A Conditional Perspective of Belief Revision

Dissertation

zur Erlangung des Grades eines

Doktor der Naturwissenschaften (Dr.rer.nat.)

der Technischen Universität Dortmund
an der Fakultät für Informatik

vorgelegt von
Meliha Sezgin

Dortmund

2023

Tag der mündlichen Prüfung: 18.07.2023

Dekan: Prof. Dr.-Ing. Gernot A. Fink

Gutachterin: Prof. Dr. Gabriele Kern-Isberner
Gutachter: Prof. Dr. Eduardo Fermé

Acknowledgements

I look back fondly at the beginning of my time as a Ph.D. student. I had many questions and was joyfully excited. Now that this time is ending answering these questions was mainly a gateway to new ones. However, the joy always remained due to the great working environment, and I want to thank everyone who was part of it.

I want to thank Gabriele Kern-Isberner for allowing me to do my Ph.D. at LS1 and for being a dedicated advisor in many ways. She listened to me, discussed ideas, gave feedback, and taught me to write scientific papers. She introduced me to the research world and helped me become a scientist. Thank you, Gaby!

Sincere thanks go to Eduardo Fermé, who inspired me to work on parameterized belief change. He helped and encouraged me along the winding paths of research, and his enthusiasm for Belief Revision was more than once contagious.

The dissertation itself is the result of successful collaboration and careful guidance. Special thanks go to Christoph Beierle, who always gave me valuable feedback and whose considerate and dedicated manner made the collaboration pleasant and fruitful. Many thanks go to my colleagues Andre, Marco, and Diana of the Information Engineering group for exchanging ideas and discussions and to everyone from the Lehrgebiet Wissensbasierte Systeme of the Fernuniversität Hagen for the helpful feedback during our various meetings online, in-person and at conferences. Also, I want to thank the whole team at LS1 for providing a friendly and productive research environment – special thanks to Christopher for the walks together and to Anja for her reliability. Also, I want to thank the inspiring women from the "Frauennetzwerktreffen" all over the faculty.

If I can say today that I have completed my time as a Ph.D. student, this success can be credited to my hard work and discipline. But that would only be part of the truth. I am a woman with an immigrant background from parents of Turkish descent, the fact that I am where I am today is mainly due to the love and support my parents, and whole family gave me. Thank you, thank you so much. Especially to my powerful, wise, funny, [...] sister Selen, thank you for everything. So, I want to shout out to the whole "crew" and my great-great friends: *We made it!*

Finally, one person remains to thank, someone who never allowed any doubts about me, even if the biggest ones came from myself, and whose love I could always be sure of. Thorsten, you are my piece of happiness, my team- and my soulmate.

Zusammenfassung

Die Wissensrevision ist ein Teilbereich der Wissensrepräsentation und Wissensverarbeitung, in dem untersucht wird, wie die Überzeugungen eines intelligenten Agenten in Reaktion auf neue Informationen rational revidiert werden können. Es gibt verschiedene Ansätze zur Wissensrevision, aber ein bekannter Ansatz ist das AGM-Modell, das auf die Arbeit von Alchourrón, Gärdenfors und Makinson zurückgeht. Dieses Modell bietet Axiome, die wünschenswerte Eigenschaften von Wissensrevisionsoperatoren definieren, die die Wissensmenge des Agenten, d.h. eine Menge von propositionalen Formeln, manipulieren. Eine berühmte Erweiterung des klassischen AGM-Rahmens der Wissensrevision ist der Ansatz von Darwiche und Pearl zur iterierten Wissensrevision. Sie entdeckten, dass der Schlüssel zu rationalem Verhalten unter Iteration in der adäquaten Erhaltung von konditionalem Wissen liegt, d.h. von Überzeugungen, die der Akteur bereit ist, im Lichte (hypothetischer) neuer Informationen zu akzeptieren. Daher führten sie Operatoren zur Wissensrevision ein, die auf den Wissenszustand des Agenten einwirken und aus konditionalem Wissen aufgebaut sind. Kern-Isberner axiomatisierte ein Prinzip der konditionalen Erhaltung für die Wissensrevision, das den Kern der angemessenen Behandlung von konditionalem Wissen während der Revision erfasst. Dieses mächtige Axiom bietet den notwendigen konzeptionellen Rahmen für die Revision von Wissenszuständen mit Mengen von Konditionalen als Input und zeigt, dass konditionales Wissen subtil, aber wesentlich für die Untersuchung der Prozesse bei der Wissensrevision ist.

Diese Arbeit zeigt eine konditionale Perspektive der Wissensrevision für verschiedene Szenarien der Wissensrevision auf. Im ersten Teil führen wir einen Begriff der Lokalität für Wissensrevisionsoperatoren auf semantischer Ebene ein und untersuchen ihn. Damit nutzen wir die Eigenschaften von Konditionalen, die es uns erlauben, lokale Fälle aufzustellen und entsprechend dieser Fälle zu revidieren, d.h. die Komplexität der Revisionsaufgabe erheblich zu reduzieren. Im zweiten Teil betrachten wir Wissensrevision im Hinblick auf zusätzliche Metainformationen, welche die Eingabedaten begleiten. Wir demonstrieren die Vielseitigkeit und Flexibilität von Konditionalen für Wissensrevision, indem wir die parametrisierte Eingabe für zwei bekannte parametrisierte Revisionsoperatoren auf eine konditionale reduzieren. Unsere Ergebnisse zeigen, dass die Berücksichtigung von konditionalem Wissen bei der Revision neue Einsichten in die Dynamik von Wissensrevisionsprozessen bietet.

Abstract

Belief Revision is a subarea of Knowledge Representation and Reasoning (KRR) that investigates how to rationally revise an intelligent agent's beliefs in response to new information. There are several approaches to belief revision, but one well-known approach is the AGM model, which is rooted in work by Alchourrón, Gärdenfors, and Makinson. This model provides a set of axioms defining desirable properties of belief revision operators, which manipulate the agent's belief set represented as a set of propositional formulas.

A famous extension to the classical AGM framework of Belief Revision is Darwiche and Pearl's approach to iterated belief revision. They uncovered that the key to rational behavior under iteration is adequate preservation of conditional beliefs, i.e., beliefs the agent is willing to accept in light of (hypothetical) new information. Therefore, they introduced belief revision operators modifying the agent's belief state, built from conditional beliefs. Kern-Isberner fully axiomatized a principle of conditional preservation for belief revision, which captures the core of adequate treatment of conditional beliefs during the revision. This powerful axiom provides the necessary conceptual framework for revising belief states with sets of conditionals as input, and it shows that conditional beliefs are subtle but essential for studying the process of belief revision.

This thesis provides a conditional perspective of Belief Revision for different belief revision scenarios. In the first part, we introduce and investigate a notion of locality for belief revision operators on the semantic level. Hence, we exploit the unique features of conditionals, which allow us to set up local cases and revise according to these cases, s.t., the complexity of the revision task is reduced significantly. In the second part, we consider the general setting of belief revision with respect to additional meta-information accompanying the input information. We demonstrate the versatility and flexibility of conditionals as input for belief revision operators by reducing the parameterized input to a conditional one for two well-known parameterized belief revision operators who are similarly motivated but very different in their technical execution.

Our results show that considering conditional beliefs as input for belief revision operators provides a gateway to new insights into the dynamics of belief revision.

Contents

Chapter 1

Introduction

The first chapter of this thesis provides an introduction to the results presented in the course of the following investigations. The first section is dedicated to the broader research context and the general motivation of this thesis. In Section 1.2, we outline this work's goals by addressing research questions that provide deeper insights into the main contributions presented in this thesis. Also, we give an overview of the organization and overall structure of the following text. The last section presents publications related to this thesis authored or co-authored by the author.

1.1 Research Context and Motivation

This thesis is located in the area of Knowledge Representation and Reasoning (KRR), a major subfield of Artificial Intelligence (AI) that is concerned with the fundamental issues of representing knowledge in a way that a computer system can use it to reason about the world [101]. This involves creating models of the world that can be manipulated and queried to derive new information or make decisions. Some common approaches to KRR include logic-based approaches, such as propositional and first-order logic, rule-based systems, and semantic networks. KRR has applications in various fields, including natural language processing, robotics, and expert systems [19]. It is also used in several practical applications, such as in medical diagnosis and treatment or automated planning and scheduling [8].

Reasoning, i.e., the process of drawing conclusions from existing knowledge, plays a significant role in today's research about AI [19]. *Non-monotonic Reasoning* is a subfield of Knowledge Representation and Reasoning that enables intelligent systems to adapt their beliefs and conclusions in response to new information. Unlike classical reasoning, it allows for modifying or retracting previously made inferences based on new or conflicting evidence [21]. This is crucial for applications such as expert systems, decision support systems, and automated reasoning.

In this thesis, we examine essential aspects of how an *intelligent agent* should change her beliefs in light of new information. The branch of research that answers this question is *Belief Revision*, a subfield of Non-monotonic Reasoning. Note that, in the context of this thesis, an intelligent agent is an AI system capable of perceiving its environment, processing information, and taking actions to achieve specific goals. One of the key features of intelligent agents is their ability to reason about the information they receive and make decisions based on that reasoning. Dealing with change in the world is an essential aspect of these agents, so understanding the phenomenon of change is a concern for designing intelligent systems in AI [20].

Theory of Belief Revision. *Belief revision* is the process of revising one's beliefs in response to new information [33]. In a typical belief revision scenario, an agent receives new information that makes her change her beliefs. In the principal case, where the new information contradicts her initial beliefs, the agent must withdraw some of the old beliefs to accommodate the new information and maintain a consistent worldview. This might lead to interactions between old beliefs and new information, which have to be monitored [90]. The study of the process of belief revision gave rise to the exciting research area *Belief Revision*, which can be traced back to the early 1980s. However, the article widely considered to mark the birth of the field is the seminal work [2] of Alchourrón, Gärdenfors, and Makinson, where rationality postulates that constrain the outcome of belief change operations are introduced. These postulates lay the AGM framework's foundation, named after its three founders' initials. The rationality postulates for belief revision that are crucial for *AGM revision operators* have proven to provide valuable guidelines for the belief revision of intelligent agents in many cases, and the AGM framework is to this date the dominant framework in Belief Revision [96]. Formally, the AGM approach employs representations of knowledge in so-called *belief sets* that are affected by *belief* *change operators* [33]. A set of prior beliefs of an agent is changed in the light of a new belief via a belief revision operator, and we receive a posterior belief set. One of the key findings in Belief Revision was presented by Katsuno and Mendelzon in [60], who showed that *total preorders* (TPOs) over possible worlds with specific features are essential for AGM revision operators, i.e., the representation of beliefs via TPOs is a necessary and sufficient condition for the definition of AGM revision operators.

Even though the AGM framework has often provided valuable guidelines to define a theory of belief change, there are a few shortcomings. One of the main problems is that the AGM paradigm lacks guidelines for handling repeated, so-called *iterated belief revision*. There were several approaches to address the problem [117, 79], one of the most influential proposals is the work of Darwiche and Pearl (DP for short) in [29], who proposed, in extending the original AGM framework, to consider not only certain beliefs in the revision process but also underlying preferential information, i.e., enrich the former *belief sets* and consider the whole *belief state*, also called *epistemic states* [48], as complex representations of an agent's cognitive state. Building upon the central result on AGM revision operators from Katsuno and Mendelzon [60], Darwiche and Pearl showed that TPOs are the essential meta-structure for representations of belief states which provide the necessary means to define rational belief revision on epistemic states. Before we discuss the crucial role of conditionals in Belief Revision, we shortly give an overview of representations of belief states that are important within the scope of this thesis.

There are several ways to represent the TPOs defining an agent's belief state [117, 45, 42]; we focus in this thesis on qualitative frameworks that encode relative plausibility via total preorders on possible worlds [60, 29] and on the semi-quantitative framework of *ranking functions* or *ordinal conditionals functions* (OCFs) firstly introduced by Spohn [117, 119] which enrich the qualitative TPOs by assigning numbers to plausibility levels. The employment of numbers makes it easier to compare different plausibility levels within the belief representation framework and equips the corresponding belief revision operators with a powerful arithmetic[1].

[1] Note that various works recognized the expressiveness and benefits of representing belief states by ranking functions in the context of belief revision [119, 64, 106].

Conditionals in Belief Revision. One of the main observations that led to the now famous *DP framework for iterated Belief Revision* [29] was that the reasonable minimization of changes in conditional beliefs, i.e., beliefs an agent is prepared to adopt conditioned on any hypothetical evidence [15, 50], is a crucial feature when it comes to designing rules for rational (iterable) changes of beliefs. Conditionals $(B|A)$, to be read as *"if A, the premise, is true, then we can conclude the consequent B (plausibly)"*, extend the narrow framework of propositional logic and can be seen as a gateway to the revision of belief states as they provide the necessary means to observe the very process of belief revision by making it possible to compare worlds outside an agent's belief set according to her preference. More precisely, the relation between belief revision and conditionals is expressed by the so-called *Ramsey Test*, which states that the acceptance of a conditional is equivalent to the acceptance of its consequent after a revision with the conditional's premise. Note here that the agent's belief state is taken explicitly into account as it is necessary to define the acceptance of a conditional belief [48]. The necessity to employ TPOs as belief representation frameworks, which arises from the characterizations of AGM revision operators in [60, 29], corresponds naturally to the fundamental challenge to preserve conditional beliefs rationally during the revision process posed by Darwiche and Pearl because both frameworks rely on the notion of comparison according to preference. Thus, conditional beliefs are not merely logical artifacts but play an essential role by guiding the belief revision process naturally by considering an agent's cognitive state at a given time, which places them at the core of Belief Revision. Moreover, conditionals are *defeasible*, i.e., the consequent of a conditional may be overruled in the light of new information. This makes them fundamentally different from material implications $A \Rightarrow B$ from classical logic [76, 84, 44].

Darwiche and Pearl acknowledged the relevance of conditional beliefs and their preservation. In [29], they proposed additional postulates that regulate changes in conditional beliefs to provide coherence under successive belief changes. The need for preserving conditional beliefs was widely recognized in the field of Belief Revision (for example, in [17, 79, 121]). Here the work [63] from Kern-Isberner stands out since it goes beyond the DP postulates for preserving conditional beliefs by providing a complete formalization of a rational *principle of conditional preservation* (PCP) [63, 65] which implies the DP postulates and extends them for conditional revision, i.e., the process of revising one's beliefs not only in the light of new propositional but also conditional information. While Darwiche and Pearl explicitly took conditional beliefs into account, their approach lacks guidelines for changing beliefs if the new information itself is conditional. The (PCP) here provides the whole picture and constitutes an essential paradigm for guiding (iterated) revision of an agent's belief state. The (PCP) leads to the definition of so-called *c-revisions*, conditional revision operators capable of revising ranking functions with sets of conditionals rationally by obeying the (PCP). To the best of our knowledge, c-revisions are the only revision operator capable of revising with sets of conditional simultaneously while respecting the guidelines for iterated Belief Revision.

Before presenting the aim and structure of this thesis in the next section, we motivate and explain the conditional perspective of Belief Revision, which drives our research. Note that, for the remainder of this Chapter, we consider belief revision on belief states in the sense of Darwiche and Pearl.

Conditional Perspective of Belief Revision. In everyday life, we often encounter situations where the outcome depends on certain conditions, and we need to make decisions or draw conclusions based on the information available to us [58]. Conditionals allow us to represent and reason about these situations in a logical and structured way by specifying the conditions under which a specific outcome or event is likely to occur [31]. Their vital role in reasoning about uncertainty has been acknowledged many times [87, 23, 102] Thus, they can be seen as one of the hallmarks of human reasoning because they enable us to reason about uncertain or hypothetical situations. Moreover, understanding conditionals is also essential for effective communication, as they are a common feature of natural language and play a significant role in conveying meaning and expressing uncertainty [22, 104]. Overall, looking at the general significance of conditional beliefs, it is not surprising how essential conditionals are also for the area of Belief Revision.

Belief revision operators, broken down quite simply, are functions that take as input two kinds of objects, a representation of an agent's initial belief state and a piece of new input information. Thus, the central objects of interest within the area of (iterated) Belief Revision are epistemic states on the one hand and input information on the other. It holds that conditionals play a subtle but crucial role for both. Darwiche and Pearl [29] pointed out their importance by connecting them to belief states, which encode an agent's conditional beliefs, and providing hints on the importance of preserving them. Then, Kern-Isberner completed the picture by axiomatizing a principle of conditional preservation that leads to a versatile revision mechanism, the c-revision mentioned above, which takes sets of conditionals as input and treats them appropriately, i.e., in such a way that conditional structures and interactions within the initial belief state and also, within the input are respected and preserved as far as possible by obeying the (PCP). In this thesis, taking on a conditional perspective in the context of Belief Revision means paying close attention to and exploiting the expressiveness and internal strengths of conditionals. Thus, revealing the underlying interplay of conditional beliefs in the process of belief revision and profiting from their powerful semantics is a promising path that motivates the results presented in this thesis.

In the following chapter, we discuss the aim and structure of this work and allocate our results based on research questions.

1.2 Aim and Structure

In this section, we present the general aim of this thesis and phrase major research questions that lead to the investigations and results presented in the following parts. We conclude this section by giving an overview of the thesis structure.

Research Questions and Contributions. This thesis aims to illustrate how conditional information as input for revision operators advances the fundamental understanding of belief revision processes in various scenarios. Thus, our investigations are guided by what we call a conditional perspective of Belief Revision, i.e., a perspective where we use conditionals as logical entities that allow us to directly influence changes in the agent's belief state during the belief revision process without taking the usual detour via propositional input information, that is typical for most operators. This leads to insightful investigations presented in the course of this thesis.

Note that while the importance of conditional information in the representation of an agent's belief state is beyond doubt and has been widely recognized in the Belief Revision community [1, 81, 63] leading to many propositional revision operators [15, 82, 13], there has been surprisingly little research on operators that are capable of revising with conditional information[2]. Here, c-revisions by Kern-Isberner stand out since they are capable of revising with sets of conditionals while respecting a fundamental principle of conditional preservation. Most of the results in this thesis can be realized via c-revisions. Thus, they display a fundamental cornerstone of this work and are helpful in many different ways to illustrate our investigations and provide deeper insights into the interplay of conditional beliefs.

In the context of our research, we advocate the employment of conditionals as input for belief revision operators, and we believe the advantages are manifold. In this thesis, we focus on two different scenarios of belief revision and thus discuss them in two separate parts. In the first part, we focus on specific dynamics when revising with a set of conditionals and investigate a *Kinematics principle for Belief Revision*. In the second part, we take propositional revision with some additional input and transfer it to the framework of conditional revision investigating so-called a *Parameterized Belief Revision*. We begin each part by presenting the specific research setting and introducing an (advanced) belief revision problem for which this thesis provides novel contributions. Moreover, we pose specific research questions, which guide the thorough investigations presented in the corresponding part. In the following, we give an overview of our main contributions and address general research questions that lead to our diverse lines of research.

Part I: *Revising w.r.t. Local Cases – A Kinematics Principle for Belief Revision*

One of the characteristic properties of conditionals, which distinguishes them fundamentally from propositional sentences, is that they are able to deal with information that is linked to a specific context. In the general setting of belief revision, a conditional translates to statements of the form *"if A is true, then we can conclude B (plausibly)"*, i.e., we can plausibly assume that B holds in the context A, a notion of conditionals that is already present in the famous Ramsey Test [95, 120]. It is evident that including the specific context given by conditional information

[2]Some notable extensions that investigate revision with a single conditional can be found in [18, 27].

simplifies the revision process since it enables us to exclude parts of our previous beliefs unrelated to the new input. Especially in cases where we have to process information from different mutually exclusive contexts at the same time, this leads to a reduction in the complexity of the revision task, which is typical for human reasoners [88]. This gives rise to the following research questions:

> *How does the specific context of information affect the revision task, and how can we benefit from the inclusion of exclusive contexts in the revision process?*

In probability theory, the concept of conditionalization to a specific context is well-established. There have been several approaches in probabilistic Belief Revision to connect the revision of beliefs with conditionalization [88, 57]. In this part of the thesis, we present and investigate a Kinematics principle for Belief Revision in the context of qualitative and semi-quantitative frameworks, which states an invariance property connecting belief revision and conditionalization. We show that in the context of ranking functions, under particular prerequisites, c-revisions are capable of dealing with information from exclusive contexts independently and investigate the advantages of our previously proposed Kinematics principle in the context of propositional revision. In the qualitative context, the Kinematics principle represents a target for qualitative revision operators, which provides ground for our investigation of a qualitative conditionalization operator and a suitable transformation schema between ranking functions and purely qualitative representations of epistemic states, leading to a novel conditional revision operator taking sets of conditionals as input in the qualitative framework which is driven by c-revisions. We conclude our investigations of the Kinematics principle by drawing connections between our results and the conditional revision operator provided by Chandler and Booth [27].

