ABSTRACT

Title of Dissertation: NEW PERSPECTIVES ON INQUISITIVE SEMANTICS

Yichi Zhang, Doctor of Philosophy, 2022

Dissertation directed by: Associate Professor Eric Pacuit
– Department of Philosophy
Associate Professor Paolo Santorio
– Department of Philosophy

Inquisitive semantics offers a unified analysis of declarative and interrogative sentences by construing information exchange as a process of raising and resolving issues. In this dissertation, I apply and extend inquisitive semantics in various new ways. On the one hand, I build upon the theoretical insight of inquisitive semantics and explore the prospect of incorporating other types of content into our conception of information exchange. On the other hand, the logical framework underlying inquisitive semantics is also of great interest in itself as it enjoys certain unique properties and is thus worth further investigation. In the first paper, I provide an account of live possibilities and model the dynamics of bringing a possibility to salience using inquisitive semantics. This account gives rise to a new dynamic analysis of conditionals, which is capable of capturing what I call the *Extended Sobel Inference*. In the second paper, drawing on the fact that disjunction in inquisitive semantics is understood as introducing a set of alternative answers to a question, I propose a Questions-Under-Discussion-based account of informational redundancy to tackle various Hurford sentences. In the third paper, I explore the prospect of cashing out the theoretical intuition behind inquisitive semantics using a non-bivalent framework. I develop a new logic

which invalidates the Law of Excluded Middle just like inquisitive logic, but unlike inquisitive logic, it employs a negation that vindicates Double Negation Elimination.

NEW PERSPECTIVES ON INQUISITIVE SEMANTICS

by

Yichi (Raven) Zhang
The Department of Philosophy

Dissertation submitted to the Faculty of the Graduate School of the University of Maryland, College Park in partial fulfillment of the requirements for the degree of
Doctor of Philosophy
June, 2022

Advisory Committee:

Associate Professor Eric Pacuit, Chair
Associate Professor Paolo Santorio, Co-Chair
Associate Professor Fabrizio Cariani
Professor Valentine Hacquard
Associate Professor Alexander Williams, Dean's Representative

Acknowledgement

First and foremost, I would like to extend my deepest gratitude to my supervisors: Eric Pacuit and Paolo Santorio. Without their guidance, extensive feedback, and unwavering support, this dissertation would not have been possible. I am also immensely grateful to my other committee members: Fabrizio Cariani, Alexander Williams, and Valentine, for their feedback, encouragement and support.

Additionally, I would like to thank Steven Kuhn for providing me with very helpful comments, and members of the logic group and the MEANING group at the University of Maryland for valuable discussion.

Lastly, I am thankful to my colleagues, faulty and staff at the department of philosophy for their support.

Table of Contents

Chapter 1

Introduction

We use language for a multitude of purposes, and one common use is to exchange information. Among the types of information we normally exchange, there are those that concern what the world is like and the ways things are. In asserting "Alice is at the party," the speaker proposes to make it a piece of common ground information that what the sentence says is true (cf. Stalnaker, 1978). Whereas the focus on truth-conditional information reflects historical emphasis on declarative sentences and assertions, linguistic investigation into questions (Hamblin, 1973; Karttunen, 1977; Groenendijk & Stokhof, 1984) and, in particular, the recent development of inquisitive semantics (Ciardelli & Roelofsen, 2011; Ciardelli et al., 2013; 2018) have contributed to expanding the conception of information exchange. According to inquisitive semantics, in uttering an interrogative sentence, the speaker is taken not as suggesting that the world is in any particular way but rather as raising an issue by putting forward alternative ways for it to be resolved, thereby orienting future exchange in a certain direction. In uttering "Is Alice at the party?" for instance, the speaker proposes two different ways to update the common ground information: either Alice is at the party, or she is not. It then invites the interlocutor to resolve this issue by settling on one of the two alternatives. Information exchange, under inquisitive semantics, can thus be viewed as a process of raising and resolving issues.

The main aim of this dissertation is to apply and extend inquisitive semantics in various new ways. On the one hand, I build upon the theoretical insight of inquisitive semantics and explore the prospect of incorporating other types of content into our conception of information exchange. On the other hand, the logical framework underlying inquisitive semantics is also of great interest in itself as it enjoys certain unique properties and is thus worth further investigation.

Before describing my work in some more detail, let me provide a brief overview of inquisitive semantics. Formally, inquisitive semantics construes the semantic content of a question as a set of Stalnakerian propositions (i.e., sets of possible worlds) each of which represents a body of information that is capable of resolving the question under discussion. The idea that the semantic content of a question is identified as a set of propositions is not unique to inquisitive semantics. For instance, Hamblin (1973) takes a question to denote a set of propositions each of which represents a possible answer to the question; similarly for Karttunen (1977), each proposition represents a *true* answer to the question. While these other accounts spell out the semantics of a question in terms of its answerhood, inquisitive semantics does so in terms of resolution conditions. A set of worlds can contain enough information so as to resolve a question even if the corresponding proposition is not normally viewed as an answer to the question.

To elucidate, consider a universe that contains four worlds: $AB, A\bar{B}, \bar{A}B, \bar{A}\bar{B}$. In AB, both Alice and Bob are at the party; in $\bar{A}\bar{B}$, neither of them is; in $A\bar{B}$ and $\bar{A}B$, one but not the other is. Under the Hamblin-style alternative semantics, the question "Is Alice at the party?" denotes the set $\{\{AB, A\bar{B}\}, \{\bar{A}B, \bar{A}\bar{B}\}\}$. This set contains two propositions: the proposition that Alice is at the party (i.e., $\{AB, A\bar{B}\}$), and the proposition that she isn't (i.e., $\{\bar{A}B, \bar{A}\bar{B}\}$). By contrast, under inquisitive semantics, in addition to the two sets $\{AB, A\bar{B}\}$ and $\{\bar{A}B, \bar{A}\bar{B}\}$, the set of propositions the question denotes also contains all the subsets of these two sets such as $\{AB\}$, $\{A\bar{B}\}$, and so on. Given that the set $\{AB\}$ embodies the information that Alice and Bob are both at the party, the information contained is enough to resolve the question "Is Alice at the party?". Hence, this set is included in the denotation of the question, even if "Alice and Bob are both at the party" is not normally conceived of as an answer to "Is Alice at the party?" under alternative semantics. By taking the resolution condition instead of answerhood as its central notion, inquisitive semantics places additional constraints on what kind of set-theoretic entities can serve as the proper denotation of questions: only sets of sets of worlds that are closed under the subset relation can serve this purpose.

As a consequence of making the semantic denotation satisfy this closure property, inquisitive semantics imparts a well-behaved algebraic structure that facilitates finding a suitable notion of entailment as well as defining other logical operations. In inquisitive semantics, semantic entailment can simply be defined via the subset relation: $\phi \vDash \psi$ iff $[\![\phi]\!] \subseteq [\![\psi]\!]$, where $[\![\phi]\!]$ and $[\![\psi]\!]$ are the denotation of ϕ and ψ. For example, consider the question "Who (of Alice and Bob) is at the party?". With respect to the aforementioned universe consisting of the four worlds $AB, A\bar{B}, \bar{A}B$, and $\bar{A}\bar{B}$, this question denotes the set $\{\{AB\}, \{A\bar{B}\}, \{\bar{A}B\}, \{\bar{A}\bar{B}\}, \varnothing\}$.[1] Now recall that the question "Is Alice at the party?" denotes the set that contains $\{AB, A\bar{B}\}$, $\{\bar{A}B, \bar{A}\bar{B}\}$, and all of their subsets. Given this, the denotation of "Who (of Alice and Bob) is at the party?" is indeed a subset of "Is Alice at the party?". Thus, under inquisitive semantics, the former question entails the latter. This is a desirable result since intuitively what this entailment conveys is that every complete answer to "Who (of Alice and Bob) is at the party?" also provides an answer to "Is Alice at the party?" (cf. Roberts, 1996).

In a similar vein, inquisitive semantics manages to provide simple set-theoretic definitions for logical operations on pairs of questions. For instance, conjunction is simply defined as set-intersection. The denotation of "Who (of Alice and Bob) is at the party?" can be derived from taking the intersection of the denotation of "Is Alice at the party?" and that of "Is Bob at the party?"; intuitively, the former *wh*-question is equivalent to conjoining the two latter polar questions. Analogously, disjunction is defined as set-union. As such, to resolve a disjunctive question such as "Where can we rent a car, or who might have one that we could borrow?" (Ciardelli et al., 2018, p.16.), it is enough to resolve one of the disjoined questions. This ability to properly define entailment and other logical operations is one main advantage of inquisitive semantics over other Hamblin-style alternative-based semantics (see Ciardelli et al. 2017 for further discussion).

What the formal architecture of inquisitive semantics further enables is a unified analysis of interrogative and declarative sentences. Both interrogative and declarative sentences now denote a set of sets of worlds. For declarative sentences, this amounts to taking the

[1] The empty set is included because it is a subset of every set. Conceptually, the empty set represents a body of information that is inconsistent. Assuming that the principle of explosion holds, an inconsistent body of information will entail everything and as such will resolve every question.

Stalnakerian proposition that the sentence standardly denotes and then forming the set that contains it and all of its subsets. By unifying the semantic denotation of declarative and interrogative sentences, inquisitive semantics offers a straightforward way to capture embeddings of interrogatives under propositional attitudes (Ciardelli & Roelofsen, 2015; Ciardelli et al., 2018), as in (1) and (2):

(1) John knows who is at the party.

(2) John wonders whether Alice is at the party.

Furthermore, a unified account allows inquisitive semantics to supply an enriched notion of conversational contexts. By construing conversational contexts as a set of sets of worlds instead of a single set of worlds as under the more traditional analysis (e.g., Stalnaker, 1978; Veltman, 1996), inquisitive semantics enables contexts to encode not only what facts have been settled but also what issues have been brought up. As such, inquisitive semantics can model information exchange that goes beyond the exchange of truth-conditional content.

The current project of exploring new perspectives on inquisitive semantics has yielded three independent papers. In ***Weak Persistence***, I articulate a notion of live possibilities with the help of inquisitive semantics. In addition to conveying truth-conditional information and raising issues, sentences in natural language can also draw attention to certain possibilities. As a prominent example, epistemic *might*-claims, such as "Alice might be at the party", can be conceived of as having the discourse function to highlight a possibility or entertain a new one. Conversational contexts can likewise be enriched to represent what possibilities have become salient in discourse. While similar ideas have been explored in the past (e.g., Willer, 2013; see also, Westera, 2017), my current aim is to cash out this dynamics of bringing a possibility to salience in an inquisitive framework in the hope of better integrating this notion of live possibilities with inquisitive content in the future.

In this paper, I present a framework that is capable of modeling the dynamics of bringing a possibility to salience with the aim of capturing what I call the *Extended Sobel Inference* (ESI) as illustrated by (3):

(3) If Alice comes, the party will be fun. But if Alice and Bob both come, the party won't be fun. $\Rightarrow$ Therefore, if Alice but not Bob comes, the party will be fun.

Despite the intuitiveness of this inference, ESI poses a challenge for two leading accounts of conditionals: the variably strict analysis (Lewis, 1973; Stalnaker, 1968) and the dynamic strict analysis (von Fintel, 2001; Gillies, 2007; Willer, 2017). While the latter fails to vindicate ESI, the former licenses the additional inference to "If Alice comes, then Bob won't come", which does not seem to be a natural inference that people will normally draw. By contrast, my account is able to validate ESI without validating the additional inference. I adopt an enriched notion of conversational contexts that is capable of distinguishing live possibilities that have become salient from plain possibilities. I then postulate the following principle governing the dynamics of raising a possibility to salience:

> **Weak Persistence:** When a plain possibility ϕ is brought to salience, past information is preserved either in all the ϕ-possibilities or in all the $\neg\phi$-possibilities.

In a nutshell, when the possibility of Alice and Bob both coming to the party is brought to salience by the second conditional in the sequence—i.e., "If Alice and Bob both come, the party won't be fun"—the past information embodied by "If Alice comes, the party will be fun" is preserved in all the possibilities where Alice but not Bob comes. This means that in all the possibilities where Alice but not Bob comes, the party will still be fun, which thereby vindicates ESI. Moreover, it does not validate the further inference to "If Alice comes, then Bob won't come" and thus avoids the drawback of the variably strict analysis.

The second paper ***A Question Under Discussion Based Account of Redundancy*** is focused on providing an account of informational redundancy that can adequately predict the infelicity of various Hurford sentences. Hurford (1974) observed that disjunctions where one disjunct entails the other are generally infelicitous:

(4) #John was born in Paris, or he was born in France.

To explain their infelicity, one common approach is to take Hurford disjunctions as involving a truth-conditionally redundant constituent whose deletion has no effect on the truth-conditional content of the whole sentence. Such a redundancy account has been offered within an inquisitive framework (Ciardelli & Roelofsen, 2017; Anvari, 2021). To explain in brief, recall that under inquisitive semantics, declarative and interrogative sentences uniformly denote sets of Stalnakerian propositions that are closed under the subset relation. Given that disjunction is viewed as a device to introduce alternatives and defined via set-theoretic union, (4) denotes the set that contains the two Stalnakerian propositions ⟦John was born in Paris⟧ and ⟦John was born in France⟧ as well as all of their subsets. But give that ⟦John was born in Paris⟧ is itself a subset of ⟦John was born in France⟧, what (4) denotes in fact is just the set that contains ⟦John was born in France⟧ and all of its subsets. In other words, (4) is equivalent to the simplification "John was born in France". As a result, the first disjunct "John was born in Paris" turns out to be redundant.

However, this analysis fails to explain why the disjunction in (5), which is an instance of what I call *conjunctive Hurford disjunctions*, is perceived defective:

(5) #John was born in Paris, or he was born in France and Mary was born in London.

Indeed, no theory on the market can adequately capture the infelicity of conjunctive Hurford disjunctions. To address this challenge, I follow Simons (2001) and propose an analysis that utilizes the notion of questions under discussion (van Kuppevelt 1995; Roberts, 1996; Büring, 2003). Similar to the inquisitive approach sketched above, I take disjunction as a device to introduce alternatives in the sense that each disjunct is supposed to provide an answer, and moreover a distinct answer, to some discourse question. But disjunction also embodies additional information that is not covered by the standard inquisitive approach, namely that each disjunct should also answer the discourse question "in the same way". I will show how this analysis can capture the infelicity (and felicity) of a wide range of Hurford sentences.

Finally, as previously mentioned, the formal architecture and algebraic structure underlying inquisitive semantics is also of great interest in itself. In ***A Non-Bivalent Approach to Inquisitive Logic***, I aim to extend the study of inquisitiveness in this direction. The central question here is whether there are other formal frameworks that can cash out the theoretical intuition behind inquisitive semantics. One essential feature of inquisitive semantics is failure of the Law of Excluded Middle (LEM)—that is, $A \vee \neg A$ is no longer valid. In inquisitive semantics, $A \vee \neg A$ serves as the logical form of a polar question like "Is Alice at the party?". In the standard setting, what it means for a sentence to be valid is for it to be completely informationally trivial: an utterance of it will have no effect on the conversational context. But since $A \vee \neg A$ raises a question, it is not completely informationally trivial. Consequently, LEM is not valid in inquisitive semantics.

Standard inquisitive semantics invalidates LEM by employing an intuitionistic negation, which as a consequence also invalidates Double Negation Elimination (DNE)—that is, $\neg\neg A$ no longer entails A. By contrast, I show that there is a different route adopters of the inquisitive approach can take. I develop a new logic for modeling inquisitiveness which employs a negation that validates DNE but not LEM. The framework is non-bivalent in the sense that it simultaneously utilizes two non-complementary notions when defining semantic satisfaction conditions: *support* and *rejection*. For instance, given a body of information, an assertion is supported if it is settled true and rejected if it is settled false; a question is supported if it is resolved by the body of information and rejected if it is incompatible with the body of information, and in the latter case, asking the question will be deemed infelicitous in the first place. For example, if it is already common ground that neither Alice nor Bob is at the party, then the following alternative question will be rejected and deemed infelicitous.

(6) Is Alice at the party↑, or is Bob at the party↓?[2]

The two notions above are non-complementary because there can be sentences that are neither supported nor rejected by a given body of information. I present an algebraic semantics for this logic via the so-called twist-structures. As a generalization, I show how this method of constructing twist-structures can be systematically employed to convert bivalent systems to many-valued ones. As such, it serves as a valuable tool in a philosopher's toolbox.

[2]The upward and downward arrows are used to indicate intonation. An interrogative sentence can receive different interpretations depending on the intonation pattern (see e.g., Roelofsen, 2015)

Chapter 2

Weak Persistence

This paper investigates an inference pattern that is intuitively valid in our ordinary reasoning with conditionals. I call this inference the *Extended Sobel Inference* (ESI). I show that the validity of ESI poses a problem for both the standard variably strict analysis of conditionals (VSA) and the dynamic strict analysis of conditionals (DSA). Whereas DSA fails to vindicate ESI, VSA, albeit validating ESI, validates another inference which does not appear to hold under all circumstances. In response, I propose a new dynamic analysis of conditionals by first devising a novel dynamic inquisitive framework that enables the modeling of information dynamics associated with entertaining a new possibility. In particular, I formulate a notion of *weak persistence*, which captures how past information is preserved in light of new possibilities. Combined with a notion of postsuppositions, the current framework manages to vindicate ESI without incurring the drawback of VSA.

Keywords: Conditionals, Sobel sequences, Dynamic strict analysis, Inquisitive semantics, Update semantics, Postsupposition

2.1 Introduction

Consider the following inferences:

$$
\text{(1)}\quad \begin{array}{l}
\text{(a) If Alice comes, the party will be fun.} \\
\text{(b) But if Alice and Bob both come, the party won't be fun.} \\
\hline
\text{(c) If Alice but not Bob comes, the party will be fun.}
\end{array}
$$

$$
\text{(2)}\quad \begin{array}{l}
\text{(a) If Alice had come, the party would have been fun.} \\
\text{(b) But if Alice and Bob both had come, the party wouldn't have been fun.} \\
\hline
\text{(c) If Alice but not Bob had come, the party would have been fun.}
\end{array}
$$

Intuitively, the first two conditionals in each sequence entail the last one. Moreover, the naturalness of such inferences appears unaffected by whether the conditionals involved are indicatives, as in the case of (1), or counterfactuals, as in the case of (2). Let us call this general form of inferences in our conditional reasoning the *Extended Sobel Inference* (ESI), which we can schematically represent as $\phi > \psi, (\phi \wedge \chi) > \neg\psi \vDash (\phi \wedge \neg\chi) > \psi$, where $\ulcorner \phi > \psi \urcorner$ abbreviates $\ulcorner$if ϕ, then (would) $\psi\urcorner$.

The inference is so called because its two premises constitute a Sobel sequence, that is, a sequence consisting of, in the following order, two conditionals $\phi > \psi$ and $(\phi \wedge \chi) > \neg\psi$, as exemplified by the two premises from (1) and (2). The fact that Sobel sequences are felicitous suggests that antecedent strengthening—that is, the inference from $\phi > \psi$ to $(\phi \wedge \chi) > \psi$—fails as a natural inference for conditionals (Lewis, 1973; Stalnaker, 1968; Willer, 2017).

The current observation concerning ESI suggests that although, given the failure of antecedent strengthening, $\phi > \psi$ alone entails neither $(\phi \wedge \chi) > \psi$ nor $(\phi \wedge \neg\chi) > \psi$, nonetheless, by extending the premise into a Sobel sequence, we are warranted to infer either $(\phi \wedge \chi) > \psi$ or $(\phi \wedge \neg\chi) > \psi$ depending on how $\phi > \psi$ is extended. For instance, if we follow up $\phi > \psi$ with $(\phi \wedge \chi) > \neg\psi$, as in the case of (1) and (2), we can infer $(\phi \wedge \neg\chi) > \psi$. Alternatively, if we follow up $\phi > \psi$ with $(\phi \wedge \neg\chi) > \neg\psi$, as in the case of (3) below, we can infer $(\phi \wedge \chi) > \psi$.

(3) (a) If Alice comes, the party will be fun. (b) If Alice comes but she isn't in a good mood, the party won't be fun. $\Rightarrow$ (c) If Alice comes and she is in a good mood, the party will be fun.

As it turns out, vindicating ESI is not an easy task. It causes different problems for two leading accounts of conditionals, namely the variably strict analysis (VSA) and the dynamic strict analysis (DSA). After briefly introducing these two accounts in §2.2, I argue in §2.3 that, on the one hand, DSA fails to validate ESI; on the other hand, although VSA validates ESI, it further licenses an additional inference—i.e., the inference from $\phi > \psi$ and $(\phi \wedge \chi) > \neg\psi$ to $\phi > \neg\chi$—which, as (4) demonstrates, does not appear to hold universally:

(4) (a) If Alice comes, the party will be fun. (b) But if Alice and Bob both come, the party won't be fun. $\not\Rightarrow$ (c) If Alice comes, Bob won't come.

Drawing inspiration from DSA, inquisitive semantics, and update semantics, I propose a new dynamic account of conditionals that vindicates ESI without incurring the drawback of VSA. The details of this framework need to wait for the second half of this paper, but here is the general idea.

Standardly, the effect of assertions is to remove certain possibilities from a stock of shared information (Stalnaker, 1978), and whenever a possible world is removed, it can never be brought back. By contrast, one distinguishing feature of this new framework is that it allows updates to be genuinely non-eliminative. This is meant in the following sense: when we entertain a new possibility, worlds that have been previously eliminated can be resurrected. For instance, in the case of (1), the utterance of the first conditional (1a) $A > F$ eliminates all worlds where Alice comes to the party and the party is not fun. As we entertain the new possibility that Alice and Bob both come to the party with (1b), some of those previously eliminated worlds can be resurrected. As a result, it is no longer the case that in all worlds wherein Alice comes to the party, the party is fun. This in turn allows us to capture failure of antecedent strengthening.

At the same time, the revival of worlds needs to be constrained in a way such that ESI comes out valid. To this end, I postulate a constraint called *weak persistence*, which, for the time being, can be characterized rather informally as follows:

> **Weak Persistence:** When a hitherto unentertained possibility ϕ is introduced, past information is preserved either in all the ϕ-worlds or in all the $\neg\phi$-worlds.

Essentially, what weak persistence does is afford two alternative ways for past information (in particular, the information embodied by the first conditional in a Sobel sequence) to be preserved in light of new possibilities. Thus, when the second conditional in a Sobel sequence subsequently eliminates one of the two alternatives, we are warranted to either infer $(\phi \wedge \chi) > \psi$ or infer $(\phi \wedge \neg\chi) > \psi$. In the case of (1), since the second conditional "if Alice and Bob both come, the party won't be fun" rules out the option of preserving the information embodied by (1a) (viz., that all worlds where Alice comes are worlds where the party is fun) in all worlds where Alice and Bob both come, the information must be preserved in the other alternative, thereby licensing the inference to (1c).

2.2 Preliminaries

2.2.1 Variably Strict Analysis of Conditionals

The observation that Sobel sequences are felicitous causes trouble for any analyses of conditionals that validate antecedent strengthening. One such prominent theory is the strict material analysis of conditionals, according to which $\phi > \psi$ is interpreted as $\Box(\phi \supset \psi)$ and is thus evaluated true at a world w iff in all accessible worlds from w where ϕ is true, ψ is true. But since all the worlds where ϕ and χ are both true must be worlds where ϕ is true, antecedent strengthening holds. Hence, as the folklore goes, conditionals cannot receive a strict interpretation.

By contrast, antecedent strengthening fails under the Lewis-Stalnaker style variably strict analysis (Lewis, 1973; Stalnaker, 1968, 1981). Under VSA, roughly, $\phi > \psi$ is true at w iff ψ is true at all the closest accessible world(s) to w where ϕ is true. Following Lewis, we use a system of similarity spheres to model closeness among worlds. A system of spheres is a collection of nested sets of possible worlds centered on a world w at which the utterance is evaluated. One world w' is closer to w than another world w'' iff w' is located in a sphere that is a proper subset of the sphere in which w'' is located. As Figure 2.1 shows, since it is possible that the closest worlds where Alice comes to the party (i.e., the A-worlds in S_2) are distinct from the closest worlds where Alice and Bob both come to the party (i.e., the $(A \wedge B)$-worlds in S_3), we do not have the entailment from $A > F$ to $(A \wedge B) > F$. Hence, antecedent strengthening fails under VSA, as desired.

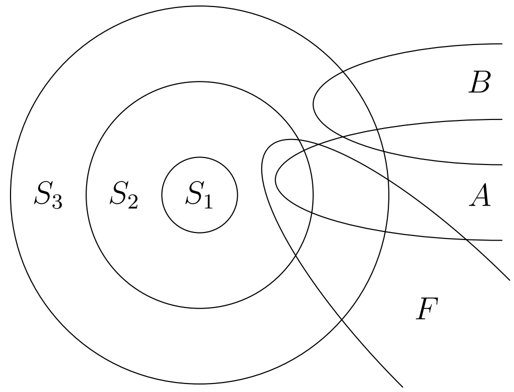

Figure 2.1: A system of spheres

2.2.2 Dynamic Strict Analysis of Conditionals

Within the dynamic tradition, conversation is understood as a process of updating common ground information. The meaning of a sentence is taken to be its update potential on information states, and hence the name *context change potential* (CCP). A context or information state is standardly construed as a set of worlds, and uttering a sentence has the effect of eliminating from it worlds that are incompatible with the information embodied by the sentence. For instance, uttering "Alice is at the party" will eliminate every world where Alice is not at the party from its input information state.

Under DSA, conditionals receive a strict interpretation with respect to a constantly evolving modal domain (von Fintel, 2001; Gillies, 2007; Willer, 2017, 2018). Gillies represents this background modal domain via a system of spheres—just like the one depicted in Figure 2.1—and he calls it a *hyper-domain*. Spheres in a hyper-domain represent information states, and the innermost sphere represents the current *modal center* over which the universal quantifier of the strict conditional ranges.[1]

The CCP of a conditional $\ulcorner \phi > \psi \urcorner$ comprises two parts: the CCP of its entertainability presupposition $\ulcorner \Diamond \phi \urcorner$, and that of its asserted content $\ulcorner \Box(\phi \supset \psi) \urcorner$. The entertainability presupposition $\Diamond \phi$ requires that the conditional's antecedent is possible with respect to the current modal center, i.e., the innermost sphere of a hyper-domain. As such, updating with

[1]The types of ordering used to generate systems of spheres depend on the type of conditionals under evaluation: counterfactuals come with a standard similarity ordering whereas indicatives come with some kind of relevance ordering (cf. Willer, 2017).

the entertainability presupposition expands the innermost sphere to one that contains a ϕ-world. If the input modal center already contains a ϕ-world, then the update has no effect, or as we will say *idles*. On the other hand, if the whole hyper-domain does not contain any ϕ-world, then the update returns the empty set, thereby signaling presupposition failure.

Next, updating on this new hyper-domain with the asserted content of $\phi > \psi$, namely with the strict conditional $\Box(\phi \supset \psi)$, boils down to checking whether all the ϕ-worlds in the current modal center are ψ-worlds. If so, then the update idles; otherwise, the update returns the empty set, thereby indicating discourse anomaly.

To illustrate how antecedent strengthening fails under DSA, consider the hyper-domain depicted in Figure 2.2(a). Suppose the initial modal center is S_1. When $A > F$ is produced, given that S_1 does not contain any A-world, the modal center expands to S_2 so as to satisfy the entertainability presupposition $\Diamond A$. Since all the A-worlds in S_2 are F-worlds, updating with the asserted content $\Box(A \supset F)$ idles. Since S_2 does not contain any $(A \wedge B)$-world, the modal center needs to expand further upon the utterance of $(A \wedge B) > \neg F$, thereby establishing S_3 as the new center as shown in Figure 2.2(b). (The dashed lines represent spheres that have been eliminated.) Given that all the $(A \wedge B)$-worlds in S_3 are $\neg F$-worlds, updating with the asserted content $\Box((A \wedge B) \supset \neg F)$ idles. The Sobel sequence is thus predicted to be consistent, which means antecedent strengthening also fails under DSA.[2]

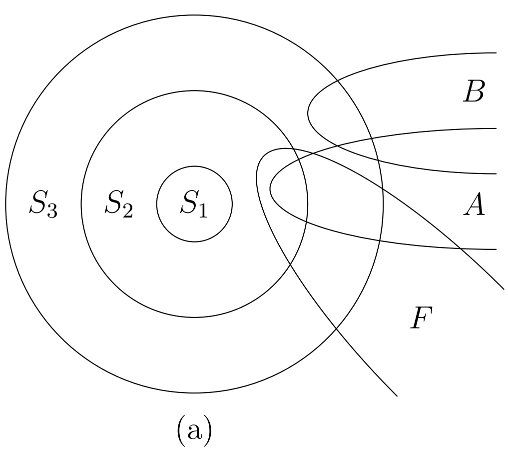

Figure 2.2: Domain expansion under DSA

[2]Additionally, DSA accounts for the infelicity of the so-called *reverse Sobel sequence* (von Fintel 2001):

#If Alice and Bob both come, the party won't be fun. But if Alice comes, the party will be fun.

2.3 The Extended Sobel Inference

2.3.1 The Problem of Unwanted Worlds

Recall that the extended Sobel inference, as exemplified by (5), is of the form:

$$\phi > \psi, (\phi \wedge \chi) > \neg\psi \vDash (\phi \wedge \neg\chi) > \psi$$

(5) (a) If Alice comes, the party will be fun. (b) But if Alice and Bob both come, the party won't be fun. ⇒ (c) If Alice but not Bob comes, the party will be fun.

The validity of ESI poses a serious challenge to DSA as it fails to vindicate this inference. Consider the hyper-domain depicted in Figure 2.3 below, which, as we have seen previously, is the output hyper-domain from the update with the first two conditionals in (5), i.e., $A > F$ and $(A \wedge B) > \neg F$.

The problem concerns the second update with (5b). When the modal center expands from S_2 to S_3 so as to satisfy the entertainability presupposition of (5b), some unwanted worlds are absorbed into the modal center. In particular, there are worlds in S_3, designated by the shaded area in Figure 2.3, that make A true but both B and F false. Given the presence of these worlds, updating with the asserted content of (5c), i.e., $\Box((A \wedge \neg B) \supset F)$,

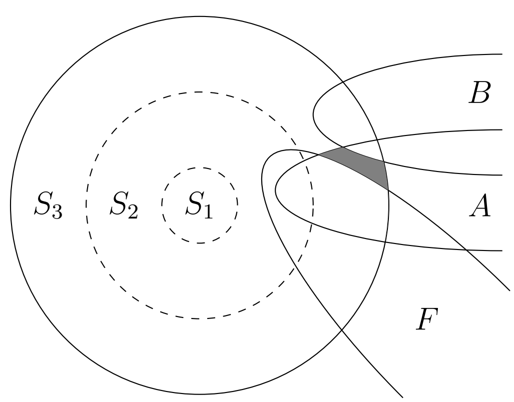

Figure 2.3: A countermodel to ESI under DSA

When $(A \wedge B) > \neg F$ becomes the first conditional, we instantly expand to S_3 to accommodate the entertainability presupposition $\Diamond(A \wedge B)$. But since S_3 contains some $(A \wedge \neg F)$-worlds, $A > F$ can longer be satisfied. Hence, reverse Sobel sequences are predicted to be inconsistent under DSA. For some other approaches to reverse Sobel sequences, see, e.g., Moss, 2012; Lewis, 2018; Ippolito, 2020.

will eliminate S_3, thereby yielding the empty set as its output. Given the existence of a countermodel, ESI is not valid under DSA.[3]

That being said, one may wonder whether there is an easy fix for the existing dynamic account. For instance, one could suggest that when the modal center expands to accommodate (5b)'s entertainability presupposition $\Diamond(A \wedge B)$, we do not simply expand from S_2 to S_3 but rather to $S_2 \cup (S_3 \cap [\![A \wedge B]\!])$. Perhaps, the thought is that the modal center should expand in a minimal fashion such that only those worlds that contribute to fulfilling the entertainability presupposition are added to the center. And since no new $(A \wedge \neg B)$-worlds are included during expansion, all the $(A \wedge \neg B)$-worlds in the modal domain are still F-worlds, thereby vindicating ESI.

This solution, however, immediately encounters another difficulty. If the modal center expands in the way just described, then we should expect (5a) and (5b) together to entail the conditional "if Bob comes to the party, then Alice will come" given that all the B-worlds in $S_2 \cup (S_3 \cap [\![A \wedge B]\!])$ are A-worlds. But this further inference is hardly intuitive. To avoid this, perhaps one could further suggest that we should also add some $(\neg A \wedge B)$-worlds in addition to the $(A \wedge B)$-worlds to the expanded modal center. But now it becomes rather unclear what the criterion is for deciding which worlds should get added. Why is it the case that only some $(\neg A \wedge B)$-worlds but no $(A \wedge \neg B)$-worlds are added? There does not seem to be a straightforward answer where one can simply read off the logical form of the conditional to decide how the modal center expands.

Another rescue strategy is to attempt at capturing the ESI in (5) via some pragmatic means, for example, as a scalar implicature. Since a conjunction and its conjuncts form a scale (cf. Sauerland, 2004), the weaker alternative "Alice comes to the party" can be exhaustified to convey that the stronger alternative "Alice and Bob both come" is false. In other words, "Alice comes to the party" can be strengthened to mean "Alice but not Bob comes to the party". Now, if we postulate that exhaustification can occur locally (e.g.,

[3] This problem of unwanted worlds is related to the *problem of intermediate worlds* pointed out by Nichols (2017). Nichols's criticism is based on the observation that DSA has difficulty rendering two counterfactuals of the form $(A \wedge B) > \neg F$ and $(A \wedge \neg B) > F$ simultaneously true (or assertible) given a particular setup, that is, when "more departure from actuality is required to make the antecedent of an earlier counterfactual true than to make a subsequent counterfactual's antecedent true but its consequent false" (p. 627). By contrast, my current focus is on DSA's inability to vindicate a desirable general inference with conditionals.

Chierchia, 2012)—more specifically, within the antecedent of a conditional—then we can view the inference to "If Alice but not Bob comes, the party will be fun" as a case scalar inference.

The problem with this strategy is that it fails to generalize to ESIs that invoke two alternatives where the weaker one cannot be exhaustified to mean the denial of the stronger alternative. For example, consider (6):

(6) (a) If Alice is from France, the party will be fun. (b) But if she is from Paris, the party won't be fun. $\Rightarrow$ (c) If Alice is from France but not Paris, the party will be fun.

Although the inference still holds, the weaker alternative "Alice is from France" cannot be exhaustified to mean "Alice is from France but not Paris". This is attested by the Hurford disjunction in (7a):

(7) (a) #Either Alice is from France, or she is from Paris.
(b) Either Alice comes to the party, or Alice and Bob both come to the party.

Compare (7a) to (7b). Disjunctions where one disjunct contextually entails the other are generally infelicitous (Hurford, 1974). But when the weaker disjunct can be exhaustified so that it no longer entails the stronger alternative, the disjunction becomes felicity (Chierchia, 2009). This is the case for (7b), but not for (7a). Hence, resorting to scalar implicature alone cannot explain why ESI holds in full generality.

Here is a brief diagnosis of the problem of unwanted worlds encountered by the existing dynamic account. DSA models incorporation of new possibilities via expansion of the modal center. But given a lack of constraint on expansion, when the center expands, it may absorb unwanted worlds, and as a result, too much old information becomes lost. One the one hand, allowing for discarding some old necessity enables DSA to capture failure of antecedent strengthening. As in the case of (5), discarding the old necessity $A \supset F$ (viz., the necessity that all A-worlds are F-worlds) enacted by (5a) leaves room for the subsequent utterance of (5b) to make $(A \wedge B) \supset \neg F$ a new necessity. On the other hand, the validity of ESI shows that when the center expands, it is hardly the case that all past information is lost. The old necessity $A \supset F$ should somewhat persist even in light of the newly entertained possibility. When the antecedent of (5b) introduces the new possibility $A \wedge B$, although the previously established necessity $A \supset F$ no longer holds in all the $(A \wedge B)$-worlds, it should nevertheless still persist in all the $\neg(A \wedge B)$-worlds. And to say that $A \supset F$ holds in all the $\neg(A \wedge B)$-worlds, given basic propositional reasoning, is to say that all the $(A \wedge \neg B)$-worlds are F-worlds,[4] thereby licensing the inference to (5c).

