes out false.

5.1.2. Why do we need a new theory? Although the two accounts assume different semantics and propose different pragmatics, they tackle the problem almost in the same way. Simply put, they both imply that the focus closure can restrict the set of relevant counterfactuals situations. But this way of explaining the focus effect crucially relies on the fact that the two counterfactuals have negative antecedents. However, counterfactuals with positive antecedents can also exhibit the focus effect, but the two accounts above do not explain it. This point is illustrated in (88) and (89), uttered in the context [Alcohol].

[Alcohol] Beer makes people burp but whisky does not. Being a fussy drinker, Clyde likes the taste of beer but only the odor of whisky. Having a glass of beer and a shot of whisky on his table, Clyde drank the beer but only sniffed at the whisky. After a while, he burped.

(88) If Clyde had drunk [the whisky]$_F$, he wouldn't have burped.

(89) If Clyde had [drunk]$_F$ the whisky, he wouldn't have burped.

This pair of counterfactuals displays a difference in truth-value, similar to (82) and (83). In particular, (88) appears true, as the relevant counterfactual situations seem to be situations where Clyde drank the whisky but not the beer. In contrast, (89) seems false: in the relevant counterfactual situations, Clyde drank the whisky rather than sniffed at it, but presumably, Clyde still drank the beer as he actually did.

But this difference is not explained by either Rooth's or von Fintel's account. The two accounts both propose that the focal closure triggered in the antecedent should restrict the

118

set of relevant counterfactual situations. However, in this case, such a restriction is vacuous, as it does not make a difference to the truth-value of the above sentences. Specifically, the focal closures triggered by the two antecedents are the following propositions.

FC(88): Clyde drank something.

FC(88): Clyde did something to the whisky.

We can see that the two focal closures are entailed, respectively, by the antecedents in (88) and (89). Thus, no matter whether the set of counterfactual situations is restricted by such a focal closure, the antecedent is to select only worlds where the focal closure is true. For (88), the antecedent guarantees that all the counterfactual situations we use to evaluate the consequent are such that Clyde drank the whisky—and thus Clyde drank something. For (88), the antecedent guarantees that all the counterfactual situations are situations where Clyde drank the whisky—and thus Clyde did something to the whisky. Therefore, the requirement imposed by the focal closure, in this example, is vacuous and does not have an effect on the truth-conditions. But given that focus indeed has an effect on the truth-conditions of this pair of counterfactuals, we need a new theory.

5.2. GQA and Focused Counterfactuals

The new theory is provided by the pragmatic principles raised in GQA, together with a very general assumption about the semantics of counterfactuals. In what follows, I will introduce the semantic assumption first. Since the GQA-based pragmatics concerns the interaction between a clause's CQ and its local context, I will then clarify what the local context is for the antecedent of a counterfactual. Thereafter, we can see that the pragmatics of Good Question and Good Answer explains why focus is capable of affecting such a local context, which in turn results in an effect on the truth-value of a counterfactual.

5.2.1. Counterfactuals: semantics and local context. Regarding the semantics of counterfactuals, I intend to assume something relatively uncontroversial without having

119

too much commitment to particular semantic theories, so that the pragmatic explanation provided by GQA can be compatible with various different semantics. The main idea here is that a counterfactual $\varphi > \psi$ is true just in case ψ is true at every φ-world in a set R of *relevant* worlds. This intuitively matches how a counterfactual is evaluated: we entertain a set R of possibilities, and use the antecedent to select all the possibilities where φ is true; then, if ψ is true in all the selected possibilities, the counterfactual is true, and it is false otherwise. To make the idea more precise, let's represent this set of relevant worlds by R, hence the semantic assumption:

> **Semantics of Counterfactuals.** A counterfactual $\varphi > \psi$, uttered in a context c, is true at w iff $[\![\varphi]\!]^O \cap R_{\langle c,w,\varphi \rangle} \subseteq [\![\psi]\!]^O$.

Formulated in this way, the semantics appears to assume the strict conditional analysis—that a counterfactual is true just in case its corresponding material conditional is true at every world within a domain of quantification. If that were the case, then the above semantics would only be compatible with semantic theories following the strict conditional analysis, e.g. von Fintel (2001). However, in the above, the set of relevant worlds is allowed to be determined by the context c of utterance, the world w of evaluation, and/or the antecedent φ. Given that, the semantic assumption is compatible with various particular semantic theories, in the sense that a particular theory can be taken as a specification of how $R_{\langle c,w,\varphi \rangle}$ is derived from the parameters among $\langle c, w, \varphi \rangle$. For example, the premise semantics, *per* Kratzer (1981a,b), holds that the context c determines a set H of propositions, i.e. an ordering source, for a world of evaluation w, and $R_{\langle c,w,\varphi \rangle}$ collects all the worlds that are compatible with $[\![\varphi]\!]$ as well as all the propositions in a maximal subset of H that is consistent with $[\![\varphi]\!]$. For the Stalnaker-Lewis semantics in Stalnaker (1968) and Lewis (1973), the context c determines a closeness ordering centered on the world w, and $R_{\langle c,w,\varphi \rangle}$ contains all the worlds that are at least as close to w as the closest φ-worlds to w. For yet another example, as von Fintel (2001) treats counterfactuals as strict conditionals, $R_{\langle c,w,\varphi \rangle}$

for him is just the contextually determined set of worlds over which *would* quantifies. Although the semantic assumption above does not specify in what way the three parameters determine the set R and leaves such specifications to particular semantic theories, what is important here is that R is at least *partially* determined by the context of utterance, which allows the set to be adjusted by pragmatic reasoning.

To apply GQA, it is not enough to have only a semantic assumption about the truth-condition of counterfactuals. Since GQA is a pragmatic theory about the appropriate relations between a clause and its local context, we need to know, for a counterfactual, what the local context of its antecedent is. According to the notion of local context proposed in §2.3, the local context of a clause is the minimized set of possible worlds that is subject to a Stalnakerian update by the clause. That is, in processing the clause, the interpreter takes the proposition expressed by the clause and adds it as a piece of information to the local context, generating a subsequent information state that serves in processing further clauses. Hence, according to the above semantic assumption, the local context of the antecedent of a counterfactual should be the set of relevant worlds. In processing a counterfactual $\varphi > \psi$, the interpreter focuses on a set of relevant worlds, picks out those at which φ is true, and then checks if ψ is true at all the possible worlds in the subsequent set. Thus, the Stalnakerian update performed by φ is on the set of relevant possible worlds—this set of relevant possible worlds is the local context of φ, which is, using the current notation, $R_{\langle c,w,\varphi \rangle}$.

With the local context clear, GQA directly applies as a pragmatic theory of how the CQ of φ relates to the local context $R_{\langle c,w,\varphi \rangle}$. According to it, the CQ of φ, call it $CQ(\varphi)$, has to be a good question in $R_{\langle c,w,\varphi \rangle}$, and the proposition expressed by φ, i.e. $[\![\varphi]\!]$, has to be a good answer to $CQ(\varphi)$ with respect to $R_{\langle c,w,\varphi \rangle}$. Specifically, $CQ(\varphi)$ has to Cover $R_{\langle c,w,\varphi \rangle}$, and $[\![\varphi]\!]$ has to be a Complete and Logically Strong answer to $CQ(\varphi)$ with respect to $R_{\langle c,w,\varphi \rangle}$. These pragmatic constraints, as we have seen in the earlier chapters, enable an upstream effect

121

of clauses on their local contexts. In the case of counterfactuals, an interpreter can make adjustments to the context upon hearing an utterance of $\varphi > \psi$, to the effect that the local context $R_{\langle c,w,\varphi \rangle}$ stands in a relation to φ in the way that all the pragmatic constraints are satisfied. Consequently, when a pair of counterfactuals are otherwise identical but trigger different CQs by their antecedents, they may bring about different upstream effects so that the local context for the antecedent is resolved differently. Furthermore, as the truth-value judgments depend on what the local context is, such a pair of counterfactuals may be interpreted as having different truth-values. This, as will see next, explains the data above, i.e. the contrast in truth-value between (88) and (89).

5.2.2. The focus effect explained. The aim now is to explain why, when uttered in the context [Alcohol], (88), repeated as (90), sounds true while (89), repeated as (91), sounds false.

(90) If Clyde had drunk [the whisky]$_F$, he wouldn't have burped.

(91) If Clyde had [drunk]$_F$ the whisky, he wouldn't have burped.

In what follows, I will show how GQA is able to explain this phenomenon.

As a preliminary step, notice that the antecedents (92) and (94), having different components focused, have different CQs, as specified in (93) and (95), respectively.

(92) Clyde had drunk [the whisky]$_F$.

(93) $[\![(92)]\!]^F = \{p : p$ is a proposition of the form *Clyde had drunk x* $\}$

(94) Clyde had [drunk]$_F$ the whisky.

(95) $[\![(94)]\!]^F = \{p : p$ is a proposition of the form *Clyde had x'ed the whisky* $\}$

Then, the explanation of the focus effect comes in three steps. First, it is to specify what pragmatic constraints there are, according to GQA, that govern the relationship between the antecedents and their local contexts. Second, it is to figure out what it takes to for these local contexts to satisfy such constraints. In particular, I will argue that as the two antecedents have distinct CQs, the pragmatic requirements imposed on their local

122

contexts are different as well. Finally, the difference between the pair of counterfactuals in their truth-values is explained in terms of the different results of accommodating these pragmatic constraints.

According to GQA, each of (92) and (94) has to express a good answer to its CQ, with respect to the local context.[3] This means that $[\![(92)]\!]^O$ must be a Complete and Logically Strong answer to $[\![(92)]\!]^F$ with respect to the context that (92) updates, and similar for (94).

Second, it is to specify what it takes to satisfy the good answer constraints. By Logical Strongness, the propositional content of (92) should not be logically weaker than any of its contextually equivalent answers to the same CQ. That is, for any other proposition in (93), if it is logically stronger than (92), it must not be equivalent to (92) with respect to the local context. On the other hand, Completeness requires the local context to be such that (92) contextually entails either the truth or falsity of every other answer to the CQ. That is, for any proposition in (93), either (92) entails it or (92) entails its negation with respect to the local context. These two conditions combined, as will see, require the local context to be a set of possible worlds such that if Clyde drank the whisky at a possible world in the set, then at that world, Clyde did not drink the beer.

To see this result, note that the following three propositions are answers to the CQ triggered by (92). In other words, they are members of (93).

(96) p: Clyde had drunk the whisky ($p = [\![(92)]\!]^O$)

(97) q: Clyde had drunk the beer

---

[3]Although GQA also implies that the CQs must be good questions with respect to the local context, this good question requirement is not relevant in this case. In particular, to satisfy the good question constraint, the CQ has to be Non-Trivial in the local context and Cover the local context. This means that, on the one hand, the local context of (92) has to only include possible worlds where Clyde drank something, and among them, there must be some possible worlds where Clyde drank something other than the whisky. On the other hand, the local context of (94) has to only include possible worlds where Clyde did something (drink, sniff, etc.) to the whisky, and among them, there must be some possible worlds where Clyde didn't drink the whisky but did something else to it. However, satisfying these conditions does not contribute to the difference in truth values between (90) and (91): the conditions say nothing about whether Clyde drank the beer (and as a result, whether he burped) in the relevant whisky-drinking worlds.