Part II: *Revising w.r.t. Meta-Information – Parameterized Belief Revision*

In most real-life settings, meta-information accompanying new information impacts how we change our beliefs [125, 116, 93]. The reliability of the source and our prior beliefs can affect how we revise our beliefs based on new information. Additionally, if the new information appears highly implausible in light of our prior beliefs, accepting it without reservation may be difficult. To deal with such additional information, we need to enrich the standard framework of belief revision operators so that they can take parameterized information into account. The parameterized belief revision operators we focus on in this thesis, *Revision by Comparision* by Fermé and Rott [34] and *Bounded Revision* by Rott [100], are able to revise with propositional inputs that are accompanied by meta-information in the general form of a propositional reference sentence. Although their corresponding mechanisms differ at a crucial point, both revision operators are motivated by employing a reference sentence as a parameter that impacts the change between the prior and posterior belief state. Yet, the parameter remains on a vague meta-level, and it is difficult to see how it affects the revision mechanism. This gives rise to the following research question:

> *How can we incorporate the parameterized information into the framework of (conditional) Belief Revision so that the relation between input and reference information during the change process is evident?*

In Part II, we identify elegant reformulations of Revision by Comparison and Bounded Revision as conditional c-revisions and discuss their advantages, thereby illustrating the versatility and expressiveness of conditionals as input information, which transfer the parameterized input information to the directly usable object level, at least for c-revisions. On our way to achieving this goal, we thoroughly investigate both belief revision mechanisms, Revision by Comparison, and Bounded Revision, clarify their corresponding change mechanism, and provide new elegant representation theorems. In particular, we demonstrate how the use of ranking functions can fully accommodate each of the distinct features of the operations. We conclude this part by comparing both mechanisms, where we investigate similarities and differences based on the insights we gained during our thorough investigations.

Overview. This work is divided into two parts in 11 chapters, where some chapters depend on others. The overall outline of the thesis is illustrated in Figure 1.1, where the arrows represent dependencies among the consecutive chapters. We start with a chapter that recalls the basics for both parts:

Formal Preliminaries of This Thesis. We start with stating essential formal preliminaries that introduce epistemic states, the area of Belief Revision, and in particular, conditionals in Belief Revision in Chapter 2. Basic definitions and notations

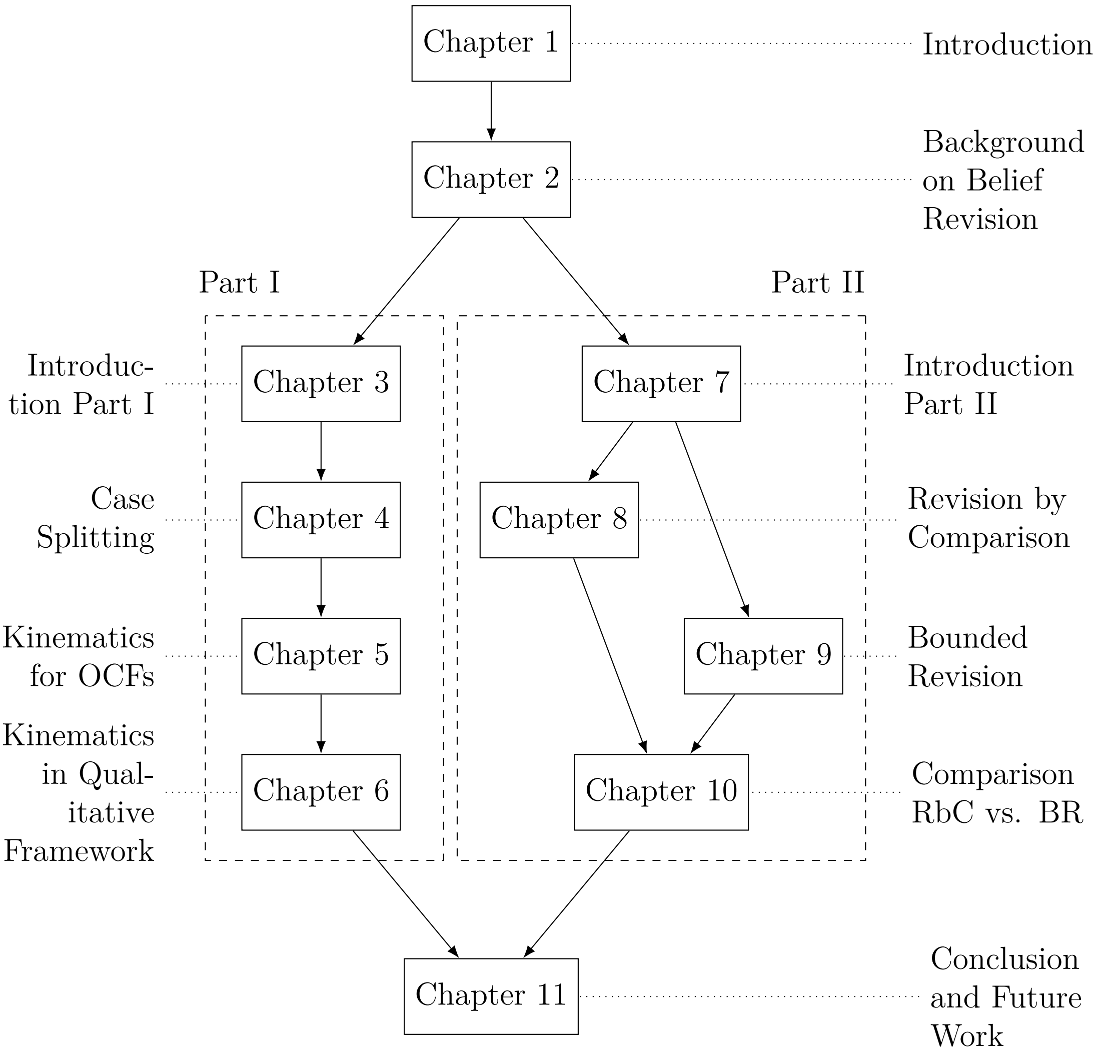

Figure 1.1: Organization of the thesis.

from propositional and conditional logic are presented in Section 2.1. Moreover, we present an extension to conditional logic, so-called *weak conditionals*, which are similar to conditionals to a certain extent. Section 2.2 recalls the basics of one-step belief revision. The extension of one-step belief revision to the framework of iterated belief revision in the style of Darwiche and Pearl [29] is presented in Section 2.3. In Section 2.4, we discuss conditionals in the context of epistemic states and present the two main frameworks used in this thesis to represent an agent's belief state. Section 2.5 concludes this chapter and presents the fundamental principle of conditional preservation, which guides the revision with sets of conditionals and eventually leads to the definition of c-revisions. We also examine the connection between c-revision with sets of weak conditionals and iterated contraction operators.

Main Body of This Thesis. The main body of this thesis is organized into two parts. Every part starts with an introduction and an overview of the part's contents. Moreover, we summarize formal preliminaries solely relevant to the corresponding part and embed our results in a broader research context by presenting related work. We conclude each part with an intermediate summary.

The first part presents and investigates the Kinematics principle for belief revision in qualitative and semi-quantitative frameworks of belief representation. We start with thoroughly investigating *case splittings*, a crucial prerequisite of the Kinematics principle in Chapter 4. Then we show in Chapter 5 that c-revisions satisfy this advanced principle of belief revision and discuss its impacts for the case of propositional belief revision and the corresponding c-revision with the complete set of conditionals. Transferring the Kinematics principle to the qualitative framework leaves us with the need for a concept of qualitative conditionalization and qualitative revision mechanisms for sets of conditionals. In Chapter 6, we define and investigate both concepts mentioned above, eventually leading to qualitative c-revisions, which rely crucially on a suitable transformation scheme between the different frameworks of belief representation. We conclude this chapter by comparing qualitative c-revisions with another conditional revision operator from [27] and show that this operator is also compatible with the Kinematics principle.

In the second part, we investigate parameterized revision operators. In Chapter 8, we investigate the mechanism and relevant properties of Revision by Comparison [34] in the context of plausibilistic TPOs and provide an elegant reformulation via a representation theorem. This provides the ground for the subsequent investigations of Revision by Comparison in the context of ranking functions. This leads us to the characterization of Revision by Comparison as a c-revision with a set of weak conditionals. The following Chapter 9 deals with Bounded Revision [100], which represents another parameterized belief revision mechanism that shares some similarities with the previously presented Revision by Comparison but also crucially differs from it in its characteristics. The investigations of Bounded Revision lead to a characterization of its change mechanism as a c-revision with a designated conditional. We close this part with a thorough comparison between parameterized belief revision mechanisms in the qualitative and semi-quantitative frameworks in Chapter 10, where we provide new insights about similarities and differences.

The thesis closes with a summary and discussion of future work.

1.3 Publications Containing Results of the Thesis

During the last four years when this thesis was written, I was involved in multiple research activities and worked jointly with various colleagues. I authored multiple publications jointly with these colleagues in the context of these research activities. In the following, I give an overview of those previous publications that contain results also presented in this thesis and point out my contributions to these publications.

Part I: The Kinematics Principle in Belief Revision.

When I first started my Ph.D. position, me and Gabriele Kern-Isberner had many fruitful discussions about quantitative approaches to reasoning, like the *principle of maximum entropy*[3], during which Gabriele Kern-Isberner gave the initial impetus to examine a (qualitative) *Kinematics principle* in the context of belief revision. The study of this principle proved to be highly insightful. It resulted in one conference and two journal papers in high-ranked and well-reputed journals in the field of Artificial Intelligence. These journal papers were joint work with Gabriele Kern-Isberner and Christoph Beierle, and I list them here chronologically.

> [107] M. Sezgin, G. Kern-Isberner, *Generalized ranking Kinematics for iterated belief revision*, in: R. Barták, E. Bell (Eds.), Proceedings of the Thirty-Third International Florida Artificial Intelligence Research Society Conference, Originally to be held in North Miami Beach, Florida, USA, May 17-20, 2020, AAAI Press, 2020, pp. 587–592.

A largely extended version of this conference paper was published in 2021 in the journal *Annals of Mathematics and Artificial Intelligence* as

[3]I investigated numerical solution processes for the principle of maximum entropy in the context of my master thesis, which Gabriele Kern-Isberner supervised.

> [111] M. Sezgin, G. Kern-Isberner, C. Beierle, *Ranking Kinematics for revising by contextual information*, Ann. Math. Artif. Intell. 89 (10-11) (2021) 1101–1131.

The initial idea and conceptual outline for *Generalized Ranking Kinematics* came from Gabriele Kern-Isberner; this encompasses the draft and design of the central postulate in [107]. I developed the technical elaborations, especially the definitions, the algorithm, and proofs of the theorems and propositions in the conference paper [107] and its journal version [111]. Moreover, [111] provides meaningful extensions of the Kinematics principle, which I designed and elaborated. These results are presented in Chapter 5. In particular, the thorough investigation of the case splittings, one of the prerequisites of the Kinematics principle, and the algorithm in [107] are my contributions. They are presented in Chapter 4.

Most of the proofs in [111] use strategies for c-revisions which were first introduced in [111]. The draft, design, and conceptual outline for these strategies came from Christoph Beierle and are presented in the preliminaries of Part I in Section 3.1.2.

After our first publications [107, 111] about the Kinematics principle in the context of ranking functions, the transfer into entirely qualitative frameworks lay close. Here, Gabriele Kern-Isberner was the driving force and gave the initial spark for further investigations. However, elaborating on the technical details in full depth proved to be more tedious and complicated than expected. Therefore, we decided to opt for a publication in the leading journal in the area of AI, the prestigious Artificial Intelligence Journal (AIJ). Together with my colleagues Gabriele Kern-Isberner and Christoph Beierle, we presented and investigated a qualitative version of the Kinematics principle in the following publication:

> [70] G. Kern-Isberner, M. Sezgin, C. Beierle, *A kinematics principle for iterated revision*, Artif. Intell. 314 (2023) 103827.

Gabriele Kern-Isberner designed and drafted the Kinematics postulate in Section 6.1 and the concept of qualitative conditionalization presented in Section 6.2. The qualitative Kinematics principle represents a target, and my specific contribution was the technical elaboration of the concepts needed for the Kinematics principle.

My idea was to take c-revisions as a blueprint for a qualitative revision operator. This led to the central transformations between the qualitative and the quantitative framework and the corresponding results concerning qualitative conditionalization resp. c-revisions presented in Section 6.3 and in Section 6.4. Also, I contributed to investigating the Kinematics principle for the conditional revision operator from [27] and the comparison to qualitative c-revisions presented in Chapter 6.5.

As in [111], the concept of strategies for c-revisions provides the technical basis for significant results presented in [70]. Christoph Beierle once again provided us with his expertise and designed the conceptual outline of strategies presented in Section 3.1.2.

Part II: Parameterized Belief Change.

In the early stages of my Ph.D., I was lucky to be part of a scientific exchange program funded by the German Academic Exchange Service (DAAD) with the title "Advanced belief change operations based on comparisons and conditionals: Towards a general framework" which was granted to Hans Rott (University of Regensburg, Germany), Gabriele Kern-Isberner and Eduardo Fermé (University of Madeira). I had the chance to meet them and have inspiring conversations and discussions about advanced belief change operations in Regensburg and Madeira. Unfortunately, the COVID-19 pandemic has prevented the possibility of further meetings, but my interest in parameterized belief change operations roots in the context of this inspiring environment of researchers. Together with Gabriele Kern-Isberner, I investigated two parameterized belief change mechanisms in qualitative and quantitative frameworks of belief representation. These research activities lead to the following publications in the two main conferences in the world of AI, the *IJCAI – International Joint Conferences on Artificial Intelligence* and the *AAAI Conference on Artificial Intelligence*:

[109] M. Sezgin, G. Kern-Isberner, *Revision by Comparison for Ranking Functions*, in: L. D. Raedt (Ed.), Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022, ijcai.org, 2022, pp. 2734–2740.

[110] M. Sezgin, G. Kern-Isberner, *Implementing Bounded Revision via Lexicographic Revision and c-Revision*, Proceedings of the Thirty-

Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, Washington, D.C., USA, 7-12 February 2023, aaai.org, 2023.

I designed and developed the concept and the technical elaboration of the papers [109] and [110]. This includes drafting and proving the results, theorems, and resulting applications. The (extended) results of these papers are presented in Chapter 8 and Chapter 9. Chapter 10 combines and compares the work presented in both papers [109] and [110] and presents novel and unpublished results. Gabriele Kern-Isberner's contribution to both of the aforementioned publications was to provide expertise in the field of Belief Revision, propose enhancements, and provide support throughout the whole process, and of course, she introduced me to Eduardo Fermé and Hans Rott in the course of the fruitful project meetings.

Chapter 2

Belief Revision, Epistemic States and Conditionals

In this chapter, we introduce the basic notions used in this thesis. We generally aim to introduce formal preliminaries and notation on a need-to-know basis throughout the thesis. However, some notions permeate this work, and thus we state them at this stage of the thesis.

2.1 Basic Definitions and Notations

This section presents basic definitions and notations from propositional and conditional logic.

Propositional Logic. Throughout this thesis, we assume that $\mathcal{L}_\Sigma$ denotes a finitely generated *propositional language* built over a non-empty set of *propositional atoms* or *variables*, which we call *signature* $\Sigma = \{a, b, c, \ldots\}$. We write $\mathcal{L}$ instead of $\mathcal{L}_\Sigma$ when Σ is clear from context or consideration of a particular Σ is not of any importance. We consider the standard logical connectives *and* $\wedge$, *or* $\vee$ and *not* $\neg$, as well as the constants $\top$ denoting *logical truths* or *tautologies* and $\bot$ for *logical fallacies* or *contradictions*. In general, the set $\mathcal{L}_\Sigma$ of *propositional formulas* over Σ is generated as the smallest set obeying the following rules:

- Every atom in Σ is an (atomic) formula, i.e., $a, b, \ldots \in \Sigma$ are formulas of $\mathcal{L}_\Sigma$.

- If $A, B \in \mathcal{L}$, then $A \wedge B$, $A \vee B$ and $\neg A$ are formulas of $\mathcal{L}_\Sigma$.

A *literal* $\dot{a}$ is either an atom a, in which case $\dot{a}$ is a positive literal, or its negation $\neg a$, in which case $\dot{a}$ is a negative literal. Note that, by lower case letters $a, b, c, \ldots$ we denote propositional variables, intended to represent issues that can be subject of reasoning, thought or deliberation. And most of the time, we denote formulas using capital letters $A, B, C, \ldots$, or if we want to emphasize the special role of designated formulas small greek letters $\alpha, \beta, \ldots$. Furthermore, we often omit the logical *and*-connector, writing AB instead of $A \wedge B$. Also, overlining formulas indicates *negation*, i.e., $\overline{A}$ means $\neg A$. We use $A \Rightarrow B$ as shorthand for $\neg A \vee B$ and $A \Leftrightarrow B$ as shorthand for $(A \Rightarrow B) \wedge (B \Rightarrow A)$.

We apply a model-theoretic approach to propositional logic. A *possible word* (alternatively *interpretation*) is a valuation function $\omega : \Sigma \to \{0, 1\}$ mapping every atom in Σ to either *true*, denoted as 1 or *false*, denoted as 0. By Ω we denote the *set of all possible worlds*, i.e., Ω is a complete set of interpretations of $\mathcal{L}$ and it holds that $|\Omega| = 2^{|\Sigma|}$. And, sometimes worlds are identified simply with their corresponding complete conjunction, s.t.

$$\omega = \bigwedge_{\substack{v \in \Sigma \\ \omega \models \dot{v}}} \dot{v}. \tag{2.1}$$

We define an *evaluation function* $[\![\cdot]\!]_\cdot : \mathcal{L}_\Sigma \times \Omega \to \{0, 1\}$ of Σ for all formulas $A, A_1, A_2 \in \mathcal{L}_\Sigma$ and $\omega \in \Omega$ as follows:

- For each $a \in \Sigma$ it holds that $[\![a]\!]_\omega = \omega(a)$
- $[\![A]\!]_\omega = 1$, if and only if $[\![\overline{A}]\!]_\omega = 0$
- $[\![A_1 \wedge A_2]\!]_\omega = 1$, if and only if $[\![A_1]\!]_\omega = 1$ an $[\![A_2]\!]_\omega = 1$
- $[\![A_1 \vee A_2]\!]_\omega = 1$, if and only if $[\![A_1]\!]_\omega = 1$ or $[\![A_2]\!]_\omega = 1$

We write $\omega \models A$ if $[\![A]\!]_\omega = 1$ and call ω a *model* of A. By $\mathrm{Mod}(A)$, we notate the set of all A-worlds, s.t.

$$\mathrm{Mod}(A) = \{\omega \in \Omega \,|\, \omega \models A\}.$$

A formula A is called *consistent* iff $\mathrm{Mod}(A) \neq \emptyset$. For logical truths $\top$, $\mathrm{Mod}(\top) = \Omega$ and for fallacies $\bot$, $\mathrm{Mod}(\bot) = \emptyset$ holds. Each symbol is a synonym for an arbitrary (unique) formula.

Given two formulas A, B, we say that B is a *logical consequence* of A, $A \models B$ if it holds that $\mathrm{Mod}(A) \subseteq \mathrm{Mod}(B)$. And formulas A and B are *equivalent*, $A \equiv B$ iff $A \models B$ and $B \models A$, i.e., $\mathrm{Mod}(A) = \mathrm{Mod}(B)$. The *classical logical consequence* operator $Cn(A) = \{B \in \mathcal{L} \mid A \models B\}$ subsumes the set of all logical consequences of A. These notions can be naturally lifted to sets of formulas $\mathcal{B} \subseteq \mathcal{L}$, s.t. $Cn(\mathcal{B}) = \{B \in \mathcal{L} \mid \mathcal{B} \models B\}$ and $\mathrm{Mod}(\mathcal{B}) = \{\omega \in \Omega \mid \omega \models B, \text{ for all } B \in \mathcal{B}\}$. We call sets $\mathcal{B}$ that are closed under logical consequences, i.e., $Cn(\mathcal{B}) = \mathcal{B}$, *deductively closed*. The deductively closed set $Th(W) = \{A \in \mathcal{L} \mid \omega \models A \text{ for all } \omega \in W\}$ which has exactly a subset $W \subseteq \Omega$ as models is called a *formal theory* of W. We call a set of formulas $\{A_i\}_{i=1,\dots,n}$ *exclusive* if $\bigwedge_{i=1,\dots,n} A_i \equiv \bot$ and *exhaustive* if $\bigvee_{i=1,\dots,n} A_i \equiv \top$.