Hence, to vindicate ESI, we need a principle governing information preservation in light of new possibilities. This principle, as we now recall, is weak persistence:

> **Weak Persistence:** When a hitherto unentertained possibility ϕ is introduced, past information is preserved either in all the ϕ-worlds or in all the $\neg\phi$-worlds.

To illustrate, consider (5) again, repeated below:

(5) (a) If Alice comes to the party, it will be fun. (b) But if Alice and Bob both come, it will not be fun. $\Rightarrow$ (c) If Alice but not Bob comes, the party will be fun.

Suppose we are speculating whether this weekend's party will be fun. We know that whenever Alice is at the party, the party is almost always fun. Hence, we establish (5a). So far, the possibility that Alice and Bob both come to the party remains not salient to us, since, say, Bob travels a lot. But as we remind ourselves that Bob is in town this week and become aware that he might also come, we no longer have to *fully* endorse (5a) given that we can establish (5b) as a new necessity. However, this is not to say that the past information, namely, the necessity established by (5a), is completely lost. By weak persistence, when the possibility that Alice and Bob both come becomes salient, our previous conclusion that the party will be fun if Alice comes should still hold either in all the worlds where Alice and Bob both come or in all the worlds where that is not the case. But since (5b) eliminates the first option, it must be that in all the $\neg(A \wedge B)$-worlds, the party will be fun if Alice comes. Hence, we can draw the conclusion that if Alice but not Bob comes, the party will be fun.

2.3.2 The Problem of Unwarranted Inference

The variably strict analysis does not face the problem of unwanted worlds which plagues DSA, as the former only quantifies over the closest worlds where the antecedent is true.

[4] To see this, note that $\neg(A \wedge B) \supset (A \supset F)$ is equivalent to $(A \wedge \neg B) \supset F$.

Consider Figure 2.3 again. Since the closest $(A \wedge \neg B)$-worlds are all F-worlds, ESI holds under VSA. Crucially, what renders ESI valid under VSA is the fact that the truth of the two conditionals in a Sobel sequence forces a particular ordering: any system of spheres that makes $A > F$ and $(A \wedge B) > \neg F$ true must force the closest A-and-B-worlds to be closer than the closest B-worlds. Consequently, all the closest A-worlds must be $\neg B$-worlds. And since all the closest A-worlds are also F-worlds, VSA predicts the entailment to $(A \wedge \neg B) > F$. However, the particular ordering invoked by VSA also happens to vindicate the following additional inference:

$$\phi > \psi, (\phi \wedge \chi) > \neg\psi \vDash \phi > \neg\chi$$

This additional inference is good for some counterfactuals, though not for all (see Klecha 2015, 2021). Moreover, it is clearly too strong for indicative conditionals. As I have mentioned in §2.1, the inference from (4a) and (4b) to (4c) seems quite bad.

(4) (a) If Alice comes, the party will be fun. (b) But if Alice and Bob both come, the party won't be fun. $\not\Rightarrow$ (c) If Alice comes, Bob won't come.

It has been suggested by Klecha (2015, 2021) that there are two types of Sobel sequences: those that do license the inference to $\phi > \neg\chi$, which Klecha deems as genuine Sobel sequences, and those that do not, which he names *Lewis sequences*. For Klecha, all indicative versions of Sobel sequences are indeed Lewis sequences as they do not license the further entailment to $\phi > \neg\chi$.[5]

The current observation is that the validity of ESI does not hinge on whether the Sobel sequences under discussion are genuine Sobel sequences or are in fact Lewis sequences; ESI is valid regardless. If so, then a unified explanation of the validity of ESI should not rest on the particular ordering that is induced by the first two conditionals as in a genuine Sobel sequence.[6] The main aim of this paper is thus to explore such a unified explanation.

[5] As a side note, Klecha (2015, 2018) views Lewis sequences as instances of a more general phenomenon of precisification. It remains to be seen whether his account can adequately capture the validity of ESI.

[6] This is not to say that I wish to completely dispense with any ordering, since as we shall see in §2.5.4, ordering can help in supplying a reservoir context from which previously eliminated worlds can be selected to be revived. My contention is that an adequate analysis that captures ESI should not merely appeal to some particular ordering alone.

2.3.3 Interim Conclusion

To take stock, the standard dynamic strict account fails to validate ESI because modal expansion can potentially assimilate unwanted worlds; the variably strict analysis, albeit validating ESI, also vindicates the additional inference to $\phi > \neg\chi$, which intuitively does not hold under all circumstances. In response, I aim to validate ESI by developing a more sophisticated update system that can capture bringing a hitherto unentertained possibility to salience without explicitly relying on any particular ordering. This update framework enables us to formally cash out weak persistence, which in turn allows us to capture the validity of ESI.

A brief note on presentation of examples in what follows: since my goal is to validate ESI in cases where the additional inference to $\phi > \neg\chi$ cannot be drawn, I will focus my discussion on the indicative versions of ESI. I leave it open how to extend this analysis to counterfactuals, which requires a detailed discussion of how similarity ordering can be incorporated into my update framework.

2.4 Weak Persistence: An Informal Sketch

Before delving into details of my framework, I informally sketch how weak persistence will be realized. To capture weak persistence, the formal apparatus needs to do two things:

Requirement A: It should afford a principled way to distinguish *live possibilities* that have become salient from *plain possibilities* that have yet to be actively entertained;

Requirement B: It should be able to encode a set of alternatives, each of which provides a complete description of the space of possibilities at the current stage of discourse. Such a possibility space contains information about what has been settled, what possibilities have become live, and how live possibilities are related.

The first requirement is needed for modeling the process of bringing a hitherto unentertained possibility to salience, while the second is needed for codifying alternative ways to preserve past information in light of new possibilities.

A formal framework that fulfills both requirements can then capture weak persistence. Consider the party example again. Let Figure 2.4(a) represent the total body of information after the utterance of the first conditional "if Alice comes to the party, it will be fun". The outermost rectangle with rounded corners represents the total logical space. The inner rectangle, which is partitioned into some A-worlds and some $\neg A$-worlds, represents the current space of possibilities. The dotted region represents the set of A-worlds that are also F-worlds.[7] After the utterance of the first conditional $A > F$, it becomes common ground that all worlds where Alice comes are worlds where the party is fun. Hence, all the A-worlds in 4(a) are in the dotted region.

As the possibility of Alice and Bob both coming to the party has yet to become salient in 4(a), accommodating the entertainability presupposition of the conditional "if Alice and Bob both come, the party won't be fun" will resurrect some $(A \wedge \neg F)$-worlds. This is reflected in Figure 2.4 by the fact that in both 4(b) and 4(c), the total possibility space is extended with an additional rectangle which comprises A-worlds but is nonetheless not contained in the dotted region. Also notice that by making a new partition, both 4(b) and 4(c) are now

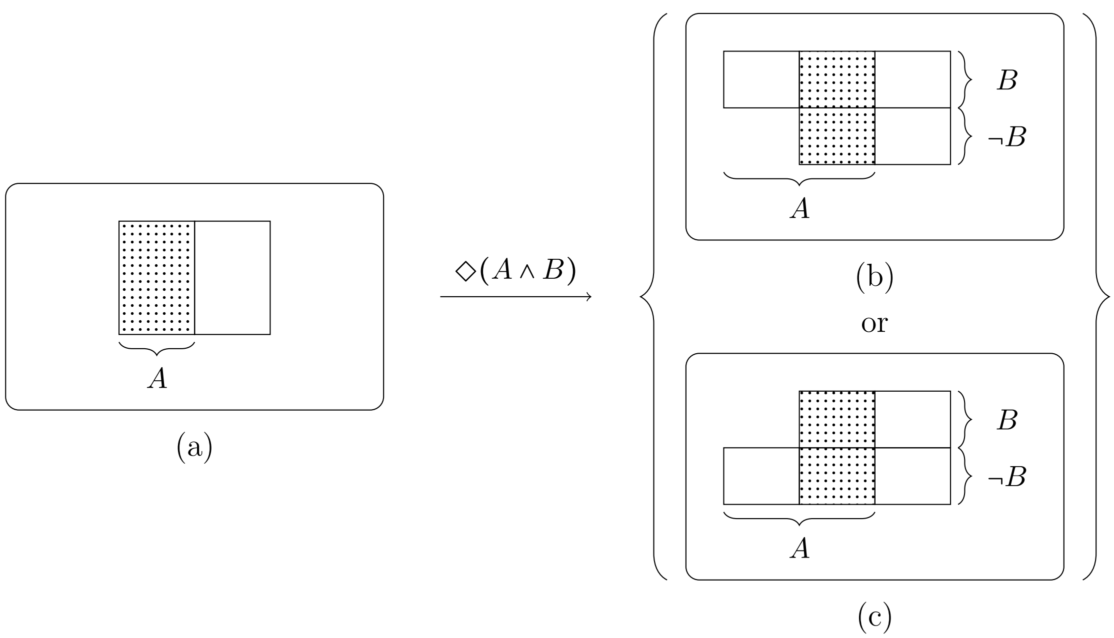

Figure 2.4: A sketch of weak persistence

[7] To clarify, the dotted region does not represent the set of all F-worlds.

actively distinguishing between B and $\neg B$. The difference between 4(b) and 4(c) concerns whether the resurrected worlds are B-worlds or $\neg B$-worlds. In 4(b), entertaining the new possibility that Alice and Bob both come brings back some worlds where Alice comes but the party is not fun, all of which are worlds where Bob also comes. By contrast, in 4(c), entertaining the same antecedent again brings back some worlds where Alice comes but the party is not fun, but this time, all of them are worlds where Bob does not come. As such, we are presented with two alternative ways to preserve the information that all A-worlds are F-worlds: the information is preserved either in all the B-worlds or in all the $\neg B$-worlds. This is how weak persistence can be realized.

Given weak persistence, here is how we can vindicate ESI without also vindicating the additional inference from VSA. As it turns out, 4(c) is indeed incompatible with the asserted content of the second conditional "if Alice and Bob both come, the party won't be fun" because all the worlds where Alice and Bob both come in 4(c) are in the dotted region and are thus worlds where the party is fun. Consequently, this alternative is discarded, which leaves us with 4(b). Given that in 4(b), all the worlds where Alice comes to the party but Bob does not are worlds where the party is fun, we successfully predict the entailment to $(A \wedge \neg B) > F$. On the other hand, since 4(b) still contains some worlds where Alice and Bob both come to the party, we do not get the additional entailment to $A > \neg B$.

Now, both Requirement A and Requirement B can be fulfilled in a framework that enables representation of alternatives. The connection to the second requirement is obvious. As for the first requirement, we can avail ourselves of alternatives to differentiate between live and plain possibilities in the following way. We can construe ϕ as a live possibility with respect to a given body of information just in case the body of information makes an active distinction between two alternatives: namely, ϕ and $\neg\phi$.

One such framework is inquisitive semantics (Ciardelli & Roelofsen, 2011; Ciardelli et al., 2015, 2018). In inquisitive semantics, disjunction functions to supply a set of alternatives. While inquisitive semantics was originally designed to provide a uniform analysis for both declarative and interrogative sentences rather than modeling live and plain possibilities, we can readapt the framework for our current purposes. More specifically, given that the Law of Excluded Middle ($\phi \vee \neg\phi$) is not valid under inquisitive semantics, we can exploit this feature to distinguish a body of information where ϕ has become salient from one where ϕ has yet to be entertained.

Be that as it may, we still need to make some significant modifications to the existing inquisitive framework for present purposes. In particular, since alternatives will be used for the dual purpose of representing live possibilities and providing alternative ways to preserve past information, we need to employ alternatives at two different levels without conflating them. This will then require us to move one level up on the set-theoretic hierarchy.

2.5 Weak Persistence Semantics

Here is a quick breakdown of this section. §2.5.1 introduces some basic notions from standard inquisitive semantics. Each of the subsequent subsections motivates and introduces one piece of the machinery, which eventually leads us to a dynamic inquisitive framework capable of capturing weak persistence. First of all, §2.5.2 introduces dynamic modals alongside a notion of refinement. Together, they allow us to model the dynamics of bringing a plain possibility alive. In §2.5.3, I ascend one level up on the set-theoretic hierarchy and introduce hyper-contexts. Finally, §2.5.4 defines a refinement operation that captures weak persistence at the level of hyper-contexts. Throughout §2.5, my presentation will be semi-formal and rely heavily on diagrams. The formal details are left to the appendix. In §2.6, I present a new dynamic strict analysis of conditionals within this framework and elucidate how ESI is validated.

2.5.1 Basic Inquisitive Semantics

We begin by introducing some basic notions from inquisitive semantics particularly pertinent to us. First, we have the familiar notion of *information states* (e.g., Veltman, 1996). An information state s is a set of possible worlds compatible with a certain body of information. In standard possible world semantics, formulas are evaluated at worlds in terms of their truth and falsity. In inquisitive semantics, formulas are evaluated at information states in terms of *support*.

An information state s supports an atomic formula p iff p is settled true in s, that is,

p is true at every world in s. As an example, suppose our logical space W contains four worlds: one where A and B are both true, one where A is true but B is false, one where B is true but A is false, and one where neither is true. Among the four information states depicted in Figure 2.5, only 5(a) supports A as well as B, since both formulas are true at every world in the state $\{AB\}$; by contrast, 5(b) supports A but does not support B whereas 5(c) and 5(d) support neither A nor B. Additionally, the empty set which represents the absurd information state is taken to support everything.

A state s supports a negation $\neg\phi$ iff no subset of s except for the empty set supports ϕ. For example, 5(b) supports $\neg A$, but it does not supports $\neg B$ since it contains a non-empty subset that supports B, namely $\{AB\}$.

A state supports a conjunction iff it supports both of its conjuncts, and a disjunction iff it supports either of its disjuncts. For example, 5(b) supports $A \vee \neg A$ by supporting A, but it does not support $B \vee \neg B$ since it supports neither B nor $\neg B$ individually. As for implication, a state s supports $A \to B$ iff every subset of s that supports A also supports B; hence, only 5(a) and 5(c) support $A \to B$ since only for these two states do all of their subsets that support A also support B.

Next, let us define a technical notion of *contexts*. Standardly, contexts are often construed as sets of worlds (Stalnaker, 1978). However, in order to separate a context where A is a live possibility from one where A is a mere plain possibility, additional structure is needed. In inquisitive semantics, a context C is defined as a non-empty *downward closed* set of information states—that is, for any information state s that belongs to C, any subsets of s must also belong to C. We call a maximal element in C (i.e., an information state that is not

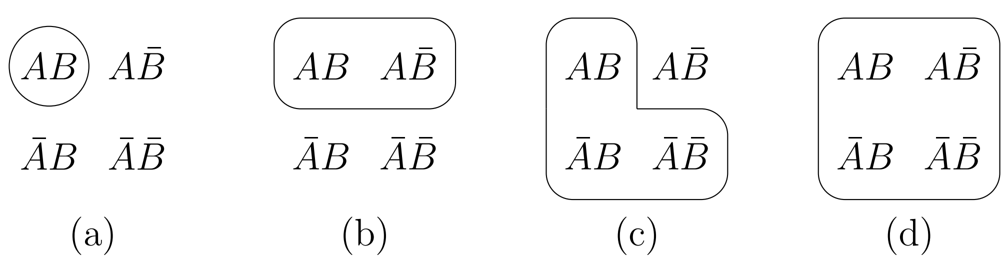

Figure 2.5: Some examples of information states

a proper subset of any other states in C) an *alternative* in C. We can then use alternatives to represent salient possibilities in a context.

Consider the four contexts depicted in Figure 2.6. For each context, only its alternatives are explicitly drawn: for example, given that contexts are downward closed, 6(c) represents the context $\{\{AB\}, \{A\bar{B}\}, \varnothing\}$, and 6(b) represents the context $\{\{AB, A\bar{B}\}, \{AB\}, \{A\bar{B}\}, \varnothing\}$. We say that a context is *inquisitive* iff it contains more than one alternative; 6(a) and 6(c) are inquisitive whereas the other two are not. In standard inquisitive semantics, an inquisitive context raises an issue by putting forth alternative ways for the issue to be resolved. Under the current setting, an inquisitive context indicates awareness of the distinctions among different alternative possibilities. Henceforth I will just describe the context itself as being aware of these possibilities. Compare 6(b) and 6(c): although both contexts are made up of the same two worlds, 6(c) is actively aware of the distinction between B and $\neg B$ whereas 6(b) is oblivious to this distinction.

We highlight two special contexts: $C_\top$ and $C_\perp$. For any given logical space W, the initial context $C_\top$ which represents a total lack of information and awareness is given by the power set $\wp(W)$. The absurd context $C_\perp$ which serves to signal discourse anomaly is identified with $\{\varnothing\}$.

Contexts are connected via updates. I write $C[\phi]_u$ to denote the context resulting from updating C with the formula ϕ.[8] Updating C with ϕ amounts to collecting only those information states in C that support ϕ, or equivalently, eliminating all states in C that do not support ϕ. For example, updating the above context 6(b) with $[B \vee \neg B]_u$ eliminates the

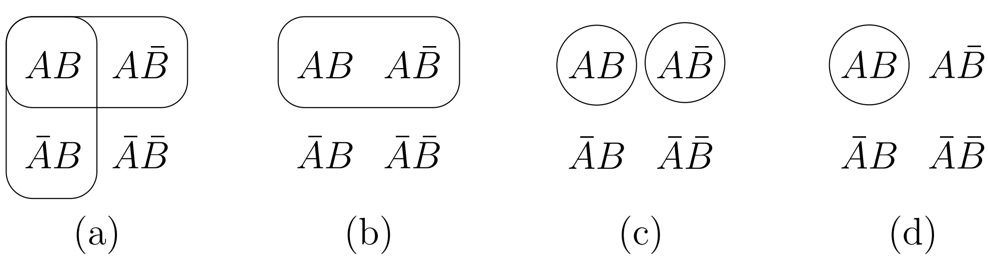

Figure 2.6: Some examples of contexts

[8]The subscript u is used to differentiate the present type of proper-updates from the type of refinement-updates to be introduced in §2.5.2.

only state in it that does not support $B \vee \neg B$, namely $\{AB, A\bar{B}\}$, thereby outputting 6(c) wherein the distinction between B and $\neg B$ now becomes salient.

2.5.2 Introduce Dynamic Modals and Refinement

We introduce dynamic modals and a notion of refinement next. Together, they enable us to model the dynamics of transforming a plain possibility into a live possibility. In accordance with the existing dynamic approach to modals (cf. Veltman 1996, 2005), we construe possibility and necessity modals (i.e., $\Diamond$ and $\Box$) as tests that examine whether certain global conditions are satisfied by the current body of information. We define the support condition of a modal formula by an information state s relative to a context to which s belongs. Given a context C, a state s in it supports $\Diamond\phi$ iff there exists an alternative in C that supports ϕ; analogously, s supports $\Box\phi$ iff every alternative in C supports ϕ. Since alternatives serve to represent live possibilities, what the above definitions amount to is simply this: s supports $\Diamond\phi$ iff ϕ is a live possibility in C; s supports $\Box\phi$ iff every live possibility in C is a ϕ-possibility.

To illustrate, let s be the state $\{AB\}$. Consider whether s supports $\Diamond B$ with respect to the following two contexts from Figure 2.6: 6(b), i.e., $\{\{AB, A\bar{B}\}\}$, and 6(c), i.e., $\{\{AB\}, \{A\bar{B}\}\}$.[9] Given that there exists an alternative in (6c), namely $\{AB\}$, such that it supports B, s supports $\Diamond B$ with respect to 6(c); given that the only alternative in 6(b), namely $\{AB, A\bar{B}\}$, does not support B, s does not support $\Diamond B$ with respect to 6(c). On the other hand, s supports $\Box A$ with respect to both 6(b) and 6(c) since every alternative in them supports A.

To elucidate why we need a new type of update operation which I call refinement, let us first try to define a notion of support at the level of contexts. As our first pass, we say that a context C supports ϕ iff $C[\phi]_u = C$, or equivalently, iff every state s in C supports ϕ. The problem with this definition is that it fails to satisfy the following intuitively desirable principle governing awareness.

Principle of Uncertainty: For any context C and formula ϕ, if C supports $\Diamond\phi$ but

[9] Here and hereafter, I will abbreviate a context using its set of alternatives.

does not support ϕ, then it must support $\Diamond\neg\phi$.

To put it in plain language, if we deem ϕ possible but do not have enough information to establish ϕ, then we must also deem $\neg\phi$ possible. However, the principle of uncertainty is violated given the present setup. Let C be the following context: $\{\{AB\}, \{A\bar{B}, \bar{A}B, \bar{A}\bar{B}\}\}$. It is easy to verify that C supports $\Diamond A$ but not A, yet it fails to support $\Diamond\neg A$.

To modify the framework so as to satisfy the principle of uncertainty, I introduce a notion of general updates $[\phi]$ and conceive them as proceeding in two separate steps: $C[\phi] := C[\phi]_r[\phi]_u$. Updating C with $[\phi]$ first *refines* the input context with $[\phi]_r$, which boils down to updating C successively with $[p \vee \neg p]_u$ for every atomic proposition p in ϕ. In inquisitive semantics, formulas of the form $\ulcorner p \vee \neg p \urcorner$, often abbreviated as $\ulcorner ?p \urcorner$, are understood as asking the question about whether or not p. Analogously, under the current setting, refining C with p amounts to entertaining whether or not p, thereby making a distinction between p and $\neg p$. Refining C with $[\phi]_r$ is then tantamount to bringing every atomic proposition in ϕ to salience in the post-refinement context.[10] This post-refinement context is in turn updated with $[\phi]_u$ as usual. As a consequence, we now define context-level support in terms of general updates: C supports ϕ iff $C[\phi] = C$.

To illustrate, consider updating the first context in Figure 2.7 with $\neg(A \wedge B)$. The general update with $[\neg(A \wedge B)]$ is divided into two steps: first, the refinement on 7(a) with $[\neg(A \wedge B)]_r$, and then the update on the post-refinement context with $[\neg(A \wedge B)]_u$. The refinement with $[\neg(A \wedge B)]_r$ equates to the sequential refinement with $[A]_r$ and then $[B]_r$, each of which is then reduced to updating with a disjunction of the form $[p \vee \neg p]_u$. Updating

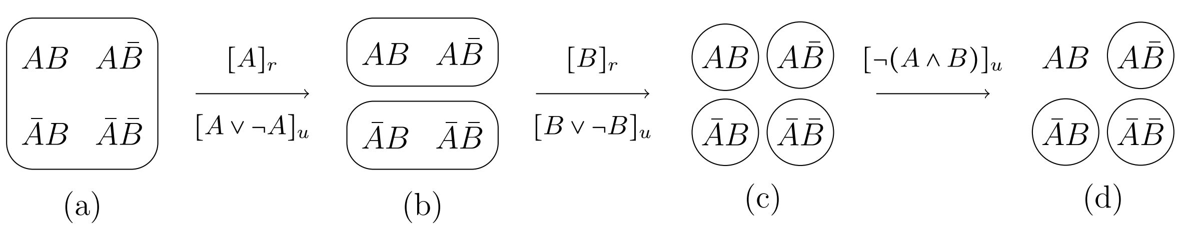

Figure 2.7: An example of the general update procedure

[10]The version of refinement proposed here provides one rather simple way to satisfy the principle of uncertainty, and as such it may need further fine-tuning. I shall leave exploring a more sophisticated definition for refinement to another occasion.

the post-refinement context 7(c) with $[\neg(A \wedge B)]_u$ eliminates the world AB and returns 7(d).

As an upshot of this modification, not all non-empty downward closed sets of information states will be considered as proper contexts. For example, there is no update procedure that can derive the purported context $\{\{AB\}, \{A\bar{B}, \bar{A}B, \bar{A}\bar{B}\}\}$, which, as we have seen, violates the principle of uncertainty. To amend this, we revise the definition of contexts as follows:

> **Contexts:** Given a logical space W, a context C is a non-empty downward closed set of information states that can be derived from performing certain general updates on the initial context $C_\top = \wp(W)$.

At the level of contexts, we can now distinguish live possibilities from plain possibilities: ϕ is a live possibility in C if $C[\Diamond\phi] = C$;[11] ϕ is a plain possibility if $C[\Diamond\phi] \neq C$ and $C[\Diamond\phi] \neq C_\perp$. Consider once more these two contexts from Figure 2.6: 6(b), i.e., $\{\{AB, A\bar{B}\}\}$, and 6(c), i.e., $\{\{AB\}, \{A\bar{B}\}\}$. While A is a live possibility (and indeed a necessity) in both of them given that the update with $[\Diamond A]$ idles for both contexts, B is only a live possibility in 6(c). Updating 6(c) with $[\Diamond B]$ idles; by contrast, updating 6(b) with $[\Diamond B]$ returns 6(c). Hence, B is merely a plain possibility in 6(b).

The current framework's ability to represent live possibilities further addresses a puzzle typically associated with the traditional dynamic approach to modals. As noted by Yalcin (2007), if updating with $\Diamond A$ only returns either the input state or the absurd state, then an utterance of $\Diamond A$ can never be truly informative, since either it does not provide any new information as the update produces no effect, or the information carried by the sentence cannot be consistently integrated. However, on many occasions, a *might*-claim is *prima facie* non-trivial and appears to communicate new information, as the following examples demonstrate (Yalcin, 2007, p. 1012):

(6) Cheerios may reduce the risk of heart disease.

(7) Late Antarctic spring might be caused by ozone depletion.

Whereas Yalcin proffers a pragmatic analysis for the seeming informativeness of these sentences, the current framework is able to supply a semantic story (see also Willer, 2013),

[11] We could further impose the condition that $C[\Diamond\phi] \neq C_\perp$ if we want to avoid considering every sentence a live possibility in the absurd context.

that is, sentences with an epistemic possibility modal like (6) and (7) can convey non-trivial information by raising some hitherto overlooked possibility to salience and thereby refining the common ground. Before the utterance of (6), the context is oblivious to the possibility that Cheerios may reduce the risk of heart disease. The possibility is a plain possibility as the common ground does not contain an alternative that supports "Cheerios reduce the risk of heart disease", nor an alternative that supports its negation. After the utterance, the possibility at issue becomes live as the updated context now contains both alternatives.

The refinement operation considered thus far is strongly persistent in the sense that it can only raise a possibility to salience but can never bring back any previously eliminated worlds. To capture failure of antecedent strengthening and weak persistence, our next step is to allow refinement to revive "dead" worlds. To do so, we will need to move to hyper-contexts.

2.5.3 The Lift to Hyper-Contexts

Recall our preliminary definition of weak persistence, which states that when a hitherto unentertained possibility ϕ is introduced, past information is preserved either in all ϕ-possibilities or in all $\neg\phi$-possibilities. To preserve past information, under the current setting, is for the refinement operation *not* to reintroduce worlds that have been previously eliminated. Suppose ψ has already been established; then while entertaining ϕ can resurrect some $\neg\psi$-worlds, either all of them need to be ϕ-worlds, or they all need to be $\neg\phi$-worlds. Hence, entertaining a new possibility should create two alternatives, each of which embodies information about how live possibilities are related—that is, either all ϕ-possibilities are still ψ-possibilities, or all $\neg\phi$-possibilities are still ψ-possibilities.

At the level of contexts, however, we cannot adequately model relations between live possibilities, since alternatives in contexts are already used to encode information about what possibilities have become live. Say we wish to model a case where either all B-possibilities are A-possibilities or all $\neg B$-possibilities are A-possibilities but not both. However, no contexts can encode such information. For instance, the context $\{\{AB\}, \{A\bar{B}\}, \{\bar{A}\bar{B}\}\}$ is such that all the B-possibilities in it are A-possibilities, and the context $\{\{AB\}, \{A\bar{B}\}, \{\bar{A}B\}\}$ is such that all the $\neg B$-possibilities in it are A-possibilities; but if we simply take the union of the two sets, the resulting context, i.e., $\{\{AB\}, \{A\bar{B}\}, \{\bar{A}B\}, \{\bar{A}\bar{B}\}\}$, is such that neither all the

B-possibilities nor all the $\neg B$-possibilities are A-possibilities. What we want instead is a set where $\{\{AB\}, \{A\bar{B}\}, \{\bar{A}\bar{B}\}\}$ and $\{\{AB\}, \{A\bar{B}\}, \{\bar{A}B\}\}$ are its two alternatives. Therefore, to enable live possibilities to be related in multiple ways, we need to move one level up on the set-theoretic hierarchy from contexts to what I call *hyper-contexts*.

A hyper-context Σ is a non-empty *restrictively* downward closed set of contexts, that is, for any context C in Σ, if x is a subset of C and x *is itself also a context*, then x is also in Σ. Furthermore, as with (proper) contexts, not all non-empty restrictively downward closed sets of contexts will be regarded as (proper) hyper-contexts because the update procedure, to be defined below, will not always produce every such set. Hence, we define hyper-contexts as follows, analogous to the revised definition of contexts:

> **Hyper-contexts** : A hyper-context Σ is a non-empty restrictively downward closed set of contexts that can be derived from performing certain updates on the initial hyper-context $\Sigma_\top$.[12]

We define alternatives in a hyper-context as maximal elements in it. While an alternative in a context represents a single salient possibility, an alternative in a hyper-context represents a set of salient possibilities, viz., a possibility space. Hence, a hyper-context that contains more than one alternative indicates that there is more than one way for live possibilities to be related.

As with contexts, general updates on hyper contexts proceed in two steps: $\Sigma[\phi] = \Sigma[\phi]_r[\phi]_u$. The (strongly persistent) refinement operation is the same as before: refining with $[\phi]_r$ amounts to refining with every atomic formula p in ϕ, which has the effect of bringing every p to salience. As for proper updates, updating Σ with $[\phi]_u$ amounts to updating each individual context C in Σ with $[\phi]_u$ and then taking the restrictive downward closure of the resulting set. We say a hyper-context Σ supports ϕ iff $\Sigma[\phi] = \Sigma$.

To adduce two concrete examples, first consider the update with $[\Diamond A]$ on Σ_1 as shown in Figure 2.8.[13] This update consists of first refining Σ_1 with $[\Diamond A]_r$—which is reduced to

[12] The initial hyper-context $\Sigma_\top$ is defined as the restrictive downward closure of the initial context $C_\top$ with respect to a given logical space W (see Appendix, Definition 13).

[13] The outermost circle in the diagram represents the hyper-context; the middle circles/ellipses represent alternatives in the hyper-context; and the innermost circles represent alternatives in their respective contexts. For example, Σ_1 contains C_1 as its sole alternative, which in turn contains $\{A\}$ and $\{\bar{A}\}$ as its two alternatives.

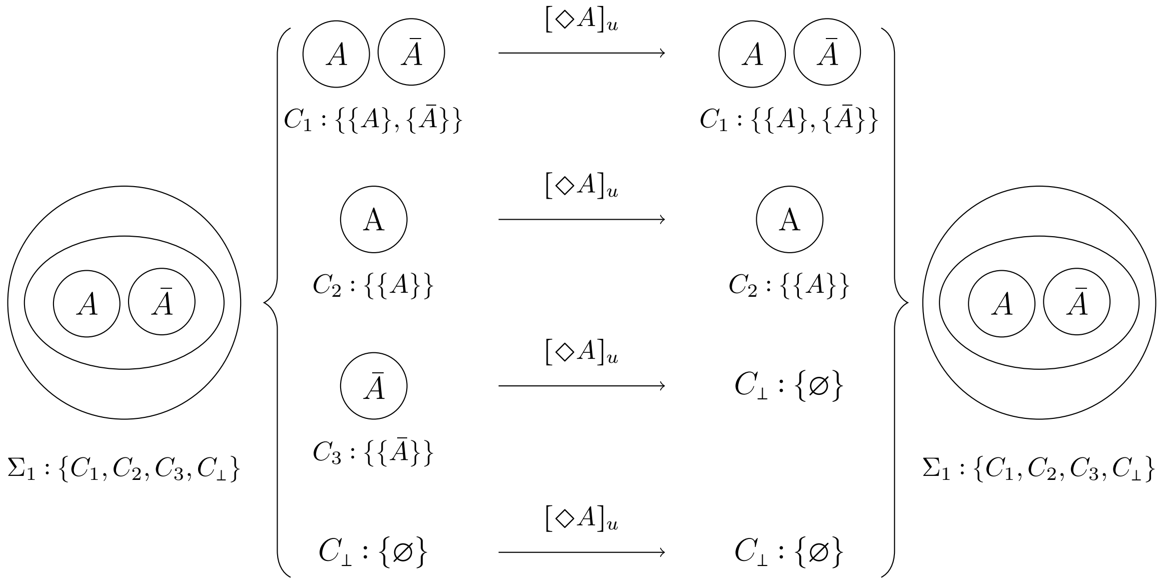

Figure 2.8: Updating Σ_1 with $[\Diamond A]_u$

updating it with $[?A]_u$—and then updating the post-refinement hyper-context with $[\Diamond A]_u$. Since Σ_1 already distinguishes between A and $\neg A$, the refinement idles. Figure 2.8 then details what the subsequent proper update looks like. Since only C_1 and C_2 contains A as a live possibility, the updates on C_1 and C_2 idle whereas the updates on the other two contexts return the absurd context. Collective the resulting contexts $\{C_1, C_2, C_\perp\}$ and then taking the restrictive downward closure give back Σ_1. Therefore, since $\Sigma_1[\Diamond A] = \Sigma_1$, Σ_1 already supports $\Diamond A$.

Next, consider the update with $[\Box A]$ on the same hyper-context Σ_1. As before, the refinement with $[\Box A]_r$ idles. Figure 2.9 depicts the subsequent update with $[\Box A]_u$. Among the four contexts contained in Σ_1, only the update on C_2 returns a non-absurd context. Collecting all the output contexts returns the posterior hyper-context Σ_3.[14]

[14]Note that we do not output the absurd hyper-context $\Sigma_\perp = \{\{\varnothing\}\}$ in this case. As such, the current framework manages to address one irregularity commonly associated with the dynamic approach to necessity modals. Since updating with $[\Box A]$ on any contexts wherein some but not all information states support A always returns the absurd state, it means that, whenever we are uncertain about the truth of some proposition A, uttering something like "*must* A" would be regarded as inconsistent. But this can hardly be the case. Suppose we are uncertain about whether Alice is at the party. In this scenario, we should not expect "Alice must be at the party" to be inconsistent with the common ground information such that the update winds up outputting the absurd state. On my account, we can thus accommodate this intuition at the level of hyper-contexts.

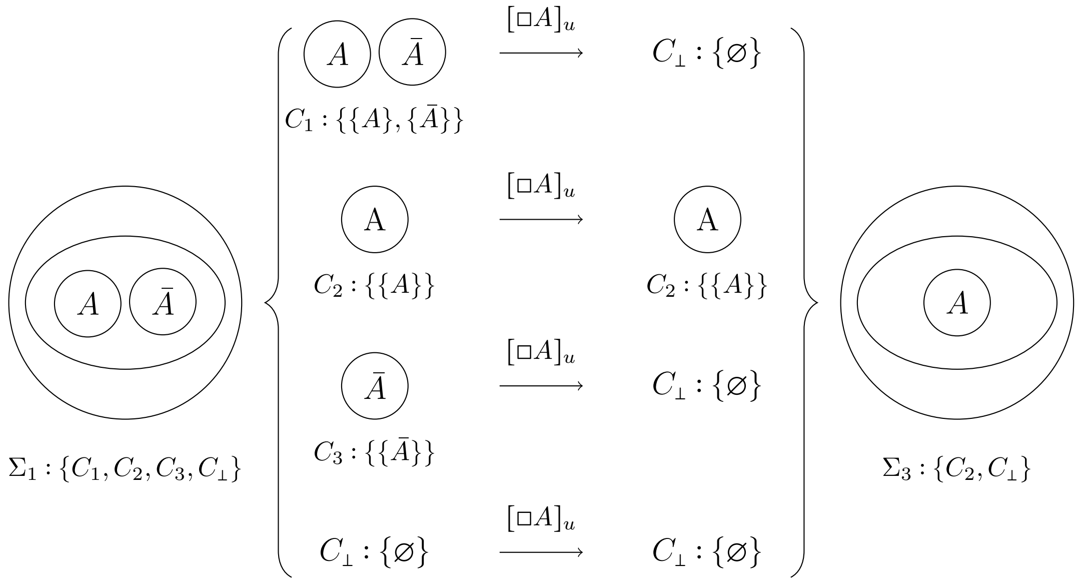

Figure 2.9: Updating Σ_1 with $[\Box A]_u$

2.5.4 Weakly Persistent Refinement

At the level of hyper-contexts, we can now cash out weak persistence induced by the conditional's antecedent in terms of weakly persistent refinement. The refinement associated with the entertainability presupposition of a conditional $[\Diamond\phi]_r$ becomes weakly persistent when the input hyper-context Σ does not actively distinguishes ϕ from $\neg\phi$, that is, when Σ does not support $?\phi$. When this happens, certain previously eliminated possibilities can be reintroduced from a reservoir context C_R, which I shall return to shortly; meanwhile, for any alternative C in Σ—which, as we recall, represents a particular way of how live possibilities are related in Σ—if C supports ψ, then refining with $[\Diamond\phi]_r$ will split C into two separate possibility spaces: one where all the ϕ-possibilities are still ψ-possibilities, thereby supporting $\phi \to \psi$, and the other where all the $\neg\phi$-possibilities are still ψ-possibilities, thereby supporting $\neg\phi \to \psi$ (see Appendix, Definition 19 for details).