(98) $p \wedge q$: Clyde had drunk both the whisky and the beer[4]

Among them, p is the OSV of (92), which is supposed to be a good (Logically Strong and Complete) answer.

Let R be a set of possible worlds. Suppose it is the local context of (92) that satisfies the good answer conditions. We can see how such an R must be like by appeal to the following inference.

(P1) $\llbracket(92)\rrbracket^O$ (i.e. p) is a Logically Strong answer to $\llbracket(92)\rrbracket^F$ w.r.t. R assumption

(P2) $\llbracket(92)\rrbracket^O$ (i.e. p) is a Complete answer to $\llbracket(92)\rrbracket^F$ w.r.t. R assumption

(P3) $p \in \llbracket(92)\rrbracket^F$, $q \in \llbracket(92)\rrbracket^F$, and $p \wedge q \in \llbracket(92)\rrbracket^F$ def. of FSV/CQ

(P4) $p \wedge q$ logically entails p but not *vice versa*   logic

(C1) p and $p \wedge q$ are not contextually equivalent w.r.t. R P1, P3, P4

(C2) p does not contextually entail q w.r.t. R C1

(C3) p contextually entails either q or $\neg q$ w.r.t. R P2, P3

(C) p contextually entails $\neg q$ w.r.t. R   C2, C3

As explained above, the Logical Strongness condition in effect requires that every answer to the CQ that is logically stronger than p must not be contextually equivalent to p in the local context R. $p \wedge q$ is such a logically stronger answer. Thus, R must not make p and $p \wedge q$ contextually equivalent. The only way to achieve this is to include in R some possible worlds where p is true but q is false. Hence the intermediate conclusion (C2). In words, it says that R includes some possible worlds where Clyde drank the whisky but not the beer.

---

[4]The proposition $p \wedge q$ is derivable as a member in the FSV of (92) if we treat 'the whisky and the beer' as denoting the mereological sum of the whisky and the beer. Similar treatment can be found in Rooth (1992) when he explains the effect of focus on scalar implicatures.

On the other hand, Completeness requires that p contextually entail the truth or falsity of every answer to the CQ. Regarding the answer q, it means that either all the p-worlds in R are q-worlds or all the p-worlds are $\neg q$-worlds. Hence the intermediate conclusion (C3). In words, in the set R, either all the whisky-drinking worlds are also beer-drinking worlds, or all the whisky-drinking worlds are not beer-drinking worlds.

Combining the two intermediate conclusions, we can see that (C3) gives two options while (C2) rejects one of them. That is, according to (C2), there are some possible worlds in R where Clyde drank the whisky but not the beer; thus, it is not the case that all the whisky-drinking worlds in R are beer-drinking worlds. Therefore, by (C3), it has to be the case that all the whisky-drinking worlds in R are not beer-drinking worlds. Hence the final conclusion (C).

According to the conclusion (C), if Clyde drank the whisky at a world $w \in R$, he didn't drink the beer at that w—there is *no drinking-both world* in R. Thus, for the counterfactual (90), in order to make the good answer conditions satisfied, the local context of the antecedent must not include any drinking-both-world. This explains why the counterfactual (90) is judged to be true. Since the pragmatic constraints imposed by the good answer conditions require the set of relevant worlds not to include any drinking-both world, the interpreter tends to accommodate this requirement by adjusting the local context accordingly. Thus, as there are no drinking-both worlds in the local context, the possible worlds picked out by the antecedent, that Clyde had drunk the whisky, will be worlds where Clyde *only* drank the whisky. Presumably, Clyde would not burp at those worlds. The counterfactual is therefore judged to be true.

In contrast, this is not the case for (91). For the antecedent of (91), as the focus is on "drunk", the CQ it triggers is a different one, i.e. the set of propositions of the form *Clyde x'ed the whisky*. By reasoning in the above manner, we can see that Logically Strongness and Completeness require the local context of the antecedent not to include any worlds

125

where Clyde both sniffed at and drank the whisky. However, this does not exclude drinking-both worlds from the local context. Furthermore, there seems no reason, given the context [Alcohol], to exclude such worlds as irrelevant. If the relevant worlds are selected by closeness to the actual world, there might well be some drinking-both worlds that are closer to the actual world than any of the worlds where Clyde only drank the whisky. Since the pragmatic reason does not make the recommendation of excluding such drinking-both worlds from the local context, it seems that when (91) is uttered in a normal context, some drinking-both worlds will be included. Presumably, Clyde would burp at those worlds because of the beer. The counterfactual is therefore judged to be false.

To sum up, the contrast between (90) and (91) is explained by appeal to GQA (especially the good answer conditions). To render *Clyde drank [the whisky]$_F$* as a good answer to *What did Clyde drink?*, the local context has to exclude drinking-both worlds, and thus drinking the whisky, with respect to such a local context, contextually entails not burping. However, as this effect does not follow from the other antecedent, *Clyde [drank]$_F$ the whisky*, the local context may include some drinking-both worlds, where Clyde would burp.

Let us turn to the old data point, i.e. the contrast between (82) and (83), repeated as follows.

(99) If Clyde hadn't [married]$_F$ Bertha, he wouldn't have inherited the money.

(100) If Clyde hadn't married [Bertha]$_F$, he wouldn't have inherited the money.

I will show that GQA applies to this example as well, though in a slightly different way as the antecedents here involves negation.

The GQA-based account I propose is in spirit very similar to the theories given by Rooth and von Fintel. In both of those previous accounts, the basic idea is that the local context (using our current terminology) has to entail the focal closure triggered by the embedded clause. First, consider the antecedent of (100). Since the CQ triggered by it is the set of all the propositions of the form *Clyde married x*, the focal closure, being the

126

disjunction of all those propositions, is the proposition that Clyde married someone. By making the local context entail this focal closure, all the possible worlds in the local context are worlds where Clyde married someone. Then, the antecedent picks out worlds where Clyde didn't marry Bertha from such a local context, and the result is a set of worlds where Clyde married someone other than Bertha. As he would still inherit the money in these worlds, the consequent is false at those worlds, which explains why (100) sounds false.

On the other hand, the CQ triggered by the antecedent of (99) is the set of all the propositions of the form *Clyde x'ed Bertha*. It follows that the focal closure is the proposition that Clyde did something to Bertha. By making the local context entail the focal closure, it is not guaranteed that in any of the worlds in the local context, Clyde married someone else. Indeed, according to the context [Marriage], it sounds likely that Clyde, as a dedicated bachelor, may only want to marry Bertha, if necessary. Thus, the local context for the antecedent may only include worlds where Clyde didn't marry anyone else. Among all those worlds, the antecedent picks out the worlds where Clyde didn't marry Bertha, and the result is a set of worlds where Clyde didn't marry anyone; thus, he wouldn't inherit the money—the consequent is true at all those worlds, which explains why (99) sounds true.

We can see that the crux in such an account is the pragmatic constraint that the local context of the antecedent has to entail the corresponding focal closure, which is an existential proposition. This is exactly what GQA is to achieve by one of its good question conditions, i.e. Cover. According to GQA, the CQ triggered by a clause has to Cover the local context of the clause, which means that at every possible world in the local context, there must be one of the answers to CQ that is true. This is to say that the local context has to entail the disjunction of all the propositions in the CQ. Thus, in explaining the contrast between (99) and (100), what GQA offers is essentially the same solution: the focal closure

127

of the clause, i.e. the disjunction of all the propositions in the clause's CQ, has to be entailed by the local context. Therefore, by appeal to the Good Question condition in GQA, the pair is explained in a similar manner as in the theories of Rooth and von Fintel.

5.3. Theoretical Conservativity

The above account certainly achieves a kind of theoretical conservativity: it adopts a semantic assumption that is compatible with mainstream semantic theories of counterfactuals, while explaining the effect of focus by appeal to some pragmatic principles. If this is correct, the relevant phenomena will not motivate a semantic theory that is committed to the semantic sensitivity of counterfactuals to focus. However, it should also be noticed that in the literature on counterfactuals, predictions on truth-values are often made directly by semantics. The current approach, in contrast, indicates that pragmatic reasoning involving CQs is integral to our judgments on the truth-values of counterfactuals. Therefore, a potential worry follows: for the examples where the traditional semantic theories make correct predictions on their truth-values, will it undermine these good results if we take GQA-based pragmatics into account? This section is intended to reassure the reader that it will not—GQA does not only protect the traditional semantic theories from the puzzling examples above but also agrees with them when they succeed.

Towards that end, I shall show that GQA preserves the predictions on truth-value made by various existing semantic theories of counterfactuals. First, note that the majority of the literature on counterfactuals does not specify the focus patterns of example sentences. Usually, a counterfactual is given in the written form, with no indication of which component is focused. That leads to a potential difficulty for GQA. On the one hand, according to GQA, our judgments on counterfactuals' truth-values can be heavily influenced by pragmatic considerations based upon particular CQs triggered in the antecedents. On the other hand, traditional semantic theories have been successful in predicting truth-values in alignment with ordinary speakers' verdicts, even when the

128

relevant example sentences do not specify focus or any other CQ triggers. Therefore, it seems that I owe an explanation of why the CQ-based pragmatic reasoning does not make a difference in those examples where truth-value judgments directly provided by the semantics have been successful.

To make the difficulty more concrete, consider the following inference from Stalnaker (1968).

(101) If J. Edgar Hoover had been born a Russian, then he would have been a Communist.

If he had been a Communist, he would have been a traitor.

Therefore, if he had been born a Russian, he would have been a traitor.

The Stalnaker-Lewis semantics, for example, successfully predicts the invalidity of this inference, and the success is based on the agreement between the invalidity of transitivity in the semantics on the one hand, and ordinary speakers' judgment that the inference is intuitively invalid on the other. However, if speakers' judgments rely on CQs, this example might not count as a success for the Stalnaker-Lewis semantics—we don't know if our intuitive judgment on it is influenced by the CQ-based pragmatic considerations.

My view is that, in general, when we consider examples without explicit CQs specified, the pragmatic theory proposed above will agree with whatever truth-value predictions the semantic theory makes. This relies on an assumption about what CQ is triggered in such an example. It seems fair to assume that the reader in this case will produce a "neutral" reading of the sentence, reading it in a neutral intonation. I assume that the CQ of a clause in such a case is a polar question. When focus, or any other CQ trigger, is unspecified, a neutral intonation reading of a clause implicitly treats the CQ as a set of two propositions: the proposition expressed by the clause, and the negation of it.

Default CQ. When there is no particular CQ trigger involved in the presentation of a clause *S*, by default, its CQ is the polar question: *Is S the case?*. Formally, $CQ(S) = \{[\![S]\!]^O, W - [\![S]\!]^O\}$, where *W* is the set of all possible worlds.