Conditional Logic. The conditional operator $\mid$ extends the propositional language $\mathcal{L}$ to a *conditional language* $(\mathcal{L}|\mathcal{L})$, s.t.

$$(\mathcal{L}|\mathcal{L}) = \{(B|A) \mid A, B \in \mathcal{L}\}.$$

In this thesis, we consider $(\mathcal{L}|\mathcal{L})$ to be a flat conditional language, so no nesting of conditionals is allowed. We call A the *antecedent* or *premise* of $(B|A)$, and B is its *consequent*.

Weak Conditionals. In the following, conditionals $(B|A) \in (\mathcal{L}|\mathcal{L})$ are referred to as *standard conditionals* or, if there is no danger of confusion, simply conditionals. We further extend our framework of conditionals to a language with *weak conditionals* by introducing a weak conditional operator $(|\cdot|\cdot|)$, s.t.

$$(|\mathcal{L}|\mathcal{L}|) = \{(|D|C|) \mid C, D \in \mathcal{L}\}.$$

Again, we consider $(|\mathcal{L}|\mathcal{L}|)$ to be a flat conditional language, i.e., no nesting of weak conditionals is allowed. As for standard conditionals, we call C the *antecedent* or *premise* and D the *consequent* of $(|D|C|)$.

2.2 Belief Revision

In this section, we discuss the basics of one-shot belief change, i.e., scenarios that deal with the change of beliefs on one specific (but generic) belief set.

In the account of belief change originally developed by Alchourrón, Gärdenfors, and Makinson [45, 2] there are three principal types of change for (propositional) beliefs: *expansion, revision, and contraction*. Each type of belief change can be axiomatized by a set of postulates, the so-called *AGM postulates*, which define rational and desirable changes of beliefs. These basic postulates lay the foundation of the well-known and fruitful *AGM theory*, named after its founders Alchourrón, Gärdenfors, and Makinson. In the AGM theory, an agent's beliefs are represented by a *belief sets*, i.e., a deductively closed set of formulas, s.t. $K = Cn(K)$. And AGM belief change operators take as input a belief set K and a piece of new information in the form of a propositional formula A and map them onto a new belief set. In the following, we consider the three change operators named above and explain and distinguish them from each other. Since, in the context of this thesis, revision is the most important one, we investigate this operation more thoroughly.

AGM Expansion. The simplest type of belief revision is *expansion*, which occurs when the new information is consistent with an agent's current beliefs in K. This type of belief change is characterized by a unique function $+$ [45], given by

$$K{+}A = Cn(K \cup \{A\}).$$

This characterization of expansion satisfies an axiomatic approach via six postulates [45]. To shorten the matter, we do not consider them here because the definition is handy for belief sets, and we consider belief expansion solely in that context. Note that, if $K \models \overline{A}$ then $K{+}A = \mathcal{L}$, i.e., if the new information is inconsistent with K expansion does not lead to rational behavior in the change of beliefs. In the following, we deal with the more general case of belief revision, where we do not presuppose consistency of A with the agent's prior belief set K. Note that if, however, consistency holds, AGM revision reduces to AGM expansion.

AGM Revision. In their seminal paper [2], Alchourrón, Gärdenfors, and Makinson proposed six basic and two additional postulates that guide the belief change process when incorporating new information A which might contradict an agent's former belief set K. The goal is to obtain a consistent posterior belief set whenever possible while avoiding unnecessary changes. Note that, unlike for expansion, these postulates are insufficient to describe a unique optimal revision operator.

Definition 2.2.1 (AGM Revision [2])**.** *A change operator* $\star$ *is called an* AGM revision operator *for a belief set* K *and a formula* $A \in \mathcal{L}$ *if the following postulates are satisfied*

(AGM⋆1)	$K \star A$ *is a belief set*	*(Closure)*
(AGM⋆2)	$A \in K \star A$	*(Success)*
(AGM⋆3)	$K \star A \subseteq K + A$	*(Inclusion)*
(AGM⋆4)	*If* $K + A$ *is consistent, then* $K \star A = K + A$	*(Vacuity)*
(AGM⋆5)	*If* A *is consistent, then* $K \star A$ *is consistent*	*(Consistency)*
(AGM⋆6)	*If* $A \equiv B$*, then* $K \star A = K \star B$	*(Syntax Independence)*
(AGM⋆7)	$K \star (A \wedge B) \subseteq (K \star A) + B$	*(Superexpansion)*
(AGM⋆8)	*If* $\overline{B} \notin K \star A$ *then* $(K \star A) + B \subseteq K \star (A \wedge B)$	*(Subexpansion)*

The first postulate (AGM⋆1) ensures that each AGM revision results in a belief set. The second postulate (AGM⋆2) demands that revision always successfully incorporates the new belief A. Due to inclusion (AGM⋆3), an AGM revision operator never yields more beliefs than an expansion of the corresponding belief set would do. Due to vacuity (AGM⋆4), it holds that if the initial beliefs K are consistent with A, then the revision by A of K is nothing more than an expansion of $K+A$. The result of a revision with consistent new information is always consistent due to the consistency postulate (AGM⋆5). On the other hand, due to (AGM⋆2), the revision of K by an inconsistent belief A always yields an inconsistent belief set. From (AGM⋆6), we can conclude that syntactic differences in the representation of beliefs do not impact the revision process. The remaining postulates (AGM⋆7) and (AGM⋆8) are considered additional postulates for revising with conjunctive beliefs. Thus, (AGM$\star$7) displays an extension of the inclusion postulate (AGM$\star$3) when adding a conjunctive belief $A \wedge B$, demanding that we should obtain $K\star(A \wedge B) \subseteq (K \star A)+B$ in addition to $K\star(A \wedge B) \subseteq K+(A \wedge B)$ [52]. In the same way, (AGM$\star$8) corresponds to the vacuity postulate (AGM$\star$4) for conjunctive beliefs. Note that, w.l.o.g., we exclude revisions with logical fallacies since revisions with $\perp$ often lead to unintuitive results. This view is also discussed in [45] by Gärdenfors and in [75] by Levi. Also, the special case of revising one's beliefs with $\perp$ is not relevant in the context of this thesis.

In their seminal paper [60], Katsuno and Mendelzon distinguished between two fundamentally different types of modifications of an agent's belief set; besides revision, they introduced the notion of *update*. Katsuno and Mendelzon considered belief revision to be an operation concerned with a *static world*, whereas update deals with updating an agent's beliefs in an *evolving* environment. We focus on their notion of revision. They rephrased the AGM postulates for revision in [60], presupposing that instead of a (possibly infinite) set of formulas K, an agent's beliefs can be represented by a single formula ψ which describes the entirety of an agents belief set. Note that the revision operator $\star$ takes as input a belief set K and a proposition A, whereas revision operators in the framework of Katsuno and Mendelzon perform revision on propositional formulas ψ. To not overload the notation, we use the same symbol $\star$ for both cases; due to the strict delimitation of belief sets represented as deductively closed sets K resp. as formulas ϕ, no confusion arises.

(KM$\star$1) $\psi \star A \models A$

(KM$\star$2) If $\psi \wedge A$ is consistent, then $\psi \star A \equiv \psi \wedge A$

(KM$\star$3) If A is consistent, then $\psi \star A$ is consistent

(KM$\star$4) If $A \equiv B$, then $\psi \star A = \psi \star B$

(KM$\star$5) $(\psi \star A) \wedge B \models \psi \star (A \wedge B)$

(KM$\star$6) If $(\psi \star A) \wedge B$ is consistent, then $\psi \star (A \wedge B) \models (\psi \star A) \wedge B$

The relation between the KM postulates and the above-stated AGM postulates is evident, and, we can recover the AGM belief set K from ψ via $K = Cn(\psi)$ [60]. Also, it holds that (AGM$\star$3) and (AGM$\star$4) are merged into (KM$\star$3). It holds that the postulates (KM$\star$5) and (KM$\star$6) represent the condition that revision should be accomplished with minimal change [59]. This *principle of minimal change* is not explicitly mentioned by the AGM postulates, but still, they were designed in its spirit [2, 39]. The postulates above provide ground for a model-theoretic characterization of belief revision in the spirit of minimal change [59], leading to a simple yet elegant condition that each KM style revision operator must satisfy.

Katsuno and Mendelzons' representation theorem relies on *faithful assignments* $\preceq_\psi$, i.e., *total preorders* over interpretations $\omega \in \Omega$ assigned to each propositional formula ψ. Note that a *preorder* $\preceq$ (PO) over Ω is a reflexive and transitive relation over the set of possible worlds. For two worlds ω, ω' it holds that $\omega \prec \omega'$ iff $\omega \preceq \omega'$ and not $\omega' \preceq \omega$. A preorder is called *total* (TPO) if for every ω, ω', $\omega \preceq \omega'$ or $\omega' \preceq \omega$ holds.

Definition 2.2.2 (Faithful assignments for propositional formulas [60])**.** *Consider a function that assigns to each propositional formula ψ a preorder $\preceq_\psi$ over Ω. This assignment is* faithful *if the following conditions hold:*

1. *If $\omega, \omega' \in Mod(\psi)$ then $\omega \prec_\psi \omega'$ does not hold*

2. *If $\omega \in Mod(\psi)$ and $\omega' \notin Mod(\psi)$ then $\omega \prec_\psi \omega'$*

3 *If $\psi \equiv \psi'$ then $\preceq_\psi = \preceq_{\psi'}$*

An interpretation ω is called *minimal* in a subset $\Omega' \subseteq \Omega$ w.r.t. $\preceq$ if $\omega \in \Omega'$ and there is no $\omega' \in \Omega'$, s.t. $\omega' \prec \omega$. We use $\min(\Omega', \preceq)$ as shorthand for all minimal models in Ω' w.r.t. $\preceq$ and by slight abuse of notation sometimes write $\min(A, \preceq)$ instead of $\min(\text{Mod}(A), \preceq)$ for a formula A. Sometimes, it is useful to consider a designated representative of the set of minimal worlds for a formula A; then we write $\omega_A \in \min(A, \preceq_\Psi)$, where the index indicates the corresponding formula.

For a faithful assignment $\preceq_\psi$, it holds that the minimal models of $\preceq_\psi$ display exactly all worlds that satisfy ψ, i.e., $\min(\Omega, \preceq_\psi) = \text{Mod}(\psi)$, without making any difference in between these worlds. The following representation theorem presented in [60] characterizes all revision operators that satisfy the KM postulates for revision.

Theorem 2.2.1 (Representation Theorem [60])**.** *The revision operator $\star$ satisfies (KM1) – (KM6) if and only if there exists a faithful assignment that map each belief* *set represented as a propositional formula* ψ *to a TPO* $\preceq_\psi$ *such that*

$$Mod(\psi \star A) = \min(A, \preceq_\psi).$$

This theorem proves that TPOs provide the essential meta-structure to rationally revise belief sets in the sense of the AGM framework.

AGM Contraction. The third important belief change operator $\div$, the *AGM contraction* $K \div A$ in the AGM framework deals with the case when an agent wants to delete a piece of information A from her former belief set K. We state the six basic AGM postulates for contraction.

Definition 2.2.3 (AGM Contraction [2])**.** *A change operator* $\div$ *is called an* AGM contraction operator *for a belief set* K *and a formula* $A \in \mathcal{L}$ *if the following postulates are satisfied*

(AGM÷1)	$K \div A$ *is a belief set*	*(Closure)*
(AGM÷2)	$K \div A \subseteq K$	*(Inclusion)*
(AGM÷3)	*If* $A \notin K$, *then* $K \div A = K$	*(Vacuity)*
(AGM÷4)	*If* $A \not\equiv \top$, *then* $A \notin K \div A$	*(Success)*
(AGM÷5)	*If* $A \in K$, *then* $K \subseteq (K \div A) + A$	*(Recovery)*
(AGM÷6)	*If* $A \equiv B$, *then* $K \div A = K \div B$	*(Syntax Independence)*

In [71], Konieczny and Pino Pérez showed a similar rephrasement of propositional AGM contraction as for AGM revision in the style of Katsuno and Mendelzon, s.t. AGM contractions can be characterized via a representation theorem in terms of faithful assignments.

2.3 Iterated Belief Revision

In this section, we extend the previously discussed framework of one-shot belief revision to revision scenarios where change can happen iteratively, i.e., the result of a belief revision can be used as input for a subsequent revision operation.

Classical belief revision in the AGM framework is concerned with situations in which the belief set of an agent changes due to new incoming information. While features and strengths of this approach have been acknowledged quite early [45], also its flaws have not passed unnoticed [29, 15, 17]. In their seminal paper [29], Darwiche and Pearl addressed the issue of iteration, which is not treated adequately by classical AGM operators. The following example from [29] demonstrates the objections of Darwiche and Pearl[1].

Example 2.3.1 ([29, 40])**.** *We consider a murder trial with two jurors, Juror 1 and Juror 2, who believe different candidates to be the possible murderer:*

> *Juror 1:"A is the murderer, B is a remote but unbelievable possibility while C is definitely innocent."*
>
> *Juror 2:"A is the murderer, C is a remote but unbelievable possibility while B is definitely innocent"*

The two jurors share the same belief $\psi_1 = \psi_2 \equiv A =$*"A is the only murderer". Now, surprising evidence, e.g., A has produced a reliable alibi, obtains* $A' =$*"A is not the murderer". The revision by* A' *should lead to different belief sets for Jurors 1 and 2. Juror 1 should now believe that B is the murderer, whereas Juror 2 should now be convinced that C is guilty. Yet, this does not align with the AGM framework since (KM⋆4) demands equal belief sets in this case. Different results are only possible if the revision does not solely depend on the corresponding belief set.*

Darwiche and Pearl identified the restriction to belief sets in the AGM framework as the leading cause for the problem of iteration. *Epistemic states*, Ψ, often referred to as *belief states*, cover the entirety of information an agent needs (at a given time) to think and reason, in particular when beliefs are applied or modified. We assume that Ψ contains meta-information about the relevance, preferences, and plausibility of beliefs. We take belief sets as the foundation of a belief state Ψ and assume that each epistemic state Ψ is implicitly equipped with a function $Bel(\cdot)$ that yields a belief set $Bel(\Psi)$, s.t. $Bel(\Psi) \subset \mathcal{L}$ contains all propositions and logical sentences the agent deems maximally plausible, i.e., truly beliefs at a given time. In the

[1]Darwiche and Pearl adapted the Example from [40].

following Section 2.4, we discuss different formal representations of belief states more thoroughly.

Viewing belief revision as an operation on belief states allows us to consider not only an agent's most plausible beliefs but also underlying *revision policies* or *inference rules* that guide the process of incorporating new information [15]. This allows us to conclude which propositional beliefs B an agent will hold in the light of new evidence A, i.e., plausible conclusions of the form '*If A then B*'. Such rules of inference are best represented by conditional $(B|A)$ and we can say that an agent accepts the conditional $(B|A)$, if the revision with A yields belief in B, which is exactly what happens in the so-called *Ramsey Test* [95, 120]:

$$\Psi \models (B|A) \text{ iff } Bel(\Psi \bullet A) \models B, \tag{2.2}$$

where $\bullet$ displays a revision operator which maps an epistemic state Ψ and a propositional sentence A onto a revised belief state $\Psi \bullet A$. The Ramsey Test reveals that belief revision and conditional beliefs are intimately related. And it holds that belief states can be expressed as the entirety of an agent's conditional beliefs [29]. In their seminal paper [29], Darwiche and Pearl strongly linked the revision with propositional information to the preservation of conditional beliefs already present in an agent's epistemic state. They extended classical AGM revision to an operation on epistemic states and, thus, paved the way to investigate iterations of revision operations [29].

The Darwiche and Pearl Framework. In the following, we state the modified KM postulates for *iterated Belief Revision* proposed in [29]. In this thesis, we use the term iterated belief revision as a synonym for belief revision on epistemic states in the sense of Darwiche and Pearl, i.e., we focus on the *preservation of conditional beliefs*. Of course, other notions of iteration are conceivable (cf. [79, 121]). Yet, for clarity, we sometimes refer to the framework of iterated belief revision presented in this chapter as the *DP framework* as an abbreviation of its founders, Darwiche and Pearl.

We presuppose for each epistemic state Ψ that we can define an associated belief set $Bel(\Psi)$, which is either a deductively closed set of formulas, like in the

AGM framework or embedded in a propositional formula, as for KM style revision operators. As before, we denote by $\Psi \bullet A$ the revised epistemic state, not the corresponding belief set. This is the first of the two main modifications to the KM postulates.

(KM•1) $Bel(\Psi \bullet A) \models A$

(KM•2) If $Bel(\Psi) \wedge A$ is consistent, then $Bel(\Psi \bullet A) \equiv Bel(\Psi) \wedge A$

(KM•3) If A is consistent, then $Bel(\Psi \bullet A)$ is consistent

(KM•4) If $\Psi_1 = \Psi_2$ and $A \equiv B$, then $Bel(\Psi_1 \bullet A) = Bel(\Psi_2 \bullet B)$

(KM•5) $Bel(\Psi \bullet A) \wedge B \models Bel(\Psi \star (A \wedge B))$

(KM•6) If $Bel(\Psi \bullet A) \wedge B$ is consistent, then $Bel(\Psi \star (A \wedge B)) \models Bel(\Psi \bullet A) \wedge B$

Although the postulates (KM⋆1) – (KM⋆6) and (KM•1) – (KM•6) may seem pretty similar at first sight, the latter are capable of regulating changes in conditional beliefs adequately, whereas the former is not. Apart from considering epistemic states, the main modification is the weakening of postulate (KM⋆4). Instead of states Ψ_1 and Ψ_2 with the same belief set, we demand equality concerning the whole state. This is because belief sets $Bel(\Psi)$ do not characterize a belief state uniquely. For a more in-depth discussion as to why the less restrictive formulation of (KM⋆4) leads to unintuitive results in terms of iteration, see [29] and [36] and Example 2.3.1.

The modified postulates (KM•1) – (KM•6) lead to a parallel characterization of belief revision as the one presented in [60] by Katsuno and Mendelzon. First, we generalize the notion of *faithful assignments* to belief states Ψ using TPOs.

Definition 2.3.1 (Faithful assignments for epistemic states [29])**.** *Let Ψ be an epistemic state. A* faithful assignment *is a function that maps Ψ to a total preorder $\preceq_\Psi$ on possible worlds Ω, $\Psi \mapsto \preceq_\Psi$, s.t.*

1. *If $\omega, \omega' \models Bel(\Psi)$ then $\omega \approx_\Psi \omega'$*

2. *If $\omega \models Bel(\Psi)$ and $\omega' \not\models Bel(\Psi)$ then $\omega \prec_\Psi \omega'$*

3 *If $\Psi = \Psi'$ then $\preceq_\Psi = \preceq_{\Psi'}$*

Note that, in contrast to faithful assignments for formulas $\preceq_\psi$, $Bel(\Psi) = Bel(\Psi')$ is not enough to yield equal TPOs $\preceq_\Psi = \preceq_{\Psi'}$.

The following notion of compatibility links faithful assignments with belief revision operators and parallels the KM style Representation Theorem 2.2.1. Note that $\mathrm{Mod}(\Psi)$ is an abbreviation of $\mathrm{Mod}(Bel(\Psi))$ and therefore, models of $\Psi \bullet A$ are precisely all minimal A-worlds in $\preceq_\Psi$.