The reservoir context C_R is a contextually supplied set that may contain worlds that have already been removed from Σ. For the time being, I remain rather agnostic about how C_R is determined. But what C_R essentially does is to supply a reservoir of worlds to be reintroduced upon entertaining a new possibility. The intuitive appeal behind postulating such a fallback position is that raising a new possibility should only allow reviving "dead" worlds that are in some sense *relevant* to the new possibility being entertained. In the party example, for instance, while entertaining the possibility of Bob's coming to the party allows us to bring back some worlds wherein the party is not fun, it should not bring back worlds wherein the basic laws of physics are different.[15]

Now, to illustrate this dynamics of raising a possibility to salience, consider the refinement depicted in Figure 2.10.[16] Suppose our input hyper-context is Σ_1 where it is established that Alice will come to the party but the possibility of Bob's coming to the party has yet to be entertained. Since Σ_1 does not support $?B$, the refinement with $[\Diamond B]_r$ will be weakly persistent. That is, upon accommodating the entertainability presupposition of "if Bob comes to the party", the sole alternative in Σ_1, namely C_1, will split into two contexts C_2 and C_3, each of which not only becomes aware of B but also brings back a world that was previously eliminated. Assume that the reservoir context C_R in this case is simply $\{\{AB, A\bar{B}, \bar{A}B, \bar{A}\bar{B}\}\}$; then both C_2 and C_3 will bring back a $\neg A$-world. At the same time, to satisfy weak persistence, the resurrected $\neg A$-world must be either a B-world as in C_2, or a $\neg B$-world as in C_3. Taking the restrictive downward closure of the set containing C_2 and C_3 yields the post-refinement hyper-context Σ_2. We thus obtain two ways of preserving past information as witnessed by the two alternatives of Σ_2: either all the $\neg B$-possibilities are

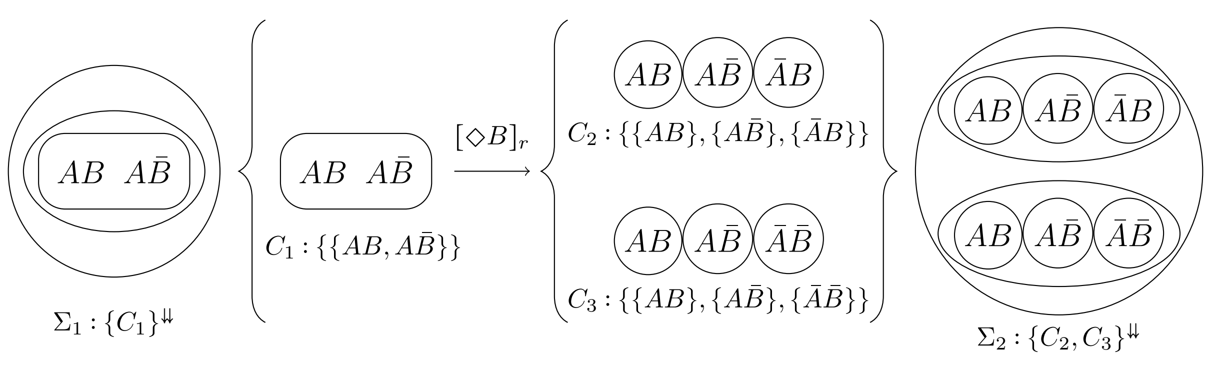

Figure 2.10: An example of weakly persistent refinement.

[15] While I do not intend to provide a detailed account of how C_R is determined in this paper, I believe this is where invoking some sort of ordering—e.g., relevance ordering as employed by Willer (2017)—can help.

[16] The symbol $\downdownarrows$ in the figure represents taking the restrictive downward closure of a set. See Appendix, Definition 12.

still A-possibilities as witnessed by C_2, or all the B-possibilities are still A-possibilities as witnessed by C_3.

As it turns out, this post-refinement hyper-context Σ_2 from the refinement with $[\Diamond B]_r$ is indeed identical to the posterior hyper-context after the general update $[\Diamond B]_r[\Diamond B]_u$. This is so because, similar to what happens with the proper update depicted in Figure 2.8, the proper update with $[\Diamond B]_u$ on each alternative in Σ_2 idles, which means applying the restrictive downward closure at the end will give back the input hyper-context.

Before ending this section, let me mention one more observation we can account for by appealing to the weakly persistent refinement induced by the antecedent of a conditional. Consider the following conversation:

(8) John: Alice will come to the party.
Mary: But what if Bob comes?

Mary appears to be neither completely agreeing nor outright disagreeing with what John just said. Instead, she seems to be asking John to reconsider his assertion in light of a new possibility without explicitly affirming or denying it. We can account for this intuition as follows: on the one hand, since entertaining a new possibility allows for revival of "dead" worlds, it explains why Mary's utterance is deemed as urging John to evaluate his claim; on the other hand, since John's claim should somewhat persist given weak persistence, it also explains why we tend to see Mary not as straightforwardly voicing a disagreement. Here, we restrict attention to the weakly persistent refinement induced by the antecedent of a conditional, but similar discourse effects can be produced by other expressions—e.g., by an epistemic *might*-claim. I shall leave exploring other triggering conditions for weakly persistent refinement to future work.

2.6 New Dynamic Strict Analysis of Conditionals

2.6.1 Introduce Postsuppositions

With a formal framework which enables us to materialize weak persistence at hand, I present a new dynamic strict analysis of conditionals in this section. In short, I construe a conditional

$\phi > \psi$ as inducing the following sequential update: $[\Diamond\phi][\Box(\phi \to \psi)][\Diamond\phi]$.[17] The first two updates resemble components from the existing dynamic strict analysis of conditionals, that is, updating with $\phi > \psi$ amounts to first updating with its entertainability presupposition $[\Diamond\phi]$ and then with its asserted content $[\Box(\phi \to \psi)]$. The last update with $[\Diamond\phi]$ involves what I call the *postsupposition* of a conditional. It functions as a delayed test whose purpose is to ensure that the posterior hyper-context after the update with the asserted content still satisfies the entertainability presupposition of the very conditional.

This idea that a sentence carries a postsupposition has been invoked in the past for various reasons to specify conditions that a relevant body of information needs to satisfy only after it has been updated with the at-issue content of the sentence. (Brasoveanu, 2012; Brasoveanu and Szabolsci, 2013; Condoravdi, 2015; Constant, 2012; Henderson, 2014). With respect to the current setting, the idea that the posterior hyper-context should still satisfy the entertainability presupposition of a conditional resonates with the idea that, in general, update procedures should be more or less idempotent (Veltman, 1996; Yalcin, 2015). More specifically, if one can utter a conditional to update a body of information consistently, then one should also be able to utter the same conditional again immediately afterwards, and the update should idle since the conditional should have already been supported. Hence, if updating a hyper-context first with the conditional's presupposition $[\Diamond\phi]$ and then with its asserted content $[\Box(\phi \to \psi)]$ makes the entertainability presupposition $\Diamond\phi$ no longer satisfiable, this suggests that the conditional cannot be felicitously uttered in the first place. To guarantee that this does not happen, we attach a postsupposition $[\Diamond\phi]$ to the first two updates.[18]

We can marshal some support for the idempotence of conditionals from the observation that conditionals of the form $A > \neg A$ (as well as $\neg A > A$) sound contradictory (cf. Wansing, 2014). This anomaly can be captured in the current framework with the aid of postsupposition. The conditional $A > \neg A$ sounds contradictory because for any hyper-context Σ,

[17] Although strictly speaking, the implication $\to$ employed in the current framework is not the material implication $\supset$, I still views this analysis as a form of dynamic strict analysis since we are still quantifying over all information states that are subsets of s when evaluating $\phi \to \psi$ at s.

[18] Additionally, it is worth mentioning that construing update idempotence as a postsupposition rather than some precondition for successful update can address the overgeneration problem pointed out by Mandelkern (2019). I leave a detailed discussion of this problem to another occasion.

$\Sigma[\Diamond A][\Box(A \to \neg A)][\Diamond A] = \Sigma_\perp$, where $\Sigma_\perp$ represents the absurd hyper-context $\{\{\varnothing\}\}$. This is so because for any Σ, the update with $[\Box(A \to \neg A)]$ necessarily eliminates every A-world from it, thereby making the subsequent update with the postsupposition $[\Diamond A]$ output the absurd hyper-context.

2.6.2 Vindicate ESI

We can finally put different pieces of the formal machinery together to show how ESI is vindicated. First, we define semantic consequence "$\vDash$" as follows:

> **Semantic Consequence:** $\phi_1, \ldots, \phi_n \vDash \psi$ iff for any hyper-context Σ, the sequential update with $[\phi_1], \ldots, [\phi_n]$ yields a hyper-context that supports ψ.[19]

To elucidate how ESI—viz., $\phi > \psi, (\phi \wedge \chi) > \neg\psi \vDash (\phi \wedge \neg\chi) > \psi$—is validated, a simple example should suffice for present purposes. Consider the following simplified variant of ESI:

$$\psi, \chi > \neg\psi \vDash \neg\chi > \psi$$

What we have above is a special instance of ESI derived from substituting a tautology $\top$ for the first antecedent ϕ. As a concrete example, (9a) and (9b) together entail (9c):

(9) (a) Alice will come to the party, (b) although she won't if Bob comes. $\Rightarrow$ (c) If Bob doesn't come, Alice will come to the party.

Here is how this inference can be captured on my account. Suppose our initial hyper-context $\Sigma_\top$, as shown in Figure 2.11, is the set $\{\{\{AB, A\bar{B}, \bar{A}B, \bar{A}\bar{B}\}\}\}$.[20] Updating $\Sigma_\top$ with (9a) eliminates all the $\neg A$-worlds in $\Sigma_\top$ and outputs Σ_1.

Since Σ_1 does not support $?B$, updating with the entertainability presupposition of (9b), namely $[\Diamond B]$, triggers the weakly persistent refinement with $[\Diamond B]_r$. This refinement is exactly the same as the one we just saw in Figure 2.10, where the only alternative in Σ_1 is split into two contexts: $\{\{AB\}, \{A\bar{B}\}, \{\bar{A}B\}\}$, and $\{\{AB\}, \{A\bar{B}\}, \{\bar{A}\bar{B}\}\}$. The post-refinement hyper-context is Σ_2, upon which the proper update with $[\Diamond B]_u$ idles.

[19] This is commonly known as the *update-to-test* consequence in the dynamic literature (cf. Veltman, 1996).

[20] I abbreviate this hyper-context using its set of alternatives.

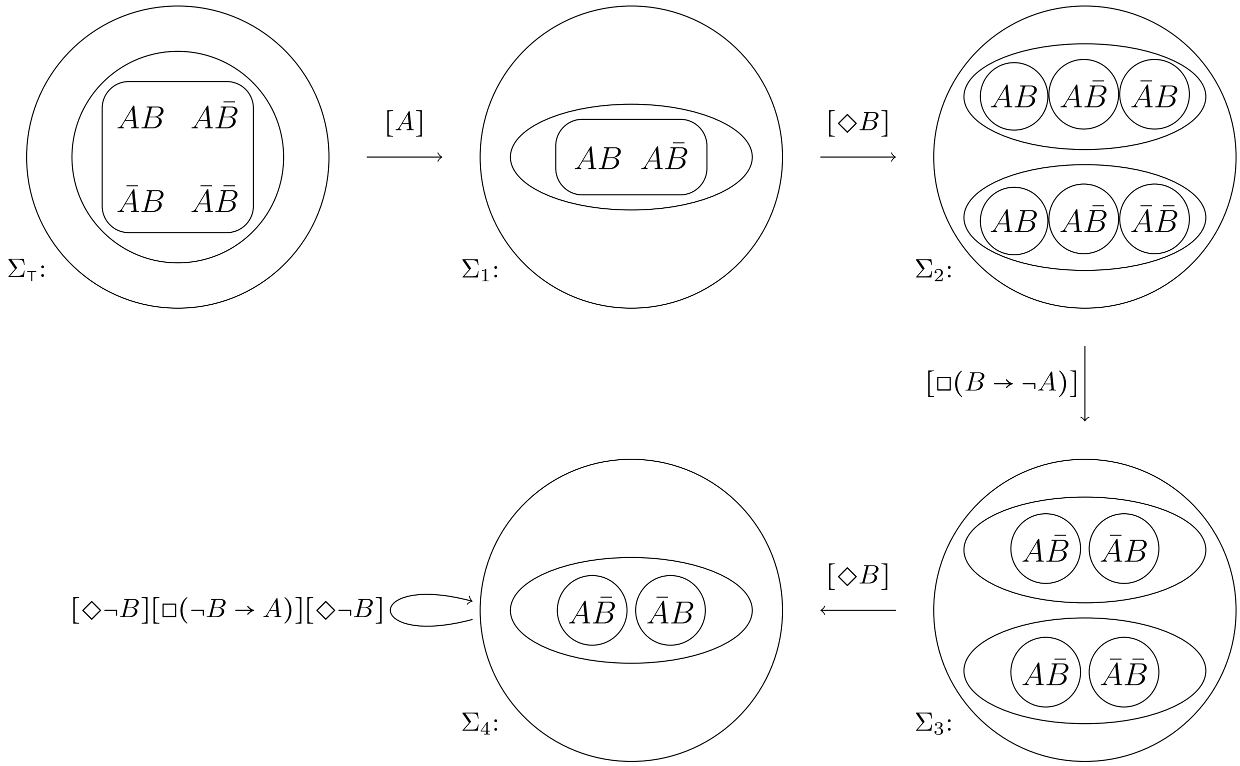

Figure 2.11: An example showcasing how ESI is vindicated

The next step is to update Σ_2 with the asserted content of (9b), namely with $[\Box(B \to \neg A)]$. Since Σ_2 is already aware of both A and B, the refinement with $[\Box(B \to \neg A)]_r$ idles; the update with $[\Box(B \to \neg A)]_u$ then eliminates every context in Σ_2 that contains the world AB. The posterior hyper-context is Σ_3.

Next, updating with the postsupposition of (9b), namely with $[\Diamond B]$, keeps only the first alternative $\{\{A\bar{B}\}, \{\bar{A}B\}\}$ but eliminates the second alternative $\{\{A\bar{B}\}, \{\bar{A}\bar{B}\}\}$. This is so because whereas B is a live possibility in the former alternative, it is deemed impossible in the latter one. Hence, we obtain Σ_4 as our posterior hyper-context after the updates with (9a) and (9b).

Finally, updating Σ_4 with the conditional (9c), that is, with the sequential update $[\Diamond\neg B][\Box(\neg B \to A)][\Diamond\neg B]$, idles. To elaborate, since Σ_4 already actively distinguishes between B and $\neg B$, the update with $[\Diamond\neg B]$ idles; and since Σ_4 does not contain any $(\neg A \wedge \neg B)$-worlds, the update with $[\Box(\neg B \to A)]$ also idles. What this means is that updating $\Sigma_\top$ with (9a) and (9b) yields a hyper-context that already supports (9c). Hence, ESI holds under the current analysis.

Moreover, since Σ_4 still contains some B-world, it does not support $\neg B$, which means the additional inference from $\top > A$ and $(\top \wedge B) > \neg A$ to $\top > \neg B$ cannot be drawn. Therefore, the current framework also manages to avoid the shortcoming of VSA.

2.7 Conclusion

I have argued that the existing dynamic strict account fails to validate the extended Sobel inference due to a lack of constraints on how the modal center expands. The traditional variably strict analysis, by directly appealing to a particular ordering, is capable of validating ESI, but in doing so, it also vindicates an additional unwarranted inference. In response, I propose a dynamic inquisitive framework which is capable of modeling the operation of raising a hitherto unentertained possibility to salience without directly relying on any particular ordering. Within this new framework, we are able to cash out a notion of weak persistence, which I consider as an important property governing how information evolves in light of new possibilities. By incorporating a notion of postsuppositions, the modified dynamic strict analysis of conditionals presented here manages to vindicate ESI without incurring the drawback associated with the variably strict account.

2.8 Appendix: Formal Definitions

Definition 2.8.1 (Information states). An information state s is a set of possible worlds and a subset of the whole logical space $s \subseteq W$.

Definition 2.8.2 (Support). We use "$\Vdash$" to denote support, specified as follows:

- $s \Vdash p$ iff $\forall w \in s : p$ is true at w, where p is atomic;
- $s \Vdash \phi \wedge \psi$ iff $s \Vdash \phi$ and $s \Vdash \psi$;
- $s \Vdash \phi \vee \psi$ iff $s \Vdash \phi$ or $s \Vdash \psi$;
- $s \Vdash \phi \to \psi$ iff $\forall t \subseteq s$: if $t \Vdash \phi$, then $t \Vdash \psi$;
- $s \Vdash \neg\phi$ iff $\forall t \subseteq s$: if $t \neq \varnothing$, then $t \nVdash \phi$.

Definition 2.8.3 (Downward closed). A set of states S is downward closed iff $\forall s \forall t$: if $s \in S$ and $t \subseteq s$, then $t \in S$.

Definition 2.8.4 (Contexts). Given a logical space W, a context C is a non-empty downward closed set of information states that can be derived from performing certain general updates on the initial context $C_\top \coloneqq \wp\{W\}$.

Definition 2.8.5 (Alternatives in a context). An alternative in a context C is a maximal element in C (i.e., an information state such that it is not a proper subset of any other states in C). We use $Alt(C)$ to denote the set of all alternatives in C.

Definition 2.8.6 (Support of modals by a state with respect to a context). Let s^c be an information state that is contained in the context C. Then,[21]

$$s^c \Vdash \Diamond\phi \quad \text{iff } \exists s' \in Alt(C) : s' \Vdash \phi$$
$$s^c \Vdash \Box\phi \quad \text{iff } \forall s' \in Alt(C) : s' \Vdash \phi$$

Definition 2.8.7 (General updates). $C[\phi] \coloneqq C[\phi]_r[\phi]_u$.

[21] Alternatively, we could define support of modals directly at contexts instead of at information states relative to a context. I adopt the above definition so that we can have a uniform definition of proper updates for both modal and non-modal formulas, as given by Definition 8.

Definition 2.8.8 (Proper updates)**.** $C[\phi]_u \coloneqq \{s \in C : s \Vdash \phi\}$.

Definition 2.8.9 (Refinement)**.** We define the refinement operation recursively in terms of the update operation as follows:

- $C[p]_r \coloneqq C[p \vee \neg p]_u$, where p is atomic;
- $C[\neg\phi]_r \coloneqq C[\phi]_r$;
- $C[\phi \wedge \psi]_r/[\phi \vee \psi]_r/[\phi \to \psi]_r \coloneqq C[\phi]_r[\psi]_r$.
- $C[\Diamond\phi]_r/[\Box\phi]_r \coloneqq C[\phi]_r$

Definition 2.8.10 (Context-level support)**.** A context C supports a sentence ϕ, notated as $C \Vdash \phi$, iff $C[\phi] = C$.

Definition 2.8.11 (Live and plain possibilities)**.**

- ϕ is a live possibility in C if $C[\Diamond\phi] = C$;
- ϕ is a plain possibility in C if $C[\Diamond\phi] \neq C$ and $C[\Diamond\phi] \neq C_\perp$ (viz., $C[\Diamond\phi] \neq \{\varnothing\}$);
- Otherwise, that is, when $C[\Diamond\phi] = C_\perp$, ϕ is an impossibility in C.

Definition 2.8.12 (Restrictive downward closed)**.** A set of contexts X is restrictively downward closed iff for all $C \in X$ and for all x, if $x \subseteq C$ and x is a context, then $x \in X$. For any given context C, $\downdownarrows C$ denotes the restrictive downward closure of C, i.e., a set of contexts containing C and every sub-context of C. For any given set of contexts X, we write $X^{\downdownarrows}$ to denote the restrictive downward closure of X, that is, $\bigcup_{C \in X} \downdownarrows C$.

Definition 2.8.13 (The initial and absurd hyper-contexts)**.** For any given logical space W,

- The initial hyper-context $\Sigma_\top \coloneqq \downdownarrows C_\top$, where $C_\top$ is the initial context w.r.t. W;
- The absurd hyper-context $\Sigma_\perp \coloneqq \{\{\varnothing\}\}$.

Definition 2.8.14 (Hyper-contexts)**.** Given a logical space W, a hyper-context Σ is a non-empty restrictively downward closed set of contexts that can be derived from performing certain updates on the initial hyper-context $\Sigma_\top$.

Definition 2.8.15 (Alternatives in a hyper-context)**.** An alternative in a hyper-context Σ is a maximal element in Σ, and we use $Alt(\Sigma)$ to denote the set of alternatives in Σ.

Definition 2.8.16 (General updates on hyper-contexts)**.** $\Sigma[\phi] = \Sigma[\phi]_r[\phi]_u$.

Definition 2.8.17 (Proper updates on hyper-contexts)**.** $\Sigma[\phi]_u = \{C[\phi]_u : C \in \Sigma\}^{\Downarrow}$.

Definition 2.8.18 (Strongly persistent refinement on hyper-contexts)**.** As before, we define refinement recursively in terms of proper updates as follows:

- $\Sigma[p]_r \coloneqq \Sigma[p \vee \neg p]_u$, where p is atomic;
- $\Sigma[\neg\phi]_r \coloneqq \Sigma[\phi]_r$;
- $\Sigma[\phi \wedge \psi]_r/[\phi \vee \psi]_r/[\phi \rightarrow \psi]_r \coloneqq \Sigma[\phi]_r[\psi]_r$;
- $\Sigma[\Diamond\phi]_r/[\Box\phi]_r \coloneqq \Sigma[\phi]_r$.

Definition 2.8.19 (Weakly persistent refinement)**.** Refining Σ with the entertainability presupposition of a conditional $[\Diamond\phi]_r$ becomes weakly persistent when $\Sigma \nVdash ?\phi$.[22] The refinement proceeds as follows:

- $\Sigma[\Diamond\phi]_r \coloneqq (\{C[\phi]_r \cup C_R[p_c]_r[\phi] : C \in Alt(\Sigma)\} \cup \{C[\phi]_r \cup C_R[p_c]_r[\neg\phi] : C \in Alt(\Sigma)\})^{\Downarrow}$, where $[p_c]_r = [p_1]_r, \ldots, [p_n]_r$ for every p such that $C \Vdash ?p$, and C_R is a contextually provided backup context.

Remark. To unpack this definition, we split every context C belonging to the set of alternatives $Alt(\Sigma)$ into two separate possibility spaces: $C[\phi]_r \cup C_R[p_c]_r[\phi]$ and $C[\phi]_r \cup C_R[p_c]_r[\neg\phi]$, each of which represents a way of how past information is preserved. Take the alternative $C[\phi]_r \cup C_R[p_c]_r[\phi]$ as an example. Its first component $C[\phi]_r$ simply brings ϕ to salience in the resulting context. The second component $C_R[p_c]_r[\phi]$ then reintroduces some previously eliminated possibilities from the reservoir context C_R. This is achieved by first refining C_R with $[p_c]_r$, which looks at every atomic proposition the input context C is actively aware of and then makes C_R also come to aware of these propositions.

[22]Here we restrict attention to weakly persistent refinement induced by the antecedent of a conditional. I leave exploring other triggering conditions for weakly persistent refinement to future work.

As such, we do not lose any distinctions the input context C is already making. The second update with $[\phi]$ ensures that all the worlds from C_R we are bringing back are indeed ϕ-worlds. As a consequence, we preserve past information in all the $\neg\phi$-worlds. Likewise, for the second alternative $C[\phi]_r \cup C_R[p_c]_r[\neg\phi]$, past information is preserved in all the ϕ-worlds. By conjoining the two sets and taking the restrictive downward closure $(\{C[\phi]_r \cup C_R[p_c]_r[\phi] : C \in Alt(\Sigma)\} \cup \{C[\phi]_r \cup C_R[p_c]_r[\neg\phi] : C \in Alt(\Sigma)\})^{\Downarrow}$, we obtain a new hyper-context such that every alternative in the original hyper-context Σ is split into two: one wherein past information is preserved in all the ϕ-worlds, thereby supporting $\phi \to \psi$, and the other wherein past information is preserved in all the $\neg\phi$-worlds, thereby supporting $\neg\phi \to \psi$.

Definition 2.8.20 (Hyper-context support). $\Sigma \Vdash \phi$ iff $\Sigma[\phi] = \Sigma$.

Definition 2.8.21 (Dynamic consistency). An update sequence $[\phi_1], \ldots, [\phi_n]$ is dynamically consistent iff there exists a hyper-context Σ such that $\Sigma[\phi_1], \ldots, [\phi_n] \neq \Sigma_\perp$.

Definition 2.8.22 (Dynamic consequence). $\phi_1, \ldots, \phi_n \vDash \psi$ iff $\forall\Sigma : \Sigma[\phi_1], \ldots, [\phi_n] \Vdash \psi$.

Chapter 3

A Questions-Under-Discussion-Based Account of Redundancy

Hurford (1974) observed that disjunctions where one disjunct entails the other (e.g., #*John was born in Paris or in France*) are generally infelicitous. To explain their infelicity, one common approach is to take Hurford disjunctions as involving a truth-conditionally redundant constituent whose deletion has no effect on the truth-conditional content of the whole sentence. However, this approach faces difficulty in light of other variants of Hurford disjunctions (Marty & Romoli, 2022). Drawing on Simons (2001), I present a new account of redundancy which utilizes questions under discussion (Roberts, 1996) and discourse trees (Büring, 2003). By postulating general constraints on the structure of discourse trees, I show how the (in)felicity of various Hurford sentences can be accounted for.

Keywords: Hurford disjunction, redundancy, question under discussion, discourse tree

3.1 Introduction

Consider an example of so-called Hurford Disjunctions (HDs):

(1) #John was born in Paris, or he was born in France.

In a context where the interlocutors know that Paris is located in France, "John was born in Paris" contextually entails " John was born in France", and as (1) demonstrates, disjunctions where one disjunct contextually entails the other sound odd. One explanation pins this oddness down on the apparent redundancy HDs involve. Since "John was born in Paris" contextually entails "John was born in France", (1) is truth-conditionally equivalent to "John was born in France". As a result, the first disjunct of (1) comes out as truth-conditionally redundant: deleting it would have no effect on the sentence's truth condition. Consequently, by appealing to Grice's (1975) maxim of manner—to be more specific, the submaxim of brevity—we can attribute the oddness of HDs to the presence of apparently redundant material.

A general notion of redundancy also helps to explain why certain HDs are nonetheless felicitous. Such examples were first noted by Gazdar (1979). For example, consider (2):

(2) John read some or all of the books .

In (2), although "John read all of the books" entails "John read some of the books", the disjunction is nonetheless felicitous. One common explanation as to why disjunctions like (2) remain unmarked is that the first disjunct is locally exhaustified to receive a strengthened reading so that it is no longer entailed by the second disjunct (cf. Chierchia et al., 2012). For example, "John read some of the books" in (2) is exhaustified to mean that John read *some but not all* of the books, and as such it is no longer entailed by "John read all of the books". Given that the first disjunct is no longer entailed by the second, its deletion would have an impact on the sentence's truth condition. Consequently, disjunctions like (2) are no longer regarded as containing any redundant material and are thus deemed felicitous.

Although the suggestion of invoking redundancy to explain the oddness of HDs appears promising, working out the details is not an easy task. A general theory of redundancy needs to capture not only the infelicity of HDs but also the infelicity (or felicity) of other Hurford sentences in the vicinity. To illustrate, consider an example of what Marty and

Romoli (2022) call *Quasi Hurford Disjunctions* (QHDs):

(3) John was born in Paris, or he was born in France but not in Paris.

For some concreteness, let us evaluate (3) against the following constituent-based non-redundancy account (Katzir & Singh, 2013; Marty & Romoli, 2022).

> **Constituent-based non-redundancy** A sentence S cannot be used in context c if there is any constituent X in S that is contextually equivalent to one of X's subconstituents.

On this account, (3) is predicted to be just as infelicitous as (1) because both sentences contain "John was born in France" as a subconstituent which is contextually equivalent to the original disjunction. Hence, both (1) and (3) are considered containing redundant material and thus deemed infelicitous. This prediction is contrary to our intuitive judgment: unlike (1), (3) sounds perfectly natural. The constituent-based non-redundancy account fails to distinguish the felicitous QHDs from the infelicitous HDs.

Moreover, as we shall see in §3.2, other accounts on the market also run into trouble in light of a wide variety of Hurford sentences. In particular, I will draw attention to disjunctions like (4) which I call *Conjunctive Hurford Disjunctions* (CHDs):

(4) #John was born in Paris, or he was born in France and Mary was born in London.

Since (4) is not contextually equivalent to any of its subconstituents and nor does it contain any constituent that is contextually equivalent to any of its subconstituents, it is not ruled out on the constituent-based non-redundancy account. Thus, it is predicted to be felicitous, contrary to our intuitive judgement. In fact, as we shall see, CHDs pose a thorny challenge for almost all existing theories of redundancy.

To account for the (in)felicity of various types of Hurford sentences, I present an analysis that utilizes the notion of Questions Under Discussion (QUDs) (van Kuppevelt 1995; Roberts, 1996; Asher and Lascarides, 2003) while still retaining the spirit of the redundancy account. Under this approach, we understand assertions as responding to some discourse question. For example, the assertion "John was born in France$_F$" with a focus marking on "France" is normally understood as responding to the question "Where was John born?". Given this relation, we can recast the ban against redundant material as a general constraint on how one should respond to discourse questions. When a sentence purports to offer multiple answers to the same question, these answers must be distinct. In fact, an account along this line has been suggested by Simons (2001). According to her, each disjunct in a disjunction should provide a distinct answer to the discourse question so that none should "contextually entail" another. This paper aims to further develop such a QUD-based approach so that the (in)felicity of various Hurford sentences can be properly captured. In particular, I propose that the disjuncts in a disjunction are not only taken to provide distinct answers to the same question—call this the *distinctness constraint*—but also understood as engaging the question "in the same way"—call this the *uniformity constraint*. I will explicate this idea of answering a question "in the same way" using Büring's (2003) notion of discourse-trees.

Additionally, my account can be extended to explain why certain particles such as "at least" appear to repair an otherwise infelicitous HD, as (5) demonstrates:

(5) John was born in Paris, or at least he was born in France.

Past literature (e.g., Singh, 2008) often sets aside disjunctions like (5) as special cases that do not contain a genuine disjunction and thus falls outside the explananda to be accounted for. In this paper, I argue for a more unified analysis and show how such an account is possible under the current QUD-based approach.

This paper proceeds as follows. In §3.2, I review the literature on Hurford disjunction, lay out the data points I plan to capture, and highlight the difficulties faced by the existing redundancy accounts. I then offer my QUD-based analysis in §3.3. By postulating general constraints on the structure of discourse-trees, I shall demonstrate how the (in)felicity of various Hurford sentences can be properly captured: §3.3.1 focuses on the standard HDs and other infelicitous variants, §3.3.2 discusses the felicitous ones, namely QHDs, and §3.3.3 extends the analysis to accommodate the repairing effect of scalar particles such as "at least". §3.4 briefly discusses one open issue concerning Hurford questions. §3.5 concludes.

3.2 Empirical Landscape

3.2.1 Quasi Hurford Disjunctions

Recall that according to the constituent-based account of redundancy, a sentence contains redundant material if it contains a constituent that is contextually equivalent to one of its subconstituents. This account, though simple and intuitively appealing, fails to explain the contrast between the infelicitous HD in (1) and the felicitous QHD in (3). Both sentences contain as a subconstituent "John was born in France" which is contextually equivalent to the original Hurford sentence. To explain this contrast, Marty and Romoli (2022) develop a more sophisticated redundancy account which takes a sentence's exhaustified meaning into consideration. More specifically, they make use of the exhaustified interpretation a disjunction "p or q" can receive, namely the exclusive reading "p or q, but not both". To explain why QHDs are felicitous with the help of exhausitification, they draw on Mayr and Romoli's (2016) explanation as to why a disjunction like (6) is felicitous even though it also has a seemingly truth-conditionally redundant component and thus could potentially be blocked by its simplification in (7):

(6) Mary isn't pregnant, or she is (pregnant) and she is happy.

(7) Mary isn't pregnant, or she is happy.

According to Mayr and Romoli, in order to determine redundancy, we should also compare the exhaustified meaning of a sentence containing an allegedly redundant component to that of its simplification without it, and only when the allegedly redundant component contributes nothing to altering the sentence's exhaustified meaning can it be truly deemed redundant.

Since the exhaustified interpretation of a disjunction $(A \vee B)$ is given by the exclusive disjunction $(A \vee B) \wedge \neg(A \wedge B)$, (6) receives the following strengthened reading:

(8) (Mary isn't pregnant, or she is pregnant and she is happy) $\wedge\neg$(Mary isn't pregnant, AND she is pregnant and she is happy) $\approx$ (6).

Given that the conjunction "Mary isn't pregnant, AND she is pregnant and she is happy" is a contradiction and thus its negation a tautology, the exhaustified interpretation given by (8) turns out to be equivalent to (6). In other words, exhaustification is vacuous in the case of (6). By contrast, exhaustifying (7) yields:

(9) (Mary isn't pregnant, or she is happy) $\wedge\neg$(Mary isn't pregnant, AND she is happy) $\approx$ Either Mary isn't pregnant, or she is happy, but not both.

What (9) rules out is the possibility that Mary is happy but not pregnant. This possibility is not explicitly ruled out by the more complex disjunction in (6). As (6) and (7) differ in their exhaustified meaning, the constituent "she is (pregnant)" in (6) is deemed non-redundant.

Now, since "Mary isn't pregnant, or she is and she is happy" shares a similar structure with the QHD "John was born in Paris, or he was born in France but not in Paris", Marty and Romoli contend that we can as well appeal to exhaustification to explain why QHDs are felicitous. Compare the exhaustified reading of the QHD in (10) to that of its simplification in (11):

(10) John was born in Paris, or he was born in France but not in Paris.

(11) John was born in Paris, or he was born in France.

As with (6), exhaustifying (10) has no effect since the conjunction "John was born in Paris, AND he was born in France but not in Paris" is a contradiction. On the other hand, the exclusive interpretation of (11) is given by (12):

(12) (John was born in Paris, or he was born in France) $\wedge\neg$(John was born in Paris, AND he was born in France) $\approx$ John was born in France but not in Paris.

Since (10) and (11) differ in their exhaustified meaning, the additional conjunct in (10) is not redundant, and (10) is predicted to be unmarked.

However, Marty and Romoli's solution faces several problems. First, unlike "Mary isn't pregnant, or she is happy" which can be strengthened to mean "Mary isn't pregnant, or she is happy, but not both", "John was born in Paris, or he was born in France" cannot be taken to mean "John was born in France but not in Paris". On one hand, it is highly doubtful that the strengthened interpretation is empirically attested; on the other hand, the strengthened interpretation would render the standard HD "John was born in Paris, or he was born in France" no longer infelicitous on Mayr and Romoli's modified redundancy account. To see this latter point, note that if we were to allow "John was born in Paris, or he was born in France" to mean "John was born in France but not in Paris", then "John was born in Paris, or he was born in France" and its simplification "John was born in France" would differ in their exhaustified meanings. As a result, "John was born in Paris, or he was born in France" would no longer be predicted to be defective.