For example, the antecedent of the first premise in (101), under such a default reading, has the CQ *Was J. Edgar Hoover born a Russian?*, which is a set of two propositions, that he was, and that he wasn't.

Based on this assumption of default CQ, the pragmatic theory GQA is only idle and not adding anything to the interpretation of counterfactuals. In other words, when a counterfactual's antecedent triggers such a polar CQ, accommodating the GQA conditions will not make a difference to the truth-value judgments. This can be shown by considering each condition in GQA. First, the two good answer conditions, Logcally Strongness and Completeness, are immediately satisfied when the relevant CQ is a polar one: $[\![\varphi]\!]$ is a Logically Strong answer to $\{[\![\varphi]\!], W - [\![\varphi]\!]\}$ in any context just because it is not contextually equivalent to any other answer to the question (unless the context is absurd, i.e. the empty set); $[\![\varphi]\!]$ is Complete in any context because it logically entails the falsity of the other answer. Furthermore, Cover is also satisfied, as the two answers to $\{[\![\varphi]\!], W - [\![\varphi]\!]\}$ jointly Cover *W*, the set of all possible worlds. Lastly, turning to Triviality, which requires the local context the be compatible with at least two possible answers to the CQ—in the current case, with both φ and $\neg\varphi$. Not all local contexts satisfy Triviality, but accommodating it normally does not make a difference to the truth-value judgments. Recall that the local context for the antecedent φ is described as the set of relevant possible worlds, which should be a set of possible worlds that is compatible with φ whenever possible (that is, whenever φ is not a necessary falsehood). Suppose such a local context is compatible only with φ but not its negation. Then, to accommodate Triviality, some $\neg\varphi$-worlds need to be added into the local context. However, this does not influence the truth-value judgment

on the counterfactual $\varphi > \psi$, as the $\neg\varphi$-worlds in the local context will not survive the update by φ anyway.

To summarize, assuming that the default CQ of a clause is a polar question, GQA should not worry anyone who has the conviction that a certain semantic theory is successful in making truth-value predictions in many cases. Agreeing with the predictions in such "default CQ" cases, what GQA achieves is that we are now able to explain the cases of non-default CQs in terms of pragmatics so that we can avoid a drastic change to the semantics due to the effect of CQ triggers.

CHAPTER 6

More Puzzles Solved Pragmatically

At the end of the preceding chapter, we have seen that many counterfactuals discussed in the literature can be taken as involving "default CQs", i.e. polar questions that have only a positive and a negative answer. In those cases, GQA is by and large idle and does not change the truth-value predictions made by semantic theories. However, in this chapter, I will argue that sometimes there could be a non-polar CQ that is not made explicit in the written form but nevertheless plays a non-trivial role in our interpretation of certain counterfactuals. In examples of this sort, we often mistakenly ignore the relevant CQs and take it for granted that whenever our truth-value judgments disagree with what a semantic theory predicts, it is a puzzle that deserves a semantic solution, and thereby a semantic revolution should ensue. However, once we become aware of the CQs as well as the pragmatics of good questions and answers, these puzzles do not have to be semantic puzzles, as they can be solved by the pragmatic theory of GQA without overthrowing the reign of the standing semantic theories.

I shall discuss two puzzles of this sort. The first is the puzzle of *true-true* counterfactuals: counterfactuals with a true antecedent and a true consequent are predicted to be invariably true by Lewis's semantics, while many of them do not sound true to ordinary speakers. The second is the puzzle of *Simplification of Disjunctive Antecedent* and *Substitution of Equivalent Antecedents*, which are two inference patterns of counterfactuals. The problem here is that both inference patterns seem somewhat plausible to ordinary speakers, but they cannot be both valid. For each puzzle, I will show that once we take the relevant CQs into account, GQA provides a plausible solution that i) explains ordinary

speakers' judgments and ii) is compatible with the standing semantic theories, such as Lewis's.

6.1. True Antecedent

Call counterfactuals with both true antecedents and true consequents "true-true counterfactuals".[1] Some true-true counterfactuals seem false. For example, (102) sounds false when uttered in the following intended context.

(102) [A coin is to be tossed twice. I bet on "two heads" and I win. I say:] If at least one head had come up, I would have won.(McDermott, 2007)

This datum appears to impose a constraint on the semantics for counterfactuals: it should not predict that all true-true counterfactuals are true. In other words, an adequate semantics should not validate *Conjunction Conditionalization* (CC).

CC. $\varphi \wedge \psi \vdash \varphi > \psi$

On the other hand, it is not easy to invalidate CC while keeping the semantics plausible as a whole. To start with, CC is valid in the standard similarity-based semantics such as Stalnaker (1968) and Lewis (1973).[2] Even if the standard semantics mistakenly validates CC, it is nonetheless considered successful for the most part, in the sense that most predictions it makes on (in)validity are quite desirable. Thus, it is the responsibility of the opponent of CC to either give a new semantics that invalidates CC and argue

---

[1]"True-true counterfactual" might seem like an oxymoron—how could a *counter*-factual have a true antecedent? Throughout the paper, I take the term 'counterfactual' as interchangeable with "subjunctive conditionals", as many authors on this topic did. A counterfactual is grammatical even if its antecedent is true. Also, true-true counterfactuals are not invariably unassertible, although many of them are, in many contexts. Using the example from Anderson (1951), the following counterfactual, in the given context, seems perfectly assertible and intuitively true: Investigating Jones's death, a doctor says, "If Jones had taken arsenic, he would have shown just exactly those symptoms which he does in fact show."

[2]In addition, CC is also valid in Pollock (1976) and Kratzer (1981b). Besides, in a restricted sense, CC is valid in some other semantics. In strict analyses such as von Fintel (2001) and Gillies (2007), a true-true counterfactual is guaranteed to be true if it is the first counterfactual uttered in a discourse. In causal modeling semantics, initially formulated in Galles and Pearl (1998) and later extended in Briggs (2012), all non-nested true-true counterfactuals are invariably true. Thus, putative counterexamples such as (102) are against these semantic theories alike.

for the adequacy of its other logical properties, or conservatively modify the standard semantics so that CC becomes invalid while other desirable logical properties of the standard semantics are kept intact. However, as Walters (2016) shows, many of such anti-CC semantic attempts in the literature throw the baby out with the bathwater (if CC is indeed bathwater). Furthermore, the failure of these anti-CC proposals might not be a mere accident. Walters and Williams (2013) prove that in order to invalidate CC, we would have to abandon some logical principles that appear very plausible.

Hence the dilemma: a semantics that validates CC is at odds with our linguistic intuition that some true-true counterfactuals are false, but a semantics that invalidates CC will unacceptably invalidate some compelling logical principles. This predicament makes a pragmatic approach preferable. As expected, I will use the pragmatic theory, GQA, to explain true-true counterfactuals that *appear* false, while keeping the standard semantics and thus the validity of CC.

Specifically, I will first, in §6.1.1, elaborate on the aforementioned dilemma by laying out some more putative counterexamples to CC and briefly rehearsing the results from Walters and Williams (2013), which show that CC can only be invalidated at the expense of some compelling logical principles. Then a pragmatic account of the putative counterexamples starts off. A divide and conquer strategy is proposed. First, the putative counterexamples to CC are sorted into two categories. One of the categories, I argue in §6.1.2, should be explained by a theory of conversational implicature. On the other hand, in §6.1.3, I argue that the other category is to be explained by GQA. The reason why the true-trues in the second category sound false, then, is because of the need of accommodating some GQA-based pragmatic conditions.

6.1.1. The Dilemma. The following true-true counterfactuals sound false in their respective contexts and thus are putative counterexamples to CC. (The example (102) above is repeated as (105).)

134

(103) [Casper attended a party and the party was fun, despite that Casper ruins most parties he attends and nearly this one.] If Casper had come, the party would have been fun. (Bennett, 1974)

(104) [A student worked hard but was only able to pass the exam by cheating. He did cheat and pass. Not aware of the facts, I say:] If he had worked hard, he would have passed. (Fine, 1975)

(105) [A coin is to be tossed twice. I bet on "two heads" and I win. I say:] If at least one head had come up, I would have won. (McDermott, 2007)

(106) [I'm sitting in front of the Eiffel Tower and I'm seeing it. I say:] If I were here or anywhere else in France, I'd be able to see the Eiffel Tower. (Nichols, 2017, slightly adapted)

It appears that the above examples require a revision of the semantics for counterfactuals in order to make CC invalid. In the literature, indeed, various such revisions are made, resulting in semantic theories that deviate from the standard semantics in one way or another.[3] These attempts all have their respective problems, although I cannot examine them here one by one.[4] But even if all the extant attempts in the literature are not completely successful, one might still think that we need a semantic revolution that invalidates CC anyway. This may not be true, because any such revolution, as shown in Walters and Williams (2013), will lead to non-trivial changes to the logic of counterfactuals, resulting in some arguably very plausible logical principles to be invalidated.

---

[3]Lewis (1973) himself considered removing the Strong Centering assumption from his semantics, resulting in the invalidation of CC. Besides, some authors, such as Penczek (1997), McGlynn (2012), and He (2016), pursue the idea that some sort of "connection" between antecedent and consequent is necessary for a counterfactual to be true and subsequently give a new truth-condition for counterfactuals that requires such a connection, so that the falsity of some true-trues is explained by the lack of the connection.

[4]For those who are interested, McDermott (2007) convincingly argues against the attempt of abandoning Strong Centering. Besides, the proposal in Penczek (1997) that connection is necessary for a counterfactual's being true suffers from the fact that some counterfactuals (in particular, false-true counterfactuals) are meant to assert the lack of such connections. See McGlynn (2012) and He (2016) for more discussions on this problem.

In particular, most of the following, if not all, are plausible logical principles that should be taken as valid.

(a) MP: $(\varphi > \psi) \supset (\varphi \supset \psi)$

(b) Necessary Consequent: if $\vdash \psi$, then $\varphi > \psi$

(c) SEA: if $\vdash \varphi \equiv \psi$, then $(\varphi > \chi) \equiv (\psi > \chi)$

(d) Disjunction: $((\varphi > \chi) \wedge (\psi > \chi)) \supset ((\varphi \vee \psi) > \chi)$

(e) LAS: $((\varphi > \chi) \wedge \neg(\varphi > \neg\psi)) \supset ((\varphi \wedge \psi) > \chi)$

(f) VLAS: $((\varphi > \chi) \wedge (\varphi > \psi)) \supset ((\varphi \wedge \psi) > \chi)$

(g) Substitution: $((\varphi > \psi) \wedge (\psi > \varphi) \wedge (\varphi > \chi)) \supset (\psi > \chi)$

In addition, (S) is a semantic assumption which is true of a lot of counterfactuals: intuitively, it is the statement that for $\varphi > \psi$, there is a proposition that is relevant to neither a counterfactual's antecedent nor its consequent:

(S) For $\varphi > \psi$, there is a χ such that $\chi > \varphi$, $\neg\chi > \varphi$, $\chi > \psi$, and $\neg\chi > \psi$ are all true.