Theorem 2.3.1 ([29])**.** *The revision operator $\bullet$ satisifies (KM$\bullet$1) – (KM$\bullet$6) if and only if there exists a faithful assignment that maps each epistemic state Ψ to a TPO $\preceq_\Psi$ such that*

$$Mod(\Psi \ \bullet \ A) = \min(A, \preceq_\Psi).$$

In general, $\preceq_\Psi$ is a *plausibility (pre)ordering* of possible worlds in Ω, which represents the belief state Ψ as a TPO on Ω. So, $\omega \preceq_\Psi \omega'$ means that the agent with belief state Ψ deems the world ω at least as plausible as ω'. It holds that $\preceq_\Psi$ satisfies the conditions 1 and 2 from Definition 2.3.1 and $\mathrm{Mod}(\Psi) = \min(\Omega, \preceq_\Psi)$ [15]. For a propositional formula A and Ψ equipped with $\preceq_\Psi$, it holds that

$$\Psi \models A \text{ iff } A \in Bel(\Psi) \text{ iff } \omega \models A \text{ for all } \omega \in \min(\Omega, \preceq_\Psi)$$

holds. Hence, each plausibility ordering on possible worlds $\preceq_\Psi$, representing a belief state Ψ, induces a plausibility relation on formulas via:

$$A \preceq_\Psi B \ \text{ iff } \omega_A \preceq_\Psi \omega_B \text{ for all } \omega_A \in \min(A, \preceq_\Psi) \text{ and } \omega_B \in \min(B, \preceq_\Psi) \qquad (2.3)$$

and therefore, displays a plausibility relation in the sense of Grove [42, 16]. Note that, since $\mathrm{Mod}(\bot) = \emptyset$, we get that

$$\omega \prec_\Psi \bot \text{ for all } \omega \in \Omega, \qquad (2.4)$$

since the minimum of the empty set corresponds to the supremum within the respective ordering. Note that, in this case, we refer to the representation of possible worlds as full conjunctions over all variables in the corresponding signature, i.e., ω is represented as the corresponding formula.

In this thesis, we call epistemic states that can be identified via plausibility orderings $\preceq_\Psi$ *plausibilistic TPOs*[2] and refer to them either via Ψ or, exploiting the

[2]Such epistemic states are called *revision models* in [15].

direct correspondence between Ψ and the underlying plausibilistic TPO, just use $\preceq_\Psi$ as identification of epistemic states. In cases where the corresponding epistemic state Ψ for a plausibilistic TPO $\preceq_\Psi$ is clear from the context or the specific epistemic state Ψ is irrelevant we omit the index Ψ and write $\preceq$ instead of $\preceq_\Psi$. Note that this direct correspondence between belief states Ψ and the corresponding TPOs $\preceq_\Psi$ is established by Theorem 2.3.1, which lays the foundation of all iterated belief revision operators considered in this thesis. Thus, we assume that iterated belief revision operators in the sense of Darwiche and Pearl map plausibilistic TPOs and propositional sentences onto (revised) plausibilistic TPOs.

We now state some basic observations which connect plausibilistic TPOs and conditional beliefs.

Lemma 2.3.2. *A conditional $(B|A)$ is accepted in an epistemic state Ψ represented as a plausibilistic TPO $\preceq_\Psi$ iff all minimal models of A satisfy B,*

$$\Psi \models (B|A) \textit{ iff for each } \omega_A \in \min(A, \preceq_\Psi) \textit{ it holds that } \omega_A \models B$$

This follows immediately from the Ramsey Test (2.2) and Theorem 2.3.1:

$$\begin{aligned}\Psi \models (B|A) \text{ iff } & Bel(\Psi \bullet A) \models B,\\ & \text{i.e., } \mathrm{Mod}(\Psi \bullet A) = \min(A, \preceq_\Psi) \subseteq \mathrm{Mod}(B).\end{aligned}$$

The investigation of iterated revision led Darwiche and Pearl to conclude that iteration is inherently linked to preserving conditional beliefs. They proposed four additional postulates that rationally ensure conditional preservation and, thus, define *iterated belief revision operator.*

(C1) If $C \models B$ then $\Psi \models (D|C)$ iff $\Psi \bullet B \models (D|C)$

(C2) If $C \models \overline{B}$ then $\Psi \models (D|C)$ iff $\Psi \bullet B \models (D|C)$

(C3) If $\Psi \models (B|A)$ then $\Psi \bullet B \models (B|A)$

(C4) If $\Psi \not\models (\overline{B}|A)$ then $\Psi \bullet B \not\models (\overline{B}|A)$

These postulates highlight the claim in minimization of changes in conditional beliefs for iterated revision operators within the framework proposed in [29][3]. Also, it

[3] For a discussion of different notions of conditional preservation, see the original paper [29].

holds that none of (C1) – (C4) are derivable from the AGM postulates [29]. (C1) and (C2) state that no accommodating information B, independent of whether it specifies or contradicts C, should destroy any conditional belief with C as a premise. Postulate (C3) states that conditional beliefs, which lead to the conclusion that B is plausible, should not be given up after revising with B. On the other hand, (C4) states that revision with B should not lead to the acceptance of $(\overline{B}|A)$.

Via the regulation of conditional beliefs in postulates (C1) – (C4), it is possible to constrain the relationship between the TPOs $\preceq_\Psi$ and $\preceq_{\Psi \bullet A}$. Darwiche and Pearl state the following representation theorem, which characterizes iterated belief revision operators for faithful assignments.

Theorem 2.3.3 ([29])**.** *Let Ψ be an epistemic state and $\preceq_\Psi$ its corresponding faithful assignment. Suppose that a revision operator satisfies (KM•1) – (KM•6). The operator under a faithful assignment satisfies postulates (C1) – (C4) iff $\Psi \bullet A$ and its corresponding faithful assignment $\preceq_{\Psi \bullet A}$ satisfy:*

(DP1) *If $\omega_1, \omega_2 \in Mod(A)$, then $\omega_1 \preceq_\Psi \omega_2$ iff $\omega_1 \preceq_{\Psi \bullet A} \omega_2$*

(DP2) *If $\omega_1, \omega_2 \notin Mod(A)$, then $\omega_1 \preceq_\Psi \omega_2$ iff $\omega_1 \preceq_{\Psi \bullet A} \omega_2$*

(DP3) *If $\omega_1 \in Mod(A)$ and $\omega_2 \notin Mod(A)$, then $\omega_1 \prec_\Psi \omega_2$ if $\omega_1 \prec_{\Psi \bullet A} \omega_2$*

(DP4) *If $\omega_1 \in Mod(A)$ and $\omega_2 \notin Mod(A)$, then $\omega_1 \preceq_\Psi \omega_2$ if $\omega_1 \preceq_{\Psi \bullet A} \omega_2$*

Each postulate (DP1) – (DP4) leads to the preservation of some part of the prior ordering of worlds in the posterior ordering. Yet, some crucial parts are not preserved. For $\omega_1 \models \overline{A}$ and $\omega_2 \models A$, if $\omega_1 \preceq_\Psi \omega_2$ (or $\omega_1 \prec_\Psi \omega_2$) then the postulates do not insist on $\omega_1 \preceq_{\Psi \bullet A} \omega_2$ (or $\omega_1 \prec_{\Psi \bullet A} \omega_2$). An in-depth discussion of the rationale behind this is discussed in [29]. In general, the addition of postulates that lead to the complete minimization of changes in conditional beliefs is not desirable. In [29], the authors provide examples showing that some changes are legitimate.

Furthermore, we state a common principle for propositional belief revision [63], expressing that revisions with tautologies should not change the prior belief state.

Tautological Vacuity (TV) A revision operator • satisfies (TV) if it holds for any epistemic state Ψ that

$$\textbf{(TV)} \qquad \Psi \bullet \top = \Psi. \tag{2.5}$$

This principle can be seen as a reasonable additional property to all aforementioned revision operators[4].

Iterated Belief Revision Operators for Propositions. The DP postulates for conditional preservation (C1) – (C4) and the resulting characterization for TPOs $\preceq_\Psi$ in Theorem 2.3.3 were widely endorsed in the Belief Revision community [17, 13, 82]. Yet, they do not fully characterize iterated belief revision operators, and there have been several approaches to strengthen the spirit of conditional preservation by providing concrete belief revision operators that obey the DP postulates, among them operators presented in the original paper on iterated belief revision by Darwiche and Pearl [29].

In [26], Chandler and Booth argue in favor of three well-known iterated belief revision operators *natural revision* $\bullet^{\mathrm{n}}$ [15], *restrained revision* $\bullet^{\mathrm{r}}$ [13] and *lexicographic revision* $\bullet^{\ell}$ [82]. The authors show in [26] that these operators fit within the DP framework but also satisfy some other desirable properties of iterated revision. They call these operators *elementary revision operators*. For a more detailed discussion of the additional postulates, see [26]. The following definition subsumes the semantic definitions of each elementary operator.

Definition 2.3.2 (Elementary Revision Operators [26, 17, 13, 82])**.** *Let Ψ be an epistemic state associated with the plausibilistic TPO $\preceq_\Psi$ and $A \in \mathcal{L}$. The operators* natural revision $\bullet^{\mathrm{n}}$ *[17],* restrained revision $\bullet^{\mathrm{r}}$ *[13] and* lexicographic revision $\bullet^{\ell}$ *[82] map $\preceq_\Psi$ onto the corresponding revised TPO $\preceq_{\Psi\bullet^{\mathrm{n}}A}$ for natural revision, $\preceq_{\Psi\bullet^{\mathrm{r}}A}$ for restrained revision and $\preceq_{\Psi\bullet^{\ell}A}$ for lexicographic revision are defined as follows:*

- $x \preceq_{\Psi\bullet^{\mathrm{n}}A} y$ *iff (1)* $x \in \min(A, \preceq_\Psi)$*, or (2)* $x, y \notin \min(A, \preceq_\Psi)$ *and* $x \preceq_\Psi y$
- $x \preceq_{\Psi\bullet^{\mathrm{r}}A} y$ *iff (1)* $x \in \min(A, \preceq_\Psi)$*, or (2)* $x, y \notin \min(A, \preceq_\Psi)$ *and either (a)* $x \prec_\Psi y$ *or (b)* $x \sim_\Psi y$ *and* $(x \in \mathit{Mod}(A)$ *or* $y \in \mathit{Mod}(\overline{A}))$
- $x \preceq_{\Psi\bullet^{\ell}A} y$ *iff (1)* $x \in \mathit{Mod}(A)$ *and* $y \in \mathit{Mod}(\overline{A})$*, or (2)* $(x \in \mathit{Mod}(A)$ *iff* $y \in \mathit{Mod}(A))$ *and* $x \preceq_\Psi y$*.*

[4] Note that, for the classical AGM resp. KM revision operators, (TV) applies solely to the underlying belief sets.

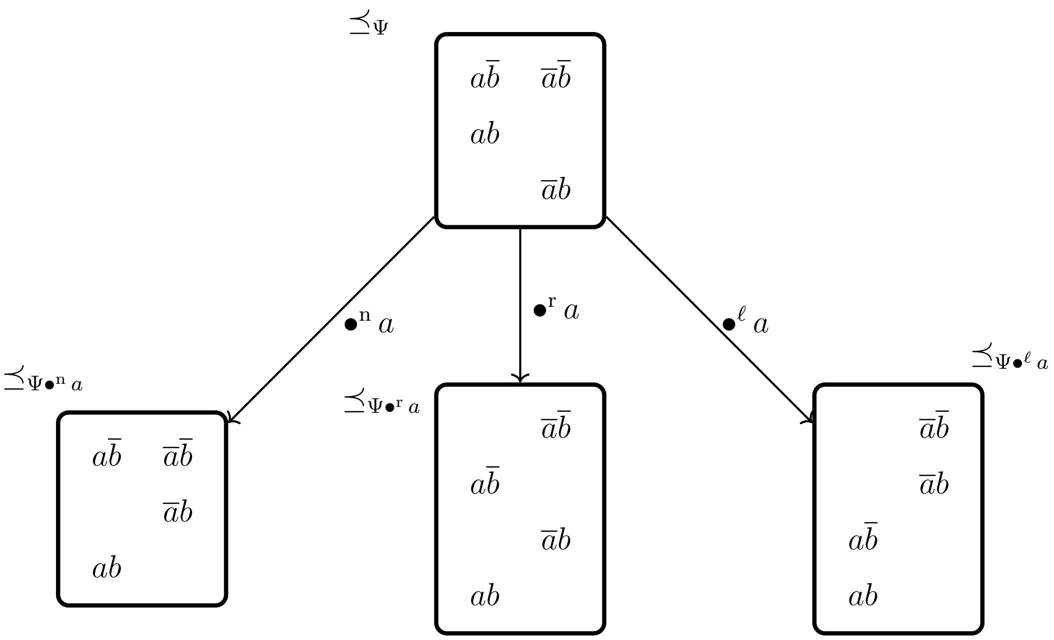

Figure 2.1: Revision of Ψ by A via elementary revision operators $\bullet^{\mathrm{n}}$, $\bullet^{\mathrm{r}}$ and $\bullet^{\ell}$ from Definition 2.3.2.

The natural revision operator $\bullet^{\mathrm{n}}$ by Boutilier shifts the minimal A-worlds to the lowermost level of plausibility while the relations between the remaining worlds are kept. This is the most conservative change that still obeys the DP postulates [13, 29]. Lexicographic Revision $\bullet^{\ell}$ [82] on the other hand, makes all A-worlds more plausible than the $\overline{A}$-worlds while keeping the corresponding plausibility relations between worlds in $\mathrm{Mod}(A)$ and $\mathrm{Mod}(\overline{A})$, i.e., the change is quite rigorous, and it is the least conservative of the DP operators [13]. The restrained revision operator $\bullet^{\ell}$, introduced by Booth and Meyer [13], shifts all minimal A-worlds to the lowermost layer of the revised TPO. The relative ordering is kept for the remaining worlds, except for equally plausible A-worlds and $\overline{A}$-worlds; those are split, s.t. the A-worlds become strictly more plausible. The following example illustrates the revision mechanism of the elementary revision operators from Definition 2.3.2.

Example 2.3.2. *The boxes in Figure 2.1 represent epistemic states and the associated TPOs over the signature $\Sigma = \{a, b\}$. It holds that the lower the world is arranged within the box, the more plausible it is. We revise the TPO $\preceq_\Psi$ by a via* *natural, restrained, and lexicographic revision. The posterior TPOs are depicted in Figure 2.1. The natural revised TPO* $\preceq_{\Psi \bullet^{\mathrm{n}} a}$ *is similar to the prior plausibilistic TPO* $\preceq_{\Psi}$*, except that the minimal a-world ab is shifted to the lowermost level. Among the three posterior TPOs, the lexicographic revised TPO* $\preceq_{\Psi \bullet^{\ell} a}$ *displays the most rigorous change; here, all a-worlds are less plausible than all* $\overline{a}$*-worlds. For the restrained revised TPO* $\preceq_{\Psi \bullet^{\mathrm{r}} a}$*, it holds that the minimal a-world is shifted to the lowermost level and for* $a\overline{b}$ *and* $\overline{a}\overline{b}$*, which are equally plausible in the prior TPO, it holds that* $a\overline{b}$ *is now strictly more plausible than* $\overline{a}\overline{b}$*. The example illustrates (some) gradations in the change mechanism possible within the DP framework.*

2.4 Conditionals and Epistemic States

In the previous section, we have seen that conditional beliefs, their modification, their preservation, and the general task of belief revision are closely related to each other. Now, we want to investigate the vital role of conditionals in the broader context of different representations of epistemic states.

Conditionals. There have been several approaches to explain the notion of conditionals (cf. [1, 120, 77, 84]) trying to define reasonable ways to connect antecedent and consequent of a conditional $(B|A)$. Probably the most famous account leads to the conclusion that conditionals $(B|A)$ express the notion that '*If A holds, then it follows plausibly that B holds*' [83]. Thus, they describe plausible relationships between antecedents and consequents. The Ramsey Test illustrates how conditionals can be seen as inference rules which guide the process of belief revision; that is why we sometimes refer to conditionals as *conditional rules*.

Most of the approaches, like e.g., [77, 83, 41], to defining a logic of conditionals were driven by the *paradox of material implication*, which was firstly described by C.I. Lewis in [76]. It concerns the fact that for any sentences A and B, the material implication '*If A, then B*' follows from '*not A*' but also from 'B', and therefore we can create true conditionals via true or false sentences irrespective of the content. Thus, identifying conditionals with the corresponding material implication leads to unintuitive results; for more, see Ramsey's argumentation in [95]. Therefore, we consider them no longer classical valued for classical input, i.e., non-boolean objects. In [30], de Finetti put forward the idea that conditionals are three-valued logical entities which can be defined w.r.t. to an interpretation ω as follows:

$$[\![(B|A)]\!]_\omega = \begin{cases} 1 & \omega \models AB \\ 0 & \omega \models A\overline{B} \\ \mathrm{u} & \omega \models \overline{A} \end{cases}, \tag{2.6}$$

where u stands for undefined. Thus, conditionals can be seen as *generalized indicator functions* $(\cdot|\cdot) : \Omega \to \{0, 1, u\}$ on worlds. We call the formula AB the *verification* and $A\overline{B}$ the *falsification* of the conditional $(B|A)$. According to [30, 63], we say that two conditionals $(B|A)$ and $(B'|A')$ are *conditionally equivalent*, denoted by $(B|A) \equiv (B'|A')$, if they have the same verification and the same falsification behavior, i.e.,

$$AB \equiv A'B' \text{ and } A\overline{B} \equiv A'\overline{B'}.$$

Note that the conditional language is taken to include the propositional language $\mathcal{L}$ by identifying a proposition A with the conditional $(A|\top)$ with a tautological premise[5].

Furthermore, it holds that for a conditional $(B|A)$, the conditional $(\overline{B}|A)$ is considered to be the *strict negation* of it since verification and negation are swapped. At the same time, both remain undefined for precisely the same worlds.

Epistemic States. Via the generalized indicator function $[\![\cdot]\!]_\omega$, we are able to evaluate conditionals with respect to possible worlds. However, we have seen in Lemma 2.3.2 that to decide whether a conditional, as an entity, is accepted, we need richer semantic structures like epistemic states. Epistemic states Ψ in the sense of Halpern [48] enable us to give appropriate semantics to conditionals.

In [72], Kraus, Lehmann and Magidor used *preferential relations* to perform non-monotonic reasoning. To model plausibility relations over possible worlds, preferential relations have proven useful [35, 37]. However, in [72], the authors use a

[5]This identification is sufficient in the context of this thesis and widely accepted in the research area [63, 119].

broad definition of preferential relations, e.g., they do not exclude cyclic relations. In this thesis, we focus on faithful total preorders $\preceq_\Psi$ from Definition 2.3.1 to represent Ψ semantically, and assume that each Ψ is equipped with such faithful assignments $\preceq_\Psi$ according to Definition 2.3.1. We have seen in Theorem 2.3.3 that such assignments $\preceq_\Psi$ are necessary and sufficient to guarantee the postulates (KM$\star$1) – (KM$\star$6) in the context of belief revision, i.e., they display the fundamental meta-structures for belief revision, which places them at the core of Belief Revision. For evaluating conditionals and (iterated) belief revision operators in the context of this thesis, we focus on qualitative and semi-quantitative preferential relations, which we describe in the following.

Plausibilistic TPOs. Plausibilistic TPOs $\preceq_\Psi$ are a common qualitative representation for iterated belief revision and have already been discussed in the context of the DP framework in Section 2.3. They order worlds according to their plausibility, where $\omega \preceq_\Psi \omega'$ means that the agent with belief state Ψ deems ω to be at least as plausible as ω', and as usual $\omega \prec_\Psi \omega'$ holds if $\omega \preceq_\Psi \omega'$ but not $\omega' \preceq_\Psi \omega$, and by $\omega \approx_\Psi \omega'$, we abbreviate that $\omega \preceq_\Psi \omega'$ and $\omega' \preceq_\Psi \omega$ holds. Each plausibilistic TPO $\preceq_\Psi$ on possible worlds induces a preference ordering on formulas as follows

$$A \preceq_\Psi B \text{ iff for all } \omega_A \in \min(A, \preceq_\Psi) \text{ and } \omega_B \in \min(B, \preceq_\Psi)\ \omega_A \preceq \omega_B.$$

TPOs over possible worlds display a qualitative approach to belief representation, as they do not use numbers to indicate plausibility or preference.

Technically, $\preceq_\Psi$ is a binary relation on Ω and the belief set $Bel(\preceq_\Psi)$ is defined via the set of minimal worlds in $\preceq_\Psi$, s.t.

$$Bel(\preceq_\Psi) = Th(\{\omega \,|\, \omega \preceq_\Psi \omega' \text{ for all } \omega \in \Omega\}).$$

The belief set of a plausibilistic TPO is defined by the minimal worlds in $\preceq_\Psi$. Thus, the plausibility of a world ω is expressed as closeness to the belief set $Bel(\preceq_\Psi)$. We define the set of models $\mathrm{Mod}(\preceq_\Psi) = Bel(\preceq_\Psi)$, s.t. $\preceq_\Psi \models A$ iff $A \in Bel(\preceq_\Psi)$ which is coherent with the definition of $\Psi \models A$ on page 28.