Furthermore, even if we grant that the constituent "but not in Paris" in (10) is not redundant, we still have not adequately explained why (10) is felicitous. This is so because even if (10) and (11) differ in their exhaustified meaning, (10) still has the same exhaustified meaning as the further simplification "John was born in France". Even after exhaustification, (10) and "John was born in France" still share the same truth condition, that is, *true* iff John was born in France. Hence, Marty and Romoli's account would still predict (10) to be infelicitous because of the availability of a simpler truth-conditionally equivalent sentence.

In short, while appealing to exhaustification may offer a solution to the puzzle concerning why disjunctions such as "Mary isn't pregnant, or she is (pregnant) and she is happy" do not involve redundancy, the same strategy cannot be replicated so as to capture the contrast between felicitous QHDs and infelicitous HDs.

3.2.2 Long Distance Hurford Disjunction

A different way to modify the redundancy account is to check redundancy not at the global level but at the molecular level (Katzir & Singh, 2013; Ciardelli & Roelofsen, 2017):

> **Molecular Non-Redundancy:** A sentence S is deviant in a context c if S's logical form contains a node $O(X, Y)$ where O is a binary connective and X and Y are its two arguments such that $O(X, Y)$ is contextually equivalent to either X or Y.[1]

On this account, HDs like (1) are still correctly predicted to be deviant: "John was born in Paris or in France" is contextually equivalent to "John was born in France". On the other hand, it also predicts the QHD in (2) to be felicitous as desired: the whole disjunction "John was born in Paris, or he was born in France but not in Paris" is not contextually equivalent to either of its disjuncts; likewise, the embedded conjunction "he was born in France but not in Paris" is not contextually equivalent to either of its conjuncts.

[1]Here, I follow Ciardelli and Roelofsen (2017) in formulating the constraint at the level of logical forms. And to further simplify, let us assume that X and Y are both sentential/propositional constituents of S.

The molecular non-redundancy account, however, fails to explain why (13), which is an example of what Marty and Romoli call *Long Distance Hurford Disjunctions* (LDHDs), is infelicitous.

(13) #John was born in London or in Paris, or he was born in France.

Since none of the binary nodes can be replaced with one of its immediate constituents, the molecular non-redundancy account incorrectly predicts it to be felicitous. One immediate response is to introduce n-ary connectives:

> **Non-Redundancy with n-ary Connectives:** A sentence S is deviant in a context c if S's logical form contains a node $O(X_1, X_2, \ldots, X_n)$ where O is an n-ary connective ($n \geq 2$) such that $O(X_1, X_2, \ldots, X_n)$ is contextually equivalent to a sentence derived from deleting any constituent from $X_1, X_2, \ldots, X_n$.

Under the n-ary version, for example, since (13) is truth-conditionally equivalent to the simpler "John was born in London, or he was born in France", the extra disjunct "John was born in Paris" now becomes redundant. One way to formally realize such an account is via inquisitive semantics (Ciardelli & Roelofsen, 2017; Anvari, 2021; see also Ciardelli et al., 2018). In brief, under inquisitive semantics, disjunction introduces a set of alternatives. The LDHD in (13), when taken as a whole, introduces three alternatives: [John was born in London], [John was born in Paris], and [John was born in France]. But since "John was born in Paris" entails "John was born in France", the alternative [John was born in Paris] collapses into [John was born in France] and thus becomes redundant. That being said, simply making disjunction n-ary and adopting inquisitive disjunction is still inadequate because similar oddness resurfaces in sentences that contain a combination of disjunction and conjunction, as in the case of conjunctive Hurford disjunctions.

3.2.3 Conjunctive Hurford Disjunctions

Recall the CHD from (4), repeated below as (14):

(14) #John was born in Paris, or he was born in France and Mary was born in London.

The infelicity of CHDs is problematic because, unlike HDs and LDHDs, there does not appear to be any constituents that are truth-conditionally redundant at first glance. None of the constituents in (14) appears deletable without affecting the truth condition of the whole sentence. Nevertheless, the sentence is still perceived as defective.

To my knowledge, the only existing account that addresses this issue to some degree comes from Simons (2001). She discusses the following example:

(15) Q: What kind of car does Jane drive?

A: #Either she drives a Subaru station wagon, or George drives a Toyota and she drives a Subaru.

The answer takes the form of a CHD. According to Simons, although (15A) is not truth-conditionally equivalent to any of its simplifications, it still in a sense contains redundant material when viewed as an answer to the question in (15Q). Since the question concerns the kind of car that Jane drives, the subconstituent "George drives a Toyota" in (15A) contributes nothing to answering this question. Therefore, with respect to the question (15Q), (15A) is indeed as informative as its simplification "Jane drives a Subaru".[2] Consequently, (15A) is deemed infelicitous as an answer to (15Q).

To put it another way, what Simons essentially postulates is a distinctness constraint on disjunctive answers. That is, each disjunct in a disjunction should provide a unique answer to the topic question. Since the answers provided by the two disjuncts in (15A) with respect to the question in (15Q) are not non-entailing, infelicity ensues.

However, Simons's account only partially addresses the challenge posed by CHDs because it fails to generalize to other discourse questions. For example, suppose that the discourse question in (15) is "Who drives what?" instead. Then, since "George drives a Toyota" does contribute as a partial answer to this question, (15A) is no longer as informative as the simplification "Jane drives a Subaru" with respect to this question. As a result, contrary to our intuitive judgement, (15A) would be regarded as a perfectly natural answer. Similarly, on Simons's account, although the CHD in (14) is deemed infelicitous as an answer to the question "Where was John born?", it is again incorrectly predicted as felicitous with respect to the question "Who was born where?".

To address this inadequacy, I propose to modify Simons's account in the following way.

[2]This is cashed out formally in Simons (2001) using Groenendijk and Stokhof's (1984) partition semantics. The detail does not concern us too much here.

The inability to account for the infelicity of (15A) with respect to the discourse question "Who drives what?" stems from the fact that, on Simons's account, the distinctness constraint only applies at the level of discourse questions. Given this, one solution is thus to make the constraint also apply at a more local level. More specifically, "Jane drives a Subaru" and "Jane drives a Subaru station wagon" can still be regarded as responding to the same QUD (e.g., "What does Jane drive?") even if this QUD is not the discourse question. As such, we can still obtain a violation of the distinctness constraint. Moreover, the reason why distinctness can be checked with respect to the question "What does Jane drive?" even if the discourse question is "Who drives what?" instead is that, assuming Jane is a relevant individual, the question "What does Jane drive?" is inherently related to the discourse question given that an answer to the former provides a partial answer to the latter. Hence, answering the question "What does Jane drive?" can be regarded as part of the inquiry strategy to address the discourse question "Who drives what?". On my account, the distinctness constraint applies not only at the level of the discourse question but also at the level of subquestions contained in the strategy of inquiry that is used to tackle the discourse question. To materialize this picture, I will make use of Büring's (2003) notion of discourse trees and postulate an additional constraint called *uniformity*.

To take stock, existing accounts of redundancy fail to capture the whole range of data concerning Hurford disjunctions. The table below summarizes the relevant data points discussed so far alongside how well the existing redundancy accounts handle these data.

Types of Hurford Disjunctions	Are they felicitous?	Constituent	Molecular	N-ary
Simple Hurford	No	✓	✓	✓
Quasi Hurford	Yes	✗	✓	✓
Long Distance Hurford	No	✗	✗	✓
Conjunctive Hurford	No	✗	✗	✗

As the table shows, none of the accounts on the market can adequately handle the full range of Hurford disjunctions; in particular, they all fail to explain why conjunctive Hurford disjunctions are infelicitous. Building upon Simons (2001), I devise a more sophisticated QUD-based account that manages to capture these data. As we shall see in §3.3.1, I explain why simple Hurford, long-distance Hurford and conjunctive Hurford disjunctions are all infelicitous by appealing to two constraints on the structure of discourse trees: uniformity and distinctness. These two constraints together help to cash out the thought that a disjunction's disjuncts should tackle a given QUD in the same way and each provide a distinct answer to it. The felicity of quasi Hurford disjunctions will be dealt with in §3.3.2 which requires the postulation of additional constraints.

3.2.4 Hurford Disjunction with a Scalar Particle

I conclude §3.2 by introducing another piece of data that I wish to capture. As mentioned, it has been observed that certain scalar particles such as "at least" (Singh, 2008) and "even" (Westera, 2020) can repair an otherwise infelicitous Hurford disjunction:

(16) John was born in Paris, or at least in France.

(17) John was born in France, or even in Paris.

Sentences like (16) and (17) have often been set aside in the literature on Hurford disjunction. For example, in his brief explanation as to why "at least" fixes an otherwise infelicitous HD, Singh (2008) comments that the "or" in "or at least" is not a genuine disjunction; rather, "or at least" should be better viewed as a form of correction: the speaker retracts what had just been said proceeding it.

I disagree with this diagnosis as I think the "or" in (16), as well as in (17), retains its disjunctive force. For one thing, note that we can embed (16) in the antecedent of a conditional and the resulting conditional still licenses the inference commonly known as the simplification of disjunctive antecedent:

(18) Mary will date John if he was born in Paris or at least in France. $\Rightarrow$

(19) Mary will date John if he was born in Paris.

(20) Mary will date John if he was born in France.

The inferences from (18) to (19) and (20) seem natural. However, if we were to construe "or at least" simply as a device of retraction which renders (18) come out expressing the same as (20), we would have difficulty explaining the inference from (18) to (19). This is so because conditionals are not monotonic in the antecedent place as attested by the so-called

Sobel sequences (Lewis, 1973; Stalnaker, 1968; Willer, 2017). For instance, (21) does not feel inconsistent:

(21) Mary will date John if he was born in France. But she won't if he was born in Paris.

Since we cannot view (19) as being entailed by (20), the inference from (18) to (19) would be left unexplained on the retraction account.

A more plausible approach, I think, is to treat the apparent disjunction in "or at least" as genuine disjunction with its disjunctive force intact. I will provide such an account in §3.3.3. To explain the reason why the presence of "at least" repairs what would otherwise be an infelicitous Hurford disjunction, I shall once more invoke QUDs. In a nutshell, scalar particles such as "at least" and "even" introduce alternative strategies of inquiry such that the two answers "John was born in Pairs" and "John was born in France" will no longer respond to the same QUD. Hence, even though "John was born in Paris" entails "John was born in France", no violation of the distinctness constraint occurs as they do not answer the same QUD. Consequently, (16) and (17) are judged as felicitous.

3.3 A New QUD-Based Analysis

3.3.1 D-Trees and Core Constraints

In this section, I will show how the infelicity of HDs, LDHDs, and CHDs can be captured on a QUD-based account of redundancy. As mentioned in §3.2.3, in order to explain why CHDs are infelicitous, I will modify Simons's account by having the distinctness constraint checked not only at the level of discourse questions but also locally at the level of subquestions. According to Roberts (2012), discourse participants devise strategies to answer a given topic question by breaking it down to more specific subquestions. Let us call Q' a subquestion of Q iff Q entails Q', that is, iff every complete answer to Q also provides a complete answer to Q'.[3] A subquestion can be broken down into further subquestions. Thus, QUDs enjoy a hierarchical structure, and one convenient way to visually represent such a hierarchical structure is to organize QUDs into a d(iscourse)-tree (Büring, 2003).

[3] Alternatively, we can view Q' as a subquestion of Q iff a complete answer to Q' provides a partial answer to Q. The detail does not concern us greatly here.

We define a d-tree as a partially ordered set of nodes, each representing a declarative or an interrogative sentence. For present purposes, we may simplify by assuming that these nodes stand in one of the two relations: (a) the question-and-subquestion relation, which relates two interrogative sentences, and (b) the question-and-answer relation, which relates an interrogative sentence to a declarative sentence.[4] If an interrogative sentence dominates a declarative sentence, then the latter must answer the former. In other words, we shall enforce the question-answer congruence constraint as standardly postulated (Roberts, 1996; Onea & Zimmermann 2019). Next, let us call any sub-tree of a d-tree that consists exclusively of questions a *QUD-tree*. A QUD-tree represents a strategy of inquiry which breaks down a superquestion into more specific subquestions, thereby laying out a particular way of tackling the superquestion. To predict the infelicity of various Hurford sentences, I posit the following core constraints (as for capturing the felicity of QHDs, additional constraints are needed which will be laid out in §3.3.2):

Uniformity: A disjunction's disjuncts must evoke the same strategy of inquiry with respect to the QUD answered by the whole disjunction.

Distinctness: Answers to the same question must be distinct in terms of non-entailment.

Consider the uniformity constraint first. It is supposed to capture the idea that the disjuncts of a disjunction should answer a question in the same way. Given that strategies of inquiry are represented by QUD-trees, uniformity says that the root question Q of a disjunction should branch into two identical QUD-trees each of which is dominated by a copy of Q. To illustrate with a simple example, consider the HD "John was born in Paris, or he was born in France" in (1). Suppose that this disjunction is uttered as a response to the question "Where was John born?". We can then depict its corresponding d-tree as shown in Figure 3.1. The root question to which the disjunction responds in this case is "Where was John born?". By uniformity, it branches into two copies of itself each of which immediately dominates one disjunct of the whole disjunction. Since the two disjuncts fail to provide non-entailing answers to the same question, distinctness is violated, and thus we

[4] Recent work on d-trees have postulated other types of relations such as dependent questions and potential questions (see, e.g., Onea & Zimmermann 2019). It is possible to modify the current framework to allow for other types of relations between nodes.

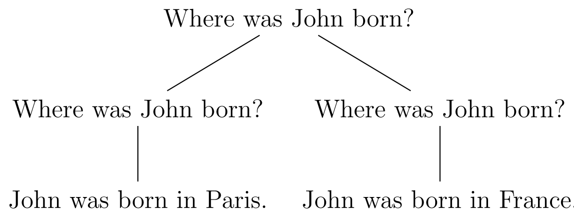

Figure 3.1: D-tree for the Hurford disjunction in (1).

predict (1) to be infelicitous. In this example, each disjunct does not invoke complex inquiry strategies; hence, only the root question itself is copied on each branch. In other cases where complex strategies are evoked as in the case of the CHD shown in Figure 3.4 below, the entire QUD-tree will be duplicated on each branch to observe uniformity.

To adduce a different example where the uniformity constraint is violated, consider (22A) which is infelicitous as a response to (22Q).

(22) Q: Where was John born? And where was Mary born?
A: #John was born in Paris or Mary was born in London.

Uniformity straightforwardly explains why (22A) feels odd: since the first disjunct purports to answer the first question whereas the second disjunct a completely different one, uniformity is violated.

The current implementation of d-trees also provides a convenient way to tease apart trees that correspond to disjunction and conjunction. Conjunctions such as (23) and (24) where the two conjuncts differ in only one dimension are conceived of as generating a d-tree where the two conjuncts are immediately dominated by the root question.

(23) John likes raspberries and strawberries.

(24) #John likes raspberries and berries.

For example, suppose that (24) is uttered as a response to the discourse question “what fruit does John like?”; we can then draw the d-tree for (24) as shown in Figure 3.2. Intuitively, when the two conjuncts of a conjunction differs only in one dimension, they each purport to provide a partial answer to the QUD answered by the whole conjunction (cf. Roberts, 1996; Jasinskaja & Zee, 2008; Riester, 2019). But since one conjunct in (24) entails the other, the two answers fail to be distinct, and thus (24) is regarded as defective. While this simple

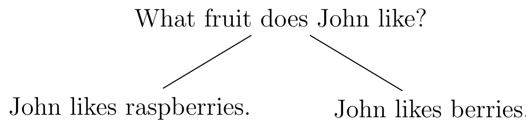

Figure 3.2: D-tree for the Hurford conjunction in (24).

treatment of conjunctive answers is not without any criticism (e.g., Onea, 2019), I shall stick to it in this paper for it offers an elegant way to reflect the difference between disjunction and conjunction in d-trees.

One last caveat regarding uniformity is that a disjunction will not always require branching. When branching is involved, each disjunct is considered as a separate answer to the QUD addressed by the whole disjunction. But branching does not always occur. For example, when the QUD is a polar question such as "Does John like raspberries or strawberries↑?", a positive answer "yes" to this question will be immediately dominated by it without the polar question splitting into further questions.

Now, as for the distinctness constraint, it is modeled after Simons (2001) and requires that answers to the same question be distinct and non-entailing. The difference is that distinctness is no longer just a constraint on disjunction with respect to the topic question but intended as a general constraint on all question-answer pairs in a given d-tree. It dictates that no two question-answer pairs in a d-tree can share the same question but have two answers where one entails the other. This includes cases where the two answers are dominated by the same node as in the tree from Figure 3.2 as well as cases where the two answers are dominated by two different nodes where the two nodes just happen to denote the same question as in the tree from Figure 3.1. Since both of them violate distinctness, we correctly predict the Hurford disjunction in (1) and the Hurford conjunction in (24) to be infelicitous.

A similar story can be told as to why long-distance Hurford disjunctions are infelicitous. This is where a QUD-based analysis starts to pay off. Consider the LDHD in (13), repeated below as (25):

(25) #John was born in Paris or in London, or he was born in France.

Since the first disjunct of (25) is itself a disjunction, the question to which it responds will split once more into two copies of itself. By uniformity, the QUD-tree on the left branch

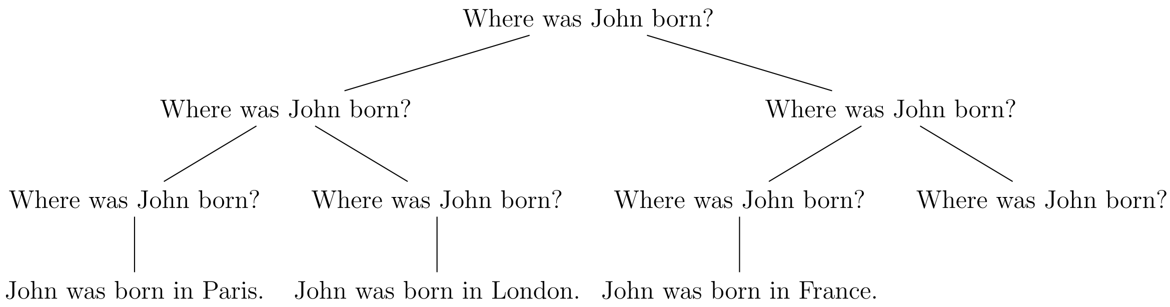

Figure 3.3: One tentative d-tree for (25).

is again duplicated on the right, thereby yielding the QUD-tree in Figure 3.3. One small wrinkle here concerns how we should connect the second disjunct “John was born in France” to the QUD-tree on the right branch, since there are two question nodes but only one answer node. A full resolution of this issue is not required for our current purposes as it does not play a significant role in explaining the infelicity of (25). For example, as Figure 3.3 shows, we could leave the last question node unanswered, which in turn can be used to signify that there isn’t a fourth possibility regarding John’s birthplace. Alternatively, if we consider disjunction as n-ary rather than binary, we can split the root question into three copies of itself and thereby circumvent the problem. In any case, since “John was born in Paris” and “John was born in France” respond to the same question, (25) is predicted to be infelicitous as desired.

Move on to conjunctive Hurford disjunctions. We can associate different d-trees with the CHD in (26) depending on how the discourse question is construed.

(26) #John was born in Paris, or he was born in France and Mary was born in London.

Suppose that (26) is uttered as a response to the question “Who was born where?”. In regard to this kind of multiple *wh*-question, it is natural to treat the second disjunct “John was born in France and Mary was born in London” where the two conjuncts differ in two dimensions as invoking contrastive topics. Following Büring’s (2003) analysis of contrastive topics, we can treat the two conjuncts of “John was born in France and Mary was born in London” as each responding to a different subquestion of “Who was born where?”. The resulting d-tree is depicted in Figure 3.4. On the right branch “John was born in France” and “Mary was born in London” are dominated by two separate subquestions of “Who was born where?”,

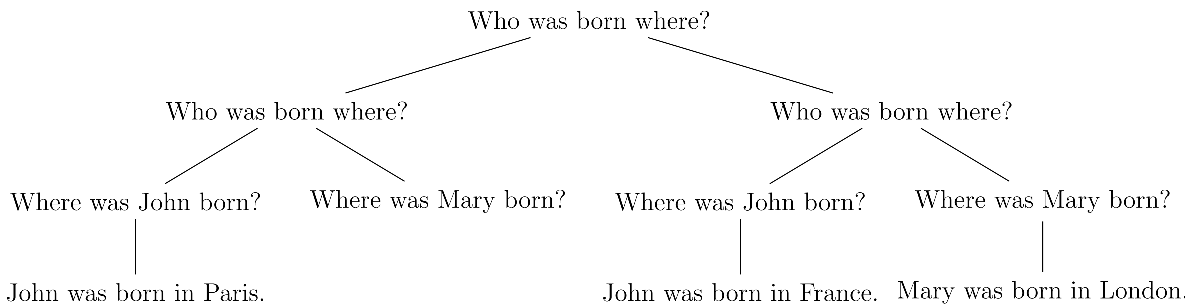

Figure 3.4: A d-tree for (26) with contrastive topics; note that the QUD-tree on the right is copied to the left branch.

namely "Where was John born?" and "Where was Mary born?" respectively. By uniformity, the QUD-tree on the right branch is copied to the left. Since the two occurrences of "Where was John born?" are associated with two answers where one entails the other, distinctness is violated and (26) is predicted to be infelicitous.

On the other hand, we can also construct an alternative d-tree for (26) that does not invoke contrastive topics. This is more natural when (26) is used as an answer to a single *wh*-question such as "Under what conditions do I win the bet?". This time, as Figure 3.5 shows, the question node "Under what conditions do I win the bet?" that appears on the right branch no longer further splits into additional subquestions but instead immediately dominates the two conjuncts of the conjunctive answer "John was born in France and Mary was born in London". Nonetheless, since "John was born in Paris" and "John was born in France" are still two non-distinct answers to the same question, we correctly predict (26) to be infelicitous with respect to this new discourse question as well.

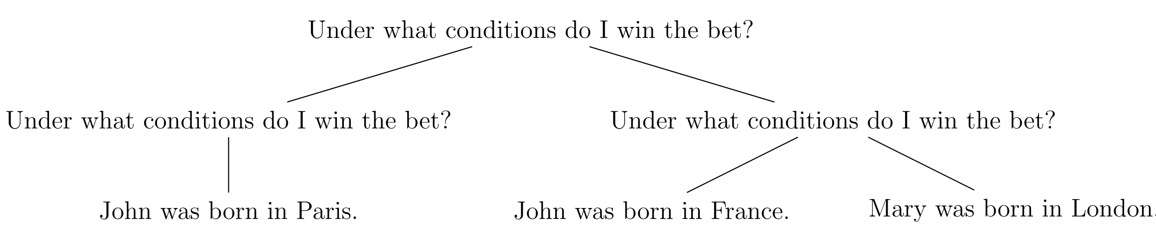

Figure 3.5: A d-tree for (26) with a different discourse question.

3.3.2 Additional Constraints

The QUD-based analysis so far successfully captures the infelicity of simple Hurford disjunctions, long-distance Hurford disjunctions, and conjunctive Hurford disjunctions. However, it does not immediately explain why quasi Hurford disjunctions are felicitous. Consider again the felicitous QHD in (3), repeated below as (27):

(27) John was born in Paris, or he was born in France but not in Paris.

To correctly predict (27) as felicitous, we need to block the d-tree depicted in Figure 3.6.

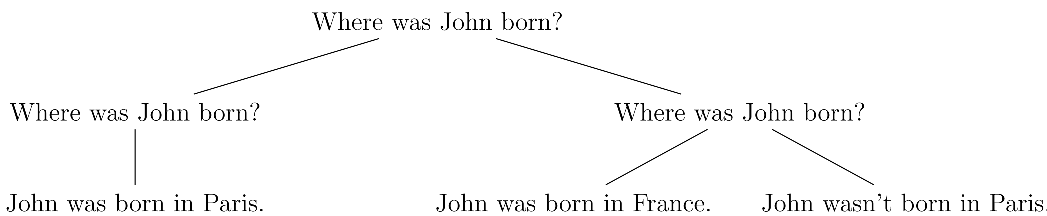

Figure 3.6: A potential d-tree for (27) that needs to be ruled out.

Since this d-tree violates distinctness, were it to be allowed, we would incorrectly predict (27) as infelicitous. One way to rule it out as an admissible representation for (27) is to constrain how implicit QUDs can be reconstructed. Although the issue about how to reconstruct the implicit QUD for a target sentence remains rather elusive, some headway has been made (e.g., Riester 2019; Onea 2019). One constraint that I will appeal to in particular is the following one postulated by Riester (2019):

Maximize-Q-Anaphoricity: Implicit QUDs should contain as much given (or salient) material as possible.

The relevant notion of givenness is from Schwarzschild (1999) (see also Rochemont, 2016). To give an example from Riester, among the three candidates QUDs in (29), (29c) is preferred as the implicit QUD for (28) over (29a) and (29b) because it maximizes the given/salient material in (28).

(28) He literally suffocated.

(29) (a) What happened?

(b) What about him?

(c) How bad was his condition?

The underlying intuition behind this constraint is that an implicit QUD reconstructed from an answer should enjoy a certain degree of specificity with respect to its answer: an implicit QUD should be maximally specific granted that it does not contain any new non-salient material. To provide another example, between the two sentences in (31), (31b) seems to serve as a better answer to (30) than (31a).

(30) John was born in Paris. What about Mary?

(31) (a) She was born in the UK.
(b) She was born in London.

This is so because the implicit QUD answered by "John was born in Paris" can be readily reconstructed as "Which city was John born in?". As such, the *what-about* question will be interpreted as asking "Which city was Mary born in?", thereby making (31b) the more suitable answer. Of course, this is not to say that a suitable answer to (30) will always specifies a city. When the speaker cannot specify the city or when the city in which Mary was born is not well-known, an answer that specifies a country would be more appropriate due to other general Gricean considerations. Setting these cases aside, it does seem that a proper answer to "What about Mary?" should provide the same degree of specificity as that of "John was born in Paris".

Imposing Maximize-Q-Anaphoricity now blocks the problematic d-tree in Figure 3.6 because "Where was John born?" fails to maximize salient material when taken as an implicit QUD for "John was born in France" and "John was born in Paris". As the first answer mentions a country while the second one a city, a better d-tree that accommodates this is shown in Figure 3.7. The question "Where was John born" is divided into "Which country?" and "Which city?", and by uniformity, this happens on both branches. As a result, the two entailing answers "John was born in Paris" and "John was born in France" no longer respond to the same QUD, which means distinctness is no longer violated. Hence, we correctly predict the QUD in (27) to be felicitous.

One small complication concerns the leftmost answer node in Figure 3.7. It is enclosed in a pair of parentheses because although the question node that dominates it is not explicitly

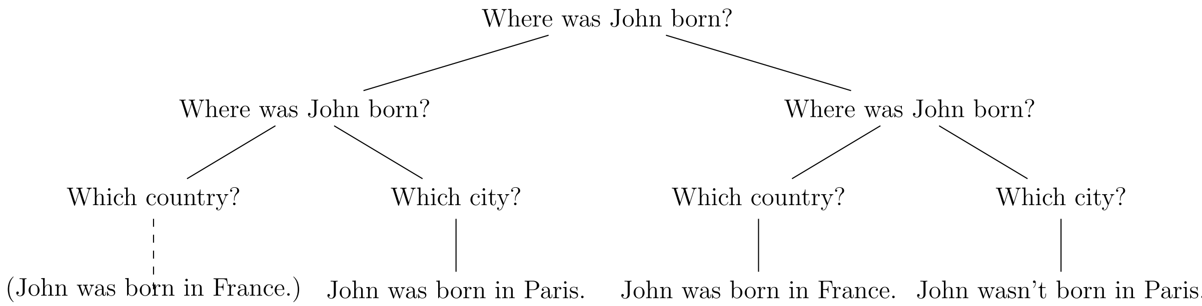

Figure 3.7: A d-tree for (27). The "John was born in France" on the left is enclosed in parentheses as this answer is not explicitly given but only contextually entailed by "John was born in Paris".

answered by the first disjunct of (27), an answer to this question, namely that John was born in France, is entailed by "John was born in Paris". We could have alternatively left the "Which country?" node on the left unanswered without affecting the core prediction of the current analysis, but enclosing a contextually entailed yet not explicitly mentioned answer in parentheses enables us to maximize the information embodied by d-trees. Also with this change, answer nodes that are not explicitly mentioned do not figure into deciding whether a violation of distinctness has occurred.

It is worth noting that the d-tree in Figure 3.7 is not the only available representation for (27). For example, this tree can be further refined to be made compatible with existing accounts regarding the implicit QUDs introduced by negated sentences. According to Tian et al. (2010, 2016), a simple negated sentence such as "John wasn't born in Paris" assumes a QUD that takes the form of a positive polar question, i.e., "Was John born in Paris?", in absence of other contextual information. By contrast, a cleft negation sentence such as "It is John who wasn't born in Paris" assumes a negative *wh*-question, i.e., "Who wasn't born in Paris". They appeal to this difference to explain why the positive component of a negative sentence appears to be represented only for simple negations but not for cleft negations. The simple negative sentence but not the cleft sentence triggers the representation of the positive component because only the former evokes a QUD that takes the form of a positive polar question.

Assuming that a simple negative sentence invokes a positive polar question as its immediate QUD, we can provide a corresponding d-tree as shown in Figure 3.8. Here, the

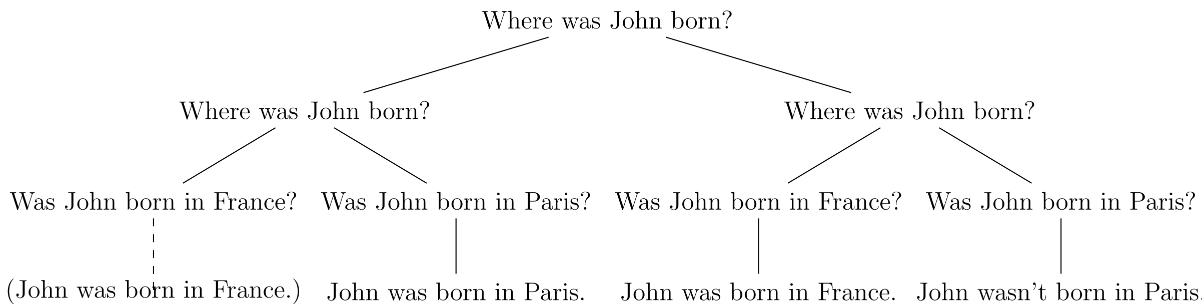

Figure 3.8: A different d-tree for (27) where the QUD that dominates the negated sentence "John wasn't born in Paris" is construed as the polar question "Was John born in Paris".

QUD that dominates "John wasn't born in Paris" is now the polar question "Was John born in Paris?". The resulting d-tree still observes Maximize-Q-Anaphoricity. Additionally, the implicit QUDs are made more specific compared to the ones from Figure 3.7. Despite this difference in details, since "John was born in Paris" and "John was born in France" still respond to two different QUDs, this new d-tree also predicts (27) to be felicitous.

Now, allowing the superquestion "Where was John born?" to be divided into the subquestions "Which country?" and "Which city?" as in Figure 3.7 raises an additional problem. If such an inquiry strategy is viable, then why doesn't the vanilla HD "John was born in Paris or he was born in France" in (1) evoke this strategy? After all, if we were to adopt this strategy and construct a tree accordingly as in Figure 3.9, we would incorrectly predict (1) to be felicitous since distinctness would no longer be violated as "John was born in Paris" and "John was born in France" now respond to different questions.

In order to rule out the d-tree in Figure 3.9 as an admissible representation for (1), I

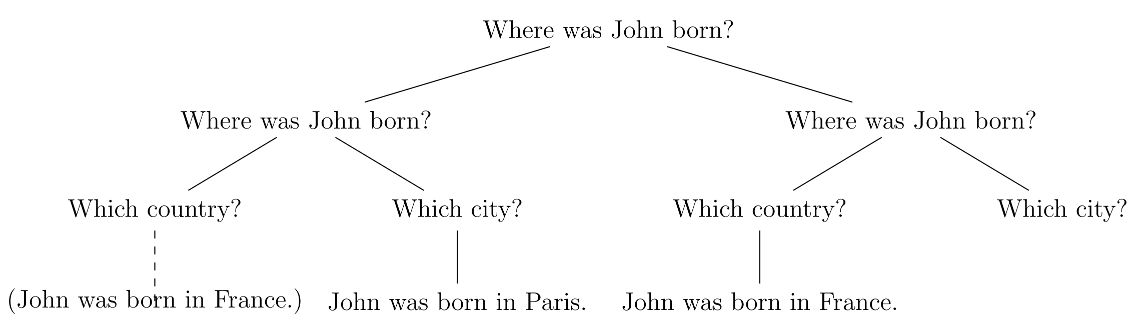

Figure 3.9: A d-tree for (1) which would incorrectly render the Hurford disjunction felicitous.

postulate the following constraint:

> **Q-Quantity:** Implicit QUDs should make the answer(s) appear as informative as possible. In other words, a d-tree that has fewer unanswered questions is preferable to one that has more.

Q-quantity is reminiscent of Grice's (1975) maxim of quantity but is now applied to implicit QUDs. The maxim of quantity (or to be more precise, its first submaxim) demands conversational participants be as informative as possible. But unlike in Grice's case, given that we do not begin with an already fixed QUD, we do not infer based on what is not said in addressing some conversational goal to the conclusion that the speaker is not in a position to assert it. Rather, according to Q-quantity, assuming that the speaker is being maximally informative, when we try to select a QUD-tree among its competing alternatives, we will take the lack of an answer to a particular question to indicate that the question is not part of the current QUDs. For example, an utterance of "John was at the party last night" in an out of the blue context will invoke "Was John at the party last night?" rather than "Was John at the party, and was Mary at the party last night?" as its implicit QUD. The lack of an answer to Mary's presence at the party is used to signify that the corresponding question is not part of the current QUD.

This is not to say that any d-trees that contain an unanswered question are automatically discarded. Q-quantity helps to select an optimal d-tree from a set of candidates only after other constraints on d-trees such as those imposed by contrastive topics, negations, and other discourse markers have been taken into account. For example, the d-tree for "John was born in Paris, or he was born in France and Mary was born in London" in Figure 3.4 contains an unanswered question given that the first disjunct does not provide an answer to "Where was Mary born?". Nonetheless, this d-tree is deemed admissible since the existence of the unanswered question is largely due to the use of contrastive topics. By contrast, since the d-tree in Figure 3.9 has a proper competitor, namely the d-tree in Figure 3.1, which does not contain any unanswered questions, it is dispreferred.

On the other hand, what indeed are ruled out automatically are cases where the reconstruction of implicit QUDs includes a question that is left completely unaddressed. To use an earlier example, with respect to "John was at the party last night", the QUD "Was John at the party, and was Mary at the party last night?" leaves the subquestion "Was Mary at the party last night?" completely unaddressed. Hence, this QUD is automatically eliminated. By contrast, the question "Where was Mary born?" in Figure 3.4 does receive an answer—i.e., "Mary was born in London"—on the right branch. Hence, there is no hard violation of Q-quantity. For a ban against this type of hard violation of Q-quantity, it is crucial that we are only concerned with the reconstruction of implicit QUDs. Q-quantity is by no means intended as a general constraint on explicit discourse questions, since otherwise we would be prohibited from constructing d-trees where the interlocutor is unable to provide an answer to some discourse question.

One final comment regarding the two constraints discussed in this section: We can attempt to unify Maximize-Q-anaphoricity and Q-quantity since they can be viewed as placing an upper and lower bound on how specific implicit QUDs should be:

Q-Specificity: Implicit QUDs should be maximally specific (as per Maximize-Q-anaphoricity) but not overly specific so that they are left unanswered (as per Q-quantity).

The reconstruction of implicit QUDs can then be viewed as a matter of finding the optimal candidate that manages to balance conflicting constraints. As such, future research may explore the prospect of further theorizing QUD-reconstruction from the perspective of Optimality Theory (Prince & Smolensky, 1993; Blutner, 2000, 2013).

3.3.3 Markers of Alternative Strategies of Inquiry

In §3.2.4, I argued in favor of a uniform analysis under which the apparent disjunction appearing in "or at least" retains its disjunctive force. I shall provide an analysis along this line here. I submit that scalar particles such as "at least" and "even" can repair an otherwise infelicitous HD because they function to evoke alternative inquiry strategies so that the resulting d-trees no longer violate the distinctness constraint.