Walters and Williams (2013) prove that whenever (S) is true of a counterfactual $\varphi > \psi$, $\varphi > \psi$ is derivable from $\varphi \wedge \psi$, assuming any one of {(a), (c), (d), (e)}, {(b), (d), (g)}, and {(c), (d), (f)} is a collection of valid principles.

To appreciate the significance of this result, note that all of (a)–(g) are theorems in Lewis's counterfactual logic VW (which is his VC without CC). Even though some of them might be controversial, the above three combinations of them respectively validate CC, making the rejection of CC inevitably at the same time a rejection of at least one principle in each of the combinations. This seems a great logical cost that few would like to pay. Although each of the three combinations, together with (S), can validate CC, they share a common component, i.e. (d), so rejecting (d) might be a tempting way to jettison CC without costing too much deviation from VW. Here I am not defending (d), or any other principle above. It suffices for my purposes to say that (d) is endorsed by many weaker counterfactual logics as well, e.g. the logic W in Nute (1980, 23), and

136

thus abandoning even this single principle would be a great logical cost—one that CC's opponents may not be willing to pay.

Second, even if one considers none of the three combinations compelling, it is not an easy task to argue that we should reject them all. Indeed, the result at the very least shows that there is no genuinely *conservative modification* of Lewis's semantics: we cannot subtract CC from VC and be satisfied with the subsequent logic VW, as VW entails CC given the semantic assumption (S). As shown in Walters (2016), semantics approaches invalidating CC in the literature all fail to preserve the whole of VW.

Finally, one who wants to conservatively invalidate CC might try to resist (S). However, (S), the claim that there is a proposition that is irrelevant to both components of a counterfactual, is true of vastly many counterfactuals, including all the putative counterexamples presented above. To illustrate, take (103) to be our $\varphi > \psi$. The proposition that (χ) there are even number of trees in New Zealand satisfies (S): Casper would come regardless of whether χ were true, and the party would be fun regardless of whether χ were true. Given that, although it is theoretically possible to conservatively invalidate CC by providing examples of which S does not hold, the aforementioned examples do not work. Therefore, the result at least shows that any example of which (S) is true cannot adequately motivate a conservative modification of the standard semantics that invalidates CC. Unfortunately, all the above examples, as well as most natural language counterfactuals that might interest us, are such examples.

6.1.2. A Pluralistic Pragmatic Account. A dilemma follows immediately. On the one hand, there is strong evidence that CC should not be invalidated, since in doing so one has to abandon some logical principles that are innocent and compelling. On the other hand, having CC as a valid inference in the counterfactual logic seems to leave our intuitions on the above putative counterexamples unexplained, which is apparently unsatisfying. However, this latter horn in the dilemma relies on the dubious assumption that if a sentence

is *felt* to be false, the only possible way to explain is to give a semantics predicting that it *is* false. The falsity of this assumption is obvious to anyone who is committed to pragmatics as part of what determines the felt truth-values of sentences. In short, it is plausible to assume that ordinary speakers normally cannot separate semantic content from inference-based, pragmatic content, so that their intuitive truth-value judgments are determined by not only sentences' semantic contents but also speakers' pragmatic inferences. Then, the view I am to propose is that, on the one hand, the seemingly false true-trues have true semantic contents, and, on the other, our pragmatic reasoning generates certain inferentially derived contents that are false. To avoid confusion, in what follows, I will call a sentence/utterance false just in case its semantic content is false. Otherwise, if its semantic content is true but some pragmatically derived content is false, it is merely infelicitous.

To explain the data pragmatically, it is worth considering first whether they deserve a uniform treatment. The data show that some true-true counterfactuals are infelicitous, but it might well be the case that some true-true counterfactuals are infelicitous due to one pragmatic reason, while a distinct pragmatic reason is responsible for other cases. In this spirit, what I argue is that the above true-true counterfactuals are infelicitous due to distinct pragmatic reasons and thus deserve distinct pragmatic accounts. In particular, I will argue for the following dichotomy in the rest of this section:

Category 1. A true-true with an antecedent and a consequent that are irrelevant to each other, such as (103) and (104), belongs to Category 1, and its infelicity should be explained in terms of its false conversational implicature that the antecedent and the consequent are relevant in some way.

Category 2. Other true-trues, such as (105) and (106), belong to Category 2, and their infelicity should be explained by some distinct pragmatic account.

138

To start with, it is worth noting that some authors do not believe that all the examples introduced above work on par against CC. McDermott thinks that true-true counterfactuals with antecedents and consequents that are irrelevant to each other, i.e. Category 1 true-trues such as (103) and (104), are *intuitively true*, and "[m]any ordinary speakers agree." (McDermott, 2007, 333) Furthermore, he observes that pairing such a Category 1 true-true with its false-true counterpart "is a natural way to *say*" (334) that there is no dependence between the antecedent and the consequent. For example, (103) sounds true in the following sequence.

(107) If Casper had come, the party would have been fun. If he hadn't, the party would still have been fun.

(103) is infelicitous when asserted alone, as it seems to falsely suggest that Casper's presence is relevant, in a way, to the party's being fun. But it sounds better in (107), intuitively because the second counterfactual serves to clarify that in fact there is no such relevance.

I argue that this phenomenon is evidence that the infelicity of Category 1 true-trues should be explained by a Gricean conversational implicature account. The core idea is that (103), uttered in a normal context, will trigger the conversational implicature (108):

(108) Casper's coming was *relevant* to the party's being fun.

If this is right, why (103) is infelicitous when standing alone can be explained, since it conversationally implicates something false. In addition, the fact that the true-true sounds better in (107) supports this explanation, because it is well-expected based on the cancellability of conversational implicature. In particular, I assume that the second counterfactual in (103) expresses the content that Casper's coming was *not* relevant to the party's being fun. The plausibility of this assumption is based on our intuition that the counterfactual is normally meant to convey such a lack of relevance. However, I have no commitment on how this content of irrelevance is expressed, i.e. whether it is part of the

139

semantic content of the counterfactual or a content derived by pragmatic reasoning. Either way, if the second counterfactual expresses this content of irrelevance in the given context, it is incompatible with (108). This is evidence that (108) is a conversational implicature of (103). When (103) is uttered alone in the intended context, since (108), which is a falsehood, is conversationally implicated, (103) comes out infelicitous. But when uttered in the sequence (107), the conversational implicature (108) is incompatible with the second counterfactual and thus it is canceled, which explains why the sequence does not sound as bad as (103) uttered alone.

The above account appeals to a rudimentary inference to the best explanation: If (108) were conversationally implicated, the contrast between (107) and (103) (uttered alone) would ensue, which is actually the case. Whether it is the *best* explanation depends on whether there is a better alternative. Further investigations are required in that regard, but at the very least, the phenomenon that (103) sounds better in the sequence (107) than alone shows that the source of its infelicity comes from some cancelable content it expresses, even if the relevance content (108) is something other than a Gricean conversational implicature. It follows, then, that a pragmatic theory is needed for Category 1 true-trues, as semantic contents are uncancelable. Hence, although the above account needs refinements, it should be on the right track in attributing the infelicity of (103) to some cancelable, implicature-like content.

Besides, what also makes the account plausible is that it explains the phenomenon with an antecedently established pragmatic theory that is independently motivated. The gist of the account, that the contrast between (103) and (107) is due to the presence of the false conversational implicature of relevance in the former and the cancellation of it in the latter, is just a mundane application of some paradigm pragmatic explanation oft-used elsewhere. Indeed, parallel examples are easy to find. If it is commonly known that I got home sober, (109) sounds infelicitous.

140

(109) I got drunk and drove home.

Perhaps most people agree that the intuitive judgment of (109) should be explained by a false conversational implicature it has, that I got drunk *first* and *then* drove home, as it will turn out to be felicitous once the false implicature is canceled—or less infelicitous, as (110) still seems infelicitous but due to a distinct reason, i.e. its verbosity:

(110) I got drunk and drove home, but not in that order.

Although this account is a plausible candidate for the best explanation of Category 1 true-trues, it is only a programmatic proposal and details are still pending. For one thing, the term "relevance" used in (108) is admittedly loose and no proper analysis or definition is offered here. Consequently, the content (108) itself is left vague. Relatedly, I do not attempt to give an explanation of how this conversational implicature is triggered. Some Gricean maxims might be responsible, but without making the implicated content explicit, a precise derivation of it from Gricean maxims is not available.[5]

But I am to leave the account as is. As the main goal of this subsection is to separate the two categories and to give a novel explanation of Category 2, for now, it suffices to just give a promising direction for Category 1 and then set it aside. In the rest of this section, I am to argue that the implicature-based account does not apply to Category 2—examples such as (105) cannot be explained by it. To see the point, note that (105) does not sound good even if paired with its false-true counterpart.

(111) If at least one head had come up, I would have won. If no head had come up, I would still have won.

If the above account were applicable, we would expect the same contrast between (103) uttered alone and the sequence (107) to reappear here between (105) and (111). However, it

---

[5]Grice's Maxim of Quantity seems to be responsible, as Walters (2016) suggests. $\varphi \wedge \psi$ is logically stronger than $\varphi > \psi$, given the validity of CC, which makes uttering $\varphi > \psi$ a violation of Quantity whenever $\varphi \wedge \psi$ is known to be true. This ostensible violation can be interpreted charitably as implicating that $\varphi > \psi$ is true in a way independent of $\varphi \wedge \psi$, i.e. it is true in virtue of the relevance between φ and ψ. Hence the conversational implicature of relevance.

is not the case, since (105) does not sound better when uttered in the sequence (111). Then, what is responsible for the infelicity of (105) cannot be its conversationally implicating some content of relevance—a distinct account is required.

However, there is a loophole in the above argument. Some might worry that the second counterfactual in (111) is a confounder. In particular, while the second counterfactual in (107) is true, the second counterfactual here is false, because I bet on "two heads" and I would lose had no head come up. Thus, the infelicity of the sequence could be due to the known falsity of the second sentence, *despite that the sequence might be indeed a case of implicature cancellation.* This is a fair challenge. After all, if one wants to show that (111) is not a case of implicature cancellation, based on the fact that the sequence is infelicitous, she must guarantee that the would-be implicature-canceling sentence, i.e. the second counterfactual in (111), does not introduce any new infelicity-inducing factor. But as the second counterfactual in (111) is false, it is not surprising that it does not make (105) sound better, as its incapability of canceling implicature might just be a consequence of its known falsity.

But this concern can be avoided because even if the second counterfactual in (111) were true, the sequence would still sound infelicitous, and it is this fact that really shows the genuine difference between the two categories. In order to make the second counterfactual in (111) true, consider the following variation of the context of (105).

> [A coin is to be tossed twice. I bet on "two heads" and I win. According to some rules unbeknownst to me, whenever I bet on "two heads", I will win if either two heads or no head comes up, and I will lose otherwise.]