Ranking Functions. Spohn introduced in [117, 119] *ordinal conditional functions* (OCFs), also called *ranking functions*, which substantiate preference orderings on worlds with numerical ranks.

Definition 2.4.1 ([119])**.** Ordinal conditional functions *(OCFs) or* ranking functions *are functions κ that maps each possible world ω a plausibility rank, s.t. $\kappa : \Omega \to \mathbb{N}_0$ with $\kappa^{-1}(\{0\}) \neq \emptyset$.*

It holds that κ assigns to each world ω and (im)plausibility rank $\kappa(\omega) \geq 0$, s.t. the higher ω is ranked by κ, the less plausible it is. The normalization constraint $\kappa^{-1}(\{0\})$ stipulates that there exist worlds with maximal plausibility, i.e., with rank zero. For a propositional formula $A \in \mathcal{L}$, the OCF ranking $\kappa(A)$ is defined as

$$\kappa(A) = \min\{\kappa(\omega) \,|\, \omega \models A\}, \tag{2.7}$$

i.e., the plausibility of A is defined via minimal worlds satisfying A. Together with the normalization constraint, this implies that for each $A \in \mathcal{L}$, it holds that $\kappa(A) = 0$ or $\kappa(\overline{A}) = 0$. For the contradiction $\perp$, it holds that $\kappa(\perp) = \min\{\kappa(\omega) \,|\, \omega \models \perp\} = \infty$, since the minimum over the empty set equals ∞ in $\mathbb{N}$. Furthermore, it holds that

$$\kappa(A \vee B) = \min\{\kappa(\omega) \,|\, \omega \models A \text{ or } \omega \models B\} = \min\{\kappa(A), \kappa(B)\}. \tag{2.8}$$

OCFs κ display a TPO on possible worlds, i.e., the notion of minimal worlds can also be applied to OCFs. We notate sets of minimal worlds for $\Omega' \subseteq \Omega$ as follows:

$$\min(\Omega', \kappa) = \{\omega \in \Omega' \,|\, \kappa(\omega) \leqslant \kappa(\omega') \text{ for all } \omega' \in \Omega'\}.$$

OCFs κ are suitable representations of an agent's belief state, and the numerical ranks correspond to degrees of disbelief [119]. Worlds an agent deems most plausible, i.e., that are ranked zero, constitute κ's belief set

$$Bel(\kappa) = Th(\kappa^{-1}(\{0\})) = \{A \in \mathcal{L} \,|\, \omega \models A \text{ for all } \omega \text{ with } \kappa(\omega) = 0\}.$$

An OCF κ accepts a formula A, $\kappa \models A$, iff $\kappa(\overline{A}) > 0$ which is equivalent $A \in Bel(\kappa)$. Hence, A is believed iff $\overline{A}$ is disbelieved to some positive degree [119]. We define $\mathrm{Mod}(\kappa) = Bel(\kappa)$, and it holds that $\kappa \models A$ iff $A \in Bel(\kappa)$.

Moreover, ranking functions are suitable representations of epistemic states [119] and each OCF κ induces a plausibilistic TPO $\preceq_\kappa$ on Ω via

$$\omega \preceq_\kappa \omega' \text{ iff } \kappa(\omega) \leqslant \kappa(\omega'). \tag{2.9}$$

Note, that $\preceq_\kappa$ is a faithful assignment in sense of Darwiche and Pearl, i.e., satisfies Definition 2.3.1 on page 27.

In [119], Spohn defined a conditionalization for OCFs analog to the classic conditionalization of probability distributions. For $A \in \mathcal{L}$, the conditionalized OCF $\kappa|A$ displays a ranking function on all models of A, s.t.

$$\kappa|A : \mathrm{Mod}(A) \to \mathbb{N}_0 \text{ with } \kappa|A\,(\omega) = \kappa(\omega) - \kappa(A). \tag{2.10}$$

Note that, for a minimal A-model $\omega_A \in \min(A, \kappa)$, it holds that $\kappa|A\,(\omega_A) = \kappa(\omega_A) - \kappa(A) = \kappa(A) - \kappa(A) = 0$ due to the minimality of ranks. Hence, each $\kappa|A$ satisfies Definition 2.4.1, i.e., displays an OCF. Ranking functions display a more concrete approach to belief representation than abstract orderings of worlds like plausibilistic TPOs. Yet, they are less burdened with many numbers than the most famous fully quantitative approach to representing an agent's epistemic state, probability distributions, which are considered the most adequate representations of belief states [45, 117]. Ranking functions provide a qualitative abstraction of probabilities and, therefore, can be seen as *semi-quantitative* representations of epistemic states.

Moreover, OCFs provide us with the powerful arithmetic of natural numbers to express plausibility or beliefs. Thus, they not only allow us to compare worlds resp. formulas according to their plausibility but also investigate their relative distances. Note that OCFs can have empty layers, i.e., $\kappa(\omega_1) < r < \kappa(\omega_2)$, s.t. there exists no world ω' with $\kappa(\omega') = r$. Empty layers are a distinguishing feature between OCFs and plausibilistic TPOs and contribute to the flexibility of their semi-quantitative approach to belief representation. Although it is useful to introduce empty layers within the plausibility ordering of ranking functions, it can sometimes be useful to exclude them[6].

Definition 2.4.2. *We call an OCF κ* convex *if it has no empty layers, i.e., if for*

[6]Some of the results in Part I rely on convex OCFs. Also, in [69] and [55], it is useful to consider ranking functions that do not allow for empty layers.

each rank r with $0 \leqslant r \leqslant \max_{\omega \in \Omega} \kappa(\omega)$, there exists a world ω with $\kappa(\omega) = r$.

If two OCFs differ only concerning the position of their empty layers, we call them equivalent. Formally, two OCFs κ_1, κ_2 are *equivalent* if for all $\omega_1, \omega_2 \in \Omega$, it holds that $\kappa_1(\omega_1) \leqslant \kappa_1(\omega_2)$ iff $\kappa_2(\omega_1) \leqslant \kappa_2(\omega_2)$.

Conditionals and Epistemic States Conditional beliefs and TPOs as representations of epistemic states for belief revision operators are closely connected since both provide the necessary means to compare beliefs on a semantic level according to their plausibility. Lemma 2.3.2 states that a plausibilistic TPO accepts or satisfis a conditional $(B|A)$, $\preceq_\Psi \models (B|A)$, iff all minimal models of A satisfy B. This provides us with the following equivalent acceptance condition for conditionals in $\preceq_\Psi$

$$\preceq_\Psi \models (B|A) \text{ iff } \omega_A \models B \text{ for } \omega_A \in \min(A, \preceq_\Psi) \text{ iff } AB \prec_\Psi A\overline{B}.$$

Thus, $\preceq_\Psi \models (B|A)$ holds if the verification of $(B|A)$ is strictly more plausible than the falsification. This translates directly to the OCF framework and it holds that

$$\kappa \models (B|A) \text{ iff } \omega_A \models B \text{ for } \omega_A \in \min(A, \kappa) \text{ iff } \kappa(AB) < \kappa(A\overline{B}). \tag{2.11}$$

We call a finite set of conditionals $\Delta = \{(B_i|A_i)\}_{i=1,\dots,n}$ a *conditional belief base.* The study of conditional belief bases has been a fruitful field of research in the past (e.g., in [1, 89, 73, 63, 55]). It holds that a plausibilistic TPO satisfies or accepts Δ, i.e., $\preceq_\Psi \models \Delta$, iff $\preceq_\Psi \models (B|A)$ for all $i = 1, \dots, n$. And this notion of acceptance can be easily applied to OCFs, s.t. $\kappa \models \Delta$ iff $\kappa \models (B_i|A_i)$ for all $i = 1, \dots, n$. We have seen before that the acceptance of a conditional imposes restrictions on the underlying epistemic state, and therefore, it is not trivial to check whether a set of conditionals is consistent, i.e., whether there exists an epistemic state Ψ, s.t. $\Psi \models \Delta$. Adams defined in [1] a notion of *tolerance* for conditionals and sets Δ, and Goldszmidt and Pearl presented a *tolerance test* in [40] based on this notion, which ultimately leads to a *consistency-test algorithm.* In this thesis, we investigate conditional belief bases in light of epistemic states and belief revision and rely on a handy (model-based) definition of consistency for conditional belief bases. We say that Δ is consistent iff there exists an epistemic state Ψ, s.t. Ψ is a model of Δ, i.e., $\Psi \models \Delta$. This definition is equivalent to the one presented in [1] resp. [40] according to Pearl [89]. Also, it corresponds to the consistency of formulas in classical logic.

Throughout this thesis, we stick to the following common assumptions about conditional belief bases $\Delta = \{(B_1|A_1), \ldots, (B_n|A_n)\}$ to avoid easy, but possibly lengthy, case distinctions.

Uniqueness We assume that Δ does not contain conditionally equivalent conditionals, i.e., for all $i, j \in \{1, \ldots, n\}$, $(B_i|A_i) \equiv (B_j|A_j)$ implies $i = j$.

Global Consistency Each Δ is consistent unless stated explicitly otherwise.

Next, we further investigate weak conditionals in the context of epistemic states.

Weak Conditionals. In Part II of this thesis, we make use of *weak conditionals*, which have been considered a valuable extension to the framework of conditional logic [32, 112, 109].

In general, $(|D|C|)$ expresses the notion '*If C, then D might be the case but $\overline{D}$ is not plausible*', i.e., the acceptance of D is not guaranteed if C is accepted but might be possible. So, it holds that weak conditionals implement *negative conditional information*, expressing that the corresponding negated standard conditional $(\overline{D}|C)$ does not hold [77]. This notion of negative information has been investigated in [14] in the light of *rational consequence relations* defined by Gabbay and Makinson [38, 78]. Note that declaring that $(\overline{D}|C)$ does not hold is certainly not the same as declaring that the (contrary) conditional $(D|C)$ holds. The acceptance of a weak conditional expresses the notion that an agent does *not* believe in a conditional. Weak conditionals as negative information are crucial in [32, 103]. And in [112], we investigate a notion of relevance inherent to conditionals which can be expressed using a set of standard and weak conditionals.

The evaluation of a weak conditional $(|D|C|)$ corresponds to the three-valued response behavior of interpretations ω in (2.6) of standard conditionals $(D|C)$. We distinguish between verification $\omega \models CD$, falsification $\omega \models C\overline{D}$ and neutrality $\omega \models \overline{C}$. As for standard conditionals, to validate and further investigate (weak and standard) conditionals, we need richer epistemic structures than plain propositional interpretations [83, 16].

A weak conditionals $(|D|C|)$ is accepted in an epistemic state Ψ, represented by $\preceq_\Psi$ or an OCF κ, written as $\Psi \models (|D|C|)$, if and only if $\Psi \not\models (\overline{D}|C)$, i.e.,

$$CD \preceq_\Psi C\overline{D} \text{ resp. } \kappa(CD) \leqslant \kappa(C\overline{D}). \tag{2.12}$$

For weak conditionals $(|D|C|)$, the acceptance condition of standard conditionals is weakened, and that is where the term actually comes from. The acceptance of $(|D|C|)$ allows for indifference between verification CD and falsification $C\overline{D}$. In this case both $(D|C)$ and $(\overline{D}|C)$ fail to be accepted. As for standard conditional belief bases, it holds that weak conditional belief bases $\Delta^{\mathrm{w}} = \{(|D_i|C_i|)\}_{i=1,\ldots,n}$ are consistent iff there exists a belief state Ψ which accepts it $\Psi \models \Delta^{\mathrm{w}}$, i.e., $\Psi \models (|D_i|C_i|)$ for $i = 1, \ldots, n$. Due to the weakened acceptance condition that allows for indifference between verification and falsification, it holds for a uniform OCF κ_u with $\kappa_u(\omega) = 0$ for all $\omega \in \Omega$, that $\kappa_u \models \Delta^{\mathrm{w}}$, since $\kappa(C_iD_i) = 0 = \kappa(C_i\overline{D_i})$ for all $i \in \{1, \ldots, n\}$. The OCF κ_u corresponds to a plausibilistic TPO $\preceq_\Psi$ with just a single layer.

Proposition 2.4.1. *Each set* $\Delta^{\mathrm{w}} = \{(|Y_i|X_i|)\}_{i=1,\ldots,n}$ *is consistent.*

For mixed sets of conditionals, i.e., conditional belief bases consisting of both standard and weak conditionals, we provided a consistency test algorithm in [107]. Weak conditionals introduce interesting dynamics to general conditional belief bases, yet, the paper in [107] is somewhat technical, and we do not use mixed sets of conditionals in this thesis. Therefore, we skip the in-depth discussion of [107].

2.5 Conditional Belief Revision

In this section, we motivate revision operators on epistemic states that take (sets of) conditionals as input and discuss the fundamental *principle of conditional preservation* (PCP), which ensures that conditional beliefs are preserved in a rational way. Then, we define c-revisions first for sets of (standard) conditionals and then for sets of weak conditionals in the context of OCFs. C-revisions from [63] obey the (PCP) and serve as a proof of concept in different ways in the course of this thesis.

2.5.1 Motivation for Conditional Belief Revision

Starting from the classical propositional AGM revision, which takes deductively closed belief sets K as a basis, Darwiche and Pearl recognized the need to broaden the scope of Belief Revision to epistemic states Ψ to reflect an agent's complex attitudes, thoughts and preferences. Their account on iterated belief revision for epistemic states has proven to align with not only the preservation of propositional beliefs, as required by the principle of minimal change, but also with the preservation of conditional beliefs to some extent [29, 15]. Conditional beliefs guide the revision process implicitly, via the representation of epistemic states as preference relations $\preceq_\Psi$, or explicitly, when given as a set of conditional beliefs, and strike down as a paradigm of *preserving conditional beliefs*. In the DP framework, the preservation of conditional beliefs for propositional revision is axiomatized via four postulates (C1) – (C4) (cf. Section 2.3). This paradigm has been widely recognized as one of the key features for designing rational revision operators (cf. [29, 13, 82, 15, 26]). However, it is not particularly far to assume that new information might not only appear in the form of a propositional sentence but also as a conditional belief itself, such that the agent needs to revise with a new conditional belief, i.e., the investigation of a rational *revision operator for conditional revision* displays a meaningful and essential extension to the framework of Belief Revision.

Most approaches to conditional revision operators presented in the past (cf. [81, 18, 27]) are somewhat based on propositional revision. An exception is the highly general framework of *c-revisions* presented by Kern-Isberner in [63, 64]. C-revisions provide a framework of belief revision on epistemic states represented as ranking functions that enable us to incorporate new information in the form of (a set of) conditionals, thus reflecting a newly acquired revision policy[7]. Note that, in general, propositions A can be identified with conditionals $(A|\top)$ with tautological antecedents (cf. Section 2.4). Thus conditional revision can be seen as a generalization of revision with propositional beliefs.

[7] The notion of conditionals as revision policies are discussed in several works (cf. e.g., [15, 95]) and roots in philosophy, see [31] for an overview.

2.5.2 The Principle of Conditional Preservation

To incorporate new conditional beliefs adequately Kern-Isberner proposed in [62] several *postulates for conditional revision*, which subsume the DP postulates (C1) – (C4) for the induced propositional operator. These postulates for conditional revision axiomatize a fundamental *principle of conditional preservation* (PCP) which was first uncovered by Kern-Isberner in [63] and fully axiomatized for Belief Revision in [64]. And in [65], a *qualitative principle of conditional preservation* $(\mathrm{PCP})^{\bullet}$ for iterated belief change was introduced. We recall *the principle of conditional preservation for OCFs* $(\mathrm{PCP})^{*}$ [67] and $(\mathrm{PCP})^{\bullet}$ from [65] via a less algebraic but equivalent presentation from [70]. Note that the operator $\bullet$ refers to revision operators for qualitative epistemic states, like plausibilistic TPOs, whereas we use $*$ as revision operators that can be applied to OCFs. Since the revision with (sets of) conditionals can be seen as a generalization of propositional revision, we use the same symbols to indicate belief revision with propositional resp. conditional input.

Let $\Delta = \{(B_1|A_1), \ldots, (B_n|A_n)\}$ be a set of conditionals, and let $\tilde{\Omega}_1 = \{\omega_1, \ldots, \omega_k\}$ be a (multi)set of worlds, i.e., the worlds need not all be distinct. For each conditional $(B_j|A_j) \in \Delta$, we count the number of verifications and falsifications in the multiset $\tilde{\Omega}_1$:

$$\begin{aligned}\#\mathit{Ver}_{(B_j|A_j)}(\tilde{\Omega}_1) &= ||\{\omega \in \tilde{\Omega}_1 : \omega \models A_j B_j\}|| \\ \#\mathit{Fals}_{(B_j|A_j)}(\tilde{\Omega}_1) &= ||\{\omega \in \tilde{\Omega}_1 : \omega \models A_j \overline{B_j}\}||,\end{aligned}$$

where $||\cdot||$ counts the elements of multisets with multiplicity. With this notation, the (PCP) for conditional OCF revision reads as follows [64, 67]:

(PCP)* Let $*$ be a revision operation for OCFs κ and $\Delta = \{(B_1|A_1), \ldots, (B_n|A_n)\}$ a conditional belief set, s.t. $\kappa^* = \kappa * \Delta$. If two multisets of possible worlds $\tilde{\Omega}_1 = \{\omega_1, \ldots, \omega_k\}$ and $\tilde{\Omega}_2 = \{\omega'_1, \ldots, \omega'_k\}$ with the same cardinality satisfy

$$\#\mathit{Ver}_{(B_j|A_j)}(\tilde{\Omega}_1) = \#\mathit{Ver}_{(B_j|A_j)}(\tilde{\Omega}_2) \text{ and } \#\mathit{Fals}_{(B_j|A_j)}(\tilde{\Omega}_1) = \#\mathit{Fals}_{(B_j|A_j)}(\tilde{\Omega}_2)$$

for each conditional $(B_j|A_j) \in \Delta$ then prior κ and posterior $\kappa^* = \kappa * \Delta$ are related by

$$(\kappa(\omega_1) - \kappa(\omega_1')) + \ldots + (\kappa(\omega_k) - (\kappa(\omega_k')) \\ = (\kappa^*(\omega_1) - \kappa^*(\omega_1')) + \ldots + (\kappa^*(\omega_k) - \kappa^*(\omega_k')). \tag{2.13}$$

The (PCP)* relies on the arithmetic OCFs are equipped with, but for qualitative representations of epistemic states, like TPOs, we do not have addition and subtraction. Yet, to define a qualitative principle of conditional preservation, we can make use of an immediate high-level consequence of (2.13): if the left-hand side of (2.13) (corresponding to the prior differences) is negative/positive, then its right-hand side (corresponding to posterior differences) must also be negative/positive. From this, we can derive a qualitative PCP for conditional revision of TPOs [65]:

(PCP)$^\bullet$ Let $\bullet$ be a revision operator for (qualitative) epistemic states $\Psi = (\Omega, \preceq_\Psi)$ and conditional belief sets $\Delta = \{(B_1|A_1), \ldots, (B_n|A_n)\}$. If two multisets of possible worlds $\tilde{\Omega}_1 = \{\omega_1, \ldots, \omega_k\}$ and $\tilde{\Omega}_2 = \{\omega_1', \ldots, \omega_k'\}$ with the same cardinality satisfy

$$\# Ver_{(B_j|A_j)}(\tilde{\Omega}_1) = \# Ver_{(B_j|A_j)}(\tilde{\Omega}_2) \text{ and } \# Fals_{(B_j|A_j)}(\tilde{\Omega}_1) = \# Fals_{(B_j|A_j)}(\tilde{\Omega}_2)$$

for each conditional $(B_j|A_j) \in \Delta$, then prior Ψ and posterior $\Psi^\bullet = \Psi \bullet \Delta$ satisfy the following two conditions:
(1) If for all $i, 1 \leqslant i \leqslant k$, it holds that $\omega_i \preceq_\Psi \omega_i'$, and there is at least one $i, 1 \leqslant i \leqslant k$ such that $\omega_i \prec_\Psi \omega_i'$ holds, then there is $j, 1 \leqslant j \leqslant k$, such that $\omega_j \prec_{\Psi^\bullet} \omega_j'$ holds.
(2) If for all $i, 1 \leqslant i \leqslant k$, it holds that $\omega_i \preceq_{\Psi^\bullet} \omega_i'$, and there is at least one $i, 1 \leqslant i \leqslant k$ such that $\omega_i \prec_{\Psi^\bullet} \omega_i'$ holds, then there is $j, 1 \leqslant j \leqslant k$, such that $\omega_j \prec_\Psi \omega_j'$ holds.