For concreteness, let us focus on "at least". Although the exact lexical entry for "at least" is still up for debate and I do not intend to settle the issue here, several recent proposals have argued for the idea that "at least" invokes an ordering on the set of alternative answers to some current QUD, ranked by the pragmatic strength of the answers (Coppock & Beaver, 2011, 2013; Biezma, 2013). For instance, consider the following lexical entry taken from Coppock and Beaver (2011):

- $[\![at\ least]\!] = \lambda p.\lambda w.\exists p' \in CQ_S[p'(w) = 1 \wedge p' \geq_S p]$.

In words, "at least p" states that there exists an proposition p' in the set of answers to the current QUD (abbreviated as CQ) such that p' is true and is ranked not lower than p in terms of pragmatic strength. Pragmatic strength subsumes semantic entailment but also goes beyond it. For example, consider (32), uttered as a response to an inquiry about the academic position that John holds:

(32) John is at least an assistant professor.

The set of answers to the current QUD includes propositions such as "John is an assistant professor" and "John is an associate professor", and they are ordered by pragmatic strength such that "John is an associate professor" is ranked higher than "John is an assistant professor" even though the former does not entail the latter. Given such an ordering, (32) is true just in case either "John is an assistant professor" is true or a higher-ranked alternative is true.

Now, when trying to explain the repairing effect of "at least" in HDs, if we were to similarly treat the ordering induced by "at least" as placed on the set of answers to the current QUD, we would run into a problem. Consider again the use of "at least" in (33A):

(33) Q: Where was John born?

A: He was born in Paris or at least in France.

What are the stronger alternatives to the prejacent of "at least" in (33A)? Since the prejacent of "at least" is "John was born in France", one could plausibly take the stronger alternatives as consisting of answers such as "John was born in Paris", "John was born in Marseille", and so on. This would in turn force "John was born in France" and "John was born in Paris" to respond to the same QUD. But this is exactly what renders the standard HD in (1) infelicitous in the first place. In other words, if "at least" requires the two disjuncts in (33A) to be viewed as addressing the same QUD, then the reconstructed d-tree for (33A) will inevitably violate distinctness, and as a result we will incorrectly judge it as infelicitous. On the other hand, if we choose not to treat answers such as "John was born in Paris" and "John was born in Marseille" as stronger alternatives to "John was born in France", it becomes rather unclear what other options remain.

In response, I propose that the use of "at least" in (33A) does not invoke an ordering on the set of alternative answers to the current QUD addressed by the prejacent but instead on a set of alternative QUDs among which the current QUD belongs. More specifically, assuming that the set of alternative QUDs is not vacuous in the sense that it contains at least one QUD that is different from the current QUD, we can treat "at least" as signaling the existence of at least one alternative QUD in the d-tree such that it is a sister of the current QUD and is pragmatically stronger than it.

The use of "at least" in (33A) signifies that the QUD answered by the prejacent "John was born in France" has a sister node that is stronger. This requirement rules out the d-tree depicted back in Figure 3.1 where the question "Where was John born?" on the right branch immediately dominates the answer "John was born in France" which is the prejacent of "at least". Instead, the question "Where was John born?" on the right should be further divided into (at least) two subquestions such that one is pragmatically stronger than the other. One straightforward way to cash out strength for the current purpose is simply to use the entailment relation. For example, consider the d-tree in Figure 3.10. Assuming that the name of the country where a given city is located is common ground knowledge, then the question "Which city?" contextually entails "Which country?" since an answer to the first question also answers the second question. Hence, "Which city?" serves as a stronger alternative QUD to "Which country?", thereby satisfying the demand from the use of "at

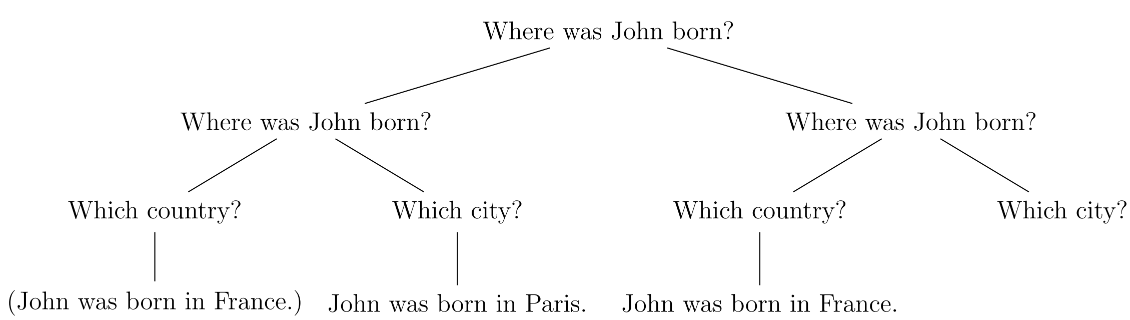

Figure 3.10: A d-tree for (33A). The presence of "at least" signals that the QUD "Which country?" on the right branch has "Which city?" as a stronger sister.

least”. And by uniformity, the same QUD-tree is duplicated on the left. As a result, “John was born in Paris” and “John was born in France” now respond to two different QUDs. Since there is no violation of distinctness, (33A) is judged as felicitous.

Note that the d-tree in Figure 3.10 is in fact identical to the one from Figure 3.9. Recall that, in §3.3.2, this particular tree is ruled out as an optimal representation for the HD in (1) because there exists a competing d-tree that does not contain any unanswered question, namely the one in Figure 3.1 where “John was born in France” on the right branch is immediately dominated by the discourse question “Where was John born?”. Thus, by Q-quantity, the d-tree in Figure 3.9 is dispreferred. The situation is different here. The addition of “at least” signals the need for a more complex inquiry strategy which eliminates the simpler d-tree as a potential competitor. In addition, although the above d-tree contains an unanswered question on the right branch, the same question (i.e., “Which city?”) does receive an answer on the left branch. Hence, there is no hard violation of Q-quantity. As a result, the d-tree in Figure 3.10 now becomes admissible with respect to (33A).

If the current analysis is on the right track, then we should expect the repairing effect of “at least” to more or less diminish in cases where it is harder to find a stronger sister of the QUD that dominates the prejacent of ”at least”. For instance, the answer in (34A) still sounds odd even with the presence of “at least”.

(34) Q: Where was John born?

A: ??He was born in France, or at least he was born in France or Germany.

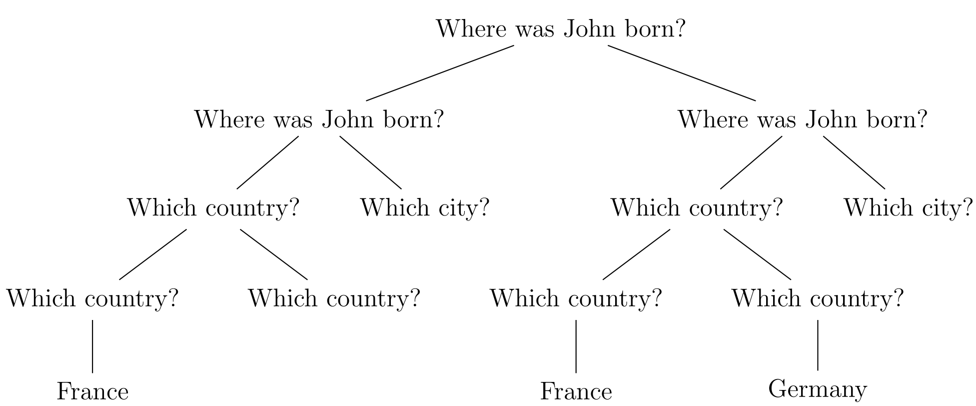

Figure 3.11: A potential d-tree for (34A) which is blocked by Q-quantity since “Which city?” does not receive an answer on either branch.

What we want to know is whether we can generate an admissible d-tree for (34A) such that the QUD that dominates the prejacent of “at least” has a stronger sister. One candidate is given in Figure 3.11, where just like in Figure 3.10, the discourse question is split into “Which country?” and “Which city?”. The “Which country?” on the right is further divided into two questions each dominating a disjunct of the prejacent of “at least” in (34A). However, unlike in Figure 3.10 where the first disjunct of (33A), namely “John was born in Paris”, supplies an answer to “Which city?”, this question is left completely unaddressed in Figure 3.11: neither of the two nodes representing this question receives an answer. As such, the d-tree involves a hard violation of Q-quantity and thus is considered inadmissible.

By contrast, the above answer does appear to fit better as a response to the following alternative question:

(35) Q: Was John born in France, Germany, or Belgium?
A: He was born in France, or at least he was born in France or Germany.

Figure 3.12 shows a possible d-tree for (35A). The QUD that dominates the prejacent of “at least” is conceived of as the polar question “Was John born in France or Germany?” this time around. As such, the question “Which country?” constitutes a stronger sister of it, since a complete answer to the former question also provides an answer to the latter. Moreover, since none of the questions in Figure 3.12 is left completely unaddressed, there is no hard violation of Q-quantity. And as there is no violation of distinctness in Figure 3.12, we predict (35A) to be an unmarked answer to (35Q).

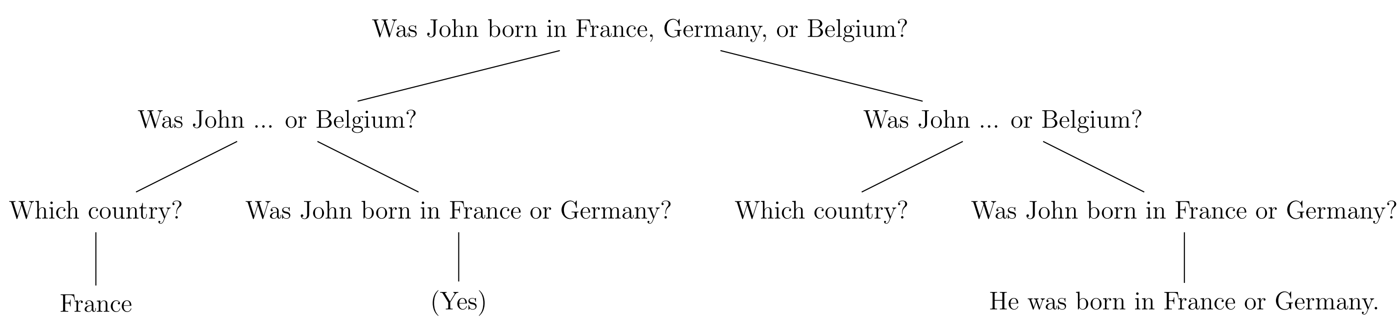

Figure 3.12: A d-tree for (35A). The prejacent of “at least” on the right branch is dominated by the polar question “Was John born in France or Germany?” which has “Which country?” as its stronger sister. (A positive answer to the polar question on the left branch is entailed by the disjunct “He was born in France”).

To recapitulate the difference between (34) and (35), the reason why (35A) feels more natural than (34A) is that given the discourse question in (35Q), it is easier to construct a QUD like "Was John born in France or Germany?". When the discourse question is simply "Where was John born?", it is hard to envision why such a polar question would figure into a reasonable inquiry strategy. By contrast, when the discourse question has already mentioned "France", "Germany", and "Belgium" as three candidate answers, the question "Was John born in France or Germany?" can be conceived of as a sensible way to narrow down and develop the original discourse question.

Besides "at least", we can explain the repairing effect of "even" along the same lines. A common recipe for the lexical analysis of "even" is an ordering on the set of focus alternatives to the prejacent of "even" (see, e.g., Rooth 1992; Chierchia 2013; Greenberg 2016). Roughly, "even" presupposes that its prejacent should be stronger than any of its alternatives. Since the focus alternatives of a proposition go hand in hand with the QUD addressed by it, we can provide a similar analysis by treating "even" as inducing an inquiry strategy which requires that the QUD immediately dominates the prejacent of "even" has a sister that is weaker than it (see Figure 3.13). Just as in the case of "at least", this more complex inquiry strategy evoked by the use of "even" yields a d-tree where "John was born in France" and "John was born in Paris" respond to two different QUDs, thereby capturing the repairing effect of "even".

Lastly, it is worth noting that not all scalar particles seem capable of fixing HDs. For example, the following disjunction still sounds odd with the addition of "at most":

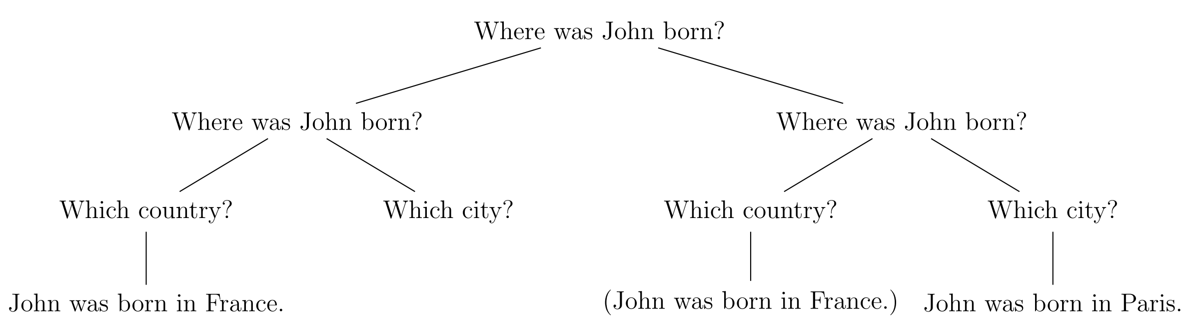

Figure 3.13: A d-tree for "John was born in France or even in Paris". The QUD "Which city?" on the right branch which dominates the prejacent of "even" has "Which country?" as its weaker sister.

(36) #John was born in France or at most in Paris.

This seems to indicate that scalar particles such as "at least" and "even" have a special status. For one thing, it has been suggested that "at least" is indeed associated with multiple distinct lexical entries (Kay, 1992; Nakanishi & Rullmann, 2009; Cohen & Krifka, 2011; Coppock & Brochhagen, 2013). On the other hand, there have also been attempts to provide a uniform lexical entry for "at least" (see, e.g., Biezma, 2013). How the current observation figures into this debate is a question to be explored in future research.

3.4 Open Issue

One remaining issue I want to briefly highlight concerns Hurford questions. It has been observed that Hurford disjunctions can also take the form of an interrogative sentence (Ciardelli & Roelofsen, 2017). For instance, just as its declarative counterpart, the alternative question in (37) is deemed infelicitous:[5]

(37) #Was John born in Paris↑, or was he born in France↓?

Ciardelli and Roelofsen (2017) combine a molecular redundancy account with inquisitive semantics to explain why questions such as (37) are infelicitous. Since (37) introduces two alternatives such that one entails the other, the alternative question is defective. However, their account is unable to account for the interrogative variants of conjunctive Hurford disjunctions. For instance, since the two alternatives introduced by (38) are non-entailing, resorting to inquisitive semantics alone cannot explain its oddness.

(38) #Was John born in Paris↑, or was he born in France and Mary in London↓?

Neither does my current analysis offer an immediate solution to the puzzle raised by Hurford questions. We could, as a makeshift solution, try to account for the data by examining the d-tree associated with the assertion that corresponds to the sentence-radical of a given question. For example, under inquisitive semantics, the alternative question in (38) shares the same sentence-radical as the assertion "John was born in Paris, or he was born in France and Mary in London". As such, we can explain why (38) is bad by appealing to the same d-tree that explain why the matching CHD is bad.

[5]The upward and downward arrows are used to indicate intonation.

Additionally, by taking a simple negative sentence as assuming a positive polar question as its QUD, we can also explain why Hurford disjunctions are infelicitous when they are embedded under negation in a similar fashion. For example, the negation in (39) is perceived as odd because the question it allegedly answers in (40) is defective in the first place.

(39) John wasn't born in Paris or in France.

(40) #Was John born in Paris or in France↑?

Although this roundabout strategy yields the right predictions, a more preferable approach would be to give a theory of how implicit QUDs can be reconstructed from explicit questions. This would enable us to generate a d-tree for the alternative question in (38) in a more straightforward manner. Such an account is worth exploring in the future.

3.5 Conclusion

This paper presents a QUD-based analysis of informational redundancy that is capable of capturing the (in)felicity of a wide range of Hurford disjunctions. In particular, it readily predicts the infelicity of conjunctive Hurford disjunctions, which is left unaccounted for under existing redundancy theories. To accomplish this, I make use of a notion of discourse trees and postulate several constraints governing their structures. More specifically, I propose that the disjuncts of a disjunction should evoke the same inquiry strategy and that answers to the same question in a d-tree should be distinct and non-entailing; additionally, the implicit QUDs reconstructed from a target sentence should be maximally but nor overly specific. I have further illustrated how this analysis can be extended to elucidate the repairing effect of scalar particles such as "at least", thereby enabling a uniform analysis of disjunction in discourse.

Chapter 4

A Non-Bivalent Approach to Inquisitive Logic

In this paper, I explore a non-bivalent approach to inquisitiveness and develop a new logic called *the logic of pseudo-complemented propositions* (LPP). Standard inquisitive logic, which takes *support* as its central notion, rejects the Law of Excluded Middle (LEM) by employing an intuitionistic negation. Under the current approach, I take *support* and *rejection* as two separate central notions and reject LEM on the ground that the two central notions, that of support and that of rejection, do not fully complement each other and thus allow gaps between them. As such, LPP incorporates a negation that does not validate LEM but does validate Double Negation Elimination. I supply an algebraic semantics for LPP via the so-called *twist-structures*. I define a new class of twist-structures, call them *pseudo-complemented twist-structures* (PTS), and prove the completeness of LPP with respect to the class of PTS. The utility of this method of constructing twist-structures will be further highlighted via exploration of some generalizations of PTS.

Keywords: Inquisitive logic, Law of Excluded Middle, Constructive logic, Algebraic semantics, Twist-structures

4.1 Introduction

Inquisitive semantics (Groenendijk, 2009; Ciardelli & Roelofsen, 2011; Ciardelli et al., 2015, 2018) offers a unified formal framework for analyzing both declarative and interrogative sentences. In doing so, it provides a new conception of information exchange that goes beyond the exchange of truth-conditional content typically associated with declarative sentences. Under inquisitive semantics, information exchange is construed as a process of raising and resolving issues: issues are raised by asking a question, which can then be subsequently resolved by a declarative sentence. As such, asking a question can be viewed as proposing a set of alternative ways for the common ground information to be updated. For example, a polar question like "Is Alice at the party?" puts forward two alternatives which correspond to the positive and negative answer to this question.

The central component of inquisitive semantics is an alternative-based account of disjunction: disjunction is interpreted as presenting a set of alternatives. For example, we can symbolize the aforementioned polar question using the disjunction $A \vee \neg A$ which introduces two alternative possibilities, namely, one where Alice is at the party and one where she is not. Given this treatment of disjunction, the Law of Excluded Middle (LEM) is no longer valid:

$$\text{(LEM)} \quad \phi \vee \neg\phi$$

The loss of LEM has an intuitive explanation. For a sentence to be valid in inquisitive semantics, it needs to be accepted or evaluated true with respect to every body of information, which means that the sentence must be informationally trivial. But since $\phi \vee \neg\phi$ symbolizes a question in inquisitive semantics, the sentence is not completely trivial as it raises the issue of whether ϕ, thereby making a significant semantic contribution in terms of its update effect on the conversational context. Given this, LEM is rejected as a validity in inquisitive semantics.

Now, if we assume a standard semantic clause for disjunction, that is, a disjunction is evaluated true whenever one of its disjuncts is evaluated true, then we can choose between two broad strategies to reject LEM: either we abandon *valuation bivalence* and adopt a non-bivalent valuation function, or we stick to a bivalent valuation function but instead resort to an account of negation (e.g., intuitionistic negation) that does not vindicate LEM.

To elaborate, let us assume a possible world semantics with the following satisfaction clause for disjunction: $w \Vdash_1 \phi \vee \psi$ iff $w \Vdash_1 \phi$ or $w \Vdash_1 \psi$. That is, $\phi \vee \psi$ is assigned the value 1 at a world w iff either ϕ is assigned 1 at w or ψ is assigned 1 at w. (We shall come back to the exact interpretation of the value 1 shortly.) To invalidate LEM, we need to make it happen that $w \nVdash_1 \phi$ and $w \nVdash_1 \neg\phi$. Here are the two general strategies mentioned above:

Non-bivalent valuation: $w \nVdash_1 \phi \not\equiv w \Vdash_0 \phi$; but $w \Vdash_0 \phi \equiv w \Vdash_1 \neg\phi$.

Non-classical negation: $w \nVdash_1 \phi \equiv w \Vdash_0 \phi$; but $w \Vdash_0 \phi \not\equiv w \Vdash_1 \neg\phi$.

Both proposals, by rejecting one of the two equivalences, allow $w \nVdash_1 \phi$ and $w \nVdash_1 \neg\phi$, thereby invalidating LEM, but they achieve this through different means.[1] According to the first approach where valuation bivalence is abandoned, ϕ can receive a value that is neither 1 nor 0 (including the possibility of not receiving a value at all in the case of partial valuation functions); consequently, not assigning 1 to ϕ is not equivalent to assigning 0 to ϕ. The second equivalence $w \Vdash_0 \phi \equiv w \Vdash_1 \neg\phi$ then construes negation as a toggle operation that connects the two values: $\neg\phi$ is assigned 1 at w just in case ϕ is assigned 0 there. As for the other option, whereas valuation bivalence is upheld, assigning 1 to $\neg\phi$ is not equivalent to assigning 0 to ϕ, and thus, in presence of the first equivalence $w \nVdash_1 \phi \equiv w \Vdash_0 \phi$, not equivalent to not assigning 1 to ϕ. As such, this second approach needs to find a suitable non-classical negation that denies the second equivalence.

Whereas the existing inquisitive semantics, laid out in Ciardelli et al. (2018), chooses the second route, I will explore the non-bivalent approach in this paper. Generally speaking, given that there are many reasons why one may want to reject LEM, which of the two approaches one choose largely comes down to one's motivation, and more specifically, how one interprets $\vDash_1$ and $\vDash_0$. This is not to say that given any particular motivation, we must always settle on one approach. In fact, for some (and perhaps many) cases, we can motivate both a bivalent and a non-bivalent approach.

[1]Of course, there could be a third option where we reject both equivalences. Moreover, even if we adopt a non-bivalent approach, we could still incorporate another negation that does not reinforce the second equivalence. But for the time being, let us focus on the two minimalist solutions given above.

To help illuminate this point, let us consider a different case. Take constructive mathematics as an example. According to constructivism, mathematical objects are mental constructs (cf. Iemhoff, 2019). Thus, for any mathematical statement ϕ to be true, there needs to be a mental construct that verifies ϕ via a proof of ϕ. Likewise, for $\neg\phi$ to be true, there needs to be a proof of not-ϕ, which under the traditional Brouwer-Heyting-Kolmogorov-interpretation, is understood as having a way to turn any proof of ϕ into a proof of a manifestly false claim such as $1 = 0$. Given this interpretation, LEM is rejected on the ground that there are certain mathematical statements, e.g., Goldbach's conjecture, that are neither provably true nor provably false.

Now, depending on how one chooses to interpret $\vDash_1$ and $\vDash_0$, we arrive at different methods to formally cash out constructivism. By taking *proof* as the sole central notion—that is, ϕ is assigned 1 at w iff there exists a proof of ϕ at w—and interpreting the other value (i.e., 0) as derivative of that of the first one, namely, as *lack of proof*, standard intuitionistic logic counts as a bivalent account where LEM is rejected via the intuitionistic negation. On the other hand, we can adopt a non-bivalent approach by bringing in another central notion, namely *counterexample* (cf. Nelson, 1949). We can then interpret 0 as standing for the existence of a counterexample. As such, LEM is rejected on the ground that it is not the case that for any mathematical statement ϕ, we can always either construct a proof of ϕ or find a counterexample to ϕ. This change of perspective results in a different logic, Nelson's three-valued logic N3, which incorporates a toggle negation $\sim$ that invalidates LEM while still vindicating Double Negation Elimination (DNE), viz., $\vDash {\sim\sim}\phi \to \phi$.[2] It demonstrates that intuitionistic logic is not the only formalism available for constructivists.

In this paper, I will apply this lesson to inquisitive semantics and explore a non-bivalent counterpart to inquisitive logic, just as N3 can serve as a non-bivalent counterpart to intuitionistic logic. Whereas the existing inquisitive semantics follows a bivalent approach by taking *support* as its sole central notion and opting for an intuitionistic negation to account for the failure of LEM, I entertain a different approach that takes *support* and *rejection* as two central notions that do not necessarily complement each other. I shall leave a more detailed discussion of how the additional notion of rejection can be understood in an inquisitive

[2]To disambiguate, I will reserve $\neg$ for the intuitionistic negation and use $\sim$ for the toggle negation.

setting to §4.3, but here is the basic idea. In standard inquisitive semantics (Ciardelli et al., 2018), supporting a question amounts to possessing enough information to be able to answer the question. Analogously, we can then understand what it means to reject a question as having enough information to determine that the question cannot be felicitously asked in the first place—for example, because the question presupposes something that is already known to be false. Since it is not the case that for any question, we are able to either answer it or decide that it cannot be felicitously asked in the first place, the two notions of support and rejection do not necessarily complement each other.

The resulting logic of the current non-bivalent approach, which I name the *Logic of Pseudo-complemented Propositions* (LPP), resembles N3 in the sense that they both invalidate LEM without simultaneously invalidating DNE. But unlike N3, the rejection conditions of complex formulas are interpreted non-constructively, and unlike N3 but like the standard inquisitive logic (Ciardelli & Roelofsen, 2011), LPP is not closed under uniform substitution. My main objective here is not to argue in favor of LPP over the existing inquisitive logic; rather the aim is to explore new possibilities by showing that the existing bivalent approach (e.g., Ciardelli & Roelofsen, 2011) is not the only option for theorizing about inquisitiveness.

I present an algebraic semantics for LPP via the so-called *twist-structures*. Twist-structures are algebras whose carrier set consists of ordered pairs which are capable of codifying two separate bodies of information—for example, information regarding the support and rejection conditions of a sentence.[3] As such, they fit well with logics that adopt a non-bivalent perspective. This method of constructing twist-structures has been previously used to devise algebraic semantics for N3 as well as other Nelson's logics (see, e.g., Vakarelov, 1977; Kracht, 1998; Odintsov, 2004). A second goal of this paper is thus to further popularize this method by showing that by varying certain parameters of the twist-operation, we can derive a wide range of twist-structures. Among these twist-structures are the pseudo-complemented twist-structures (PTS) which will supply us with an algebraic semantics for LPP.

This paper proceeds as follows. In §4.2, I review the basic inquisitive logic and the

[3]A note on terminology. The name "twist-structure" comes from Kracht (1998); Vakarelov (2005) calls this method "counterexample-construction". But since the current paper explores generalizations of this method beyond constructive mathematics, I will adopt Kracht's terminology.

logic N3. Since I intend to extend existing inquisitive logic in the same fashion as how N3 extends the standard intuitionistic logic—that is, by incorporating a toggle negation—I shall introduce the algebraic semantics for N3 based on basic twist-structures (BTS) to set up the stage for my own use of the twist-construction method. In §4.3, I first motivate a non-bivalent approach to inquisitiveness and then elucidate why the existing logic N3, albeit being non-bivalent, is unsuitable for this purpose. This prompts us to devise a new non-bivalent framework along the lines of N3. In §4.4, I present LPP together with an algebraic semantics for it via PTS and prove completeness. In §4.5, I move on to the aforementioned second objective by exploring several generalizations of the particular twist-operation proposed in §4.4. And §4.6 concludes.

4.2 Preliminaries

4.2.1 Basic Inquisitive Semantics and Inquisitive Logic

Let us begin by defining the language for basic inquisitive logic. The language is given as follows, where p is a member of a countable set of atomic propositional formulas:

Definition 4.2.1 (Logical syntax)**.** $\phi ::= p|\top|\bot|\phi \wedge \phi|\phi \vee \phi|\phi \to \phi|\neg\phi|?\phi$

The last expression $\ulcorner ?\phi \urcorner$ abbreviates $\ulcorner \phi \vee \neg\phi \urcorner$.

As mentioned previously, under inquisitive semantics, a question proposes alternative ways to update common ground information. Each alternative introduced by a question is formally represented as a set of possible worlds that are compatible with the information embodied by the corresponding answer to the question. For example, the alternative "Alice is at the party" can be represented as a set of worlds where Alice is at the party. Since a question introduces multiple alternatives, the propositional content associated with a question is taken to be a set of sets of possible worlds under inquisitive semantics. To obtain uniformity, we can then construe the propositional content of any sentence, be it declarative or interrogative, as a set of sets of possible worlds, that is, as a set of information states. We define information states, propositions, and alternatives as follows:

Definition 4.2.2 (Information states)**.** An information state s is a set of possible worlds and a subset of the total logical space (i.e., the set of all worlds) $s \subseteq W$.

Definition 4.2.3 (Propositions)**.** A proposition P is a non-empty downward closed set of information states, that is, for all information states s and t, if $s \in P$ and $t \subseteq s$, then $t \in P$.

Definition 4.2.4 (Alternatives in a proposition)**.** The maximal elements of a proposition P are the alternatives in P. We use $Alt(P)$ to denote the set of alternatives in P.

An information state is a set of worlds that are compatible with a certain body of information. The proposition expressed by a sentence is then identified with those information states that either already accept the assertion of it if the sentence is declarative, or resolve the question it raises if the sentence is interrogative. And since whenever an information state already contains enough information to accept an assertion or resolve a question, any subset of it, which contains more information by excluding certain worlds, must also be capable of doing so, propositions are construed as downward closed sets of information states.

Figure 4.1 displays some examples of propositions all generated from the same logical space $\{AB, A\bar{B}, \bar{A}B, \bar{A}\bar{B}\}$, where AB stands for a world where Alice and Bob are both at the party, $A\bar{B}$ a world where Alice but not Bob is at the party, and so on. I use $[\![\;]\!]$ to denote propositions, and for readability, I will represent a proposition using its alternatives. For instance, Figure 4.1(a) depicts the proposition $[\![A \vee \neg A]\!]$, which will be abbreviated as $[\![?A]\!]$, and as the figure shows, this proposition is represented by its two alternatives $\{AB, A\bar{B}\}$ and $\{\bar{A}B, \bar{A}\bar{B}\}$. We call a proposition that contains more than one alternative *inquisitive*. $[\![?A]\!]$ is inquisitive: it raises the question of whether or not A, with each alternative representing a way to answer the question, that is, by either affirming or denying A. Similarly, the proposition $[\![A \vee B]\!]$ depicted by 4.1(b) is also inquisitive: it corresponds to the alternative question

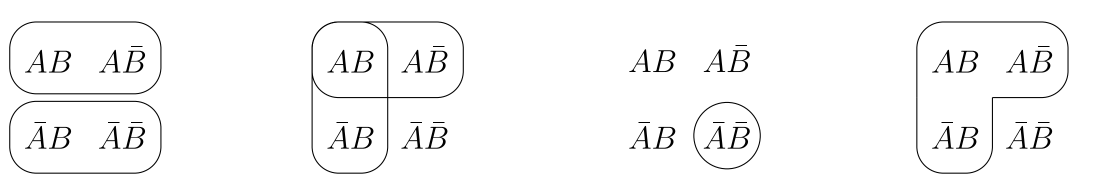

(a): $[\![?A]\!]$ (b): $[\![A \vee B]\!]$ (c): $[\![\neg(A \vee B)]\!]$ (d): $[\![\neg\neg(A \vee B)]\!]$

Figure 4.1: Some examples of propositions

"Is Alice at the party$^{\uparrow}$, or is Bob at the party$^{\downarrow}$?".[4] By contrast, the propositions depicted by 4.1(c) and 4.1(d) are not inquisitive as they contain only one alternative. Additionally, there are two special propositions for any given logical space: the minimal proposition $\top := \wp(W)$—i.e., the power set of the logical space W—and the absurd proposition $\bot := \{\varnothing\}$.

Now, given that propositions under inquisitive semantics are sets of information states, we can apply basic algebraic operations to them. More specifically, the set of all propositions, call it Σ, ordered by the subset relation $\subseteq$ forms a complete Heyting algebra $\langle \Sigma, \cap, \cup, \Rightarrow, \neg, \bot, \top \rangle$, where the set intersection $\cap$ and the set union $\cup$ denote its meet and join, respectively, and $\Rightarrow$ denotes the relative pseudo-complement of this algebra (Roelofsen, 2013; Ciardelli et al., 2018).[5] Propositions are then connected via the following algebraic operations:

$$\begin{aligned}
\top &:= \wp(W);\\
\bot &:= \{\varnothing\};\\
[\![\phi \wedge \psi]\!] &:= [\![\phi]\!] \cap [\![\psi]\!];\\
[\![\phi \vee \psi]\!] &:= [\![\phi]\!] \cup [\![\psi]\!];\\
[\![\phi \rightarrow \psi]\!] &:= [\![\phi]\!] \Rightarrow [\![\psi]\!];\\
[\![\neg\phi]\!] &:= [\![\phi]\!] \Rightarrow \bot.
\end{aligned}$$

Traditionally, sentences are evaluated at worlds, and a sentence ϕ is evaluated true at a world w iff w belongs to the proposition expressed by ϕ, that is, the set of ϕ-worlds. Under inquisitive semantics, sentences are evaluated at information states in terms of *support*. As mentioned earlier, to support a sentence is to accept its assertion when the sentence is declarative or to resolve the issue it raises when the sentence is interrogative. We then define semantic consequence in terms of preservation of support.

Definition 4.2.5 (Support). An information state s supports ϕ, notated as $s \Vdash_1 \phi$, iff

[4] The upward and downward arrow represent raising and falling intonation, respectively. Varying the intonation will also vary the type of questions raised by this sentence (cf. Ciardelli et al., 2018).

[5] We can be more specific. The set of all propositions under inquisitive semantics indeed forms a special type of Heyting algebras that satisfy certain additional properties (see Bezhanishvili et al., 2020, and §5.1 below).

$s \in [\![\phi]\!]$.[6]

Definition 4.2.6 (Consequence via support). $\Gamma \vDash \phi$ iff for all logical spaces W and all $s \subseteq W$, if $s \Vdash_1 \psi$ for every $\psi \in \Gamma$, then $s \Vdash_1 \phi$.

As an essential feature of inquisitive semantics, LEM is invalid. Figure 4.1(a) provides a counterexample to the validity of $?A$, since, given the setup, it is not the case that all information states support $?A$: the information state $\{AB, A\bar{B}, \bar{A}B, \bar{A}\bar{B}\}$ does not, as it does not belong to $[\![?A]\!]$. Also, since inquisitive semantics employs an intuitionistic negation,[7] neither is DNE valid in general: as Figure 4.1(d) and 4.1(b) demonstrate, the information state $\{AB, A\bar{B}, \bar{A}B\}$ belongs to $[\![\neg\neg(A \vee B)]\!]$ but does not belong to $[\![A \vee B]\!]$. That being said, DNE does obtain for atomic propositions.

The above semantics gives rise to inquisitive logic which can be axiomatized by a system that contains the following set of axiom schemas along with the inference rule of *modus ponens* (Ciardelli & Roelofsen, 2011; see also Ciardelli, 2016):

IPL. All axiom schemas of intuitionistic propositional logic.

KP. $(\neg\phi \to (\psi \vee \chi)) \to ((\neg\phi \to \psi) \vee (\neg\phi \to \chi))$.[8]

DNE. $\neg\neg p \to p$, for every atomic proposition p.

As a noteworthy feature of inquisitive logic, since DNE only obtains for atomic propositions, the logic is not closed under uniform substitution.

Lastly, given that propositions under inquisitive semantics form a complete Heyting algebra, we can provide an algebraic semantics for inquisitive logic via Heyting algebra (see Bezhanishvili et al., 2019). Given a Heyting algebra $\mathfrak{H}$, let us first consider a special class of valuation functions that assign every atomic proposition to an element of $\mathcal{H}_{\neg\neg} = \{\neg\neg x : x \in \mathcal{H}\}$; call such valuations *inquisitive valuations*. Since, as Johnstone (1982)

[6] Given the algebraic foundations of inquisitive propositions, this definition coincides with the recursive definition of support conditions (Ciardelli et al., 2018). Additionally, although inquisitive semantics is not originally formulated using Kripke models, we can convert it into a Kripke semantics using the following Kripke frame $\langle \wp(W) - \{\varnothing\}, \supseteq \rangle$ (Ciardelli & Roelofsen, 2011). Hence, we can use the same Kripkean satisfaction relation $\Vdash_1$ for the support relation.