In this context, the second counterfactual in (111) is true. But the whole sequence is still infelicitous. Given that the second counterfactual is true and able to cancel the relevance implicature, the infelicity of the sequence cannot be explained by the implicature-based account.

142

The methodological point underlying the above argument can be made more explicit. For any intuitively false true-true counterfactual $\varphi > \psi$, first, we check whether its false-true counterpart $\neg\varphi > \psi$ is true. If it is, then $\varphi > \psi$ belongs to Category 1 if and only if the sequence $\varphi > \psi; \neg\varphi > \psi$ is felicitous (or less infelicitous than $\varphi > \psi$ uttered alone). If $\neg\varphi > \psi$ is false, we need to make some minimal changes to the context to make $\neg\varphi > \psi$ true. Hence $\varphi > \psi$ belongs to Category 1 if and only if the sequence is felicitous (or less infelicitous compared to $\varphi > \psi$ uttered alone) in the altered context. To apply this method to (106), we consider the false-true counterpart of it.

(112) If I were not in France, I'd be able to see the Eiffel Tower.

As this false-true counterfactual is false in a normal context where people outside of France cannot see the Eiffel, we need to change the context in order to make it true. So suppose people can see the Eiffel Tower in every country other than France (through some technological assistance), and the only place to see Eiffel within France is Paris (as that technological assistance is not available in France). Then (112) is true in this context. But the sequence (113) still sounds infelicitous, which shows that (106) is a Category 2 true-true.

(113) If I were here or anywhere else in France, I'd be able to see the Eiffel Tower. If I were not in France, I'd be able to see the Eiffel Tower.

Intuitively this is an expected result, since (106) seems to be infelicitous in the same way as (105) is, and thus they should belong to the same category. Specifically, each of them seems to have an antecedent that stands for more than one counterfactual supposition. The antecedent of (105) supposes that there are *one* or *two* heads coming up, and that of (106) supposes that I am *here* or *somewhere else in France*. The infelicity or seeming falsity of such a counterfactual is because there is at least one supposition that cannot guarantee the truth of the consequent. This, in a straightforward way, is drastically different from Category 1 true-trues, the infelicity of which comes from the lack of relevance between antecedent and

consequent. Since the implicature-based account, though being programmatic, provides a promising explanation of why Category 1 true-trues are infelicitous, I will set Category 1 aside from now on, and focus on Category 2 true-trues in the next section.

6.1.3. Focused Antecedent. The starting point of my take on (105) is that it actually involves a focus on the phrase "at least one", so a more accurate representation of it is (114),

(114) If [at least one]$_F$ head had come up, I would have won.

I submit that ordinary speakers, when reading (105), naturally add a focus on "at least one" and thus the syntactic structure subject to interpretation is actually (114). If that is correct, the pragmatics of CQ is potentially relevant to our felicity or truth-value judgment in this case. In what follows, I will pursue such an account and apply it to Category 2 true-trues, explaining why they are felt to be infelicitous by appeal to pragmatic reasons based on CQs.

As argued in Chapter 5, for a counterfactual $\varphi > \psi$, the local context of the antecedent φ is a set R of relevant possible worlds, and this local context is part of what determines the truth-conditional interpretation of the counterfactual: the counterfactual is interpreted to be true just in case all the φ-worlds in R are ψ worlds. What is crucial to the pragmatic theory of GQA is that the CQ triggered by the antecedent φ has an upstream effect on the local context R: the interpreter tends to accommodate the good question and good answer conditions, to the effect that the local context R is adjusted to render the CQ of φ a good question and the proposition expressed by φ a good answer to the CQ with respect to the local context R. This pragmatic effect of accommodation in turn affects the truth-conditional interpretation of the counterfactuals, since the latter depends on what the local context is like.

144

To apply the account to (114) again, note that the CQ of the antecedent of (114) is Q, which is a set that contains ≥ 1, 1, and 2 (and perhaps more) as members. In addition, the answer that the antecedent provides is ≥ 1.

Q: the set of propositions of the form *x head(s) came up*, where *x* is any generalized quantifier such as *zero, one, two, at least one*, etc..

≥ 1: the proposition that at least one head came up.

1: the proposition that exactly one head came up

2: the proposition that exactly two heads came up

Then, according to the good answer condition of Logical Strongness, ≥ 1 has to be a Logically Strong answer to Q with respect to the set R of relevant worlds. This means that for any answer that is contextually equivalent to ≥ 1 with respect to R, it cannot be logically stronger than ≥ 1. Thus, there must be some 1-worlds and some 2-worlds in R—namely, in R, there are some worlds where exactly one head came up and some worlds where exactly two came up—otherwise either 1 or 2 would be contextually equivalent to but logically stronger than ≥ 1, violating the Logical Strongness condition. Given that in R, there are some 1-worlds and some 2-worlds, it follows that some ≥ 1-worlds in R are 1-worlds and some are 2-worlds. Thus, the counterfactual (114) is interpreted as false: as some ≥ 1-worlds in R are 1-worlds, the consequent is not true at some ≥-worlds in R—I bet on "two heads" and thus do not win in these 1-worlds.

Now we have reached the truth-value prediction that (114) is false, but what I have promised is a pragmatic theory explaining that it *seems false* but *is actually true*. To see the point, we need to consider the role played by the pragmatic principle in the above derivation of the truth-value prediction. The point here is that the falsity belongs to a content that is not identical to the semantic content of (114) but is recovered by pragmatic reasoning. Assuming that Lewis's semantics gets the truth-condition right, the set of relevant worlds should be the smallest sphere that admits the given antecedent, which in

this case is the singleton set of the actual world because the antecedent is actually true. But the pragmatic requirement imposed by the Logical Strongness condition would not be satisfied; namely, it would be infelicitous to use the antecedent, focused as such, to update such an R. Thus, upon hearing the counterfactual (114), a competent speaker should be able to reason that the actual world cannot be the only intended relevant world, and then to include more worlds in R until the pragmatic constraint is satisfied—in other words, she is to expand R until some 1-worlds and 2-worlds are included. When this is done, the counterfactual (114) is in effect re-interpreted, as the interpreter obtains a different truth-conditional content. Although she judges this new truth-conditional content as false, it is compatible with the truth of (114) itself. In conclusion, (114) itself is true, but it is felt to be false because of the interpreter's re-interpretation of it according to the GQA conditions.[6]

The same account applies to (106) as well, with one additional assumption: that disjunction is also a device that governs what CQ is triggered. In particular, I assume that a disjunction $\varphi \vee \psi$ can address an alternative question that has φ, ψ, and $\varphi \vee \psi$ (and perhaps more) as members. This assumption sounds plausible. For example, it is intuitive to say that the antecedent of (106) addresses the question *Where am I?*, which includes propositions of the form *I am at x*. Also, it seems that the antecedent proposition *I am here or anywhere else in France* should also be a member of this question, since the proposition is expressible by "I am in France." Then, the question the antecedent of (106) addresses includes at least the following propositions as members:

here: I am here (before Eiffel).

else: I am somewhere else in France.

---

[6]The diagnosis is that the truth-conditional content of (114) itself and the truth-conditional content generated by pragmatically re-interpreting (114) are distinct. Now one may wonder what the latter really is. However, I am not able to offer an answer here, for two reasons. First, although the current view claims that the re-interpretation is an expansion of R from the singleton set of the actual world into a larger one that includes some 1-worlds and some 2-worlds, I do not have a theory of exactly what worlds are to be included. Second, even if we have decided on which worlds are to be included, there may not be a natural language expression that expresses the truth-condition in question, other than something involving a semantic ascent: "In the actual world and all the newly included worlds where at least one head came up, I won."

or: I am here or somewhere else in France.

Similar to the reasoning above, the Logical Strongness condition require that there be some *here*-worlds and some *else*-worlds in R. If R satisfies this requirement, then (106) is to be evaluated as false, because in the *else*-worlds in R I cannot see the Eiffel Tower.

So far the above two Category 2 true-true counterfactuals are explained by the pragmatic theory alike. The overall diagnosis is that the question-answer property of an antecedent, determined by focus or disjunction, has to be related to its set of relevant worlds in the way specified by the GQA conditions. To satisfy the pragmatic constraints, the two Category 2 true-trues are re-interpreted, resulting in some pragmatically derived truth-conditions that are different from their original semantic contents, which explains why the two true-trues *are* true but *seem* false in many contexts. That being said, I shall not take a stance on whether all Category 2 true-trues can be explained in the same way. I am inclined to think that true-trues might be infelicitous due to even more reasons and those two addressed above are probably not exhaustive.

6.1.4. Summary. Motivated by ordinary speakers' intuition that some true-true counterfactuals seem false and also the proofs by Walters and Williams showing that invalidating CC would result in overthrowing the largely successful standard semantics for counterfactuals, the current account attempts to preserve the validity of CC while pragmatically explaining away the intuition against it. Furthermore, the pragmatic account is pluralistic, in the sense that it separates two categories of seemingly false true-trues and uses a conversational implicature theory and a pragmatic theory of focus to explain them respectively. There might be examples of intuitively false true-trues that are not covered by either account, but all the examples discussed by the researchers on this topic, to my knowledge, can be explained in either of the two ways. Therefore, at the very least, this two-fold account can be a starting point of a complete pragmatic account of the intuitive

invalidity of CC, which, if successful, would release logicians and semanticists from the burden of overhauling conditional logic and semantics just in order to invalidate CC.

6.2. Disjunctive Antecedent vs. Equivalent Antecedents

In this section, I shall explore another application of GQA. The topic here is the debate about whether *Simplification of Disjunctive Antecedent* (SDA) is a valid inference pattern.

SDA. $\varphi \vee \psi > \chi \vdash (\varphi > \chi) \wedge (\psi > \chi)$

On the one hand, this inference pattern seems to work well in many examples, so it has some intuitive appeal. On the other hand, validating SDA leads to a notoriously invalid inference, Antecedent Strengthening (AS), when Substitution of Equivalent Antecedents (SEA) is assumed.

SEA. $\varphi > \psi \vdash \varphi' > \psi$, if φ and φ' are logically equivalent.

AS. $\varphi > \psi \vdash \varphi \wedge \chi > \psi$

According to the standard Stalnaker-Lewis semantics, the choice is to abandon SDA but embrace SEA. However, it seems that we need to explain why SDA sounds very appealing while being invalid. In §6.2.1, I will show that GQA provides such an account: a disjunctive antecedent often invokes an upstream effect on its local context due to the GQA conditions, which leads to a pragmatics-affected interpretation of the counterfactuals that render SDA plausible. Thereafter, in §6.2.2, I will further support the choice of validating SEA, as opposed to SDA, by using GQA to explain some putative counterexamples to SEA.

6.2.1. Disjunctive antecedent. There are three main desiderata concerning the validity of SDA. First, a lot of instances of SDA seem to show that it is valid. For example, both (116) and (117) seem to follow from (115).

(115) If Ann had come to the party or Bill had come to the party, it would have been fun.

(116) If Ann had come to the party, it would have been fun.