The (PCP) is based solely on observing conditional structures like the number of verifying resp. falsifying worlds without taking acceptance conditions of conditionals into account. This focus on structural aspects of conditional preservation is the reason for this principle's broad applicability and flexibility within the context of iterated belief revision.

2.5.3 C-Revisions

C-revisions were introduced in [63] as a general class of revision operators for sets of conditionals Δ which obey the principle of conditional preservation while simultaneously satisfying the classical success condition for belief revision [64].

In this thesis, we develop a qualitative version of c-revisions capable of revising plausibilistic TPOs with sets of conditional rules (cf. Section 6.4). But for now, we focus on defining c-revisions for ranking functions and sets of conditionals, which map a ranking function κ onto a c-revised ranking function $\kappa^{\mathrm{c}} = \kappa *^{\mathrm{c}} \Delta$, s.t. the posterior OCF $\kappa *^{\mathrm{c}} \Delta \models \Delta$. It holds that each c-revision is an iterated revision operator for OCFs with sets of conditionals in the sense of Darwiche and Pearl, i.e., they satisfy the postulates (C1) – (C4) since (PCP)* implies the DP postulates [63].

Definition 2.5.1 (C-revisions for OCFs [63])**.** *Let κ be a ranking function and $\Delta = \{(B_1|A_1), \dots, (B_n|A_n)\}$ a set of conditionals. Then a* c-revision of κ by Δ *is an OCF $\kappa^{\mathrm{c}} = \kappa *^{\mathrm{c}} \Delta$ constructed from non-negative impact factors η_i assigned to each $(B_i|A_i)$ and an integer κ_0 such that κ^{c} accepts Δ and is given by:*

$$\kappa^{\mathrm{c}}(\omega) = \kappa_0 + \kappa(\omega) + \sum_{\substack{1 \leqslant i \leqslant m \\ \omega \models A_i \overline{B}_i}} \eta_i \tag{2.14}$$

$$\textit{with } \kappa_0 = -\min_{\omega \in \Omega} \{\kappa(\omega) + \sum_{1 \leqslant i \leqslant m,\, \omega \models A_i \overline{B}_i} \eta_i\}. \tag{2.15}$$

It holds that κ_0 given by (2.15) is a normalization factor ensuring that $\kappa^{\mathrm{c}}(\omega) = 0$ for at least one world ω, s.t. κ^{c} is a ranking function [64].

The definition of c-revisions in (2.14) displays a basic version of the one presented in [64] since it solely punishes worlds falsifying conditionals in Δ by adding an impact factor η_i. The more sophisticated version from [64] also considers impact factors rewarding worlds which verify conditionals in Δ. The full version of c-revisions presented in [63], which considers not only impact factors for falsifying worlds but also rewarding impact factors for verifying them, is fully characterized by (PCP)* and the success condition $\kappa^{\mathrm{c}} \models \Delta$. However, (2.14) from Definition 2.5.1 is enough to provide a highly general framework for revising OCFs with sets of conditionals Δ and thus is used more frequently in the context of conditional revision [70, 107, 112].

It holds that the impact factors η_i, assigned to each $(B_i|A_i)$, have to be chosen so as to satisfy the following *Success* postulate for conditional revision:

$$\textbf{(Success)}\ \ \kappa * \Delta \models \Delta, \text{ i.e., } \kappa * \Delta(A_iB_i) < \kappa * \Delta(A_i\overline{B_i}) \text{ for all } (B_i|A_i) \in \Delta \quad (2.16)$$

Starting from the acceptance condition for conditionals in (Success), we get for the c-revision $\kappa *^{c} \Delta = \kappa^{c}$ inequalities constraining the impact factors $\eta_i \in \mathbb{N}_0$ [64]. More precisely, via the definition of OCF-ranks for formulas via minimal worlds and equation (2.14), the constraints $\kappa^{c}(A_iB_i) < \kappa^{c}(A_i\overline{B}_i)$ for $1 \leqslant i \leqslant n$ expand to

$$\min_{\omega \models A_iB_i} \underbrace{\{\kappa_0 + \kappa(\omega) + \sum_{\omega \models A_k\overline{B}_k} \eta_k\}}_{(2.17a)} < \min_{\omega \models A_i\overline{B}_i} \underbrace{\{\kappa_0 + \kappa(\omega) + \sum_{\omega \models A_k\overline{B}_k} \eta_k\}}_{(2.17b)} \tag{2.17}$$

The left minimum ranges over models of A_iB_i, so the conditional $(B_i|A_i)$ is not falsified by any considered world. Thus η_i is no element of any sum (2.17a). As opposed to this, the right minimum ranges over the models of $A_i\overline{B}_i$, so the conditional $(B_i|A_i)$ is falsified by every considered world. Thus η_i is an element of every sum in (2.17b). With these deliberations, we can rewrite the inequalities to

$$\min_{\omega \models A_iB_i} \{\kappa_0 + \kappa(\omega) + \sum_{\substack{\omega \models A_k\overline{B}_k \\ i \neq k}} \eta_k\} < \eta_i + \min_{\omega \models A_i\overline{B}_i} \{\kappa_0 + \kappa(\omega) + \sum_{\substack{\omega \models A_k\overline{B}_k \\ i \neq k}} \eta_k\} \tag{2.18}$$

and therefore, we get

$$\eta_i > \min_{\omega \models A_iB_i} \Big\{\kappa(\omega) + \sum_{\substack{j \neq i \\ \omega \models A_j\overline{B}_j}} \eta_j\Big\} - \min_{\omega \models A_i\overline{B}_i} \Big\{\kappa(\omega) + \sum_{\substack{j \neq i \\ \omega \models A_j\overline{B}_j}} \eta_j\Big\}. \tag{2.19}$$

for all $1 \leqslant i \leqslant n$. These inequalities regulate conditional dependencies within Δ, s.t. each conditional is treated adequately in the revision process. For the rest of this thesis, we adopt the following convention about the impact factors η_i:

Nonnegativity We assume that for each c-revision the impact factors η take non-negative values, i.e., $\eta \geq 0$.

This convention expresses the notion that the falsification of a conditional should not increase the plausibility of the corresponding world. In general, c-revisions $\kappa *^c \Delta$ are not unique since the impact factors in (2.14) are defined via inequalities, thus, cannot be chosen uniquely in general. Later, in Section 3.1.2, we deal with the matter of choosing suitable impact factors more thoroughly, but for now, we keep in mind that each c-revision $\kappa *^c \Delta$ refers to a family of suitable c-revisions, each of which is uniquely defined by the corresponding vector of impact factors [9].

Furthermore, it holds that c-revisions satisfy (TV) from page 30 since $\top$ is not falsified by any world ω, the sum in (2.14) is empty. We get that $\kappa = \kappa *^c \top$ [64].

We illustrate c-revisions with sets of conditionals via the following example. This example is an extended version from [64] of the famous *penguin example*, which is well-known in the context of belief revision.

Example 2.5.1 (Adapted from [64])**.** *We consider the ranking function κ from Table 2.1 which represents an agents beliefs over the signature $\Sigma = \{b, f, k, p, w\}$ concerning the atoms b - birds, f - flying, k - kiwis, p - penguins, and w - winged animals. Examining κ closely, we see that the agent accepts (among others) the following inference rules concerning her knowledge of birds:*

$\delta_1 : (f\|b)$	*birds fly,*	$\delta_2 : (w\|b)$	*birds have wings,*
$\delta_3 : (b\|p)$	*penguins are birds,*	$\delta_4 : (b\|k)$	*kiwis are birds*

We can conclude from κ that both representatives of birds – penguins and kiwis – inherit properties from their superclass. Thus the agents believe that penguins fly and kiwis have wings. Suppose now that she concludes that this false, i.e., penguins do not *fly and kiwis do* not *have wings. So, the agent revises her belief state κ by these new information $\Delta = \{(\overline{f}|p), (\overline{w}|k)\}$ and we compute the c-revision via* (2.14)

$$\begin{aligned}\kappa^c(\omega) &= \kappa *^c \Delta(\omega) = \kappa *^c \{(\overline{f}|p), (\overline{w}|k)\}(\omega)\\ &= \kappa_0 + \kappa(\omega) + \begin{cases}\eta_p, & \omega \models pf\\ 0, & else\end{cases} + \begin{cases}\eta_k, & \omega \models kw\\ 0, & else\end{cases}.\end{aligned}$$

Substantiating (2.19)*, we yield the following inequalities defining the impact factors*

η_p for $(\overline{f}|p)$ and η_k for $(\overline{w}|k)$:

$$\eta_p > \min_{\omega\models p\overline{f}}\{\kappa(\omega)+\begin{cases}\eta_k, & \omega\models kw\\ 0, & else\end{cases}\} - \min_{\omega\models pf}\{\kappa(\omega)+\begin{cases}\eta_k, & \omega\models kw\\ 0, & else\end{cases}\} = 1-0=1$$

$$\eta_k > \min_{\omega\models k\overline{w}}\{\kappa(\omega)+\begin{cases}\eta_p, & \omega\models pf\\ 0, & else\end{cases}\} - \min_{\omega\models kw}\{\kappa(\omega)+\begin{cases}\eta_p, & \omega\models pf\\ 0, & else\end{cases}\} = 1-0=1$$

Note that any impact factor η_p resp. η_k which satisfies these inequalities, constitutes a c-revision, but in order to keep numerical changes minimal, we choose $\eta_p^\star = 2$ and $\eta_k^\star = 2$. Note that the additional superscript indicates the specific choice of an impact factor satisfying the inequality defining it. For the normalization constant κ_0, we yield via (2.15)*, since*

$$\kappa_0 = \min_{\omega\in\Omega}\{\kappa_0+\kappa(\omega)+\begin{cases}\eta_p, & \omega\models pf\\ 0, & else\end{cases}+\begin{cases}\eta_k, & \omega\models kw\\ 0, & else\end{cases}\} = \kappa(\overline{p}bfw\overline{k}) = 0,$$

i.e., no further normalization is necessary. The schematic c-revised OCF $\kappa^{\mathrm{c}} = \kappa *^{\mathrm{c}} \Delta$ and the c-revision $\kappa^\star$ employing $\eta_p^\star$ resp. $\eta_k^\star$ is shown in Table 2.1.
It holds that κ^{c} still accepts the conditionals $(f|b),(w|b),(b|p)$ and $(w|k)$. Note that the new information that penguins do not fly resp. that kiwis do not have wings might cast doubts on both penguins and kiwis, being birds. However, the conditionals $(b|p)$ and $(w|k)$ are inscribed in κ with sufficient inferential strength, s.t. they do not get lost during revision. This illustrates how c-revisions properly deal with conditional interrelationships.

We consider the special case of c-revisions with a single conditional $(B|A)$, which is helpful in various occasions in this thesis. Compared to the general revision schema of c-revisions, a c-revision with a single conditional is pretty simple, and from (2.14), (2.15) and (2.19), we obtain

$$\kappa *^{\mathrm{c}} (B|A)(\omega) = -\kappa(\overline{A}\vee B)+\kappa(\omega)+\begin{cases}\eta, & \omega\models A\overline{B}\\ 0, & \omega\models \overline{A}\vee B\end{cases} \tag{2.20}$$

$\omega \in \Omega$	κ	$\kappa^{\mathrm{c}} = \kappa *^{\mathrm{c}} \Delta$	$\kappa^\star$	$\omega \in \Omega$	κ	$\kappa^{\mathrm{c}} = \kappa *^{\mathrm{c}} \Delta$	$\kappa^\star$
$pbfwk$	0	$\kappa_0 + \kappa(\omega) + \eta_p + \eta_k$	$= 4$	$pbfw\overline{k}$	0	$\kappa_0 + \kappa(\omega) + \eta_p$	$= 2$
$pbf\overline{w}k$	1	$\kappa_0 + \kappa(\omega) + \eta_p$	$= 3$	$pbf\overline{w}\overline{k}$	1	$\kappa_0 + \kappa(\omega) + \eta_p$	$= 3$
$pb\overline{f}wk$	1	$\kappa_0 + \kappa(\omega) + \eta_k$	$= 3$	$pb\overline{f}w\overline{k}$	1	$\kappa_0 + \kappa(\omega)$	$= 1$
$pb\overline{f}\overline{w}k$	2	$\kappa_0 + \kappa(\omega)$	$= 2$	$pb\overline{f}\overline{w}\overline{k}$	2	$\kappa_0 + \kappa(\omega)$	$= 2$
$p\overline{b}fwk$	4	$\kappa_0 + \kappa(\omega) + \eta_p + \eta_k$	$= 8$	$p\overline{b}fw\overline{k}$	2	$\kappa_0 + \kappa(\omega) + \eta_p$	$= 4$
$p\overline{b}f\overline{w}k$	4	$\kappa_0 + \kappa(\omega) + \eta_p$	$= 6$	$p\overline{b}f\overline{w}\overline{k}$	2	$\kappa_0 + \kappa(\omega) + \eta_p$	$= 4$
$p\overline{b}\,\overline{f}wk$	4	$\kappa_0 + \kappa(\omega) + \eta_k$	$= 6$	$p\overline{b}\,\overline{f}w\overline{k}$	2	$\kappa_0 + \kappa(\omega)$	$= 2$
$p\overline{b}\,\overline{f}\overline{w}k$	4	$\kappa_0 + \kappa(\omega)$	$= 4$	$p\overline{b}\,\overline{f}\overline{w}\overline{k}$	2	$\kappa_0 + \kappa(\omega)$	$= 2$
$\overline{p}bfwk$	0	$\kappa_0 + \kappa(\omega) + \eta_k$	$= 2$	$\overline{p}bfw\overline{k}$	0	$\kappa_0 + \kappa(\omega)$	$= 0$
$\overline{p}bf\overline{w}k$	1	$\kappa_0 + \kappa(\omega)$	$= 1$	$\overline{p}bf\overline{w}\overline{k}$	1	$\kappa_0 + \kappa(\omega)$	$= 1$
$\overline{p}b\overline{f}wk$	1	$\kappa_0 + \kappa(\omega) + \eta_k$	$= 3$	$\overline{p}b\overline{f}w\overline{k}$	1	$\kappa_0 + \kappa(\omega)$	$= 1$
$\overline{p}b\overline{f}\overline{w}k$	2	$\kappa_0 + \kappa(\omega)$	$= 2$	$\overline{p}b\overline{f}\overline{w}\overline{k}$	2	$\kappa_0 + \kappa(\omega)$	$= 2$
$\overline{p}\overline{b}fwk$	2	$\kappa_0 + \kappa(\omega) + \eta_k$	$= 4$	$\overline{p}\overline{b}fw\overline{k}$	0	$\kappa_0 + \kappa(\omega)$	$= 0$
$\overline{p}\overline{b}f\overline{w}k$	2	$\kappa_0 + \kappa(\omega)$	$= 2$	$\overline{p}\overline{b}f\overline{w}\overline{k}$	0	$\kappa_0 + \kappa(\omega)$	$= 0$
$\overline{p}\overline{b}\,\overline{f}wk$	2	$\kappa_0 + \kappa(\omega) + \eta_k$	$= 4$	$\overline{p}\overline{b}\,\overline{f}w\overline{k}$	0	$\kappa_0 + \kappa(\omega)$	$= 0$
$\overline{p}\overline{b}\,\overline{f}\overline{w}k$	2	$\kappa_0 + \kappa(\omega)$	$= 2$	$\overline{p}\overline{b}\,\overline{f}\overline{w}\overline{k}$	0	$\kappa_0 + \kappa(\omega)$	$= 0$

Table 2.1: Ranking function κ and c-revised ranking function κ^{c} resp. $\kappa^\star$ from Example 2.5.1.

with $\kappa_0 = -\kappa(\overline{A} \vee B)$ as normalization constant. As inequality constraining the corresponding impact factor η we get

$$\eta > \kappa(AB) - \kappa(A\overline{B}). \tag{2.21}$$

This inequality displays a reduction of (2.19) for cases when we revise with just a single conditional. Usually, for c-revisions with sets of conditionals, we must pay close attention to interactions with other conditionals, expressed in the sums within the minima defining general impact factors η_i in (2.19). For $\kappa *^{\mathrm{c}} (B|A)$ we can single out a unique minimal (in terms of the impact factors) c-revision by choosing

$$\eta_{\mathrm{m}} = \kappa(AB) - \kappa(A\overline{B}) + 1. \tag{2.22}$$

The following proposition shows that the replacement of the general normalization constant κ_0 in (2.20) by $-\kappa(\overline{A} \vee B)$ is correct.

Proposition 2.5.1. *For the c-revision in* (2.20)*, it holds that* $\kappa_0 = -\kappa(\overline{A} \vee B)$.

Proof. Via (2.15), it holds that $\kappa_0 = -\min\{\kappa(\overline{A} \vee B), \kappa(A\overline{B}) + \eta\}$ and we have to show that $\kappa(\overline{A} \vee B) < \kappa(A\overline{B}) + \eta$ holds. Due to the properties of OCFs (2.7) and (2.8), and the constraints defining η in (2.21) it holds that

$$\begin{aligned}\kappa(\overline{A} \vee B) &= \min\{\kappa(\overline{A}B), \kappa(AB), \kappa(\overline{A}\,\overline{B})\} \leqslant \kappa(AB) \\ &< \kappa(AB) + \kappa(A\overline{B}) - \kappa(A\overline{B}) + 1 \leqslant \kappa(A\overline{B}) + \eta\end{aligned}$$

Thus, it holds that $\kappa_0 = -\kappa(\overline{A} \vee B)$. □

We have already seen that c-revisions allow us to revise with sets of conditionals. Now, we broaden the scope of the c-revision operator $*^{\mathrm{c}}$ towards sets of weak conditionals. In Section 2.1, we discussed the acceptance condition for weak conditionals and standard conditionals for ranking functions κ and saw that both are pretty similar, except for the fact that the first one allows for indifference towards the plausibility of verification vs. falsification, while the latter does not.

In [63], the inequalities defining the impact factor η_i for a revision with a set of standard conditionals are derived from the (Success)-condition via the considerations we stated in (2.17) and (2.18), eventually leading to (2.19). For a set of weak conditionals $\Delta^{\mathrm{w}} = \{(|D_i|C_i|)\}$, the (Success)-condition for conditional revision is weakened in the same manner as the acceptance condition for conditionals in (2.11) vs. the acceptance condition for weak conditionals (2.12) and we get:

$$\begin{aligned}\textbf{(|Success|)}\quad & \kappa * \Delta^{\mathrm{w}} \models \Delta^{\mathrm{w}}, \text{ i.e.,} \\ & \kappa * \Delta^{\mathrm{w}}(C_iD_i) \leqslant \kappa * \Delta^{\mathrm{w}}(C_i\overline{D_i}) \text{ for all } (|D_i|C_i|) \in \Delta^{\mathrm{w}}\end{aligned} \tag{2.23}$$

Instead of demanding for strict inequalities, we allow for equality between verification and falsification. The weakened condition (|Success|) leads to the following definition of c-revision for sets of weak conditionals.

Definition 2.5.2 (C-revisions with weak conditionals for OCFs)**.** *Let* κ *be a ranking function and* $\Delta^{\mathrm{w}} = \{(|D_1|C_1|), \ldots, (|D_n|C_n|)\}$ *a set of weak conditionals. Then a* c-revision of κ by Δ^{w} *is an OCF* $\kappa^{\mathrm{c}}_{\Delta^{\mathrm{w}}} = \kappa *^{\mathrm{c}} \Delta^{\mathrm{w}}$ *constructed from non-negative integers* η_i *assigned to each* $(|D_i|C_i|)$ *and an normalization constant* κ_0 *such that the OCF* $\kappa^{\mathrm{c}}_{\Delta^{\mathrm{w}}}$ *accepts* Δ^{w} *and is given by*

$$\kappa^{\mathrm{c}}_{\Delta^{\mathrm{w}}}(\omega) = \kappa_0 + \kappa(\omega) + \sum_{\substack{1 \leqslant i \leqslant m \\ \omega \models C_i\overline{D}_i}} \eta_i \tag{2.24}$$

$$\textit{with } \kappa_0 = -\min_{\omega \in \Omega}\{\kappa(\omega) + \sum_{1 \leqslant i \leqslant m,\, \omega \models C_i\overline{D}_i} \eta_i\}. \tag{2.25}$$

and the following inequalities constraining η_i *for all* $1 \leqslant i \leqslant n$*:*

$$\eta_i \geq \min_{\omega \models C_i D_i}\Big\{\kappa(\omega) + \sum_{\substack{j \neq i \\ \omega \models C_j\overline{D}_j}} \eta_j\Big\} - \min_{\omega \models C_i\overline{D}_i}\Big\{\kappa(\omega) + \sum_{\substack{j \neq i \\ \omega \models C_j\overline{D}_j}} \eta_j\Big\}. \tag{2.26}$$

Note that the inequalities for η_i follow from (|Success|) in (2.23) in the same manner as the inequalities for standard c-revisions follow from (Success) in (2.16). Thus, c-revisions with weak conditionals correspond to standard c-revisions from (2.14), except that the inequalities constraining the impact factors are not strict. We continue the penguin Example 2.5.1 to illustrate c-revisions with weak conditionals.