[7] The name "intuitionistic negation" might be a slight misnomer in the inquisitive setting. Nonetheless, due to the lack of a catchier name, I will continue using it to refer to the negation employed in standard inquisitive semantics.

[8] KP stands for "Kreisel-Putnam".

shows, $\mathfrak{H}_{\neg\neg}$ constitutes a Boolean algebra, it follows that for all $y \in \mathcal{H}_{\neg\neg}$, the operation $\neg\neg y \to y$ returns the top element. Now, we say that a formula ϕ is inquisitively valid in $\mathfrak{H}$ iff for every inquisitive valuation v on $\mathcal{H}$: $v(\phi) = 1$; and ϕ is inquisitively valid *simpliciter* iff it is inquisitively valid in every Heyting algebra. Hence, given that atomic propositions are assigned to $\mathcal{H}_{\neg\neg}$, they come to satisfy DNE. Additionally, if we further require that the class of Heyting algebras satisfy the KP axiom, we derive the exact set of validities that are valid in inquisitive logic.

4.2.2 A case study of the non-bivalent approach: N3

In the following two subsections, I will describe the logic N3 in some detail. The reason for investigating N3 is mainly two-fold. On one hand, it helps us to see why N3, albeit being a non-bivalent framework, is unsuitable for the current purpose of modeling inquisitiveness, which I shall further explicate in §4.3.2. On the other hand, since the framework I will eventually propose employs a toggle negation in the spirit of N3 and comes with an algebraic semantics via twist-structures, an examination of N3 is conducive to understanding my proposal.

The logic of N3 originates from considering two different ways of proving the negation of ϕ: we can do so by either applying *reductio ad absurdum* or directly constructing a counterexample of ϕ (Nelson, 1949; see also Kapsner, 2014). The former corresponds to the intuitionistic negation $\neg\phi$, as $\neg\phi$ can be defined as $\phi \to \bot$ which conveys that any proof of ϕ leads to a proof of absurdity. On the other hand, standard intuitionistic logic is unable to capture the latter idea of refuting a sentence by constructing a counterexample, as the same intuitionistic negation cannot fulfill this additional purpose. To elaborate, in intuitionistic logic, $\neg(\phi \wedge \psi)$ does not entail $\neg\phi \vee \neg\psi$; but if we were to interpret negation as expressing the presence of a counterexample, we should expect the entailment to hold, since for us to be able to construct a counterexample of a conjunction, we must be able to construct a counterexample to one of its conjuncts.

The search for an alternative negation leads us to the aforementioned non-bivalent picture where we take *proof* and *counterexample* as two non-complementary central notions and use them to interpret $\vDash_1$ and $\vDash_0$, respectively.

To explain, let us first consider a Kripke semantics for N3. A Kripke model for N3 is a tuple $\langle W, \leq, v\rangle$, where W is a non-empty set of worlds (or information states), $\leq$ is as a partial order (i.e., a binary relation that is reflexive, transitive, and antisymmetric), and v is a partial function from pairs of an atomic proposition and a world to 1 and 0. Given a Kripke model, we define two separate satisfaction relations $\Vdash_1$ and $\Vdash_0$ that extend the valuation v to every formula in the language:

$w \Vdash_1 p$ iff $v(p, w) = 1$, when p is atomic;

$w \Vdash_0 p$ iff $v(p, w) = 0$, when p is atomic;

$w \Vdash_1 \phi \wedge \psi$ iff $w \Vdash_1 \phi$ and $w \Vdash_1 \psi$;

$w \Vdash_0 \phi \wedge \psi$ iff $w \Vdash_0 \phi$ or $w \Vdash_0 \psi$;

$w \Vdash_1 \phi \vee \psi$ iff $w \Vdash_1 \phi$ or $w \Vdash_1 \psi$;

$w \Vdash_0 \phi \vee \psi$ iff $w \Vdash_0 \phi$ and $w \Vdash_0 \psi$;

$w \Vdash_1 \phi \to \psi$ iff $\forall w' \geq w$: if $w' \Vdash_1 \phi$ then $w' \Vdash_1 \psi$;

$w \Vdash_0 \phi \to \psi$ iff $w \Vdash_1 \phi$ and $w \Vdash_0 \psi$;

$w \Vdash_1 {\sim}\phi$ iff $w \Vdash_0 \phi$;

$w \Vdash_0 {\sim}\phi$ iff $w \Vdash_1 \phi$.

As with the Kripke semantics for intuitionistic logic, we interpret worlds in the set W as *proof stages*: $w \leq w'$ just in case w' is a stage no earlier than w in our mathematical inquiry. Since proofs and counterexamples are supposed to be conclusive, we posit the following two heredity conditions:

For all w and w', if $w \Vdash_1 p$, where p is atomic, and $w \leq w'$, then $w' \Vdash_1 p$

For all w and w', if $w \Vdash_0 p$, where p is atomic, and $w \leq w'$, then $w' \Vdash_0 p$

That is, if p is either proved (viz., evaluated as 1) or refuted by a counterexample (viz., evaluated as 0) at an earlier stage w, then it must remain proved or refuted and receive the same value at any later stage w'.[9]

[9] The heredity constraints on atomic propositions together with the above satisfaction clauses for complex formulas ensure that the heredity conditions hold for all formulas (see Priest, 2001).

Note that since the valuation functions are allowed to be partial, p may receive neither 1 nor 0 at w, and in that case, it means that, at w, we can neither prove p nor construct a counterexample of p. The satisfaction conditions for complex formulas are defined separately for 1 and 0. The former very much resemble the satisfaction clauses under intuitionistic logic. In particular, the satisfaction clause for implication states that for there to be a proof of $\phi \to \psi$ at w, we need to be able to convert any proof of ϕ into a proof of ψ at any stage w'.

As for the satisfaction conditions for the other value 0, connectives are again defined constructively. For $\phi \wedge \psi$ to receive 0 at w—that is, for there to be a counterexample of it at w—either ϕ must receive 0 at w or ψ must receive 0 at w—that is, we must be able to construct a counterexample to at least one of the conjuncts. And for there to be a counterexample of $\phi \to \psi$ at w, we must be able to obtain, simultaneously, a proof of ϕ and a counterexample of ψ at w.

In N3, the negation ~ toggles between the satisfaction conditions associated with the two values. For $\sim\phi$ to receive 1 at w—that is, for the negation of ϕ to be provable at w—ϕ needs to receive 0 at w—that is, we need to be able to construct a counterexample of ϕ at w, and vice versa.

Lastly, we define semantic consequence in terms of preservation of the value 1 at every world in every model:

$\Gamma \vDash \phi$ iff for every model $\mathcal{M}$ and every world w, if $w \Vdash_1 \psi$ for every $\psi \in \Gamma$, then $w \Vdash_1 \phi$. And in the case Γ is empty, that is, $\vDash \phi$, we say ϕ is valid in N3.

A complete axiomatization of N3 is provided below, which contains the following axiom schemas complemented with the inference rule *modus ponens*. The biconditional $\phi \leftrightarrow \psi$ is defined as $(\phi \to \psi) \wedge (\psi \to \phi)$ as usual.

AS1. $\phi \to (\psi \to \phi)$

AS2. $(\phi \to (\psi \to \chi)) \to ((\phi \to \psi) \to (\phi \to \chi))$

AS3. $\phi \to (\phi \vee \psi)$

AS4. $\psi \to (\phi \vee \psi)$

AS5. $(\phi \to \chi) \to ((\psi \to \chi) \to (\phi \vee \psi \to \chi))$

AS6. $(\phi \wedge \psi) \to \phi$

AS7. $(\phi \wedge \psi) \to \psi$

AS8. $\phi \to (\psi \to (\phi \wedge \psi))$

AS9. $\sim\sim\phi \leftrightarrow \phi$

AS10. $\sim\phi \to (\phi \to \psi)$

AS11. $\sim(\phi \vee \psi) \leftrightarrow (\sim\phi \wedge \sim\psi)$

AS12. $\sim(\phi \wedge \psi) \leftrightarrow (\sim\phi \vee \sim\psi)$

AS13. $\sim(\phi \to \psi) \leftrightarrow (\phi \wedge \sim\psi)$

The first eight axiom schemas come from the positive fragment of intuitionistic logic. AS9 specifies how proof-construction and counterexample-construction are connected via toggle negation, that is, constructing a counterexample of $\sim\phi$ essentially amounts to having a proof of ϕ. Hence, unlike intuitionistic negation, toggle negation vindicates DNE. AS11–13 capture the constructiveness of counterexample-construction. Lastly, AS10, the principle of *ex falso quodlibet*, ensures that a proof of any contradiction leads to explosion.[10] Without it, the resulting logic becomes Nelson's paraconsistent logic N4. Additionally, given the presence of *ex falso*, we can define intuitionistic negation $\neg\phi$ in N3 as $\phi \to \sim\phi$.[11] Defined as such, the intuitionistic negation $\neg\phi$ is entailed by the toggle negation $\sim\phi$.[12] This means that in N3, we can capture the two distinct conceptions of negation delineated above, that is, negation via proving *reductio* and negation via constructing a counterexample.

4.2.3 An algebraic semantics for N3 via twist-structures

An algebraic semantic for N3 has been provided by Rasiowa (1958, 1974) using the so-called N-lattices (i.e., quasi-pseudo-Boolean algebras). An important method to generate

[10]To disambiguate, I shall use *ex falso quodlibet* to refer to the principle of explosion stated in the axiom form while using *ex contradictione quodlibet* to refer to the inference $\sim\phi, \phi \vDash \psi$. The two are nonetheless identical in presence of the deduction theorem.

[11]To spell out, we first define $\neg\phi$ as $\phi \to \bot$ as usual. In N3, $\bot$ can be defined as $\phi \wedge \sim\phi$, which means that $\neg\phi$ can be defined as $\phi \to (\phi \wedge \sim\phi)$. Given that we have $\phi \to \phi$ as a theorem, $\neg\phi$ can simply be defined as $\phi \to \sim\phi$.

[12]Due to the fact that $\sim\phi$ entails $\neg\phi$ in N3, the toggle negation $\sim$ is also often termed "strong negation" (e.g., Rasiowa, 1958; Vakarelov, 1977); another name for this negation is "constructive negation" (e.g., Priest, 2001). However, since neither the entailment to intuitionistic negation nor constructiveness are necessary features of $\sim$ (see, e.g., Vakarelov, 2005), I shall stick to the name of "toggle negation" as used by Kapsner (2014).

N-lattices is via the so-called twist-structures (Vakarelov, 1977; Kracht, 1998), as twist-structures provide a much more intuitive account of the meaning of the logical operators employed in N-lattices, which, as we shall see, is closely related to the interpretation of these operators given by the above Kripke semantics.

Since this paper will explore different generalizations of twist-structures, I shall call the particular kind of twist-structures proposed in Vakarelov (1977) *basic twist-structures* (BTS). Given a Heyting algebra (i.e., an implicative lattice with a distinguished bottom element)[13] $\mathfrak{H} = \langle \mathcal{H}, \wedge, \vee, \rightarrow, \top, \bot \rangle$, we define a BTS as $\mathfrak{H}^{\bowtie} = \langle \mathcal{H}^{\bowtie}, \wedge^{\bowtie}, \vee^{\bowtie}, \rightarrow^{\bowtie}, \sim, 1, 0 \rangle$. Its carrier set is:

$$\mathcal{H}^{\bowtie} = \{\langle x, x' \rangle \in \mathcal{H} \times \mathcal{H} |\ x \wedge x' = \bot\}$$

And we define the operators as follows:

$$\begin{aligned}
\langle x, x' \rangle \wedge^{\bowtie} \langle y, y' \rangle &:= \langle x \wedge y, x' \vee y' \rangle \\
\langle x, x' \rangle \vee^{\bowtie} \langle y, y' \rangle &:= \langle x \vee y, x' \wedge y' \rangle \\
\langle x, x' \rangle \rightarrow^{\bowtie} \langle y, y' \rangle &:= \langle x \rightarrow y, x \wedge y' \rangle \\
{\sim}\langle x, x' \rangle &:= \langle x', x \rangle \\
1 &:= \langle \top, \bot \rangle \\
0 &:= \langle \bot, \top \rangle
\end{aligned}$$

Alternatively, we can also construct the BTS of a Heyting algebra $\mathfrak{H}$ induced by the ordering $\leq$ by first generating the product lattice $\mathfrak{H} \times \mathfrak{H}$ with the following new ordering: $\langle a, b \rangle \leq \langle c, d \rangle$ iff $a \leq c$ and $d \leq b$, and then restricting the carrier set to $\mathcal{H}^{\bowtie}$. It can then be easily shown that the resulting algebra is closed under the set of operations $\{\wedge^{\bowtie}, \vee^{\bowtie}, \rightarrow^{\bowtie}, \sim\}$.

To adduce a concrete example, consider the BTS in Figure 4.2(b), which is generated from the Heyting algebra $\langle \{a, b, c\}, \wedge, \vee, \rightarrow, \top, \bot \rangle$ depicted in 2(a). For every pair $\langle x, y \rangle$ in the resulting BTS, $x \wedge y$ is always the bottom element in the original Heyting algebra. The

[13] An implicative lattice is a distributive lattice ordered by "$\leq$"such that for any two elements x and y, there exists a unique greatest element $x \rightarrow y$ such that $(x \rightarrow y) \wedge x \leq y$. We call $x \rightarrow y$ the pseudo-complement of x relative to y. And when the implicative lattice comes with a bottom element, we call $x \rightarrow \bot$ the pseudo-complement of x.

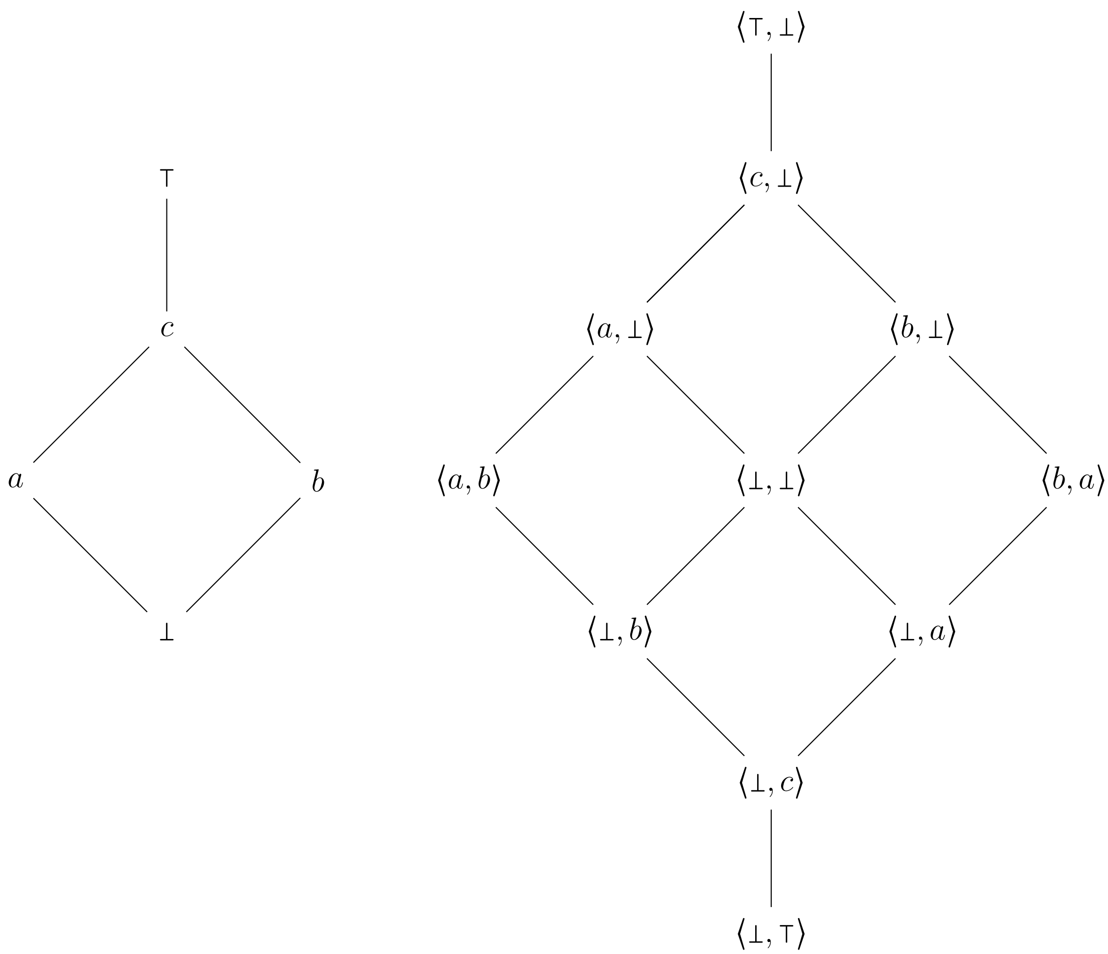

(a) (b)

Figure 4.2: An example showcasing twist-construction

new meet, join, and weak relative pseudo-complement in the resulting BTS are given by $\wedge^{\bowtie}, \vee^{\bowtie}, and \rightarrow^{\bowtie}$, respectively.

A valuation v in a twist-structure is defined in the standard algebraic way, that is, as a function from formulas to elements in the carrier of the twist-structure while observing the definitions of all the logical operators for all complex formulas.

Given this, twist-structures provide an intuitive algebraic semantics for N3. Each pair from the carrier of BTS can be construed as codifying positive and negative information about the formula assigned to the given pair.[14] In the present case, the positive information

[14]This interpretation of such order pairs is also common in the literature on bilattices (see, e.g., Fitting, 1989; Gargov, 1999).

is the formula's provability condition, and the negative information is its disprovability condition. The definitions of the logical operators, which go hand in hand with the definitions provided via Kripke semantics, then tell us how the provability and disprovability conditions of complex formulas are derived from those of their constituents. For instance, let ϕ denote the pair $\langle x, x' \rangle$ and ψ denote the pair $\langle y, y' \rangle$. The pair denoted by the conditional formula $\phi \rightarrow \psi$ is $\langle x \rightarrow y, x \wedge y' \rangle$. The first coordinate $x \rightarrow y$ tells us what it means to prove this formula, that is, by being able to convert any positive information about ϕ (i.e., a proof of ϕ) into positive information about ψ (i.e., a proof of ψ); the second coordinate $x \wedge y'$ tells us what it means to disprove it, that is, by being able to obtain both positive information about ϕ (i.e., a proof of ϕ) and negative information about ψ (i.e., a counterexample of ψ).

Lastly, validity and consequence are defined equationally in the usual way. For any given algebra $\mathfrak{A}$, ϕ is valid in $\mathfrak{A}$, notated as $\vDash_{\mathfrak{A}} \phi$, iff for all valuations v: $\mathscr{F} \rightarrow \mathcal{A}$, $v(\phi) = 1$. And for any given class of algebras $\mathbf{A}$, $\vDash_{\mathbf{A}} \phi$ iff $\vDash_{\mathfrak{A}} \phi$ for every $\mathfrak{A}$ in $\mathbf{A}$.[15]

For any given algebra $\mathfrak{A}$, ϕ is a semantic consequence of a set of sentences Γ in $\mathfrak{A}$, notated as $\Gamma \vDash_{\mathfrak{A}} \phi$, iff for all valuations v: $\mathscr{F} \rightarrow \mathcal{A}$, if $v(\psi) = 1$ for every $\psi \in \Gamma$, then $v(\phi) = 1$. And for any given class of algebras $\mathbf{A}$, $\Gamma \vDash_{\mathbf{A}} \phi$ iff $\Gamma \vDash_{\mathfrak{A}} \phi$ for every $\mathfrak{A}$ in $\mathbf{A}$.

A distinguishing feature of N3 is that $\phi \leftrightarrow \psi$ alone does not define a congruence relation (i.e., an equivalence relation that is preserved under all logical operations). Here is a counterexample using the BTS depicted in Figure 4.2(b). Put $v(\phi) = \langle b, \bot \rangle$ and $v(\psi) = \langle b, a \rangle$. Given this, $v(\phi \leftrightarrow \psi) = v((\phi \rightarrow \psi) \wedge (\psi \rightarrow \phi)) = (\langle b, \bot \rangle \rightarrow^{\bowtie} \langle b, a \rangle) \wedge^{\bowtie} (\langle b, a \rangle \rightarrow^{\bowtie} \langle b, \bot \rangle) = \langle \top, \bot \rangle$. However, $v({\sim}\phi \leftrightarrow {\sim}\psi) = v(({\sim}\phi \rightarrow {\sim}\psi) \wedge ({\sim}\psi \rightarrow {\sim}\phi)) = (\langle \bot, b \rangle \rightarrow^{\bowtie} \langle a, b \rangle) \wedge^{\bowtie} (\langle a, b \rangle \rightarrow^{\bowtie} \langle \bot, b \rangle) = \langle b, \bot \rangle \neq \langle \top, \bot \rangle$. Instead, in order for ϕ and ψ to be considered congruent, both $\phi \leftrightarrow \psi$ and ${\sim}\phi \leftrightarrow {\sim}\psi$ should obtain. As we shall see, this feature of N3 is also shared by the logic LPP to be developed in this paper.

[15] Hereafter, when the relative class of algebras is clear, I shall omit the subscript "$\mathbf{A}$".

4.3 A Non-Bivalent Approach to Inquisitiveness

4.3.1 Support and Rejection

To develop an inquisitive framework that rejects LEM from a non-bivalent perspective, first we need to spell out what the dual notion of support "$\Vdash_1$" is. There is one immediate candidate for the interpretation of "$\Vdash_0$", namely, the notion of *rejection*.

If ϕ is a declarative sentence, then for an information state s to reject ϕ is for it to already contain enough information to settle the sentence false, just as for s to support ϕ is for it to already contain enough information to settle the sentence true. Since s may not contain enough information to either settle ϕ true or settle it false, the equivalence $s \nVdash_1 \phi \equiv s \Vdash_0 \phi$ does not hold. And given that the negation $\sim$ is now interpreted as a toggle operation between the support and rejection conditions of a sentence—that is, to support $\sim\phi$ is to reject ϕ, and vice versa—LEM is rejected for the very same reason.

On the other hand, when ϕ is an interrogative sentence, one natural way to interpret the rejection of ϕ by an information state s is to consider ϕ as making a presupposition that is already settled false in s. For example, an alternative question such as "is Alice at the party$^{\uparrow}$ or is Bob at the party$^{\downarrow}$?" is often taken to presuppose that either Alice or Bob is at the party (cf. Karttunen & Peters, 1976; Biezma & Rawlins, 2012). When it is common ground that neither Alice nor Bob is at the party, the presupposition is already settled false and thus cannot be accommodated. As a result, the alternative question cannot be felicitously uttered in this context, which leads to its rejection.[16]

Now, since it is not the case that every information state must either resolve a question or directly contradict the presupposition it carries—e.g., the state $\{AB, A\bar{B}, \bar{A}B\}$ is unable to answer the question $A \vee B$ but nevertheless satisfies its presupposition—the equivalence $s \nVdash_1 \phi \equiv s \Vdash_0 \phi$ does not hold. Consequently, LEM is rejected again.

As we have seen with N3, the toggle negation $\sim$, unlike the intuitionistic negation $\neg$,

[16] Alternatively, we may hold a view according to which a question can be justifiably rejected as long as its presupposition is not settled true. This could happen in cases where accommodating a question's presupposition turns out to be difficult or costly. I will consider both interpretations of rejection as available options. And the difference here will not affect what follows. For example, if we adopt the second view, then DNE can be motivated on the ground that denying that the presupposition of a question is *not* already settled true amounts to affirming that the presupposition is already satisfied.

vindicates DNE. In N3, the intuitiveness of DNE is explained based on the thought that to find a counterexample to a counterexample of ϕ is to prove ϕ. Under the inquisitive setting, we can illuminate the validity of DNE in the following way. For declarative sentences, DNE amounts to saying that to object to a rejection of ϕ is to support ϕ. As for interrogative sentences, as mentioned above, to reject a question can be understood as saying that the presupposition of the question is already settled false, and when this happens, the question fails to make an inquiry and does not request future discourse to supply an answer to it. Given this, to object to the rejection of a question can be understood as rejecting the claim that the presupposition of the question is already settled false, thereby reinstating the question alongside its inquisitive content. Hence, after rejecting an objection to a question, the original question should still solicit an answer to it. As a paradigmatic example of this dynamics, consider cross examination in court. If the judge rules that the objection to a question having posed to the witness is "overruled", then the original question still stands, and the witness still needs to answer the question.

As such, we can make sense of the rejection of LEM from a non-bivalent perspective as well as the validity of DNE associated with the toggle negation.

4.3.2 Why N3 is Unsuitable

Since N3 is already a non-bivalent account that employs a constructive disjunction that does not vindicate LEM, we may wonder whether we can repurpose it to model inquisitiveness. However, the reason why N3 alongside BTS is unsuitable for this purpose is this. The basic twist-structure makes the rejection condition of a proposition more stringent than what we would normally expect. In particular, it requires the rejection of a conjunction to also be constructive.

For instance, let $v(A) = \langle a, a' \rangle$ and $v(B) = \langle b, b' \rangle$. Then, according to BTS, $v(A \wedge B) = \langle a \wedge b, a' \vee b' \rangle$. Under the inquisitive setting, we can interpret the two coordinates of a given pair as codifying positive and negative information about what it takes for the assigned formula to be supported and rejected. The support condition for $A \wedge B$ provided by BTS is fairly intuitive, since in order to settle "Alice and Bob are both at the party" true, we need to both know that Alice is at the party and know that Bob is at the party. The rejection condition for $A \wedge B$, on the contrary, appears too strong, since it commands that in order to reject the above assertion, we must either know that Alice is not at the party or know that Bob is not at the party. But if we know for sure that either Alice or Bob is out of town but cannot remember exactly who it is, then, arguably, we have already obtained enough information to reject the conjunction.

The above claim that the rejection condition of a sentence should not be constructive squares with the existing inquisitive semantics. Although inquisitive semantics does not employ this notion of rejection, it does construe the negation of a conjunction as non-inquisitive. Hence, unlike $\neg A \vee \neg B$, $\neg(A \wedge B)$ does not introduce any alternatives. Consequently, any information state that only eliminates all A-and-B-possibilities (e.g., $\{A\bar{B}, \bar{A}B, \bar{A}\bar{B}\}$) will come to support $\neg(A \wedge B)$.

To take stock, whereas N3 implements constructivism globally, we want, for the current purpose of devising a non-bivalent account of inquisitiveness, a framework that binds constructiveness more closely to the support condition for disjunction as disjunction uniquely plays the role of introducing alternatives. In other words, we want to introduce asymmetry between the support condition and the rejection condition of a sentence in respect of their constructiveness so as to make the support condition of a disjunction but not the rejection condition of a conjunction constructive. In the next section, I propose a framework along this line.

4.4 The Logic of Pseudo-Complemented Propositions

4.4.1 Pseudo-complemented twist-structures

We begin by exploring how we can modify BTS so as to vary the constructiveness of support and rejection conditions. In devising such a framework, I shall stick to the following two guiding principles:

Principle 1: The modified twist-structure still employs a toggle negation ~ that switches between the support and rejection conditions of a sentence.

Principle 2: The rejection condition of a sentence ϕ is given, whenever possible, by the support condition of its intuitionistic negation $\neg\phi$ (viz., $\phi \to \bot$) as provided in the original inquisitive semantics.

These two guiding principles together give rise to a class of twist-structures, i.e., *pseudo-complemented twist-structures* (PTS), that, on the one hand, incorporate a toggle negation and thus reject LEM from a non-bivalent standpoint, and on the other hand, by identifying the rejection condition of a sentence, whenever possible, with the support condition of its intuitionistic negation, render the rejection conditions of complex formulas non-constructive in a way that deviates minimally from the original inquisitive semantics. Before giving the definition for PTS, I want to emphasize that it is not my contention that PTS is the sole available non-bivalent treatment of inquisitiveness and we must adhere to the aforementioned two guiding principles. In particular, although Principle 2 affords a convenient way to specify the rejection conditions of complex formulas, we might have reasons to reject it based on other empirical motivations. Consequently, we may adopt a less uniform approach when defining the rejection condition for each connective. Be that as it may, in this paper, let us explore a framework that fulfills these two principles.

Given a Heyting algebra $\mathfrak{H} = \langle \mathcal{H}, \wedge, \vee, \to, \top, \bot \rangle$, we define a PTS, $\mathfrak{H}^{\bowtie *} = \langle \mathcal{H}^{\bowtie *}, \wedge^*_R, \vee^*_R, \to^*_R, {\sim}, 1, 0 \rangle$, as follows:

$$\mathcal{H}^{\bowtie *} = \{\langle x, x' \rangle | \text{ either } x' = x^*, \text{or } x = x'^*\}$$

where x^* abbreviates $x \to \bot$ for any x, and the logical operators are defined as follows:[17]

$$\langle x, x' \rangle \wedge^*_R \langle y, y' \rangle := \langle x \wedge y, (x \wedge y)^* \rangle$$

$$\langle x, x' \rangle \vee^*_R \langle y, y' \rangle := \langle x \vee y, (x \vee y)^* \rangle$$

$$\langle x, x' \rangle \to^*_R \langle y, y' \rangle := \langle x \to y, (x \to y)^* \rangle$$

$${\sim}\langle x, x' \rangle := \langle x', x \rangle$$

$$1 := \langle \top, \bot \rangle$$

$$0 := \langle \bot, \top \rangle$$

[17]The subscript "$_R$" indicates that the operator yields pairs whose second coordinate is the pseudo-complement of the first one. This is to be contrasted with a variation of PTS explored in §5.3 below.

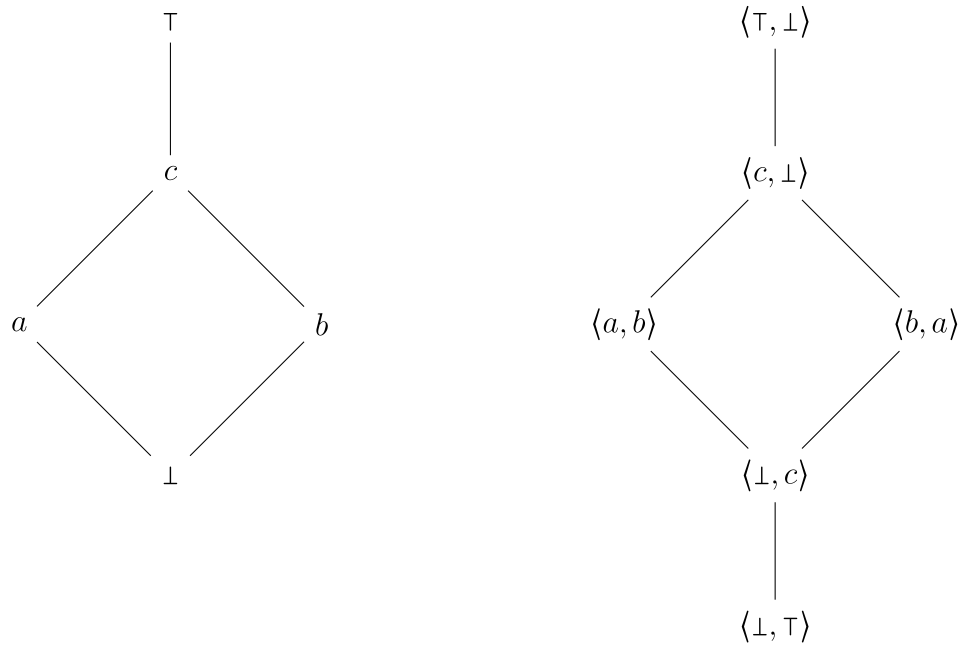

(a) (b)

Figure 4.3: An example of PTS.

Again, we define valuation functions inquisitively by mapping every atomic proposition to an element of $\mathcal{H}^{\bowtie *}_{\neg\neg} = \{((\langle x, x' \rangle \to^*_R 0) \to^*_R 0 \mid \langle x, x' \rangle \in \mathcal{H}^{\bowtie *}\}$. Semantic consequence and validity are defined in terms of inquisitive valuations as before.

Figure 4.3(b) depicts the PTS generated from the Heyting algebra as shown in 3(a). For the twist-operations associated with the three binary connectives, the operations are defined component-wise only for the first coordinate; the second coordinate is instead always the pseudo-complement of the first one. Defined as such, rejection conditions for complex formulas come to fulfill Principle 2. The definition for toggle negation remains unchanged, thereby fulfilling Principle 1. As a consequence, the carrier of PTS consists exclusively of pairs such that either the first component is the pseudo-complement of the second one, or the second component is the pseudo-complement of the first one.

As a notable feature of PTS, the two operations $\wedge^*_R$ and $\vee^*_R$ defined above do not always form the meets and joins of the new algebra. For instance, the meet of the two pairs $\langle a, b \rangle$ and $\langle b, a \rangle$ in the PTS as shown in Figure 4.3(b) is $\langle \bot, c \rangle$; however, applying the operation

$\langle a,b\rangle \wedge_R^* \langle b,a\rangle$ gives us $\langle\bot,\top\rangle$. Likewise, the join of $\langle\bot,c\rangle$ and itself is itself, whereas $\langle\bot,c\rangle \vee_R^* \langle\bot,c\rangle$ gives us $\langle\bot,\top\rangle$ again.

As it turns out, any two elements in PTS do not always have a greatest lower bound and a least upper bound; in other words, PTS does not form a lattice. To illustrate, consider the PTS generated from the Heyting algebra whose carrier set is set of all non-empty downsets in $\wp\{AB, A\bar{B}, \bar{A}B, \bar{A}\bar{B}\}$; in other words, we are considering a concrete Heyting algebra where each element from its carrier stands for some proposition understood as per inquisitive semantics. Now, let us consider the following two pairs $\langle a,a'\rangle$ and $\langle b,b'\rangle$ whose components are shown in Figure 4.4 (as with Figure 4.1, we represent each proposition using its set of alternatives). For each pair, the first component is the pseudo-complement of the second one but not vice versa.

Combining these two pairs via the old operation $\vee^{\bowtie}$ from BTS yields $\langle a\vee b, a'\wedge b'\rangle$, that is, $\langle c,d\rangle$, which is in fact the join of the two pairs in BTS. However, since neither component in $\langle c,d\rangle$ is the pseudo-complement of the other, this pair does not belong to the carrier of PTS. What we do have in PTS are the two pairs $\langle c,c'\rangle$ and $\langle d',d\rangle$, which are derived from $\langle a,a'\rangle \vee_R^* \langle b,b'\rangle$ and $\sim(\sim\langle a,a'\rangle \wedge_R^* \sim\langle b,b'\rangle)$, respectively.

Recall that PTS, as with BTS, is induced by an ordering such that for any two pairs $\langle x,x'\rangle$ and $\langle y,y'\rangle$, $\langle x,x'\rangle \le \langle y,y'\rangle$ in PTS iff $x \le x'$ and $y \ge y'$ in the original Heyting algebra. It follows that, given the absence of $\langle c,d\rangle$, the two pairs $\langle c,c'\rangle$ and $\langle d',d\rangle$ are indeed competing

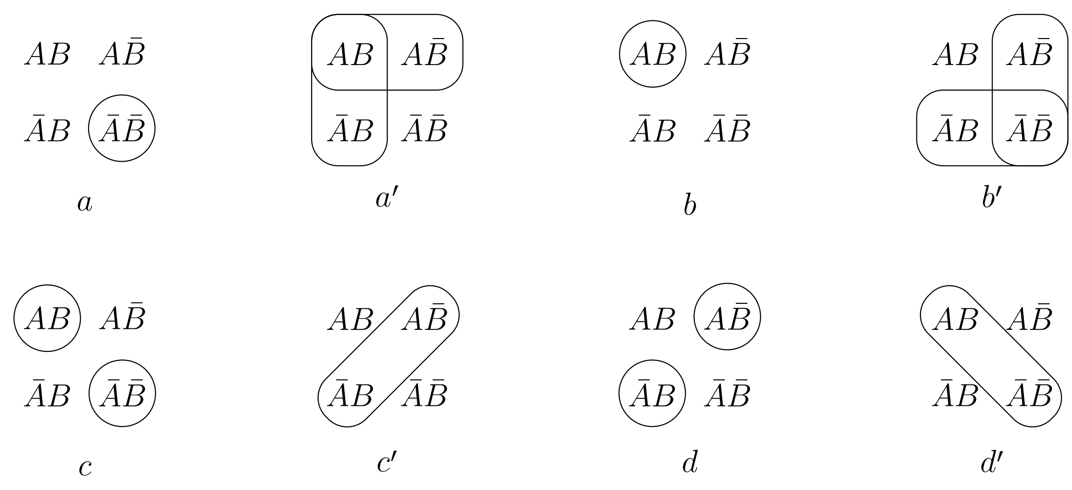

Figure 4.4: An illustration of why PTS does not form a lattice.

for the greatest lower bound of $\langle a, a' \rangle$ and $\langle b, b' \rangle$. Hence, $\langle a, a' \rangle$ and $\langle b, b' \rangle$ do not have a unique greatest lower bound. By parity of reasoning, any two elements in PTS also do not always have a least upper bound.