(117) If Bill had come to the party, it would have been fun.

Second, some peculiar examples, however, suggest the opposite. For example, (118) is true while (119) is obviously false.

(118) If the US were to spend more on defense or education, it would spend more on defense.

(119) If the US were to spend more on education, it would spend more on defense.

Third, SDA, together with the putatively valid inference pattern SEA, entails the validity of AS.

> **Proof.** Suppose SDA and SEA are both valid. Assume $\varphi > \psi$ is true. By SEA, $(\varphi \wedge \chi) \vee (\varphi \wedge \neg\chi) > \psi$ follows. By SDA, $\varphi \wedge \chi > \psi$ is true.

Since AS is unacceptable, SDA and SEA should not be both validated.

Given the three desiderata, theories about SDA can be classified into three categories, following Khoo (2018). The first includes Univocal Strong theories, which are committed to the validity of SDA. For its proponents, such as Fine (2012a,b) and Willer (2018), the burden is associated with the second and third desiderata: they have to explain away the examples where SDA seems invalid, and also to provide good reasons for rejecting SEA in order to avoid validating AS. The second kind involves Univocal Weak theories, which claim that SDA is invalid. This option is in accordance with the traditional Stalnaker-Lewis-Kratzer semantics, which has no problem with the second and third desiderata but is obligated to explain away the first. Oftentimes, Univocal Weak theorists appeal to some pragmatic mechanisms to do so, as in Bennett (2003). Finally, the third kind includes Ambiguity theories, such as Khoo (2018), which essentially claim that counterfactuals are ambiguous between a strong reading and a weak one, the former of which admits the

validity of SDA while the latter does not. In principle, this option can accommodate cases involved in both the first and second desiderata. Moreover, the third desideratum may not be a problem for Ambiguity theories either. For example, the strong reading of $\varphi \vee \psi > \chi$ which validates SDA, according to Khoo (2018), is *if φ or if ψ, then χ*, in which there are two *if* operators in the antecedent. Thus, the version of SDA that allows simplification of such double-*if* counterfactuals is not the same as in §6.2 and thus does not lead to validating AS. However, the burden on Ambiguity theories is that they have to explain, in a compositional way if possible, why the same surface form of disjunctive antecedents is ambiguous between two logical forms.

The objective here is to show that GQA is able to explain the data in line with Univocal Weak theories—that is, GQA is in the position to provide a pragmatic explanation for why SDA seems valid in a lot of cases while refusing to deviate semantically from the Stalnaker-Lewis approach. In particular, it is to explain i) that simplification of counterfactuals like (115) is acceptable given the GQA conditions are accommodated and ii) that the same pragmatic mechanism does not apply to cases like (118). Besides, since the account is not committed to the validity of SDA at all, it does not have to reject SEA and thus the third desideratum above should not be a problem.

The pragmatic explanation is actually already in place when (106) is discussed in §6.1. It crucially relies on the assumption that the CQ of a disjunction $\varphi \vee \psi$ generally involves both φ and ψ, alongside the proposition $\varphi \vee \psi$ itself, as members.[7] Given that, the antecedent of (115) has at least the following three propositions in its CQ, where *Or* is the answer to the CQ provided by the antecedent.

Or: Ann had come to the party or Bill had come to the party.

---

[7]Similar treatment of disjunctive antecedents can be found in Alonso-Ovalle (2006) and Starr (2014). However, these approaches take the set of alternatives triggered by a disjunctive antecedent as part of the input of the semantics of counterfactuals. The present view holds that such alternatives do not directly interact with the semantics of counterfactuals but rather affects the interpretation through pragmatic resolutions of the context.

Ann: Ann had come to the party.

Bill: Bill had come to the party.

Then, for the GQA conditions to be satisfied, the local context of the antecedent, i.e. the set R of relevant worlds, must include some worlds where only Ann had come. If such worlds did not exist in R, *Or* would be contextually equivalent to *Bill*. Since *Or* is logically weaker than *Bill*, the Logical Strongness condition would be violated. On the other hand, R must also include some worlds where only Bill had come because otherwise *Or* would be contextually equivalent to *Ann*, and thus Logical Strongness would be violated.

Therefore, by GQA, the local context R of the antecedent should include some only-Ann worlds and some only-Bill worlds. Hence, suppose the interpreter of (115) now follows the requirements of GQA to accommodate some only-Ann worlds and some only-Bill worlds into R. Let's also suppose the interpreter processes the counterfactual by following the Stalnaker-Lewis semantics: by default, the local context R is the smallest sphere of possible worlds that admits the truth of the antecedent *Or*, i.e. the set of all the possible worlds that are at least as close to the actual world as the closest *Or*-world to the actual world.

At this point, there are two possible cases. First, if this local context R satisfies the above pragmatic requirements, no accommodation is needed. In this case, R includes some *Ann*-worlds and some *Bill*-worlds. It follows that all the closest *Ann*-worlds as well as all the closest *Bill*-worlds (to the actual world) are in R. Therefore, if (115) is true, then the consequent is true at all the *Ann*-worlds and all the *Bill*-worlds in R, and it follows that the consequent is true at all the closest *Ann*-worlds and all the closest *Bill*-worlds. Therefore, if (115) is true, both simplifications, (116) and (117), are true.

In the second possible case, the default R, the smallest sphere that admits the truth of *Or*, does not satisfy the above pragmatic requirements; that is, either there is no only-Bill worlds or no only-Ann worlds in R. In this case, the interpreter needs to accommodate

151

the requirements by expanding R. I submit that the interpreter does not randomly add such worlds into R, but rather expands R in an orderly manner: when the requirements are not met, she adds to R the closest worlds that haven't been included in R, until the requirements are satisfied. Thus, when this process of accommodation is complete, the resulted R will include both the closest *Ann*-worlds and the closest *Bill*-worlds to the actual world. Similar to the above situation, if (115) is true, it means that the consequent is true at all the *Or*-worlds in such an R, which entails that the consequent is true at all the closest *Ann*-worlds and all the closest *Bill*-worlds—the two simplifications, therefore, are both true.

This explains the intuitive appeal of SDA in the case of (115). When (115) is processed, the interpreter does not merely consider all the closest *Or*-worlds. Rather, per the GQA conditions, she takes into consideration a local context R which includes both the closest *Ann*-worlds and the closest *Bill*-worlds. According to the Stalnaker-Lewis semantics, the invalidity of SDA is because there is no guarantee that there is any closest $\varphi \vee \psi$-world that is, say, a ψ-world. When there isn't, the fact that all the closest $\varphi \vee \psi$-worlds are χ-worlds does not guarantee that all the closest ψ-worlds are χ-worlds. However, if the above account is correct, pragmatic reasons give rise to an accommodation effect, which makes the evaluation of (115) depend on a local context R that includes all the closest *Ann*-worlds and all the closest *Bill*-worlds. Thus, when (115) is judged to be true, its simplifications are true as well.

Now we turn to the cases where SDA apparently fails, e.g. (118). It is obvious that if the same account were applicable to (118), (119) would be predicted as true, which is incorrect. The reason why the above account does not apply to this case is that the CQ of the antecedent, I argue, does *not* contain *def* or *edu* as its members. Rather, the CQ is a polar question, which only contains the proposition expressed by the antecedent, i.e. *Or*, and its negation.

152

Or: The US will spend more on defense or education.

def: The US will spend more on defense.

edu: The US will spend more on education.

As mentioned, a disjunction by default triggers an alternative question as its CQ, as in the case of (115). However, I assume that such an alternative question is not the only possible CQ that a disjunction maps to. In particular, when reading the CQ as such an alternative question is pragmatically implausible, the interpreter will switch to the polar question reading. This, I believe, is the case of (118).

To see why an alternative question reading is implausible, the crucial point here is that the consequent of (118) is identical to one of the disjuncts, i.e. *def*, and consequently it is logically stronger than the antecedent. This fact, however, leads to a conflict between the truth of the counterfactual and its pragmatic felicity. On the one hand, for the counterfactual to be true, the set of relevant possible worlds has to exclude all education-worlds. It follows that whenever the counterfactual is true, *def* is contextually equivalent to the answer provided by the antecedent, *Or*, with respect to the set of relevant possible worlds. On the other hand, if *def* were one of the members in the CQ of the antecedent, as it is logically stronger than *Or*, the counterfactual can be felicitously uttered only if *def* and *Or* are *not* contextually equivalent, because otherwise the Logical Strongness condition would be violated. Therefore, if the CQ of the disjunctive antecedent is the alternative question *The US will spend more on which area?*, the counterfactual cannot be both true and felicitous. This implies that interpreting (118) by assuming this CQ makes the felicity and truth incompatible with each other, which is implausible.

In contrast, if the CQ of the antecedent is the polar question *Is it the case that the US will spend more on defense or education?*, there is no such a problem. As we have seen in §5.3, if the CQ of the antecedent is a polar question, the GQA conditions are just idle. As a result, (118) should be interpreted by assuming that the CQ is this polar question.

153

Therefore, as GQA does not have an effect on antecedents that have polar CQs, there is no guarantee that the local context of the antecedent includes both some defense-worlds and some education-worlds. Plausibly, the closest defense-worlds are strictly closer to the actual world than all the education-worlds. Thus, the default set R of relevant worlds will not have any education-worlds in it. The truth of (118), therefore, does not guarantee the truth of its simplification (119). This explains the disparity between (115) and (118).

We have seen that the GQA-based pragmatic account successfully applies to both data points, the one where SDA seems plausible and the one where it does not. Before leaving this topic, it is worth re-emphasizing the pragmatic nature of this account. On the one hand, as it explains the seeming validity of SDA by appeal to accommodation, it is compatible with the Stalnaker-Lewis semantics, where SDA is invalid. Consequently, equipped with this pragmatic account, a Univocal Weak theory is preferred, in the sense that i) counterfactuals do not have to have two truth-conditions so that postulating ambiguity of counterfactuals is unnecessary, and ii) the truth-condition of counterfactuals does not have to validate SDA so that rejecting SEA is unnecessary. On the other hand, although counterfactuals can have a univocal truth-conditional semantics, the above account indeed suggests a different kind of ambiguity: disjunctive antecedents are ambiguous in what CQs they have. In general, the CQ of $\varphi \vee \psi$ could be either an alternative question whose members include φ, ψ, and $\varphi \vee \psi$, or a polar question consisting of $\varphi \vee \psi$ and $\neg(\varphi \vee \psi)$. Thus, in a sense, we can also classify the above account as an Ambiguity theory, but it is important that this ambiguity is *not* an ambiguity in the truth-conditional semantics of counterfactuals, but rather one in the CQ-dimension of meaning. A favorable feature of this account is that, by connecting the CQ-dimension and the truth-conditional dimension of meaning through pragmatic mechanisms, as seen in the previous chapter, we can explain the variations in perceived truth-conditional interpretations without making drastic revisions to the truth-conditional semantics.