Example 2.5.2 (Continue Example 2.5.1)**.** *We consider the c-revised OCF* κ^{c} *from Example 2.5.1, where the agent now believes that penguins and kiwis are birds, but penguins do not fly and kiwis do not have wings. Now, the agent learns something about a small island where a small population of a special class of penguins, so-called super-penguins (*s*), live. Since these birds are also penguins, the agent naturally assumes that they cannot fly. A representation of the agent's belief state via the OCF* κ_{w} *over the signature* $\Sigma_{\mathrm{w}} = \{p, f, s\}$ *can be found in Table 2.2 and it holds that* $\kappa_{\mathrm{w}} \models \{(\overline{f}|p), (p|s), (\overline{f}|s)\}$*. To shorten the matter, we neglect the beliefs about birds, kiwis, and wings.*
Now, the agent learns that super-penguins are famous because some can fly. Thus, she weakens her beliefs about the flight capability of super-penguins so that super-penguins might fly, and she no longer accepts the rule $(\overline{f}|s)$*. This corresponds to*

$\omega \in \Omega$	κ_{w}	$\kappa_{\mathrm{w}}^{\mathrm{c}}$	$\kappa_{\mathrm{w}}^{\star}$	$\omega \in \Omega$	κ_{w}	$\kappa_{\mathrm{w}}^{\mathrm{c}}$	$\kappa_{\mathrm{w}}^{\star}$
pfs	1	$\kappa_0 + \kappa(\omega)$	$= 1$	$\overline{p}fs$	2	$\kappa_0 + \kappa(\omega)$	$= 2$
$pf\overline{s}$	1	$\kappa_0 + \kappa(\omega)$	$= 1$	$\overline{p}f\overline{s}$	2	$\kappa_0 + \kappa(\omega)$	$= 2$
$p\overline{f}s$	0	$\kappa_0 + \kappa(\omega) + \eta_s$	$= 1$	$\overline{p}\overline{f}s$	1	$\kappa_0 + \kappa(\omega) + \eta_s$	$= 2$
$p\overline{f}\overline{s}$	0	$\kappa_0 + \kappa(\omega)$	$= 0$	$\overline{p}\overline{f}\overline{s}$	2	$\kappa_0 + \kappa(\omega)$	$= 2$

Table 2.2: Ranking function κ_{w} and c-revised ranking function $\kappa_{\mathrm{w}}^{\mathrm{c}}$ resp. $\kappa_{\mathrm{w}}^{\star}$ from Example 2.5.1.

the c-revision $\kappa *^{\mathrm{c}} (|f|s|)$*, which is given via* (2.24) *as follows*

$$\kappa_{\mathrm{w}}^{\mathrm{c}}(\omega) = \kappa_{\mathrm{w}} *^{\mathrm{c}} (|f|p|) = \kappa_0 + \kappa(\omega) + \begin{cases} \eta_s, & \omega \models s\overline{f} \\ 0, & else \end{cases} .$$

Since we revise with only a single weak conditional, we do not have to consider possible interactions between conditionals, when computing the impact factor η_s *for* $(|f|s|)$ *and the inequality in* (2.26) *reduces to*

$$\eta_s \geq \min_{\omega \models sf}\{\kappa(\omega)\} - \min_{\omega \models s\overline{f}}\{\kappa(\omega)\} = 1 - 0 = 1.$$

For the normalization constant κ_0*, it holds that*

$$\kappa_0 = \min_{\omega \in \Omega}\{\kappa(\omega) + \begin{cases} \eta_s, & \omega \models s\overline{f} \\ 0, & else \end{cases}\} = \kappa(p\overline{f}\overline{s}) = 0.$$

The schematic c-revised OCF $\kappa_{\mathrm{w}}^{\mathrm{c}}$ *can be found in Table 2.2 as well as the corresponding c-revision* $\kappa_{\mathrm{w}}^{\star}$ *with* $\eta_s^{\star} = 1$*. It holds that* $\kappa_{\mathrm{w}}^{\mathrm{c}}$ *still accepts that in general, penguins do not fly, i.e.,* $\kappa_{\mathrm{w}}^{\mathrm{c}} \models (\overline{f}|p)$*, and that super-penguins are penguins* $\kappa_{\mathrm{w}}^{\mathrm{c}} \models (p|s)$*, but is now indifferent towards whether super-penguins fly or not, i.e.,* $\kappa_{\mathrm{w}}^{\mathrm{c}} \not\models (f|s)$ *and* $\kappa_{\mathrm{w}}^{\mathrm{c}} \not\models (\overline{f}|s)$*.*

Note that the main difference between c-revisions with standard conditionals and c-revision with weak conditionals is that standard c-revisions satisfy the KM postulates for belief revision (KM•1) – (KM•6), i.e., display real revision operators and together with (PCP)* this suffices to make them iterated belief revision operators. For c-revision with weak conditionals, it holds that they display a revision with weaker information and therefore do not ensure the acceptance of standard conditionals, which is why the DP framework for Belief Revision is not the right choice to characterize the change mechanism of c-revisions with weak conditionals. In the following subsection, we demonstrate how c-revisions with weak conditionals can still be incorporated into the iterative belief change framework for belief contraction as opposed to belief revision.

2.5.4 C-Revisions and Iterated Contraction

In this subsection, we clarify the relation between c-revisions with (sets of) weak conditionals and belief contraction operators that can be applied to total preorders, which might be used again for contraction, so-called *iterated contraction operators*.

Compared to iterated belief revision, iterated contraction operators are less examined. Even though special aspects of iterated contraction were considered in the past (cf. [118, 80, 53, 94]), Caridroit, Konieczny, and Marquis in [24] were first to provide a set of propositional contraction postulates (KM−1) – (KM−7) for contraction operators $-$ that assign posterior epistemic states $\Psi - C$ to an initial state Ψ and $C \in \mathcal{L}$ in the style of the KM postulates for belief revision (cf. Section 2.3):

(KM−1) $Bel(\Psi) \models Bel(\Psi - C)$

(KM−2) If $Bel(\Psi) \not\models C$, then $Bel(\Psi - C) \models Bel(\Psi)$

(KM−3) If $Bel(\Psi - C) \models C$, then $C \equiv \top$

(KM−4) $Bel(\Psi - C) \wedge C \models Bel(\Psi)$

(KM−5) If $C \equiv D$, then $Bel(\Psi - C) \equiv Bel(\Psi - D)$

(KM−6) $Bel(\Psi - (C \wedge D)) \models Bel(\Psi - C) \vee Bel(\Psi - D)$

(KM−7) If $Bel(\Psi - (C \wedge D)) \not\models C$, then $Bel(\Psi - C) \models Bel(\Psi - (C \wedge D))$

For a detailed explanation of the corresponding meaning we refer to [24]. In [71] Konieczny and Pino Pérez [71] stated the following KM style characterization of contraction operators satsifying (KM−1) – (KM−7) in terms of faithful preorders:

Theorem 2.5.2 ([71])**.** *A contraction operator $-$ that assigns a posterior epistemic state $\Psi - C$ to a prior state Ψ and a proposition C fulfills (KM−1) – (KM−7) iff* *there exists a faithful assigment $\preceq_\Psi$ for Ψ s.t. for $C \in \mathcal{L}$ it holds that*

$$Bel(\Psi - C) = Bel(\Psi) \cup \min(\overline{C}, \preceq_\Psi). \tag{2.27}$$

Similar to Theorem 2.2.1 for belief revision, the above characterization is insufficient to ensure intuitive, iterative belief contraction. This problem has been addressed in [71], and the authors proposed a set of (semantic) DP style postulates for iterated contraction operators. However, in [67] Kern-Isberner, Bock, Sauerwald, and Beierle showed that the DP style postulates for iterated contraction proposed in [71] in the context of OCFs is implied by the principle of conditional preservation (PCP)* together with (KM$-$1) – (KM$-$7) and the following natural constraint for contraction

$$\kappa - C(\omega) \geq \kappa(\omega) \text{ for all } \omega \models C,$$

i.e., the intuition that under contraction by C, models of C should not be made more plausible (cf. Theorem 9 in [71]). Note that in [67], (PCP)* is formulated for a general change operator on OCFs instead of a revision operator $*$ in order to make it applicable to the case of belief contraction; apart from that, there are no differences to the formulation of (PCP)* from page 43. Thus, in the presence of an additional intuitive constraint and (KM$-$1) – (KM$-$7), a change operator satisfying (PCP)* displays an iterated contraction operator.

Now, we turn to c-revisions with sets of weak conditionals. In (2.12) on page 40, we have seen that $\kappa \models (|D|C|)$ is equivalent to $\kappa \not\models (\overline{D}|C)$, i.e., the acceptance of a weak conditional $(|D|C|)$ necessarily implies the non-acceptance of designated standard conditional $(\overline{D}|C)$[8]. So, for a revision operator $*$ for weak conditionals, which satisfies the weakened success (|Success|) on page 49, there exists a direct correspondence between the revisions with the set of weak conditionals $\Delta^{\mathrm{w}} = \{(|D_i|C_i|)\}_{i=1,\ldots,n}$ and the contraction, in the sense of non-acceptance, of the corresponding set of negated conditionals $\Delta = \{(\overline{D_i}|C_i)\}_{i=1,\ldots,n}$. C-revisions with weak conditionals satisfy (PCP)* [112]. And it is even valid that revision with weak conditionals displays iterated contraction operators. This follows immediately from a special iterated

[8] From this perspective weak conditionals can be seen as negative information, i.e., information about conditionals that do *not* hold. This notion of weak conditionals as negated information is more thoroughly discussed in [32, 108] and in a more general form in [14] for rational closure.

contraction operator, called *type β c-contraction* proposed in [67], which matches our definition of c-revisions with weak conditionals.

To summarize, while c-revisions with standard conditionals display iterated belief revision operators in the sense of Darwiche and Pearl, c-revisions with weak conditionals display iterated contraction operators. This versatility in the change mechanism is possible due to the flexible approach of c-change via impact factors, which depend on the type of conditionals and their respective interaction in the input information.

Part I

The Kinematics Principle in Belief Revision

Chapter 3

Introduction to Part I

One of the core strengths of how we revise our beliefs as human reasoners is that we can distinguish between different contexts and, in the light of new information, change our beliefs concerning this specific context. For example, suppose we receive the following information about a party next weekend: if the weather is fine, a barbecue can be expected; if it is raining, the host will serve a pasta buffet. The weather serves as the context in which we process the new information. This kind of contextual information can be modeled as conditionals, where the premise provides the context for the consequent, in this case, the weather being good or bad. Now, in the case of mutually exclusive resp. disjoint contexts, to which we refer from now on as *cases*, it makes sense to assume that the revision with the new information concerning one of these exclusive contexts should be independent of the revision of the remaining state. If the new input naturally decomposes into cases, only parts of the belief state concerning this case should be relevant for the revisions and revised accordingly. We can take this a step further and suppose that we learn that some of these disjoint scenarios are more plausible than others. The posterior plausibility of the case should not affect the revision with the conditional information related to it. More precisely, we investigate the following advanced belief revision problem[1]:

(CondCS) Let Ψ be an epistemic state represented as a total preorder $\preceq_\Psi$ on possible worlds Ω. Let $A_1, \ldots, A_n$ be exhaustive and exclusive propositions, i.e., cases, and let $\Delta = \Delta_1 \cup \ldots \cup \Delta_n$ be a set of conditionals, with subsets Δ_i

[1] The acronym CondCS abbreviates "Conditional Case Splitting".

containing conditionals whose premises imply A_i, and let $S = \bigvee_{j \in J} A_j$ with $\emptyset \neq J \subseteq \{1, \ldots, n\}$. How should Ψ be revised by Δ and S in a rational way to yield a posterior state $\Psi \bullet (\Delta \cup \{S\})$ (also represented by a total preorder) such that conditional beliefs in Ψ and Δ are treated adequately, and s.t. conditional revised beliefs given A_i are unaffected by the information provided by S?

Different cases are represented by the (exclusive and exhaustive) propositions A_i's, which induce a partition of Δ, s.t. the Δ_i's provide new information referring (only) to the individual cases. Moreover, S expresses some additional information presenting a selection of more plausible cases.

The problem (CondCS) is challenging for several reasons. First, we need to find a way to focus on the specific case in which the new information comes into play, and then, we need to revise with a set of conditionals and a propositional formula simultaneously. The Kinematics principle we propose in this part of the thesis employs the concept of conditionalization to focus on a specific part of the belief state, i.e., models of a specific case. Note that via conditionalization, we introduce a notion of locality for the epistemic state because it enables us to specify which parts of the belief state we focus on during the revision[2].

Our investigations and the results presented in the following chapters are mostly driven by the following research questions:

- *Is the set of exclusive and exhaustive formulas $A_1, \ldots, A_n$ unique for each set Δ? If not, how can we compare them, and is there a best choice for these sets which maximizes the benefits of the Kinematics principle?*

- *How do we define the Kinematics principle for OCFs, and does a revision operator for OCFs exist that is capable of satisfying it? Are there possible implications of the Kinematics principle for propositional revisions, at least in special cases? And is there a way to reconstruct the revision with the whole set Δ just using the local revisions with subsets $\Delta_i \subseteq \Delta$?*

- To provide a rational solution to (CondCS), we crucially need a concept conditionalization and a conditional revision mechanism. The semi-quantitative

[2]Note that this notion of locality we use in this thesis is based on the semantics of epistemic states, and there are several other ways to define a concept of locality for Belief Revision (cf. e.g. [28, 68])

framework of OCFs provides both of these, yet in the qualitative framework, we still need those. An intuitive idea would be to transfer those concepts from the OCF framework to the framework of plausibilistic TPOs. *So, how do we define conditionalization for plausibilistic TPOs rationally? And how do proper transformation operators from plausibilistic TPOs to OCFs and vice versa look that fully comply with conditionalization and revision even to the smallest level of detail? Can these transformation operators be used to transfer c-revisions with sets of conditionals to qualitative frameworks of belief representation?*

- After defining a concept of conditionalization for the qualitative framework and proper transformation operators, we still need to transfer the Kinematics principle from the OCF framework to qualitative TPOs. *How does a Qualitative Kinematics principle (QK) look? Do qualitative c-revisions satisfy this advanced principle?*

- *How is the Kinematics principle applicable for other conditional revision operators, like e.g., the ones presented by Chandler and Booth in [27]?*

In the following section, we present formal preliminaries for the results in this part. We start with some general notions and definitions in Section 3.1.1 and then present strategies for c-revision in Section 3.1.2, which play a significant role in the methodological realization of our results. It holds that the Kinematics principle is one solution to (CondCS) and there are several approaches to revising an agent's beliefs w.r.t. to specific contexts presented in the past; we present some of them in Section 3.2, which summarizes some related work. The following chapters of this part are structured in consecutive organized chapters:

Chapter 4 We investigate a crucial prerequisite forming the setting in which the Kinematics principle comes into play. In Section 4.1, we define sets of exclusive and exhaustive cases $A_1, \dots, A_n$, which take on a unique role in our following investigations and are useful to define cases for the conditional information given in Δ. Furthermore, we introduce a refinement and specificity property in Section 4.2, which enables us to compare different sets of exclusive and exhaustive cases for a set Δ and we present an algorithm that computes the finest splitting of premises for an arbitrary set Δ.

Chapter 5 In this chapter, we present the Kinematics principle in the context of OCFs and show that c-revisions fully realize it in Section 5.1. Then, in Section 5.2, we consider a special setting of the Kinematics principle where the premises of the conditionals in Δ are equivalent to one of the exclusive and exhaustive cases $A_1, \ldots, A_n$, which enables us to redefine revisions with conditionals as propositional revisions for conditionalized OCFs. Concluding our investigations of the Kinematics principle for OCFs, we show in Section 5.3 that via the application of an intuitive merging operator, we can set up a globally c-revised OCF from the corresponding conditionalized and c-revised local OCFs.

Chapter 6 We present the Kinematics principle in the context of plausibilistic TPOs. The following Sections 6.2 and 6.3 are dedicated to defining and investigating crucial new concepts and tools, such as a qualitative conditionalization and a fully compatible transformation between plausibilistic TPOs and OCFs, leading to qualitative c-revisions. In Section 6.4, we show that these, just like their semi-quantitative counterparts, satisfy the Kinematics principle, at least for special cases. In the last section of this part, we first analyze qualitative c-revision with a single conditional resp. with the material implication and relate both of them to each other. This provides ground for the following investigation of the (QK) in the context of the conditional revision operator defined by Chandler and Booth in [27].

Bibliographic Remark. The contents of this part are based on joint work with Gabriele Kern-Isberner and Christoph Beierle [107, 111, 70] (see Section 1.3).

3.1 Formal Preliminaries for This Part

We introduce some notations and formal preliminaries that are relevant to this part of the thesis.

3.1.1 Definitions and Notations

Throughout the following investigations in this part, we consider plausibilistic TPOs over Ω as qualitative representations of epistemic states. Especially in the light of conditionalization, we consider plausibilistic TPOs over subsets $\tilde{\Omega}$ of Ω. In these cases, all definitions and notations we stated previously apply to subsets $\tilde{\Omega}$ of Ω accordingly. Moreover, to guarantee correspondence between conditionalization for ranking functions and for TPOs to a certain degree, an additional *convexity property for TPOs* and OCFs will prove useful.[3] In contrast to TPOs, OCFs can have empty layers. However, if they do not have empty layers, we call them convex (cf. Definition 2.4.2 on page 37). For TPOs, we define convexity as follows:

Definition 3.1.1 ($\preceq_\Psi$-convex). *A subset $\tilde{\Omega} \subseteq \Omega$ is called* $\preceq_\Psi$-convex *if for all $\omega_1, \omega_2 \in \tilde{\Omega}$ with $\omega_1 \preceq_\Psi \omega_2$, the following holds: if $\omega_1 \prec_\Psi \omega' \prec_\Psi \omega_2$ with $\omega' \notin \tilde{\Omega}$, then there is $\omega_3 \in \tilde{\Omega}$ such that $\omega_3 \approx_\Psi \omega'$. A TPO Ψ is* convex *with respect to A if $Mod(A)$ is $\preceq_\Psi$-convex.*

Subsets of Ω which stretch over the layers of $\preceq_\Psi$ "without gaps", i.e., a world from $\tilde{\Omega}$ can be found in each layer, are called convex. Thus, the basic idea of convexity for TPOs and OCFs is the same. We give an example to illustrate the convexity property for TPOs resp. OCFs:

Example 3.1.1 (Convexity of TPOs and OCFs). *The TPO $\Psi : a\bar{b}c \prec \bar{a}bc, abc, \bar{a}b\bar{c}, \bar{a}\bar{b}\bar{c} \prec \bar{a}\bar{b}c, ab\bar{c}, a\bar{b}\bar{c}$ is convex with respect to formulas $A = a$ and $\overline{A} = \bar{a}$, and also with respect to $\tilde{\Omega} = \{a\bar{b}c, abc, a\bar{b}\bar{c}\}$. Hence, $\tilde{\Omega}$ is $\preceq$-convex, whereas $\Omega' = \{a\bar{b}c, \bar{a}\bar{b}c\}$ is not. The OCF κ with $\kappa(a\bar{b}c) = 0$, $\kappa(\bar{a}bc) = \kappa(abc) = \kappa(\bar{a}b\bar{c}) = \kappa(\bar{a}\bar{b}\bar{c}) = 1$ and $\kappa(\bar{a}\bar{b}c) = \kappa(ab\bar{c}) = \kappa(a\bar{b}\bar{c}) = 2$ is also convex. Note that κ arises from Ψ by numbering the layers in Ψ.*

For the consecutive section on related work, we need some basic definitions from probability theory. A probability distribution P is a full quantitative representation of an epistemic state [48]. Then, probabilistic facts are propositional formulas $A[x]$ equipped with a probability $x \in [0, 1]$. A probability distribution P satisfies $A[x]$, denoted as $P \models A[x]$, iff $P(A) = x$. Similarly, probabilistic conditionals $(Y|X)[x]$

[3] In [70], this property was defined under the name *coherence property*.

are conditionals equipped with a probability $x \in [0,1]$ and P models $(Y|X)[x]$, $P \models (Y|X)[x]$, iff $P((Y|X)) = P|X\,(Y) = x$.