Algebraically, the two operators $\wedge_R^*$ and $\vee_R^*$ satisfy commutativity (viz., $X \wedge_R^* Y = Y \wedge_R^* X$, and $X \vee_R^* Y = Y \vee_R^* X$)[18], associativity (viz., $X \wedge_R^* (Y \wedge_R^* Z) = (X \wedge_R^* Y) \wedge_R^* Z$, and $X \vee_R^* (Y \vee_R^* Z) = (X \vee_R^* Y) \vee_R^* Z$), and distributivity (viz., $X \wedge_R^* (Y \vee_R^* Z) = (X \wedge_R^* Y) \vee_R^* (X \wedge_R^* Z)$, and $X \vee_R^* (Y \wedge_R^* Z) = (X \vee_R^* Y) \wedge_R^* (X \vee_R^* Z)$)[19], but they do not satisfy the principle of absorption which is characterized by the two equations: $X \wedge_R^* (X \vee_R^* Y) = X$, and $X \vee_R^* (X \wedge_R^* Y) = X$. Take the former equation as an example, and let X be a pair wherein the first component is the pseudo-complement of the second one but not vice versa. Since $\wedge_R^*$ and $\vee_R^*$ necessarily produce a pair whose second component is the pseudo-complement of the first, the left hand side and the right hand side of the equation cannot be the same.

For the same reason, idempotence also fails. That is, the following two equations do not hold in PTS: $X \wedge_R^* X = X$, $X \vee_R^* X = X$—specifically, when the second component of X is *not* the pseudo-complement of the first one. The loss of idempotence for $\wedge_R^*$ and $\vee_R^*$ may at first appear problematic. While that being the case, ϕ and $\phi \wedge \phi$, as well as ϕ and $\phi \vee \phi$, nonetheless semantically entail each other for any formula ϕ. This is so because given how semantic consequence is defined, when the premise—be it ϕ, $\phi \wedge \phi$, or $\phi \vee \phi$—is assumed to take the top element $\langle \top, \bot \rangle$ as its value, its first and second components are forced to be the pseudo-complements of each other, thereby forcing the conclusion to also take the top element as its value. Additionally, both $\vDash (\phi \wedge \phi) \leftrightarrow \phi$ and $\vDash (\phi \vee \phi) \leftrightarrow \phi$ obtain in PTS. Therefore, potential negative impact from the loss of idempotence is largely mitigated.[20]

[18] Here and hereafter, I will use the uppercase X, Y, and Z to stand for elements from the carrier set of a twist-structure.

[19] The proofs of these results follow from the fact that, for each equation, the resulting pairs on each side are such that their first coordinates are identical given that the Heyting algebra forms a distributive lattice and thus satisfies commutativity, associativity and distributivity. Given that the second coordinate of each pair is the pseudo-complement of its first coordinate, which, as we have just said, are identical on both sides for each equation, these equations hold in PTS.

[20] What we do lose in PTS are the following: $\nvDash {\sim}(\phi \wedge \phi) \to {\sim}\phi$ and $\nvDash {\sim}(\phi \vee \phi) \to {\sim}\phi$. How troublesome this is remains to be further investigated.

4.4.2 The logic of pseudo-complemented propositions

I provide an axiomatization of the corresponding logic of PTS here and prove completeness in §4.4.3. The logic of pseudo-complemented propositions (LPP) contains, in addition to the axiom schemas AS1–10 from N3 (viz., all axiom schemas from the positive logic plus $\sim\sim\phi \leftrightarrow \phi$ and $\sim\phi \to (\phi \to \psi)$), the axiom schemas listed below complemented with the inference rule of *modus ponens*.

AS14. $\perp \to \psi$

AS15. $((p \to \perp) \to \perp) \to p$, where p is an atomic formula.

AS16. $(\alpha \to \phi) \wedge (\alpha \to \sim\phi) \to \sim\alpha$, where α is a non-negated formula, i.e., a formula whose main connective is not $\sim$, and ϕ is any arbitrary formula.

In LPP, we introduce a propositional constant $\perp$ along with AS14. Together with all axiom schemas of the positive logic, it follows that LPP is an extension of IPL. The addition of $\perp$ turns out to be useful later in characterizing certain formal results and in our proof of completeness. AS15 amounts to the restricted version of DNE found in inquisitive logic, where DNE is valid with intuitionistic negation but only for atomic propositions. Since the operations $\wedge^*_R$, $\vee^*_R$, and $\to^*_R$ do not define their second coordinates component-wise, the AS11-AS13 from N3 are not valid. In their place, we have AS16 which embodies a weakened version of negation introduction found in intuitionistic logic and classical logic with the additional restriction that α needs to be a formula whose main connective is not $\sim$. Note that this restriction does allow negation to appear when embedded as in $\phi \wedge \sim\phi$.

As a noteworthy feature of this logic, the validity of AS16 hinges on the fact that given how logical operators are defined in PTS, operations other than negation (i.e., $\wedge^*_R$, $\vee^*_R$, and $\to^*_R$) will always result in a pair wherein the second component is the pseudo-complement of the first one; additionally, since valuation functions assign atomic propositions to the regular elements of PTS, every pair that interprets an atomic formula will again be such that its second component is the pseudo-complement of the first one.

Given this, we can prove the validity of $(\alpha \to \phi) \wedge (\alpha \to \sim\phi) \to \sim\alpha$ as follows. Let v be an arbitrary valuation; put $v(\alpha) = \langle x, x' \rangle$ and $v(\phi) = \langle y, y' \rangle$. Then $v(\alpha \to \phi) = \langle x \to y, (x \to y)^* \rangle$ and $v(\alpha \to \sim\phi) = \langle x \to y', (x \to y')^* \rangle$. It follows that $v((\alpha \to \phi) \wedge (\alpha \to \sim\phi)) = \langle (x \to y) \wedge (x \to$

$y'), ((x \to y) \wedge (x \to y'))^*\rangle = \langle x \to (y \wedge y'), (x \to (y \wedge y'))^*\rangle = \langle x \to \bot, (x \to \bot)^*\rangle$. Hence, $v((\alpha \to \phi) \wedge (\alpha \to {\sim}\phi) \to {\sim}\alpha) = \langle (x \to \bot) \to x', ((x \to \bot) \to x')^*\rangle$. Now, because the pair $\langle x, x'\rangle$ that interprets α is such that x' is the pseudo-complement x, $(x \to \bot) \to x' = x' \to x' = \top$. Thus, $v((\alpha \to \phi) \wedge (\alpha \to {\sim}\phi) \to {\sim}\alpha) = \langle \top, \bot\rangle = 1$.

By contrast, when $\langle x, x'\rangle$ is such that its first component is the pseudo-complement of the second component but not vice versa, we only have $(x \to \bot) \to x' = ((x' \to \bot) \to \bot) \to x'$, which is not always equal to the top element in a Heyting algebra.

Additionally, note that by replacing AS16 with the stronger negation introduction (NI) which is just like AS16 but does not contain any additional restriction on the structure of α, we would turn LPP into classical propositional logic (CPL).[21] In order to validate the unconstrained NI, the second coordinate of any pair that interprets α needs to always be the pseudo-complement of its first coordinate, even if α is led by a single ~. To satisfy this, we would need to make the pseudo-complement relation between x and x' symmetric. This would in turn generate a new class of twist-structures, call them *complemented twist-structures* (CTS). Indeed, as we shall see in §4.5.2, the class of CTS is isomorphic to the class of Boolean algebras.

Let me highlight some other properties of LPP which we will come back to later in our proof of completeness.

Proposition 4.4.1. $\vdash_{LPP} (\alpha \to \phi) \to ({\sim}\phi \to {\sim}\alpha)$, *provided that α is a non-negated formula.*

Proof. Since we have $(\alpha \to \phi) \wedge (\alpha \to {\sim}\phi) \to {\sim}\alpha$ from AS16 and $\vdash_{LPP} {\sim}\phi \to (\alpha \to {\sim}\phi)$ from AS1, we can easily derive $\vdash_{LPP} (\alpha \to \phi) \to ({\sim}\phi \to {\sim}\alpha)$. □

Corollary 4.4.1.1. $\vdash_{LPP} (\alpha \leftrightarrow \beta) \to ({\sim}\alpha \leftrightarrow {\sim}\beta)$, *provided that both α and β are non-negated formulas.*

Remark. As with N3, since $(\alpha \leftrightarrow \beta) \to ({\sim}\alpha \leftrightarrow {\sim}\beta)$ does not hold generally, $\vdash_{LPP} \phi \leftrightarrow \psi$ alone does not define a congruence relation; instead, in order for ϕ and ψ to belong to the same congruence class, we need to have both $\vdash_{LPP} \phi \leftrightarrow \psi$ and $\vdash_{LPP} {\sim}\phi \leftrightarrow {\sim}\psi$.

Proposition 4.4.2. $\vdash_{LPP} {\sim}\alpha \leftrightarrow (\alpha \to \bot)$, *provided that α is a non-negated formula.*

[21] More specifically, AS1–8 + NI give rise to minimal logic, upon which the addition of DNE yields CPL (See Odintsov, 2008).

Proof. Left-to-right follows directly from AS10. Right-to-left follows from AS16 together with the fact that $\vdash_{LPP} \alpha \to {\sim}\bot$, which is derivable from applying contraposition to $\vdash_{LPP} \bot \to {\sim}\alpha$, given that $\bot$ is a non-negated formula. $\square$

Remark. Note that we do not have the contrapositive $\nvdash_{LPP} \alpha \leftrightarrow {\sim}(\alpha \to \bot)$ as ${\sim}\alpha$ is not a non-negated formula.

Proposition 4.4.3. *The following bi-implications are theorems of LPP, provided that α is a non-negated formula:*

$$\vdash_{LPP} (\phi \wedge {\sim}\alpha) \leftrightarrow (\phi \wedge (\alpha \to \bot))$$
$$\vdash_{LPP} {\sim}(\phi \wedge {\sim}\alpha) \leftrightarrow {\sim}(\phi \wedge (\alpha \to \bot))$$
$$\vdash_{LPP} (\phi \vee {\sim}\alpha) \leftrightarrow (\phi \vee (\alpha \to \bot))$$
$$\vdash_{LPP} {\sim}(\phi \vee {\sim}\alpha) \leftrightarrow {\sim}(\phi \vee (\alpha \to \bot))$$
$$\vdash_{LPP} (\phi \to {\sim}\alpha) \leftrightarrow (\phi \to (\alpha \to \bot))$$
$$\vdash_{LPP} {\sim}(\phi \to {\sim}\alpha) \leftrightarrow {\sim}(\phi \to (\alpha \to \bot))$$
$$\vdash_{LPP} ({\sim}\alpha \to \phi) \leftrightarrow ((\alpha \to \bot) \to \phi)$$
$$\vdash_{LPP} {\sim}({\sim}\alpha \to \phi) \leftrightarrow {\sim}((\alpha \to \bot) \to \phi)$$

Proof. The four bi-implications where both sides of $\leftrightarrow$ are non-negated formulas follow easily from Proposition 4.4.2. As for the four bi-implications of the negative form, they follow from their respective positive form by Corollary 4.4.1.1. $\square$

Remark. What Proposition 4.4.3 amounts to is that whenever a toggle negation appears alone embedded below a binary connective, it becomes equivalent to the intuitionistic negation. This allows us to reduce embedded toggle negations to intuitionistic negations. As an example, we can convert ${\sim}(A \wedge {\sim}A)$ to ${\sim}(A \wedge (A \to \bot))$.

Proposition 4.4.4 (Normal form)**.** *Formulas of LPP can be converted into either a positive formula or a positive formula prefixed by a single negation ${\sim}$, where a positive formula is defined as a formula that does not contain any ${\sim}$ (but may contain $\bot$).*

Proof. We prove this inductively. First, all literals apparently satisfy this condition. Next, we want to show that for any formulas ϕ and ψ that satisfy this condition, any complex formula that contains ϕ and ψ as its immediate constituents can be transformed into either a positive formula or a positive formula headed by a single $\sim$. Consider negation first, given the inductive hypothesis, ϕ can take the form of either α or $\sim\alpha$, where α is a positive formula. If ϕ takes the form of α, then $\sim\phi$ becomes $\sim\alpha$, which is a positive formula prefixed by a single $\sim$; if ϕ takes the form of $\sim\alpha$, then $\sim\phi$ becomes $\sim\sim\alpha$, which given DNE, is reduced to a positive formula α. For conjunction, to simplify, let us focus on the structure of ϕ by granting that ψ is a positive formula. If ϕ takes the form of α, then $\psi \wedge \phi$ is just $\psi \wedge \alpha$, which is a positive formula; if ϕ takes the form of $\sim\alpha$, then $\psi \wedge \phi$ becomes $\psi \wedge \sim\alpha$, which given Proposition 4.4.3, can be converted to $\psi \wedge (\alpha \rightarrow \bot)$, which is again a positive formula. The same would apply to ψ if it was a positive formula prefixed by a single $\sim$. Hence, conjunction preserves the above condition. By similar reasoning, disjunction and implication also produce either a positive formula or one that is prefixed by a single $\sim$. □

Remark. As it turns out, the toggle negation in LPP can only appear in one place where it makes substantial semantic contribution, that is, at a sentence's widest scope. As such, $\sim$ can be viewed as a global rejection operator under LPP. This result appears to fare well with our previous characterization of $\sim$ as an operator that toggles between the support and rejection conditions of a sentence.

Corollary 4.4.4.1. *LPP can be axiomatized alternatively by replacing AS16, i.e.,* $(\alpha \rightarrow \phi) \wedge (\alpha \rightarrow \sim\phi) \rightarrow \sim\alpha$, *with an axiom that further restricts* α *to only positive formulas.*

4.4.3 Completeness

Let me first provide a sketch of the proof. We prove completeness by establishing the following equivalences. (Different notations will be explained immediately afterwards.)

Theorem 4.4.5 (Completeness theorem)**.** *The following conditions are equivalent:*

1. $\vdash_{LPP} \phi$;
2. $\vDash_{\bowtie *} \phi$;
3. $\vDash_{\mathfrak{L}^{\bowtie *}_{LPP+}} \phi$;

4. $\vDash_{\mathfrak{L}_{LPP}} \phi$.

First, we show that (1) entails (2) (viz., ϕ is a theorem of LPP only if ϕ is valid in the class of PTS) by proving soundness. Second, let $\mathfrak{L}_{LPP}$ be the Lindenbaum-Tarski algebra of the language of LPP; we show that (4) entails (1) by proving that LPP is complete with respect to $\mathfrak{L}_{LPP}$. Third, let us define another Lindenbaum-Tarski algebra $\mathfrak{L}_{LPP+}$ using the positive fragment of LPP that does not contain ~; let $\mathfrak{L}^{\bowtie *}_{LPP+}$ be the PTS generated from the Lindenbaum-Tarski algebra $\mathfrak{L}_{LPP+}$. Since $\mathfrak{L}^{\bowtie *}_{LPP+}$ is a special PTS, (2) entails (3) immediately. Last, we show that $\mathfrak{L}^{\bowtie *}_{LPP+}$ and $\mathfrak{L}_{LPP}$ are indeed isomorphic, thereby establishing the equivalence between (3) and (4). I shall also remind the readers that in (2)-(4), validity is defined in term of the inquisitive valuation. For any atomic formula p, the valuation will assign p and $(p \to \bot) \to \bot$ to the same element.

Lemma 4.4.6 (Soundness lemma). *If* $\vdash_{LPP} \phi$, *then* $\vDash_{\bowtie *} \phi$.

Proof. We can directly show that all the axiom schemas from above are valid with respect to PTS. The proofs for AS1–8 from the positive fragment are straightforward. Take AS1 as an example. Let v be an arbitrary valuation on an arbitrary PTS, and put $v(\phi) = \langle x, x' \rangle$ and $v(\psi) = \langle y, y' \rangle$. Then, it follows that $v(\phi \to (\psi \to \phi)) = \langle x, x' \rangle \to^*_R (\langle y, y' \rangle \to^*_R \langle x, x' \rangle) = \langle x \to (y \to x), (x \to (y \to x))^* \rangle$. Since x and y are elements from a Heyting algebra, $x \to (y \to x) = \top$. Therefore, $v(\phi \to (\psi \to \phi)) = \langle \top, \bot \rangle = 1$. Likewise for AS2–7, since no negation is involved, we can show that the corresponding formulas hold in the Heyting algebra.

The validity of AS9 immediately follows from the definition of ~ in PTS.

To prove the validity of *ex falso* as stated by AS10, again let $v(\phi) = \langle x, x' \rangle$ and $v(\psi) = \langle y, y' \rangle$. It follows that $v({\sim}\phi \to (\phi \to \psi)) = {\sim}\langle x, x' \rangle \to^*_R (\langle x, x' \rangle \to^*_R \langle y, y' \rangle) = \langle x', x \rangle \to^*_R \langle x \to y, (x \to y)^* \rangle = \langle x' \to (x \to y), (x' \to (x \to y))^* \rangle$. Now, given that, $x' \to (x \to y) = (x' \wedge x) \to y = \bot \to y = \top$, $v({\sim}\phi \to (\phi \to \psi)) = \langle \top, \bot \rangle = 1$. Hence, $\vDash_{\bowtie *} {\sim}\phi \to (\phi \to \psi)$.

The validity of AS15 immediately follows from the fact that valuations on PTS map every atomic proposition p to an element of $\mathcal{H}^{\bowtie *}_{\neg\neg} = \{(\langle x, x' \rangle \to^*_R 0) \to^*_R 0 \mid \langle x, x' \rangle \in \mathcal{H}^{\bowtie *}\}$.

We have already proven the validity of AS16 in the last subsection.

Lastly, to prove that *modus ponens* holds, again let $v(\phi) = \langle x, x' \rangle$ and $v(\psi) = \langle y, y' \rangle$. Assume $v(\phi) = 1$ and $v(\phi \to \psi) = 1$, that is $\langle x, x' \rangle = \langle \top, \bot \rangle$ and $\langle x \to y, (x \to y)^* \rangle = \langle \top, \bot \rangle$. We want to show that $v(\psi) = 1$, that is, $\langle y, y' \rangle = \langle \top, \bot \rangle$. This is straightforward because for any two elements x and y in a Heyting algebra, if $x = \top$ and $x \to y = \top$, it must follow that $y = \top$. And given that $y \wedge y' = \bot$, it must follow that $y' = \bot$. Thus, $\langle y, y' \rangle = \langle \top . \bot \rangle$. □

Definition 4.4.7 (Equivalence relation $\equiv$)**.** Define $\equiv$ as follows: $\phi \equiv \psi$ iff $\vdash_{LPP} \phi \leftrightarrow \psi$ and $\vdash_{LPP} {\sim}\phi \leftrightarrow {\sim}\psi$.

Proposition 4.4.8. *The relation $\equiv$ is a congruence on the propositional formula algebra of LPP: $\mathfrak{F} = \langle \mathscr{F}, \wedge, \vee, \to, \sim, \bot \rangle$.*

Definition 4.4.9 (Lindenbaum-Tarksi algebra of LPP)**.** Let $\mathscr{F}/{\equiv}$ be the set of congruence classes $\|\phi\|$ induced by $\equiv$ on the set of formulas $\mathscr{F}$. The Lindenbaum-Tarski algebra of LPP is defined as follows:

$$\mathfrak{L}_{LPP} = \langle \mathscr{F}/{\equiv}, \wedge, \vee, \to, \sim, \top, \bot \rangle$$

where $\|\phi\| \wedge \|\psi\| := \|\phi \wedge \psi\|$, $\|\phi\| \vee \|\psi\| := \|\phi \vee \psi\|$, $\|\phi\| \to \|\psi\| := \|\phi \to \psi\|$, ${\sim}\|\phi\| := \|{\sim}\phi\|$, $\top := \|\phi \to \phi\|$, and $\bot := \|\bot\|$.

Lemma 4.4.10 (Completeness of LPP w.r.t. $\mathfrak{L}_{LPP}$)**.** *If $\vDash_{\mathfrak{L}_{LPP}} \phi$, then $\vdash_{LPP} \phi$.*

Proof. We prove its contrapositive. Suppose $\nvdash_{LPP} \phi$. Then $\|\phi\| \neq \top$. Let us define a canonical valuation $v^0 : \mathscr{F} \to \mathfrak{L}_{LPP}$ from the set of formulas to the carrier of $\mathfrak{L}_{LPP}$ as follows: for every atomic proposition p, $v^0(p) = \|p\|$. Note that given the presence of AS15, we have $\|(p \to \bot) \to \bot\| = \|p\|$ for every atomic proposition p, which in turn ensures that v^0 still qualifies as an inquisitive valuation. Now, it follows that for all formulas ψ, $v^0(\psi) = \|\psi\|$. Hence, $v^0(\phi) = \|\phi\|$. And since $\|\phi\| \neq \top$, $v^0(\phi) \neq \top$. Thus, $\nvDash_{\mathfrak{L}_{LPP}} \phi$. □

Definition 4.4.11 (Equivalence relation $\approx$)**.** Define $\approx$ as follows: $\phi \approx \psi$ iff $\vdash_{LPP} \phi \leftrightarrow \psi$.

Proposition 4.4.12. *The relation $\approx$ is a congruence on the propositional formula algebra of the negation-free fragment of LPP: $\mathfrak{F}^+ = \langle \mathscr{F}^+, \wedge, \vee, \to, \bot \rangle$.*

Definition 4.4.13 (Lindenbaum-Tarski algebra $\mathfrak{L}_{LPP+}$)**.** Let $\mathscr{F}^+/{\approx}$ be the set of congruence classes $[\phi]$ induced by $\approx$ on $\mathscr{F}^+$. The Lindenbaum-Tarski algebra $\mathfrak{L}_{LPP+}$ of the negation-free fragment of LPP is defined as follows:

$$\mathfrak{L}_{LPP+} = \langle \mathscr{F}^+/\approx, \wedge, \vee, \to, \top, \bot \rangle$$

where $[\phi] \wedge [\psi] := [\phi \wedge \psi]$, $[\phi] \vee [\psi] := [\phi \vee \psi]$, $[\phi] \to [\psi] := [\phi \to \psi]$, $\top := [\phi \to \phi]$, and $\bot := [\bot]$.

Definition 4.4.14 (PTS generated from the Lindenbaum-Tarski algebra $\mathfrak{L}_{LPP+}$)**.** We define the PTS of $\mathfrak{L}_{LPP+}$ as follows:

$$\mathfrak{L}^{\bowtie *}_{LPP+} = \langle \mathcal{L}^{\bowtie *}_{LPP+}, \wedge^*_R, \vee^*_R, \to^*_R, {\sim}, 1, 0 \rangle$$

where $\mathcal{L}^{\bowtie *}_{LPP+} = \{\langle [\phi], [\phi]^* \rangle | \ [\phi] \in \mathscr{F}^+/\approx\} \cup \{\langle [\phi]^*, [\phi] \rangle | \ [\phi] \in \mathscr{F}^+/\approx\}$, and

$$\langle [\phi], [\phi'] \rangle \wedge^*_R \langle [\psi], [\psi'] \rangle := \langle [\phi] \wedge [\psi], ([\phi] \wedge [\psi])^* \rangle$$
$$\langle [\phi], [\phi'] \rangle \vee^*_R \langle [\psi], [\psi'] \rangle := \langle [\phi] \vee [\psi], ([\phi] \vee [\psi])^* \rangle$$
$$\langle [\phi], [\phi'] \rangle \to^*_R \langle [\psi], [\psi'] \rangle := \langle [\phi] \to [\psi], ([\phi] \to [\psi])^* \rangle$$
$${\sim}\langle [\phi], [\phi'] \rangle := \langle [\phi'], [\phi] \rangle$$
$$1 := \langle \top, \bot \rangle$$
$$0 := \langle \bot, \top \rangle$$

Next, we want to show that $\mathfrak{L}_{LPP}$ and $\mathfrak{L}^{\bowtie *}_{LPP+}$ are indeed isomorphic. To prove this, we first define two maps between $\mathcal{L}_{LPP}$ and $\mathcal{L}^{\bowtie *}_{LPP+}$. We then show that they are homomorphisms that are also inverses of each other.

Definition 4.4.15. We define a map $h : \mathcal{L}_{LPP} \to \mathcal{L}^{\bowtie *}_{LPP+}$ from the carrier of $\mathfrak{L}_{LPP}$ to the carrier of $\mathfrak{L}^{\bowtie *}_{LPP+}$:

$$h(\|p\|) = \langle [p], [p]^* \rangle$$
$$h(\|\phi \wedge \psi\|) = h(\|\phi\|) \wedge^*_R h(\|\psi\|)$$
$$h(\|\phi \vee \psi\|) = h(\|\phi\|) \vee^*_R h(\|\psi\|)$$
$$h(\|\phi \to \psi\|) = h(\|\phi\|) \to^*_R h(\|\psi\|)$$
$$h(\|{\sim}\phi\|) = {\sim}h(\|\phi\|)$$

$$h(\top) = 1$$
$$h(\perp) = 0$$

Proposition 4.4.16. *h is a homomorphism.*

Proof. This is obvious given how h is defined. $\square$

Definition 4.4.17. We define a converse map $h' : \mathcal{L}^{\bowtie *}_{LPP+} \to \mathcal{L}_{LPP}$:

$$h'(\langle[\phi],[\psi]\rangle) = \begin{cases} \|\phi\| & \text{if } [\psi] = [\phi]^* \\ \|{\sim}\psi\| & \text{if } [\phi] = [\psi]^* \end{cases}$$

Example 4.4.18. To illustrate, let us consider three examples:

1. $h'(\langle[?A],[\perp]\rangle) = \|?A\|$,[22] provided that the second coordinate $[\perp]$ is the pseudo-complement of the first coordinate $[?A]$—that is, $[\perp] = [?A \to \perp]$—but not vice versa.
2. $h'(\langle[\perp],[?A]\rangle) = \|{\sim}?A\|$, provided that the first coordinate $[\perp]$ is the pseudo-complement of the second coordinate $[?A]$, but not vice versa.
3. $h'(\langle[A \to A],[\perp]\rangle) = \|A \to A\| = \|{\sim}\perp\|$, provided that $[A \to A]$ and $[\perp]$ are pseudo-complement of each other. In other words, we have that $h'(1) = \top = {\sim}\perp$.

Proposition 4.4.19. *For any $\langle[\phi],[\psi]\rangle$, if $[\psi] = [\phi]^*$ and $[\phi] = [\psi]^*$, then $h'(\langle[\phi],[\psi]\rangle = \|\phi\| = \|{\sim}\psi\|$.*

Proof. Recall again that $[\phi]^* = [\phi \to \perp]$. Since $[\psi] = [\phi]^*$ and $[\phi] = [\psi]^*$, we have $\vdash_{LPP} \psi \leftrightarrow (\phi \to \perp)$, and $\vdash_{LPP} \phi \leftrightarrow (\psi \to \perp)$. Next, since both ϕ and ψ appearing in $[\phi]$ and $[\psi]$ are formulas from the negation-free fragment of LPP, they do not contain any ~. By Proposition 4.4.2 where we establish that for any non-negated formula α, $\vdash_{LPP} {\sim}\alpha \leftrightarrow (\alpha \to \perp)$, it then follows that $\vdash_{LPP} \psi \leftrightarrow {\sim}\phi$, and $\vdash_{LPP} \phi \leftrightarrow {\sim}\psi$. And given that $\vdash_{LPP} {\sim}{\sim}\psi \leftrightarrow \psi$, we have $\vdash_{LPP} {\sim}{\sim}\psi \leftrightarrow {\sim}\phi$. Since we have obtained both $\vdash_{LPP} \phi \leftrightarrow {\sim}\psi$ and $\vdash_{LPP} {\sim}{\sim}\psi \leftrightarrow {\sim}\phi$, by the definition of the congruence relation on LPP, ϕ and ${\sim}\psi$ belong to the same congruence class, that is, $\|\phi\| = \|{\sim}\psi\|$. $\square$

[22] $?A$ abbreviates $A \vee (A \to \perp)$.

Proposition 4.4.20. *h' is a homomorphism.*

Proof. We reason by cases.

1. For negation, we have that $h'({\sim}\langle[\phi],[\psi]\rangle) = h'(\langle[\psi],[\phi]\rangle) =$

$$\begin{cases} \|\psi\| = \|{\sim}{\sim}\psi\| = {\sim}\|{\sim}\psi\| = {\sim}h'(\langle[\phi],[\psi]\rangle) & \text{if } [\phi] = [\psi]^*; \\ \|{\sim}\phi\| = {\sim}\|\phi\| = {\sim}h'(\langle[\phi],[\psi]\rangle) & \text{if } [\psi] = [\phi]^*. \end{cases}$$

   Hence, regardless of whether the first coordinate is the pseudo-complement of the second or vice versa, $h'({\sim}\langle[\phi],[\psi]\rangle) = {\sim}h'(\langle[\phi],[\psi]\rangle)$.

2. For conjunction, we establish the following equivalences first: $h'(\langle[\phi],[\phi']\rangle \wedge^*_R \langle[\psi],[\psi']\rangle) = h'(\langle[\phi] \wedge [\psi], ([\phi] \wedge [\psi])^*\rangle) = h'(\langle[\phi \wedge \psi], [\phi \wedge \psi]^*\rangle) = \|\phi \wedge \psi\|$, given that $[\phi \wedge \psi]^*$ is the pseudo-complement of $[\phi \wedge \psi]$.

   Next, we show that the following equations hold: $\|\phi \wedge \psi\| =$

$$\begin{cases} \|\phi \wedge \psi\| & \text{if } [\phi'] = [\phi]^* \text{ and } [\psi'] = [\psi]^*; \\ \|\phi \wedge {\sim}\psi'\| & \text{if } [\phi'] = [\phi]^* \text{ and } [\psi] = [\psi']^*; \\ \|{\sim}\phi' \wedge \psi\| & \text{if } [\phi] = [\phi']^* \text{ and } [\psi'] = [\psi]^*; \\ \|{\sim}\phi' \wedge {\sim}\psi'\| & \text{if } [\phi] = [\phi']^* \text{ and } [\psi] = [\psi']^*. \end{cases}$$

   The first case is obvious. It then follows that $\|\phi \wedge \psi\| = \|\phi\| \wedge \|\psi\| = h'(\langle[\phi],[\phi']\rangle) \wedge h'(\langle[\psi],[\psi']\rangle)$. Therefore, we have that when $[\phi'] = [\phi]^*$ and $[\psi'] = [\psi]^*$, $h'(\langle[\phi],[\phi']\rangle \wedge^*_R \langle[\psi],[\psi']\rangle) = h'(\langle[\phi],[\phi']\rangle) \wedge h'(\langle[\psi],[\psi']\rangle)$.

   Consider the second case. We want to show that when $[\phi'] = [\phi]^*$ and $[\psi] = [\psi']^*$, that is, when $[\phi'] = [\phi \to \bot]$ and $[\psi] = [\psi' \to \bot]$, both $\vdash_{LPP} (\phi \wedge \psi) \leftrightarrow (\phi \wedge {\sim}\psi')$ and $\vdash_{LPP} {\sim}(\phi \wedge \psi) \leftrightarrow {\sim}(\phi \wedge {\sim}\psi')$ obtain. First, note that since both $\phi \wedge \psi$ and $\phi \wedge {\sim}\psi'$ are non-negated formulas, by the restricted contraposition established in Corollary 4.4.1.1, if we can prove that $\vdash_{LPP} (\phi \wedge \psi) \leftrightarrow (\phi \wedge {\sim}\psi')$, we can prove that $\vdash_{LPP} {\sim}(\phi \wedge \psi) \leftrightarrow {\sim}(\phi \wedge {\sim}\psi')$.

   Now, consider the proof of $\vdash_{LPP} (\phi \wedge \psi) \leftrightarrow (\phi \wedge {\sim}\psi')$. Given that $[\psi] = [\psi' \to \bot]$, we have $\vdash_{LPP} \psi \leftrightarrow (\psi' \to \bot)$. We can then easily show that $\vdash_{LPP} (\phi \wedge \psi) \leftrightarrow (\phi \wedge (\psi' \to \bot))$.

Since ψ' is a formula from the negation-free fragment of LPP, it does not contain ~. Thus, by Proposition 4.4.3 which establishes the equivalence between toggle negation and intuitionistic negation when they are embedded below a binary connective, it follows that $\vdash_{LPP} (\phi \wedge \psi) \leftrightarrow (\phi \wedge {\sim}\psi')$. And since both $\phi \wedge \psi$ and $\phi \wedge {\sim}\psi'$ are non-negated formulas, by contraposition, we have $\vdash_{LPP} {\sim}(\phi \wedge \psi) \leftrightarrow {\sim}(\phi \wedge {\sim}\psi')$.

Because we have established both $\vdash_{LPP} (\phi \wedge \psi) \leftrightarrow (\phi \wedge {\sim}\psi')$ and $\vdash_{LPP} {\sim}(\phi \wedge \psi) \leftrightarrow {\sim}(\phi \wedge {\sim}\psi')$, it follows that $\|\phi \wedge \psi\| = \|\phi \wedge {\sim}\psi'\|$. Continuing with the proof that h' is a homomorphism, we have $\|\phi \wedge {\sim}\psi'\| = \|\phi\| \wedge \|{\sim}\psi'\| = h'(\langle[\phi], [\phi']\rangle) \wedge h'(\langle[\psi], [\psi']\rangle)$, provided that $[\phi'] = [\phi]^*$ and $[\psi] = [\psi']^*$.

We can apply similar reasoning to the last two cases and show that for all four cases, $h'(\langle[\phi], [\phi']\rangle \wedge^*_R \langle[\psi], [\psi']\rangle) = h'(\langle[\phi], [\phi']\rangle) \wedge h'(\langle[\psi], [\psi']\rangle)$.

3. For disjunction, by the same method, we can show that $h'(\langle[\phi], [\phi']\rangle \vee^*_R \langle[\psi], [\psi']\rangle) = h'(\langle[\phi], [\phi']\rangle) \vee h'(\langle[\psi], [\psi']\rangle)$.
4. Likewise for implication, we can show $h'(\langle[\phi], [\phi']\rangle \to^*_R \langle[\psi], [\psi']\rangle) = h'(\langle[\phi], [\phi']\rangle) \to h'(\langle[\psi], [\psi']\rangle)$.
5. For the two 0-ary operators, we immediately have $h'(1) = \top$ and $h'(0) = \bot$.

This completes our proof that h' is a homomorphism. □

Lemma 4.4.21. *$\mathfrak{L}_{LPP}$ and $\mathfrak{L}^{\bowtie *}_{LPP+}$ are isomorphic.*

Proof. Since both h and h' are homomorphisms, we can show that for all $x \in \mathcal{L}_{LPP}, h'(h(x)) = x$, and for all $X \in \mathcal{L}^{\bowtie *}_{LPP+}, h(h'(X)) = X$.

1. We prove that for all $x \in \mathcal{L}_{LPP}, h'(h(x)) = x$ via an induction.

   Base case: for any atomic proposition p, $h'(h(\|p\|)) = h'(\langle[p], [p]^*\rangle) = \|p\|$.