6.2.2. Equivalent Antecedents. We have seen how a GQA-based account explains away examples where SDA seems valid without deviating from the standard similarity-based Stalnaker-Lewis semantics for counterfactuals. As a result, although SDA, combined with SEA, entails the undesirable consequence of validating AS, as SDA is invalid, we are not pressed to abandon SEA. However, there are independent reasons questioning the validity of SEA. Thus, if there is strong evidence that SEA should be abandoned after all, we will have less motivation to invalidate SDA. In this subsection, I will address some alleged counterexample to SEA, showing that it can be explained away by GQA in a way very similar to what we have seen above.

Using an example from Ciardelli et al. (2018), given De Morgan's law, it follows from SEA that the following two counterfactuals should always have the same truth-value. However, (120) seems true but (121) seems false, when uttered in the context [Two Switches].

[Two Switches] There is a long hallway with a light in the middle and with two switches, one at each end, called switch A and switch B respectively. The light is on whenever both switches are in the same position (both up or both down); otherwise, the light is off. Right now, switch A and switch B are both up, and the light is on.

(120) If switch A or switch B were down, the light would be off.

(121) If switch A and switch B were not both up, the light would be off.

Ciardelli et al. (2018) conduct a survey that shows that most native speakers consider (120) as true but (121) false, and at the same time the subjects have no problem identifying the equivalence between the antecedents. Assuming that *not up* and *down* mean the same here, this appears to be a counterexample to SEA.

In what follows, I shall argue that we do not need to invalidate SEA after all, as GQA provides a pragmatic way of explaining away the counterexample. The account starts, again, with a specification of what CQ is triggered by the antecedent of the relevant

155

counterfactual. Consider the antecedent of (121), which says that Switch A and Switch B were not both up. Here I assume that this antecedent does not trigger a polar CQ; that is, its CQ is not the question *Were both Switch A and Switch B up?*. Rather, I assume that this antecedent actually involves a focus on "not both". This assumption is based on the intuition that the antecedent seems to be about the quantity of switches, between A and B, that were up, and the possible contrasts of it include *One of Switch A and Switch B were up*, *None of Switch A and Switch B were up*, etc.. Thus, I take this sentence as having a similar structure as a sentence discussed in §6.1.3, i.e. (114), repeated as (114):

(122) If [at least one]$_F$ head had come up, I would have won.

In (122), as the focus falls on the generalized quantifier "at least one", the antecedent is about the quantity of coins coming up. I submit that (121) is similar: "not both", which means less than two, should be treated as a generalized quantifier as well, and a focus on it would explain why it is felt that the antecedent is about the quantity of switches that were up. Therefore, the counterfactual (121), with focus taken into account, becomes

(123) If [not both]$_F$ Switch A and Switch B were up, the light would be off.

Given this focus pattern, the CQ of the antecedent is clear: it is a set of propositions of the form *x Switch A and Switch B were up*, where *x* is a generalized quantifier. In particular, the CQ has the following propositions (and perhaps more) as its members, where < 2 is the answer provided by the antecedent:

< 2: Not both A and B were up.

1: Exactly one of A and B was up.

0: None of A and B were up.

Now we can apply GQA to explain why (123) sounds false. Note that 1 is logically stronger than < 2. With respect to the local context R of the antecedent, if < 2 and 1 are contextually equivalent, the GQA condition of Logical Strongness is violated. Thus, GQA requires < 2 and 1 not to be contextually equivalent in R, which in turn means that R

should include some possible worlds where none of A and B were up, i.e. some 0-worlds. Given that the interpreter of (123) accommodates this pragmatic requirement, it follows that in some < 2-worlds in R, none of A and B were up and thus the light would be on—that is, the consequent is false at some < 2-worlds in R. Therefore, once the GQA conditions are accommodated, the interpreter will judge (123) as false.

On the other hand, the standard Stalnaker-Lewis semantics, together with a plausible closeness ordering of worlds, predicts the truth of (120). As the two switches are both up in the actual world, it is plausible to assume that the closest one-down worlds are closer to the actual world than all the both-down worlds. Thus, the light is off at all the closest A-or-B-down worlds. Furthermore, the pragmatics of GQA does not alter this result. For the antecedent of (120), as it is a disjunction, let's assume that its CQ includes the following propositions, while *Or* is the answer provided by the antecedent:

Or: A or B is not up.

not-A: A is not up.

not-B: B is not up.

Because each of *not-A* and *not-B* is logically stronger than *Or*, Logical Strongness requires *Or* not to be contextually equivalent with *not-A* or *not-B* in R. Thus, R has to include some worlds where only A is down and some worlds where only B is down. However, accommodating such worlds into R does not mean that R will include any both-down worlds, especially because both-down worlds are presumably further away from the actual world than the closest one-down worlds. Therefore, even if we bring the Logical Strongness condition into account, (120) is still predicted to be true.

We can see that the contrast between (120) and (121) is explained by the difference in what CQ is triggered by their antecedents, as CQ invokes some pragmatic effect on the local context of the antecedent, which may affect the interpreter's judgment on truth-values. This account shows that the apparent contrast between (120) and (121) does not

need to be explained by a new semantics. Rather, if the current account is successful, the contrast is because of the upstream effect of CQs, while the standard semantics, with its commitment to the validity of SEA, still stands.

6.2.3. Summary. In this section, a putative counterexample to SEA is explained away by GQA, which supports the common conviction that SEA is valid. Given the validity of SEA, we have to invalidate SDA in order to avoid validating AS. Thus, protecting SEA also indirectly motivates a pragmatic explanation of the seeming validity of SDA, as offered in §6.1. Overall, the current pragmatic theory stands with the received semantic theory and defends it by using pragmatics to explain the putative counterexamples. The crux in this account is that the CQ of the antecedent of a counterfactual is able to affect the interpretation, as the GQA conditions constrain how the CQ should be related to the local context of the antecedent. This effect is within the realm of pragmatics because what gives rise to it is the interpreter's charitable disposition of taking the expression as obeying the pragmatic principles. Explaining the counterexamples in this way, GQA can be taken as, in a broad sense, an instance of Grice's razor which resists semantic revolutions: as GQA is a pragmatic solution to some puzzles which have been thought of as semantic puzzles, it thus provides a viable alternative to semantic solutions, making deviations from the standard semantics for counterfactuals unnecessary and less attractive.

CHAPTER 7

Conclusion

In this dissertation, I introduced the GQA theory as a pragmatic theory that treats clauses as contributing question-answer pairs and imposes certain pragmatic constraints on how the question-answer pair offered by a clause should appropriately relate to its local context. This novel pragmatic theory is first motivated in Chapter 2 by appeal to two precursors: the QUD approach and theories of presupposition projection that make use of the notion of local context. If QUD theories concern how question-answer pairs contributed by entire sentences should relate to global contexts, the GQA theory is their local counterpart, as it is a pragmatic theory concerning the relationship between clauses' question-answer pairs and local contexts. With this "local pragmatics" status clear, the GQA theory is formulated in Chapter 3. In particular, I give an open list of conditions governing the relationship between a clause's question-answer pair and the local context, i.e. conditions on what questions and answers count as "good" with regard to a local context. With these conditions, we are able to explain why question-triggers, such as focus, are capable of having a truth-conditional effect in some cases. This effect on knowledge ascription sentences and counterfactuals is explained in Chapters 4 and 5. The core idea there is that according to the GQA conditions, question-triggers matter for how a local context is resolved, and thus it can further affect the truth-conditional interpretation of an entire sentence, provided the interpretation depends on how the local context is resolved. This, I argue, is exactly what happens in interpreting the knowledge ascriptions and counterfactuals that exhibit the truth-conditional effect of question-triggers. Lastly, in Chapter 6 the GQA-based account is further applied to some allegedly semantic puzzles.

159

In particular, some examples that seemingly challenge the validity of Conjunction Conditionalization and some other examples render Simplification of Disjunctive Antecedent as plausible. But I called for a re-examination of such examples that takes Congruent Questions into account. Once we stop ignoring the relevant Congruent Questions, I argue, GQA is capable of explaining away the counterexamples, making it unnecessary to change the semantics according to them.

Overall, I hope this dissertation renders GQA a conceptually plausible pragmatic theory that can be successfully applied to explaining some intriguing phenomena. As a closing remark, I would like to highlight its potential of competing with semantic theories that are meant to cope with hyper-intensionality in language. As we have seen in the previous chapters, the truth-conditional interpretation of an utterance might be sensitive to the Congruent Question involved in an embedded clause, although the Congruent Question does not make a difference in what proposition the clause expresses. To repeat one of the examples, we have seen that the antecedents of (124) and (125) have the same propositional content but the truth-conditional interpretation of the two sentences are nonetheless different due to the Congruent Questions involved.

(124) If Clyde hadn't [married]$_F$ Bertha, he wouldn't have inherited the money.

(125) If Clyde hadn't married [Bertha]$_F$, he wouldn't have inherited the money.

This phenomenon does not only shows up in counterfactuals but also in some other constructions. For example, (126) and (127), as well as (128) and (129), may have different truth-conditional interpretations, even though the propositional contents of their embedded clauses are the same.

(126) Clyde didn't get the money because he didn't [marry]$_F$ Bertha.

(127) Clyde didn't get the money because he didn't marry [Bertha]$_F$.

(128) Carl recommended Clyde to [marry]$_F$ Bertha.

(129) Carl recommended Clyde to marry [Bertha]$_F$.

160

Such examples might motivate a hyper-intensional semantics. For instance, it might be tempting to claim that the semantic value of embedded clauses is richer than their propositional contents, and operators such as *because* and *recommend* take the rich contents as arguments. However, if GQA can explain the difference between (124) and (125) in terms of the pragmatics of Congruent Questions, it seems that in principle, the same strategy might be applicable to these new examples as well. In particular, it could be the case that the relevant operators still operate on the propositional contents of the embedded clauses, while the difference in truth-conditional interpretations is due to some pragmatic effect of Congruent Questions on local contexts. If this is correct, GQA, or some augmented version of it, has a chance to be an alternative to semantic accounts of the phenomenon, and thus a hyper-intensional semantics of *because* or *recommend* might not be the only available option.

That said, I do not mean that we should reject hyper-intensional semantics, but the fact that a pragmatic account is applicable to some examples of this sort is indicative that we should not directly take the phenomenon as evidence for such a semantic solution. At the very least, a pragmatic theory of Congruent Question is a contender that needs to be addressed by those who propose a semantic account of the phenomenon. Although it is not clear, at this point, how exactly GQA, or any pragmatic theory in the vicinity, can be applied to these new cases, it suffices, for now, to point out this possibility and let further explorations decide whether it is indeed a worthy direction to go.

Bibliography

Abusch, D. (2009). Presupposition triggering from alternatives. *Journal of Semantics*, 27:37–80.

Abusch, D. (2010). Presupposition triggering from alternatives. *Journal of Semantics*, 27(1):37–80.

Aloni, M. and Égré, P. (2010). Alternative questions and knowledge attributions. *Philosophical Quarterly*, 60(238):1–27.