   Inductive step: Suppose $h'(h(\|\phi\|) = \|\phi\|$ and $h'(h(\|\psi\|) = \|\psi\|$. Then,

   (a) $h'(h(\|\phi \wedge \psi\|)) = h'(h(\|\phi\|) \wedge^*_R h(\|\psi\|)) = h'(h(\|\phi\|)) \wedge h'(h(\|\psi\|)) = \|\phi\| \wedge \|\psi\| = \|\phi \wedge \psi\|$;

   (b) By similar reasoning, $h'(h(\|\phi \vee \psi\|)) = \|\phi \vee \psi\|$, and $h'(h(\|\phi \to \psi\|)) = \|\phi \to \psi\|$;

   (c) $h'(h(\|{\sim}\phi\|) = h'({\sim}h(\|\phi\|) = {\sim}h'(h(\|\phi\|) = {\sim}\|\phi\| = \|{\sim}\phi\|$.

2. Given we have just established that for all $x \in \mathcal{L}_{LPP}, h'(h(x)) = x$, we can show for all $X \in \mathcal{L}^{\bowtie *}_{LPP+}, h(h'(X)) = X$ via a *reductio*. Suppose, on the contrary, that $h(h'(X)) \neq X$, and let $X = h(x)$. It follows that $h(h'(h(x))) \neq h(x)$. But since for all $x \in \mathcal{L}_{LPP}, h'(h(x)) = x$, it also follows that $h(h'(h(x))) = h(x)$. Therefore, by *reductio*, $h(h'(X)) = X$.

Given that h and h' are inverses of each other, $\mathfrak{L}_{LPP}$ and $\mathfrak{L}^{\bowtie *}_{LPP+}$ are indeed isomorphic. □

Corollary 4.4.21.1. $\vDash_{\mathfrak{L}_{LPP}} \phi$ *iff* $\vDash_{\mathfrak{L}^{\bowtie *}_{LPP+}} \phi$.

We have completed the proof of completeness.

4.5 Generalizations

As stated in the introduction, a second aim of this paper is to further popularize this method of twist-construction. To this end, I will explore several generalizations of PTS in this last section. To begin with, let me first make a conceptual remark regarding the use of twist-construction. As notated at the beginning of this paper, to reject LEM, we could either adopt a bivalent or a non-bivalent approach. And as we have seen, in both the case of N3 and the case of LPP, we can transform a bivalent account to a non-bivalent one via twist-operation which introduces a toggle negation that rejects LEM. Crucially, what twist-operation enables is to convert a semantic theory that employs only one central notion to one that employs two independent central notions. This is reflected in the fact that the carrier of a twist-structure consists of pairs whose components are themselves elements of the carrier of the original algebra from which the twist-structure is generated. As such, twist-structures are able to enrich the semantic interpretation a sentence receives by encoding two separate bodies of information embodied by it—e.g., its acceptance and rejection conditions.

To generalize, we can view PTS as one of the many structures that can be generated from applying the twist-operation to a certain base algebra. In using this method, we need to determine three things: first, we fix an algebra that serves as the base of the twist-operation, then we define the carrier set of the twist-structure by filtering out a subset of elements from the product algebra, and lastly we define the logical operators of the new algebra. This is not to say that these three parameters are fully independent of each other, since, for example, selecting a base that does not contain a bottom element will also limit how the carrier can be defined. In what follows, I shall examine one generalization along each dimension to illustrate the utility of twist-construction.

4.5.1 The base

The base for the twist-operation is the algebra from which the product algebra is derived. As we have seen, both PTS and BTS can be regarded as derived from Heyting algebras. Thus, a natural generalization is to explore other algebras that can serve as the base for twist-construction. One option is to explore algebras weaker than Heyting algebras. In the existing literature, twist-structures have been defined using weaker bases such as implicative lattices (Odintsov, 2004; Rivieccio, 2014) and subminimal algebras (Vakarelov, 2005).

For instance, using implicative lattices as the base gives rise to the so-called *full twist-structures* (FTS). Note that since, different from a Heyting algebra, an implicative lattice does not have to come with a unique bottom element, we cannot use the same carrier as that of BTS; consequently, the carrier of FTS is simply identified with the carrier of the whole product lattice. FTS have been used to provide an algebraic semantics for Nelson's paraconsistent logic N4, especially given that there is no requirement for the meet of the two coordinates of any given pair in FTS to be the bottom element.

Here, I wish to explore this generalization in the other direction, that is, to consider taking algebra that is stronger than the Heyting algebra as the base for twist-construction. In particular, we can consider taking the inquisitive algebra (Bezhanishvili et al., 2019) as the base of our pseudo-complemented twist-operation.

Definition 4.5.1 (Inquisitive algebra)**.** An inquisitive algebra $\mathfrak{I}$ is a Heyting algebra $\mathfrak{H}$ that is regularly generated, is well-connected, and validates the Kreisel-Putnam axiom, viz., $(\neg\phi \to (\psi \vee \chi)) \to ((\neg\phi \to \psi) \vee (\neg\phi \to \chi))$.

Definition 4.5.2 (Regularly generated)**.** A Heyting algebra $\mathfrak{H}$ is regularly generated if it is generated by $\mathcal{H}_{\neg\neg} = \{\neg\neg x | x \in \mathcal{H}\}$.

Definition 4.5.3 (Well-connected)**.** A Heyting algebra $\mathfrak{H}$ is well-connected if for every x and y in $\mathcal{H}$, if $x \vee y = 1$, then either $x = 1$ or $y = 1$.[23]

Taking inquisitive algebras as the base of the pseudo-complemented twist-operation as defined in §4.4 produces a new class of twist-structures, call them Inq-PTS. Correspondingly, we obtain an extension of LPP, call this logic Inq-LPP. In particular, since the base algebra now satisfies the KP axiom, we add the following axiom to Inq-LPP:

AS17. $((\phi \to \bot) \to (\psi \vee \chi)) \to ((\phi \to \bot) \to \psi) \vee ((\phi \to \bot) \to \chi))$

What we have above is a version of the KP schema where the relevant negation is intuitionistic. To elucidate its validity, note that any pair from the carrier of Inq-PTS that comes to interpret AS17 will be such that the pair's first coordinate is necessarily identical to the top element given that the base algebra from which the twist-structure is generated now satisfies KP. Additionally, whenever ϕ is a non-negated formula, say α, then given the conversion between intuitionistic negation and toggle negation, we can derive from AS17 the following theorem: $\vdash_{Inq-LPP} (\sim\alpha \to (\psi \vee \chi)) \to ((\sim\alpha \to \psi) \vee (\sim\alpha \to \chi))$. Hence, we also obtain a restricted version of KP schema for toggle negation.

4.5.2 The carrier

A second parameter we can vary is the twist-structure's carrier set. In the case of FTS, there is no restriction on the carrier set; hence every element in the product lattice is included in the carrier of the resulting twist-structure. In the case of BTS, the restriction is that the two components of each pair need to be such that their meet is the bottom element. And in the case of PTS, we further restrict the carrier to pairs wherein either the first component is the pseudo-complement of the other, or vice versa. In general, after we have fixed a base algebra, we can explore different ways to determine the carrier set of the twist-structure.

[23]The two properties of being regularly generated and well-connected help us to further delimit the class of Heyting algebras suitable for modeling inquisitive logic. The property of being regularly generated directly follows from the fact that, since inquisitive valuations assign every atomic proposition to a member of $\mathfrak{H}_{\neg\neg}$, all formulas are thus assigned to elements from the subalgebra generated by $\mathfrak{H}_{\neg\neg}$. On the other hand, well-connectedness follows from the fact that inquisitive logic, just like intuitionistic logic, satisfies the disjunction property, that is, if $\phi \vee \psi$ is a theorem of the logic, then either ϕ must be a theorem or ψ must be a theorem of the logic.

In this subsection, I shall explore one immediate generalization of PTS which shares the same base as PTS but impose a further restriction on the carrier. As I briefly mentioned in §4.4.1, we may further curtail the carrier set such that, for any given pair, their two coordinates are pseudo-complements of each other. Call the resulting twist-structures *complemented twist-structures* (CTS), which are generated in the following way.

Given a Heyting algebra $\mathfrak{H} = \langle \mathcal{H}, \wedge, \vee, \rightarrow, \top, \bot \rangle$, we define a CTS, $\mathfrak{H}^{\bowtie c} = \langle \mathcal{H}^{\bowtie c}, \wedge^c, \vee^c, \rightarrow^c, \sim, 1, 0 \rangle$, as follows:

$$\mathcal{H}^{\bowtie c} = \{\langle x, x' \rangle | \ x' = x^* \text{ and } x = x'^*\}$$

$$\begin{aligned}
\langle x, x' \rangle \wedge^c \langle y, y' \rangle &\coloneqq \langle (x \wedge y)^{**}, (x \wedge y)^* \rangle \\
\langle x, x' \rangle \vee^c \langle y, y' \rangle &\coloneqq \langle (x \vee y)^{**}, (x \vee y)^* \rangle \\
\langle x, x' \rangle \rightarrow^c \langle y, y' \rangle &\coloneqq \langle (x \rightarrow y)^{**}, (x \rightarrow y)^* \rangle \\
\sim\langle x, x' \rangle &\coloneqq \langle x', x \rangle \\
1 &\coloneqq \langle \top, \bot \rangle \\
0 &\coloneqq \langle \bot, \top \rangle
\end{aligned}$$

Complemented twist-structures are closed under the above operations, given that for any x in a Heyting algebra, x^{**} and x^* are necessarily pseudo-complements of each other.

Theorem 4.5.4. *Every CTS is a Boolean algebra, and every Boolean algebra is isomorphic to a suitable CTS.*

Proof. we can prove that every CTS is a Boolean algebra by showing that elements of CTS satisfy the equational definitions of Boolean algebras (cf. Rasiowa 1974). Take the equation $X \vee^c \sim X = 1$ as an example, where X is any element from the carrier of CTS. Suppose $X = \langle x, x' \rangle$. We want to show that if $x' = x^*$ and $x = x'^*$, then $\langle x, x' \rangle \vee^c \langle x', x \rangle = 1 = \langle \top, \bot \rangle$. First, by definition, $\langle x, x' \rangle \vee^c \langle x', x \rangle = \langle (x \vee x')^{**}, (x \vee x')^* \rangle$. Next, since $\mathfrak{H}$ is a Heyting algebra, $(x \vee x')^* = x^* \wedge x'^*$. And given that x and x' are pseudo-complements of each other, $x^* \wedge x'^* = x' \wedge x'^* = \bot$. Thus, we have $\langle (x \vee x')^{**}, (x \vee x')^* \rangle = \langle \top, \bot \rangle = 1$.

Conversely, every Boolean algebra can also be regarded as coming from a CTS. Since every Boolean algebra is a Heyting algebra, then given any Boolean algebra $\mathfrak{B}$, we can generate a CTS based on $\mathfrak{B}$ and show that $\mathfrak{B}$ is isomorphic to the resulting CTS. This can be done by mapping every element x in $\mathfrak{B}$ to the pair $\langle x, x^* \rangle$ in the resulting CTS. $\square$

4.5.3 The operator

After we have determined a base and a carrier, depending on how we choose to define the logical operators, we may still derive different twist-structures. Here, I shall examine a variation of PTS, call them *left-pseudo-complemented twist-structures* (LPTS), which share the same base and carrier as PTS but with their operators defined differently.

We have remarked earlier that, under the inquisitive setting, there is an asymmetry between support and rejection in terms of the amount of information required: support in general demands more information than rejection. Since disjunction under inquisitive semantics plays the role of introducing alternatives, to support a disjunction, an information state needs to contain enough information to settle on at least one of the alternatives; by contrast, a conjunction does not introduce any alternatives, and thus to reject a conjunction, it is not mandatory for an information state to be able to reject one of the conjuncts. As a consequence, in devising a suitable class of twist-structures for modeling inquisitiveness, we define the support condition constructively but define the rejection condition non-constructively, thereby taking binary connectives to produce pairs wherein the second-coordinate is always the pseudo-complement of the first coordinate.

Now, as we go beyond the quest of modeling inquisitiveness, we may consider the asymmetry in constructiveness between support and rejection as representing different global standards for acceptance and rejection of information. Under PTS, we have a relatively stringent standard on information acceptance, that is, we support fewer things, compared to our standard on information rejection. In particular, to accept a disjunction, we must be able to support one of its disjuncts.

As a natural generalization, we may entertain reversing the stringency associated with information acceptance and rejection. We may adopt a relatively stringent standard on information rejection, that is, we reject fewer things, compared to our standard on information acceptance. In particular, to reject a conjunction, we must be able to reject one of its conjuncts.

Formally, we may entertain the converse structures of PTS under which the three binary connectives produce pairs wherein the first coordinate is always the pseudo-complement of the second coordinate. Without delving too deep into how the rejection condition of a conditional should be specified, let us assume that the second coordinates are defined in accordance with the BTS as devised for the constructive logic N3. Hence, we supply the following definition for left-pseudo-complemented twist-structures.

Given a Heyting algebra $\mathfrak{H} = \langle \mathcal{H}, \wedge, \vee, \rightarrow, \top, \bot \rangle$, we define a LPTS, $\mathfrak{H}_L^{\bowtie *} = \langle \mathcal{H}^{\bowtie *}, \wedge_L^*, \vee_L^*, \rightarrow_L^*, \sim, 1, 0 \rangle$, as follows:

$$\mathcal{H}^{\bowtie *} = \{\langle x, x' \rangle | \text{ either } x' = x^*, \text{or } x = x'^*\}$$

$$\langle x, x' \rangle \wedge_L^* \langle y, y' \rangle := \langle (x \vee y)^*, x \vee y \rangle$$
$$\langle x, x' \rangle \vee_L^* \langle y, y' \rangle := \langle (x \wedge y)^*, x \wedge y \rangle$$
$$\langle x, x' \rangle \rightarrow_L^* \langle y, y' \rangle := \langle (x \wedge y')^*, x \wedge y' \rangle$$
$$\sim\langle x, x' \rangle := \langle x', x \rangle$$
$$1 := \langle \top, \bot \rangle$$
$$0 := \langle \bot, \top \rangle$$

As an abnormality of LPTS, *modus ponens* is invalid. To illustrate, under LPTS, we have $\vDash (A \wedge \sim A) \rightarrow \bot$ and $\vDash ((A \wedge \sim A) \rightarrow \bot) \rightarrow \sim(A \wedge \sim A)$, but $\nvDash \sim(A \wedge \sim A)$. To see this, put $v(A) = \langle a, a' \rangle$. Then, $v((A \wedge \sim A) \rightarrow \bot) = \langle \bot, a \vee a' \rangle \rightarrow_L^* \langle \bot, \top \rangle = \langle \top, \bot \rangle = 1$ and $v(((A \wedge \sim A) \rightarrow \bot) \rightarrow \sim(A \wedge \sim A)) = \langle \top, \bot \rangle \rightarrow_L^* \langle a \vee a', \bot \rangle = \langle \top, \bot \rangle = 1$, but $v(\sim(A \wedge \sim A)) = \langle a \vee a', \bot \rangle \neq 1$.

The failure of *modus ponens* is often a telling trait of certain non-bivalent paraconsistent logics. For instance, both Priest's (1979) Logic of Paradoxes and the paraconsistent variant of N3 (Kapsner, 2015) fail to vindicate it. The reason why *modus ponens* fails in these logics is that semantic consequence is defined not as preservation of truth or support but as preservation of non-falsity or non-rejection. That is,

$$\Gamma \vDash \phi \text{ iff in every model and at every } w, \text{ if } w \nVdash_0 \psi \text{ for every } \psi \text{ in } \Gamma, \text{ then } w \nVdash_0 \phi.$$

Consequently, if the implication $\rightarrow$ is defined so that $\phi \rightarrow \psi$ can receive a non-0 value when ϕ also receives a non-0 value even if ψ receives 0, then *modus ponens* fails because the inference does not preserve non-falsity.

However, for LPTS, the situation is rather different, as semantics consequence in LPTS is still defined in terms of preservation of support. That is, $\phi \vDash \psi$ iff for all v such that $v(\phi) = 1$, $v(\psi) = 1$. As a result, LPTS is in fact not paraconsistent, since it does vindicate *ex contradictione quodlibet*, viz., $\phi, {\sim}\phi \vDash \psi$.[24] Nevertheless, since the rejection condition for conjunction is constructive in the sense that to reject a conjunction we must be able to reject either one of its conjuncts, the Law of Non-Contradiction (LNC), viz., ${\sim}(\phi \wedge {\sim}\phi)$ is in fact invalid in LPTS. For instance, as we have just seen in the above counterexample to *modus ponens*, $\nvDash {\sim}(A \wedge {\sim}A)$.

The contrast between the validity of *ex contradictione* and the invalidity of LNC has an intuitive explanation. Given that our semantic consequence is defined as preservation of acceptance. It means if one genuinely accepts both ϕ and ${\sim}\phi$, i.e., a contradiction, then one is expected to accept everything. But this does not equate to saying that for any sentence ϕ, we are expected to either reject it or reject its negation. Since we might not have enough information to countenance either option.

Moreover, since the failure of *modus ponens* is a direct consequence of how implication is defined in LPTS, we can retain the constructiveness of the rejection condition for conjunction while at the same time reinstate *modus ponens* by replacing $\rightarrow^*_L$ with a suitable implication. One immediate candidate for this role is the previous operator $\rightarrow^*_R$ as employed in PTS, viz., $\langle x, x' \rangle \rightarrow^*_R \langle y, y' \rangle \coloneqq \langle x \rightarrow y, (x \rightarrow y)^* \rangle$. Modified in this way, the resulting twist-structure is essentially identical to the PTS, given that the only difference between the two concerns the definitions for conjunction and disjunction, and as it turns out, $\wedge^*_L$ and $\vee^*_L$ are already definable in PTS as follows:

$$\langle x, x' \rangle \wedge^*_L \langle y, y' \rangle \coloneqq {\sim}({\sim}\langle x, x' \rangle \vee^*_R {\sim}\langle y, y' \rangle)$$

[24]Here is a quick proof of this result. Given how semantic consequence is defined in LPTS, it is impossible for there to be a valuation v such that both $v(\phi) = \langle x, x' \rangle = \langle \top, \bot \rangle$ and $v({\sim}\phi) = \langle x', x \rangle = \langle \top, \bot \rangle$. Consequently, anything follows.

$$\langle x, x' \rangle \vee_L^* \langle y, y' \rangle := {\sim}({\sim}\langle x, x' \rangle \wedge_R^* {\sim}\langle y, y' \rangle)$$

This means that PTS is already capable of modeling both constructive and non-constructive acceptance for disjunction as well as both constructive and non-constructive rejection for conjunction.

4.6 Conclusion

This paper highlighted two ways to reject the Law of Excluded Middle and explored a non-bivalent approach to inquisitiveness via the construction of a suitable class of twist-structures. Just as N3 extends the standard intuitionistic logic by introducing a notion of counterexamples in addition to the existing notion of proofs, the logic of pseudo-complemented propositions (or to be more precise, the logic Inq-LPP) extends the standard inquisitive logic by introducing a notion of rejection in addition to the existing notion of support. And just like N3, LPP also incorporates a toggle negation which validates Double Negation Elimination and functions as a switch that toggles between the support and rejection conditions of a sentence. However, unlike N3, LPP embodies an asymmetry in constructiveness between support and rejection: support is interpreted constructively but rejection is not. I have presented an algebraic semantics for LPP via pseudo-complemented twist-structures and proved completeness. The general utility of twist-construction—that is, as a way to convert a bivalent semantics into a non-bivalent one—has been further highlighted via exploration of some variations of PTS. By varying twist-construction along three dimensions—i.e., the base algebra, the carrier of the twist-structure, and the definitions of the logical operators—we can generate a wide range of twist-structures which could potentially serve diverse philosophical purposes.

Bibliography

[1] Anvari, A. (2021, May 20). *Content and Redundancy.* [Presentation slides]. Institut Jean-Nicod.

[2] van Benthem, J., Bezhanishvili, N., & Holliday, W. H. (2017). A bimodal perspective on possibility semantics. *Journal of Logic and Computation, 27*(5), 1353–1389.

[3] Bezhanishvili N., Grilletti G., Holliday W.H. (2019) Algebraic and Topological Semantics for Inquisitive Logic via Choice-Free Duality. In Iemhoff R., Moortgat M., de Queiroz R. (Eds) *Logic, Language, Information, and Computation. WoLLIC 2019* (pp. 35-52). Springer, Berlin, Heidelberg.

[4] Biezma, M. (2013). Only one at least: Refining the role of discourse in building alternatives. *University of Pennsylvania working papers in linguistics, 19*(1), 3.

[5] Biezma, M. and Rawlins, K. (2012). Responding to alternative and polar questions. *Linguistics and Philosophy, 35*(5), 361–406.

[6] Blutner, R. (2000). Some aspects of optimality in natural language interpretation, *Journal of Semantics, 17*, 189–216.

[7] Blutner, R. (2013). Pragmatics in Optimality Theory. In F. Liedtke & C. Schulze (Ed.), *Beyond Words: Content, context, and inference* (pp. 33-58). Berlin, Boston: De Gruyter Mouton.

[8] Brasoveanu, A. (2012). Modified numerals as post-suppositions. *Journal of Semantics, 30*(2), 155–209.

[9] Brasoveanu, A., & Szabolsci. A. (2013). Presuppositional too, postsuppositional *too*. In M. Aloni, M. Franke, & F. Roelofsen (Eds.), *The dynamic, inquisitive, and visionary life of ϕ, $?\phi$, and $\Diamond\phi$: A festschrift for Jeroen Groenendijk, Martin Stockhof, and Frank Veltman* (pp. 55–64). Amsterdam: Institute for Logic, Language, and Computation.

[10] Büring, D. (2003). On d-trees, beans, and b-accents. *Linguistics and Philosophy, 26*, 511–545.

[11] Chierchia, G., Fox, D. & Spector, B. (2009). Hurford's constraint and the theory of scalar implicatures: Evidence for embedded implicatures. In P. Egré, & G. Magri (Eds.), *Presuppositions and implicatures* (pp. 47-62). Cambridge, MA: MIT Pre

[12] Chierchia, G., Fox, D. & Spector, B. (2012). The grammatical view of scalar implicatures and the relationship between semantics and pragmatics. In C. Maienborn, P. Portner, & K. von Heusinger (Eds.), *Semantics: An International Handbook of Natural Language Meaning, Vol. 2* (pp. 2297–2332). Mouton de Gruyter.

[13] Ciardelli, I. (2016). *Questions in logic*. ILLC Dissertation Series.

[14] Ciardelli, I. & Roelofsen, F. (2011). Inquisitive logic, *Journal of Philosophical Logic, 40*(1), 55–94.

[15] Ciardelli, I. & Roelofsen, F. (2015). Inquisitive dynamic epistemic logic. *Synthese 192*(6), 1643–1687.

[16] Ciardelli, I. & Roelofsen, F. (2017). Hurford's constraint, the semantics of disjunction, and the nature of alternatives. *Natural Language Semantics 25*, 199–222.

[17] Ciardelli, I., Groenendijk, J., & Roelofsen, F. (2009). Attention! Might in inquisitivesemantics. In S. Ito, & E. Cormany (Eds.), *Semantics and Linguistic Theory (SALT) 19*.

[18] Ciardelli, I., Groenendijk, J. & Roelofsen, F., (2015). On the semantics and logic of declaratives and interrogatives, *Synthese, 192*, 1689–1728.

[19] Ciardelli, I., Groenendijk, J., & Roelofsen, F. (2018). *Inquisitive Semantics*. Oxford University Press.

[20] Ciardelli, I., Roelofsen, F., & Theiler, N. (2017). Composing alternatives. *Linguistics and Philosophy, 40*(1), 1–36.

[21] Cohen, A. & Krifka, M. (2011). Superlative quantifiers as modifiers of meta-speech acts. *The Baltic International Yearbook of Cognition, Logic and Communication, 6*.

[22] Coppock, E., & Beaver, D. (2011). Sole sisters. *In Semantics and Linguistic Theory, 21,* 197-217.

[23] Coppock, E., & Brochhagen, T. (2013). Raising and resolving issues with scalar modifiers. *Semantics and Pragmatics, 6*(3), 1-57.

[24] Condoravdi, C. (2015). Ignorance, indifference, and individuation with wh-ever. In L. Alonso-Ovalle & P. Menéndez-Benito (Eds.), *Epistemic indefinites: Exploring modality beyond the verbal domain* (pp. 213–243). Oxford: Oxford University Press.

[25] Constant, N. (2012). English rise-fall-rise: A study in the semantics and pragmatics of intonation. *Linguistics and Philosophy, 35*(5), 407–442.

[26] Dummett, M. (1959). Truth *Proceedings of the Aristotelian Society, 59* 141–162.

[27] Dummett, M. (1969). The reality of the past. *Proceedings of the Aristotelian Society, 69*, 239–258.

[28] Dummett, M. (1976.) What is a theory of meaning? (II), In G. Evans and J. McDowell (Eds.), *Truth and Meaning*. Oxford: Clarendon Press.

[29] von Fintel, K. (2001). Counterfactuals in a dynamic context. In M. Kenstowicz (Ed.), *Ken Hale: A Life in Language* (pp.123–152). Cambridge, MA: The MIT Press.

[30] Fitting, M. (1989) Bilattices and the theory of truth. *Journal of Philosophical Logic, 18*, 225–256.

[31] Gargov, G. (1999). Knowledge, uncertainty and ignorance in logic: bilattices and beyond, *Journal of Applied Non-Classical Logics, 9*, 195–283.

[32] Gazdar, G. (1979). *Pragmatics: Implicature, Presupposition, and Logical Form.* New York: Academic Press.

[33] Gillies, A. (2007). Counterfactual scorekeeping, *Linguistics and Philosophy, 30*(3), 329–360.

[34] Greenberg, Y. (2016). A novel problem for the likelihood-based semantics of even. *Semantics and Pragmatics, 9*(2), 1.

[35] Grice, H. P. (1975). Logic and conversation. In P. Cole & J. Morgan (Eds.), *Syntax and Semantics, Vol. 3*, (pp. 41–58).

[36] Groenendijk, J. (2009). Inquisitive semantics: Two possibilities for disjunction. In P. Bosch, D. Gabelaia, & J. Lang (Eds.), *Seventh international Tbilisi symposium on language, logic, and computation*. Springer.

[37] Groenendijk, J. & Stokhof, M.(1984). *Studies on the Semantics of Questions and the Pragmatics of Answers*, Joint Ph.D. thesis, University of Amsterdam, Department of Philosophy.

[38] Groenendijk, J., & Stokhof, M. (1991). Dynamic Predicate Logic. *Linguistics and Philosophy, 14*(1), 39–100.

[39] Hamblin, C. L. (1973). Questions in Montague English. *Foundations of Language, 10*(1), 41–53.

[40] Heim, Irene, 1983, On the projection problem for presuppositions, in M. Barlow, D. Flickinger, & M. Westcoat, (Eds.), *Second Annual West Coast Conference on Formal Linguistics* (pp. 114–126). Stanford University.

[41] Henderson, R. (2014). Dependent indefinites and their post-suppositions. *Semantics and Pragmatics 7*(6), 1–58.

[42] Hurford, J. (1974). Exclusive or inclusive disjunction. *Foundations of Language, 11*(3), 409–411.

[43] Iemhoff, R. (2019). Intuitionism in the Philosophy of Mathematics, *The Stanford Encyclopedia of Philosophy* (Fall 2020 Edition).

[44] Ippolito, M. (2020). Varieties of Sobel sequences. *Linguist and Philosophy 43*, 633–671.

[45] Johnstone, P. T. (1982) *Stone Spaces.* In *Cambridge Studies in Advanced Mathematics, Vol. 3.* Cambridge University Press, Cambridge.

[46] Jung, A., & Rivieccio, U. (2012). Priestley duality for bilattices. *Studia Logica: An International Journal for Symbolic Logic, 100*(1/2), 223–252.

[47] Kapsner, A. (2014). *Logics and Falsifications: A New Perspective on Constructivist Semantics.* Springer, Berlin, Heidelberg.

[48] Karttunen, L. (1977). Syntax and Semantics of Questions, *Linguistics and Philosophy, 1*(1), 3–44.

[49] Karttunen, L. and Peters, S. (1976). What indirect questions conventionally implicate. In *CLS 12: Papers from the Twelfth Regional Meeting* (pp. 351–368). Chicago Linguistic Society.

[50] Katzir, R. & Singh, R. (2013). Hurford disjunctions: Embedded exhaustification and structural economy. *Proceedings of Sinn und Bedeuting, 18*, 201–216.

[51] Kay, P. (1992). At least. *Frames, fields, and contrasts: New essays in semantic and lexical organization,* 309-331.

[52] Klecha, P. (2015). Two kinds of sobel sequences: Precision in conditionals. In U. Steindl, T. Borer, H. Fang, A. García Pardo, P. Guekguezian, B. Hsu, C. O'Hara, & I. Chuoying Ouyang (Eds.), *Proceedings of WCCFL* (Vol. 32, pp. 131–140). Somerville, MA: Cascadilla Press.

[53] Klecha, P. (2018). On unidirectionality in precisification. *Linguistics and Philosophy, 41*(1), 87–124.

[54] Klecha, P. (2021). On the Consistency and Reversibility of Certain Sequences of Counterfactual Assertions, *Mind*, 1–33.

[55] van Kuppevelt, J. (1995). Main structure and side structure in discourse, *Linguistics 33,* 809–833.

[56] Kracht, M. (1998) On extensions of intermediate logics by strong negation, *Journal of Philosophical Logic, 27*, 49–73.

[57] Lewis, D., K. (1973). *Counterfactuals*, Cambridge, MA: Harvard University Press.

[58] Lewis, K. S. (2018). Counterfactual Discourse in Context, *Noûs, 52*(3), 481–507.

[59] Mandelkern, M. (2019). Bounded Modality. *Philosophical Review, 128*(1), 1–61.

[60] Marty, P. & Romoli, J. (2022). Varieties of Hurford disjunctions. *Semantics and Pragmatics, 15*(3).

[61] Mayr, C. & Romoli, J. (2016). A puzzle for theories of redundancy: Exhaustification, incrementality, and the notion of local context. *Semantics and Pragmatics 9*(7), 1–48.

[62] Moss, S. (2012). On the Pragmatics of Counterfactuals, *Noûs, 46*(3), 561–586.

[63] Nakanishi, K., & Rullmann, H. (2009). Epistemic and concessive interpretation of at least. Slides for a paper presented at the meeting of the Canadian Linguistics Association, Carleton University.

[64] Nelson, D. (1949). Constructible falsity, *Journal of Symbolic Logic, 14*, 16–26.

[65] Nichols, C. (2017) Strict conditional accounts of counterfactuals. *Linguist and Philosophy, 40*, 621–645.

[66] Odintsov, S.P. (2004). On the representation of N4-lattices. *Studia Logica: An International Journal for Symbolic Logic, 76*, 385–405.

[67] Odintsov, S.P. (2008). *Constructive Negations and Paraconsistency.* Springer, Dordrecht.

[68] Onea, E. (2019). Underneath Rhetorical Relations: the Case of Result. In M. Zimmermann, K. von Heusinger, & E. Onea (Eds.), *Questions in Discourse. Volume 2: Pragmatics* (pp. 194-250). Leiden: Brill.

[69] Priest, G. (1979) The Logic of Paradox, *Journal of Philosophical Logic, 8*(1), 219–241.

[70] Priest, G. (2001). *An Introduction to Non-Classical Logic*. Cambridge University Press.

[71] Prince, A. & Smolensky, P. (1993). *Optimality Theory. Constraint interaction in generative grammar*, Cambridge, MA: MIT Press.

[72] Rasiowa, H. (1958). N-lattices and constructive logic with strong negation. Fundamenta Mathematicae, 46, 61-80.

[73] Rasiowa, H. (1974). *An algebraic approach to non-classical logics.* Amsterdam: North-Holland Publishing Company.

[74] Riester, A. (2019). Constructing QUD trees. In M. Zimmermann, K. von Heusinger, & E. Onea (Eds.), *Questions in Discourse. Volume 2: Pragmatics* (pp. 164–193). Leiden: Brill.

[75] Rivieccio, U. (2014). Implicative twist-structures. *Algebra universalis, 71*, 155–186.

[76] Roberts, C. (2012). Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics, 5,* 1–69.

[77] Rochemont, M. (2016). Givenness. In C. Féry & S. Ishihara (Eds.). *The Oxford Handbook of Information Structure* (pp. 41–63). Oxford: Oxford University Press.

[78] Roelofsen, F. (2013). Algebraic foundations for the semantic treatment of inquisitive content. *Synthese, 190*(1), 79–102.

[79] Roelofsen, F. (2013). A bare bone attentive semantics for might. In M. Aloni, M. Franke, & F. Roelofsen. (Eds.), *The dynamic, inquisitive, and visionary life of ϕ $?\phi$ and $\Diamond\phi$: a festschrift for Jeroen Groenendijk, Martin Stokhof, and Frank Veltman* (pp. 190–215). ILLC publications.

[80] Roelofsen, F. (2015). The semantics of declarative and interrogative lists. Manuscript, ILLC, University of Amsterdam.

[81] Romoli, J., Santorio, P., & Wittenberg, E. (2021) Alternatives in counterfactuals: what is right and what is not. Unpublished Ms.

[82] Rooth, M. (1985). *Association with focus.* Ph.D. thesis, University of Massachusetts, Amherst.

[83] Sauerland, (2004). Scalar implicatures in complex sentences. *Linguistics and Philosophy, 27*, 367–91

[84] Simons, M. (2001). Disjunction and alternativeness. *Linguistics and Philosophy, 24,* 597–619.

[85] Singh, R. (2008). On the interpretation of disjunction: Asymmetric, incremental, and eager for inconsistency. *Linguistics and Philosophy, 31,* 245–260.

[86] Schlenker, P. (2009). Local contexts. *Semantics and Pragmatics 2*, 1–78.

[87] Schwarzschild, R. (1999). Givenness, AVOIDF and other constraints on the placement of accent. *Natural Language Semantics, 7*(2), 141–177.

[88] Stalnaker, R. C. (1968). A Theory of Conditionals. In N. Rescher (Ed.), *Studies in Logical Theory* (pp. 98–112). Oxford: Basil Blackwell.

[89] Stalnaker, R. (1978). Assertion. In P. Cole (ed.), *Syntax and semantics 9: Pragmatics,* 315–322. Academic Press.

[90] Stalnaker, R. C. (1981). A defense of conditional excluded middle. In W. Harper, R. C. Stalnaker & G. Pearce (Eds.), *Ifs* (pp. 87-104). Reidel.

[91] Tian, Y., Breheny, R., & Ferguson, H. J. (2010). Why we simulate negated information: A dynamic pragmatic account.

[92] Tian, Y., Ferguson, H. & Breheny, R. (2016). Processing negation without context – why and when we represent the positive argument. *Language, Cognition and Neuroscience 31*(5), 683–698

[93] Vakarelov, D. (1977). Notes on N-lattices and constructive logic with strong negation. *Studia Logica: An International Journal for Symbolic Logic, 36*, 109–125.

[94] Vakarelov, D. (2005). Nelson's Negation on the Base of Weaker Versions of Intuitionistic Negation. *Studia Logica: An International Journal for Symbolic Logic, 80*, 393–430.

[95] Veltman, F. (1996). Defaults in update semantics, *Journal of Philosophical Logic, 25*, 221–261.

[96] Veltman, F. (2005). Making counterfactual assumptions, *Journal of Semantics, 22*(2), 159–180.

[97] Wansing, H. (2014). *Connexive logic*. The Stanford Encyclopedia of Philosophy.

[98] Westera, M. (2017). *Exhaustivity and intonation: a unified theory.* Ph. D. thesis, ILLC, University of Amsterdam.

[99] Westera, M. (2019). Hurford disjunctions: An in-depth comparison of the grammatical and the pragmatic approach. Unpublished Ms. Leiden University.

[100] Willer, M. (2013). Dynamics of epistemic modality. *The Philosophical Review, 122*, 45–92.

[101] Willer, M. (2017). Lessons from Sobel Sequences, *Semantics and Pragmatics, 10*(4), 1–57.

[102] Willer, M. (2018). Simplifying with free choice. *Topoi 37*, 379–392.

[103] Yalcin, S. (2007). Epistemic Modals. *Mind, 116*(464), 983–1026.

[104] Yalcin, S. (2015). Epistemic Modality De Re. *Ergo: An Open Access Journal of Philosophy, 2*, 475–527.