Alonso-Ovalle, L. (2006). *Disjunction in alternative semantics*. PhD thesis, University of Massachusetts Amherst.

Anderson, A. R. (1951). A note on subjunctive and counterfactual conditionals. *Analysis*, 12(2):35–38.

Beaver, D. and Clark, B. (2003). Always and only: Why not all focus-sensitive operators are alike. *Natural Language Semantics*, 11(4):323–362.

Beaver, D. I. and Clark, B. Z. (2008). *Sense and Sensitivity: How Focus Determines Meaning*. Blackwell.

Bennett, J. (1974). Counterfactuals and possible worlds. *Canadian Journal of Philosophy*, 4(December):381–402.

Bennett, J. (2003). *A Philosophical Guide to Conditionals*. Oxford University Press.

Blome-Tillmann, M. (2009). Knowledge and presuppositions. *Mind*, 118(470):241–294.

Bonomi, A. and Casalegno, P. (1993). *Only*: Association with focus in event semantics. *Natural Language Semantics*, 2(1):1–45.

Briggs, R. (2012). Interventionist counterfactuals. *Philosophical Studies*, 160(1):139–166.

Brogaard, B. (2009). What mary did yesterday: Reflections on knowledge-wh. *Philosophy and Phenomenological Research*, 78(2):439–467.

Büring, D. (2016). *Intonation and meaning*. Oxford University Press.

Ciardelli, I., Groenendijk, J., and Roelofsen, F. (2019). *Inquisitive Semantics*. Oxford University Press.

Ciardelli, I., Zhang, L., and Champollion, L. (2018). Two switches in the theory of counterfactuals. *Linguistics and Philosophy*, 41(6):577–621.

Cohen, S. (1999). Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives*, 13(s13):57–89.

Cresswell, M. J. and Stechow, A. (1982). De re belief generalized. *Linguistics and Philosophy*, 5(4):503–535.

Dretske, F. I. (1970). Epistemic operators. *Journal of Philosophy*, 67(24):1007–1023.

Dretske, F. I. (1972). Contrastive statements. *Philosophical Review*, 81(4):411–437.

Emmanuel Chemla, P. S. (2012). Incremental vs. symmetric accounts of presupposition projection: an experimental approach. *Natural Language Semantics*, 2(20):177–226.

Fine, K. (1975). Critical notice. *Mind*, 84(335):451–458.

Fine, K. (2012a). Counterfactuals without possible worlds. *Journal of Philosophy*, 109(3):221–246.

Fine, K. (2012b). A difficulty for the possible worlds analysis of counterfactuals. *Synthese*, 189(1):29–57.

von Fintel, K. (1994). *Restrictions on Quantifier Domains*. PhD thesis, University of Massachusetts at Amherst.

von Fintel, K. (2001). Counterfactuals in a dynamic context. *Current Studies in Linguistics Series*, 36:123–152.

Galles, D. and Pearl, J. (1998). An axiomatic characterization of causal counterfactuals. *Foundations of Science*, 3(1):151–182.

Gauker, C. (1998). What is a context of utterance? *Philosophical Studies*, 91(2):149–172.

Gerken, M. and Beebe, J. R. (2016). Knowledge in and out of contrast. *Noûs*, 50(1):133–164.

Gillies, A. S. (2007). Counterfactual scorekeeping. *Linguistics and Philosophy*, 30(3):329–360.

Hamblin, C. L. (1973). Questions in Montague English. *Foundations of Language*, 10(1):41–53.

He, C. (2016). Conjunction, connection and counterfactuals. *Erkenntnis*, 81(4):705–719.

Heim, I. (1983). On the projection problem for presuppositions. In Portner, P. and Partee, B. H., editors, *Formal Semantics - the Essential Readings*, pages 249–260. Blackwell.

Herburger, E. (2000). *What counts: Focus and quantification*, volume 36. MIT Press.

Hintikka, J. (1962). *Knowledge and Belief*. Ithaca: Cornell University Press.

Hughes, M. (2013). Problems for contrastive closure: Resolved and regained. *Philosophical Studies*, 163(3):577–590.

Kadmon, N. (2001). *Formal Pragmatics: Semantics, Pragmatics, Presupposition, and Focus*.

Kadmon, N. and Sevi, A. (2011). Without 'focus'. *The Baltic International Yearbook of Cognition, Logic and Communication*, 6:18.

Kallestrup, J. (2009). Knowledge-wh and the problem of convergent knowledge. *Philosophy and Phenomenological Research*, 78(2):468–476.

Karjalainen, A. and Morton, A. (2003). Contrastive knowledge. *Philosophical Explorations*, 6(2):74–89.

Karttunen, L. (1973). Presuppositions of compound sentences. *Linguistic inquiry*, 4(2):169–193.

Karttunen, L. (1974). Presuppositions and linguistic context. *Theoretical Linguistics*, 1:181–194.

Kelp, C. (2011). A problem for contrastivist accounts of knowledge. *Philosophical Studies*, 152(2):287–292.

Khoo, J. (2018). Disjunctive antecedent conditionals. *Synthese*, 198(8):7401–7430.

Kratzer, A. (1981a). The notional category of modality. In Eikmeyer, H. and Rieser, H., editors, *Words, Worlds, and Contexts: New Approaches in Word Semantics*, pages 38–74. Walter de Gruyter.

Kratzer, A. (1981b). Partition and revision: The semantics of counterfactuals. *Journal of Philosophical Logic*, 10(2):201–216.

Krifka, M. (1992). A framework for focus-sensitive quantification. In *Semantics and Linguistic Theory*, volume 2, pages 215–236.

Langendoen, D. T. and Savin, H. (1971). The projection problem for presuppositions. In Fillmore, C. J. and Langndoen, D. T., editors, *Studies in Linguistic Semantics*, pages 54–60. Irvington.

Lewis, D. K. (1973). *Counterfactuals*. Blackwell.

Lewis, D. K. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8(1):339–359.

Lewis, D. K. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74(4):549–567.

Lewis, K. (2017). Dynamic semantics. In *Oxford Handbooks Online*. Oxford University Press.

McDermott, M. (2007). True antecedents. *Acta Analytica*, 22(4):333–335.

McGlynn, A. (2012). The problem of true-true counterfactuals. *Analysis*, 72(2):276–285.

Morton, A. (2012). Contrastive knowledge. In Blaauw, M., editor, *Contrastivism in Philosophy*, pages 101–115. New York: Routledge.

Nichols, C. (2017). *New Work on the Old World Order*. PhD thesis, Dept. of Philosophy, Princeton University.

Nute, D. (1980). *Topics in conditional logic*, volume 20. Springer Science & Business Media.

Penczek, A. (1997). Counterfactuals with true components. *Erkenntnis*, 46(1):79–85.

Pollock, J. (1976). *Subjunctive Reasoning*. Reidel.

Potts, C. (2005). *The Logic of Conventional Implicatures*. Oxford University Press UK.

Roberts, C. (2012). Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics*, 5:1–69.

Rooth, M. (1985). *Association with focus*. PhD thesis, Dept. of Linguistics, University of Massachusetts, Amherst.

Rooth, M. (1992). A theory of focus interpretation. *Natural Language Semantics*, 1(1):75–116.

Rooth, M. (1999). Association with focus or association with presupposition? In Bosch, P. and van der Sandt, R., editors, *Focus: Linguistic, cognitive, and computational perspectives*, pages 232–246. Cambridge Cambridge University Press.

Saul, J. M. (2012). *Lying, Misleading, and What is Said: An Exploration in Philosophy of Language and in Ethics*. Oxford University Press.

Schaffer, J. (2004). From contextualism to contrastivism. *Philosophical Studies*, 119(1-2):73–104.

Schaffer, J. (2005). Contrastive knowledge. In Gendler, T. S. and Hawthorne, J., editors, *Oxford Studies in Epistemology 1*, page 235. Oxford University Press.

Schaffer, J. (2007a). Closure, contrast, and answer. *Philosophical Studies*, 133(2):233–255.

Schaffer, J. (2007b). Knowing the answer. *Philosophy and Phenomenological Research*, 75(2):383–403.

Schaffer, J. (2008). The contrast-sensitivity of knowledge ascriptions. *Social Epistemology*, 22(3):235–245.

Schaffer, J. (2009). Knowing the answer redux. *Philosophy and Phenomenological Research*, 78(2):477–500.

Schaffer, J. and Knobe, J. (2012). Contrastive knowledge surveyed. *Noûs*, 46(4):675–708.

Schaffer, J. and Szabo, Z. G. (2013). Epistemic comparativism: A contextualist semantics for knowledge ascriptions. *Philosophical Studies*, (2):1–53.

Schlenker, P. (2009). Local contexts. *Semantics and pragmatics*, 2:3–1.

Schlenker, P. (2010). Presuppositions and local contexts. *Mind*, 119(474):377–391.

Schwarzschild, R. (1999). GIVENness, AvoidF and other constraints on the placement of accent. *Natural Language Semantics*, 7(2):141–177.

Simons, M. (2003). Presupposition and accommodation: Understanding the stalnakerian picture. *Philosophical Studies*, 112(3):251–278.

Simons, M., Beaver, D., Roberts, C., and Tonhauser, J. (2017). The best question: Explaining the projection behavior of factives. *Discourse Processes*, 54(3):187–206.

Sinnott-Armstrong, W. (2004). Classy pyrrhonism. In Sinnott-Armstrong, W., editor, *Pyrrhonian Skepticism*, pages 188–207. Oxford University Press.

Sinnott-Armstrong, W. (2008). A contrastivist manifesto. *Social Epistemology*, 22(3):257–270.

Stalnaker, R. (1968). A theory of conditionals. In Rescher, N., editor, *Studies in Logical Theory*, pages 98–112. Oxford: Blackwell.

Stalnaker, R. (1973). Presuppositions. *Journal of Philosophical Logic*, 2(4):447–457.

Stalnaker, R. (1974). Pragmatic presuppositions. In Stalnaker, R., editor, *Context and Content*, pages 47–62. Oxford University Press.

Stalnaker, R. (1978). Assertion. *Syntax and Semantics (New York Academic Press)*, 9:315–332.

Stalnaker, R. C. (2002). Common ground. *Linguistics and Philosophy*, 25(5-6):701–721.

Starr, W. B. (2014). What 'if'? *Philosophers' Imprint*, 14.

von Stechow, A. (1981). Topic, focus and local relevance. In *Crossing the boundaries in linguistics*, pages 95–130. Springer.

von Stechow, A. (1991). Focusing and backgrounding operators. *Discourse particles*, 6:37–84.

Steglich-Petersen, A. (2014). Knowing the answer to a loaded question. *Theoria*, 80(1):97–125.

Walters, L. (2016). Possible world semantics and true-true counterfactuals. *Pacific Philosophical Quarterly*, 97(3):322–346.

Walters, L. and Williams, R. (2013). An argument for conjunction conditionalization. *Review of Symbolic Logic*, 6(4):573–588.

Willer, M. (2018). Simplifying with free choice. *Topoi*, 37(3):379–392.