To not trivialize the revision task, we presuppose that the conditional belief sets Δ considered in this part are always consistent. Also, we do not use weak conditionals in this part and thus always refer to sets of standard conditionals Δ. Furthermore, we consider different types of revision operators for different epistemic states. For general epistemic states, in our framework represented by plausibilistic TPOs, we use $\bullet$ to notate the conditional revision $\Psi \bullet \Delta$. For quantitative and semi-quantitative belief representation frameworks, like probability distributions or ranking functions, we use $*$ to notate the (conditional) revision operator.

3.1.2 Strategies for C-Revisions

We motivate and introduce strategies for c-revision, allowing us to adapt and select specific c-revision according to the (disjoint) context of the revision[4]. A proof of concept for the Kinematics principle requires a case-sensitive revision mechanism. We show that strategies are crucial for ensuring coherence for c-revisions across different revision scenarios.

C-revisions provide a highly general framework for revising ranking functions by sets of conditionals while respecting the principle of conditional preservation (PCP)* (cf. Subsection 2.5.2) [64]. We recall the basic definition of c-revisions from (2.14) on page 44

$$\kappa *^{\mathrm{c}} \Delta(\omega) = \kappa^{\mathrm{c}}(\omega) = \kappa_0 + \kappa(\omega) + \sum_{\substack{1 \leqslant i \leqslant m \\ \omega \models X_i \overline{B}_i}} \eta_i,$$

where κ_0 is a normalization constant and η_i are impact factors for each conditional $(B_i|X_i)$ added to worlds falsifying the corresponding conditional in Δ and they have to satisfy (2.19) from page 45. We refer to Section 2.5.3 for a detailed discussion of the inequalities and their derivation. The impact factors η_i display the characteristic parameters of a c-revision. We present functions, called *strategies for c-revision*,

[4]Some concepts and results presented in this subsection are similar to those in [6] where c-inferences were examined. However, here we focus on strategies in the context of c-revision that were first introduced in [111] (cf. Section 1.3)

that map each revision problem to a single vector of impact factors. These strategies enable us to make the general concept of c-revision usable in an elegant and axiomatic way. The possible values for the impact factors η_i can be specified by a constraint satisfaction problem.

Definition 3.1.2 ($\mathit{CR}(\kappa,\Delta)$, $\mathit{cr}(\kappa,\Delta)_i$)**.** *Let κ be an OCF and $\Delta = \{(X_1|B_1),\ldots,(X_m|B_m)\}$ a set of conditionals. The* constraint satisfaction problem for c-revisions of κ by Δ, *denoted by $\mathit{CR}(\kappa,\Delta)$, is given by the set of constraints $\mathit{cr}(\kappa,\Delta)_i$, for $i \in \{1,\ldots,m\}$:*

$$(\mathit{cr}(\kappa,\Delta)_i) \qquad \eta_i > \min_{\omega \models X_i B_i} \Big\{ \kappa(\omega) + \sum_{\substack{j \neq i \\ \omega \models X_j \overline{B}_j}} \eta_j \Big\} - \min_{\omega \models X_i \overline{B}_i} \Big\{ \kappa(\omega) + \sum_{\substack{j \neq i \\ \omega \models X_j \overline{B}_j}} \eta_j \Big\} \tag{3.1}$$

where the η_i are constraint variables taking values in $\mathbb{N}$.

For a constraint satisfaction problem CSP, the set of solutions is denoted by $\mathit{Sol}(\mathit{CSP})$. A solution of $\mathit{CR}(\kappa,\Delta)$ is an m-tuple $\vec{\eta} = (\eta_1,\ldots,\eta_m) \in \mathbb{N}^m$. Thus, with $\mathit{Sol}(\mathit{CR}(\kappa,\Delta))$ we denote the set of all solutions of $\mathit{CR}(\kappa,\Delta)$. For any $\vec{\eta} \in \mathbb{N}^m$, the induced c-revision κ^{c} as defined in (2.14) is denoted by $\kappa^{\mathrm{c}}_{\vec{\eta}}$.

Proposition 3.1.1 (Soundness and completeness of $\mathit{CR}(\kappa,\Delta)$)**.** *Let κ be an OCF and $\Delta = \{(X_1|B_1),\ldots,(X_m|B_m)\}$ be a set of conditionals.*

- *If $\vec{\eta} \in \mathit{Sol}(\mathit{CR}(\kappa,\Delta))$ then $\kappa^{\mathrm{c}}_{\vec{\eta}}$ is a c-revision of κ by Δ with $\kappa^{\mathrm{c}}_{\vec{\eta}} \models \Delta$.*
- *If κ^{c} is a c-revision of κ by Δ then there is a vector $\vec{\eta} \in \mathit{Sol}(\mathit{CR}(\kappa,\Delta))$ such that $\kappa^{\mathrm{c}} = \kappa^{\mathrm{c}}_{\vec{\eta}}$.*

The proof of Proposition 3.1.1 is a direct consequence of the propositions presented in [63]. Since c-revision and the impact factors defined by (3.1) provide a general schema for revision operators, many c-revisions are possible. Nevertheless, it is useful to impose further constraints on the parameters η_i to further improve the results of c-revisions. For instance, one option is to take Pareto-minimal η_i satisfying (3.1) ensuring that the resulting OCF ranks worlds as plausible as possible, cf. [7]. To impose further restrictions on the impacts η_i determining a c-revision, we employ selection strategies similar to the ones presented in [66, 9].

Definition 3.1.3 (Selection strategy σ, strategic c-revision $*_\sigma$)**.** *A* selection strategy (for c-revisions) *is a function*

$$\sigma : (\kappa, \Delta) \mapsto \vec{\eta}$$

assigning to each pair of an OCF κ and a (consistent) set of conditionals Δ an impact vector $\vec{\eta} \in Sol(CR(\kappa, \Delta))$. If $\sigma(\kappa, \Delta) = \vec{\eta}$, the c-revision of κ by Δ determined by σ *is $\kappa^{c}_{\vec{\eta}}$, denoted by $\kappa *_\sigma \Delta = \kappa^\sigma$, and $*_\sigma$ is a* strategic c-revision.

Note that a strategic c-revision operator $*_\sigma$ selects a single c-revision for each OCF κ and each Δ. We present two useful postulates for selection strategies. First, we consider selection strategies σ that choose Pareto-minimal impact vectors in the context of c-revisions.

(PM$^\sigma$) A selection strategy σ is *Pareto-minimal* if the impact vector $\sigma(\kappa, \Delta) = \vec{\eta}$ is Pareto-minimal among all impact vectors in $Sol(CR(\kappa, \Delta))$.

Furthermore, the following property expresses that including tautologies should not change the result of c-revision with a belief base Δ.

Tautological Vacuity (cTV) A selection strategy σ satisfies (cTV) if for any OCF κ and any conditional belief base Δ, the following holds:

$$\textbf{(cTV)} \qquad \kappa *_\sigma (\Delta \cup \{\top\}) = \kappa *_\sigma \Delta \tag{3.2}$$

It can easily be seen that for each c-revision we can select an impact vector s.t. (3.2) holds since, tautologies can never be falsified in (2.14).

Now, we turn to selection strategies for c-revisions that is helpful to ensure the coherency of impact factors across different revision scenarios, like the ones presented in the general revision problem (CondCS). In [5], the notion of *elementwise equivalence* for sets of conditionals Δ, Δ' is introduced, stating essentially that for every conditional in Δ there is a conditionally equivalent conditional in Δ', and vice versa. Using our general assumptions about sets of conditionals from page 39 in Section 2.4, we employ the following slight modification of elementwise equivalence to ensure a one-to-one correspondence between the conditionals in Δ and Δ'.

Definition 3.1.4 ($\equiv_e$)**.** *Two consistent sets of conditionals* $\Delta = \{(X_1|B_1), \ldots, (X_m|B_m)\}$ *and* $\Delta' = \{(X'_1|B'_1), \ldots, (X'_m|B'_m)\}$ *are* conditionally equivalent*, denoted by* $\Delta \equiv_e \Delta'$*, if the conditionals in* Δ *(resp. in* Δ'*) are pairwise not conditionally equivalent and* $(X_i|B_i) \equiv (X'_i|B'_i)$ *for all* $i \in \{1, \ldots, m\}$*.*

In the following, we state conditions for selection strategies that aim at ensuring basic properties that we expect from belief revision implemented by c-revisions [9]. The first condition is a postulate requiring selection strategies to be dependent only on the respective constraint satisfaction problem:

(IP-EP$^\sigma$) A selection strategy σ is *impact preserving with respect to equivalent problems* if for any two equivalent constraint satisfaction problems $CR(\kappa, \Delta) = CR(\kappa', \Delta')$, we have $\sigma(\kappa, \Delta) = \sigma(\kappa', \Delta')$.

In particular, this also implies that the selection strategy σ and the resulting strategic c-revision $*_\sigma$ are *syntax independent* [9]:

(SI$^\sigma$) A selection strategy σ is *syntax independent* if for any Δ' obtained from Δ by replacing a conditional $(X|B)$ occurring in Δ by a conditional $(X'|B')$ with $(X|B) \equiv (X'|B')$, we have $\sigma(\kappa, \Delta) = \sigma(\kappa, \Delta')$.

The following observation is a direct consequence of syntax independence.

Proposition 3.1.2. *Let* $*_\sigma$ *be a strategic c-revision satisfying (SI$^\sigma$) and let* Δ, Δ' *be sets of conditionals. If* $\Delta \equiv_e \Delta'$ *holds then* $\sigma(\kappa, \Delta) = \sigma(\kappa, \Delta')$ *and thus* $\kappa *_\sigma \Delta = \kappa *_\sigma \Delta'$ *for any ranking function* κ*.*

If $\vec{\eta}$ is an impact vector with impacts corresponding to the conditionals in Δ, then for $\Delta' \subseteq \Delta$, the subvector of $\vec{\eta}$ containing only the impacts related to the conditionals in Δ' is called the *projection* of $\vec{\eta}$ to Δ' and is denoted by $\vec{\eta}_{\Delta'}$. Hence, if σ is a selection strategy, $\sigma(\kappa, \Delta)_{\Delta'}$ is the projection of $\sigma(\kappa, \Delta)$ to Δ'. The following definition extends the notion of projection to CSPs for c-revisions [9].

Definition 3.1.5 (CSP projection $CR(\kappa, \Delta)_{\Delta'}$)**.** *Let* κ *be an OCF, let* $\Delta = \{(X_1|B_1), \ldots, (X_m|B_m)\}$ *be a set of conditionals, and let* $\Delta' \subseteq \Delta$*. The* projection of $CR(\kappa, \Delta)$ *to* Δ'*, denoted by* $CR(\kappa, \Delta)_{\Delta'}$*, is the constraint satisfaction problem given by the set of constraints* $\{cr(\kappa, \Delta)_i \mid (X_i|B_i) \in \Delta'\}$*.*

W.l.o.g. let us assume that $\Delta = \{(B_1|X_1), \ldots, (B_m|X_m)\}$ and $\Delta' = \{(B_1|X_1), \ldots, (B_k|X_k)\}$ where $k \leqslant m$. Note the difference between $CR(\kappa, \Delta') = \{cr(\kappa, \Delta')_1, \ldots, cr(\kappa, \Delta')_k\}$ on the one hand and $CR(\kappa, \Delta)_{\Delta'} = \{cr(\kappa, \Delta)_1, \ldots, cr(\kappa, \Delta)_k\}$ on the other hand. While $CR(\kappa, \Delta')$ is a CSP over the constraint variables $\eta_1, \ldots, \eta_k$, $CR(\kappa, \Delta)_{\Delta'}$ is a CSP over the constraint variables $\eta_1, \ldots, \eta_k, \eta_{k+1}, \ldots, \eta_m$. Both CSP have $|\Delta'|$-many constraints, but in contrast to $cr(\kappa, \Delta)_i$, any of the constraint variables $\eta_{k+1}, \ldots, \eta_m$ in the sum in the minimizations terms given in Equation (3.1) do not occur in $cr(\kappa, \Delta')_i$.

Using projections of constraint satisfaction problems for c-revisions, we can generalize the idea of preserving impacts, as expressed by (IP-EP$^\sigma$) and (SI$^\sigma$) to equivalent subproblems. This is specified in the next axiom (cf. [111, 9]), which has far-reaching consequences.

(IP-ESP$^\sigma$) A selection strategy σ is *impact preserving with respect to equivalent subproblems* if for any two revision problems $(\kappa, \Delta), (\kappa', \Delta')$ with $\Delta_1 \subseteq \Delta, \Delta'_1 \subseteq \Delta'$ and $\Delta_1 \equiv_e \Delta'_1$ such that $CR(\kappa, \Delta)_{\Delta_1} = CR(\kappa', \Delta')_{\Delta'_1}$, we have $\sigma(\kappa, \Delta)_{\Delta_1} = \sigma(\kappa', \Delta')_{\Delta'_1}$.

As a first immediate consequence, it holds that if σ satisfies (IP-ESP$^\sigma$), then it is syntax independent.

Proposition 3.1.3 ([9]). *If a selection strategy σ satisfies (IP-ESP$^\sigma$), then σ also satisfies the axiom (SI$^\sigma$).*

It can easily be seen that (IP-ESP$^\sigma$) ensures tautological vacuity (cTV), i.e., the axiom for revision, which expresses that including tautologies should not change the revision result because tautologies can never be falsified in (3.1).

Proposition 3.1.4. *If a selection strategy σ satisfies (IP-ESP$^\sigma$), then the strategic c-revision $*_\sigma$ satisfies (cTV).*

The strategical axiom (IP-ESP$^\sigma$) will prove very helpful to substantiate the Kinematics principle with a proof of concept in the next section, and we later present an example illustrating its application for c-revisions.

3.2 Related Work

In this section, we place the Kinematics principle in a more extensive research context and discuss related work.

Probabilistic reasoning, a research branch in KRR, deals with employing probability theory and statistical inference to make predictions or decisions in situations involving uncertainty or (possibly incomplete) evidence [88]. To be emphasized in this context are two main methods of probabilistic reasoning, namely *Jeffrey's rule* [57] and *Pearl's method of virtual evidence* [88], which he proposed in the context of Bayesian networks. In [25], the authors show that for both of these methods, the axiom of *Probability Kinematics*, firstly introduced in [57], is crucial, which assumes that if we revise a probability distribution P by some uncertain evidence $\mathcal{A} = \{A_1[x_1], \ldots, A_n[x_n]\}$, s.t. $P^* = P * \mathcal{A}$, bearing on a set of exclusive and exhaustive cases $A_1, \ldots, A_n$, then the conditional probabilities given these cases should not change:

$$P^*|A_i\,(\omega) = P|A_i\,(\omega) \text{ for all } \omega \in \Omega, \tag{3.3}$$

i.e., Probability Kinematics captures the notion that even though the probabilities of the cases $A_1, \ldots, A_n$ change, their corresponding conditional probabilities w.r.t. the remaining worlds in Ω do not. This notion of invariance of probabilities under conditionalization in the light of new evidence gave the Kinematics principle we deal with in this part of the thesis its name. A more in-depth discussion about its justification can be found in [122] and [114]. For a more thorough discussion of Jeffrey's rule and Pearl's virtual evidence and how they relate to each other in the context of probabilistic Belief Revision, i.e., Belief Revision in the context of probability distributions, we refer to [25]. There have been several proposals to generalize Jeffrey's rule. For example, Wagner [123] uses an arbitrary set of propositions A_i. Smets generalized Jeffrey's rule to belief functions (see [115]), and Benferhat et al. analyzed the expressive power of possibilistic counterparts to Jeffrey's rule for modeling belief revision [10].

In this thesis, one of the most important extensions of Probability Kinematics in probabilistic Belief Revision was introduced by Shore and Johnson [113] as *Subset*

Independence. We recapture the definition from [113].

Definition 3.2.1 ([113])**.** *Let P be a probability distribution, $A_1, \ldots, A_n$ be exhaustive and exclusive formulas, P be a probability distribution and $\mathcal{R} = \mathcal{R}_1 \cup \ldots \cup \mathcal{R}_n$ be a set of probabilistic conditionals, with subsets $\mathcal{R}_i$ containing conditionals whose premises imply A_i, and $\mathcal{S} = \{A_1[x_1], \ldots, A_n[x_n]\}$ with $\sum_{i=1}^{n} x_i = 1$. The revision operator $*$ satisfies* Subset Independence *iff*

$$(P * (\mathcal{R} \cup \mathcal{S}))|A_i \;=\; P|A_i * \mathcal{R}_i. \tag{3.4}$$

As Probability Kinematics, Subset Independence deals with conditional probabilities in the light of exclusive cases induced by formulas $A_1, \ldots, A_n$. Note that Probability Kinematics follows immediately from Subset Independence if we take Δ to be the empty set [113]. Yet, Subset Independence broadens the scope of the Kinematics principle and Jeffrey's rule to revisions with a set of probabilistic facts $\mathcal{S}$, which represents new information about the probabilities of the exclusive cases, and a set of probabilistic conditionals $\mathcal{R}$, which decomposes naturally into disjoint subsets $\mathcal{R}_i$ corresponding to the cases A_i. In this setting, Subset Independence states that for the conditional probabilities in terms of each case A_i, it should not matter whether we revise with the entirety of the new information, i.e., $\mathcal{R} \cup \mathcal{S}$, or whether we revise with the separate probabilistic conditionals concerning only this case. Thus, introducing a notion of locality for revision, since only worlds talking about the corresponding case are relevant for the revision with (conditional) information that concerns precisely this case. Subset Independence is the blueprint for our Kinematics principle in the following investigations.

Chapter 4

Case Splitting for the Kinematics Principle

The advanced belief revision problem (CondCS) presented in the introduction of this part takes place in a particular setting with a set of conditional $\Delta = \Delta_1 \cup \ldots \cup \Delta_n$ where the partition $\Delta_1 \cup \ldots \cup \Delta_n$ is induced by exclusive and exhaustive formulas $A_1, \ldots, A_n$, s.t. the premises of each conditional in Δ_i imply the corresponding formula A_i. The exclusivity and exhaustiveness of the cases A_i lay the foundation for the advantages of the Kinematics principle by introducing distinguishable cases for the revision. This chapter investigates those cases, and Section 4.1 introduces the notion of a *case splitting*. Then, in Section 4.2, we introduce a refinement and specificity property for sets of cases, enabling us to define a unique (up to equivalences) finest splitting and an algorithm that computes it.

Bibliographic Remark. The contents of this part are based on joint work with Gabriele Kern-Isberner and Christoph Beierle [107, 111] (see Section 1.3).

4.1 Case Splitting

We start our investigation with the following definition of *case splittings*, i.e., sets of formulas that satisfy the prerequisites of (CondCS).

Definition 4.1.1 (Case splitting)**.** *Let* $\Delta = \{(B_1|X_1), \ldots, (B_m|X_m)\}$ *be a set of conditionals. A* case splitting $\mathcal{P}_\Delta$ of Δ *is a set* $\mathcal{P}_\Delta = \{A_1, \ldots, A_n\}$ *of exclusive and exhaustive formulas* $A_1, \ldots A_n$*, such that every premise* X_k *(*$k = 1, \ldots, m$*) implies exactly one* A_i*,* $X_k \models A_i$*.*

Each case splitting $\mathcal{P}_\Delta = \{A_1, \ldots, A_n\}$ induces a partitioning of $\Delta = \{(B_1|X_1), \ldots, (B_m|X_m)\}$ as follows:

$$(B_j|X_j) \in \Delta_i \text{ iff } X_j \models A_i \text{ for } j \in \{1, \ldots, m\} \text{ and } i \in \{1, \ldots, n\}.$$

The subsets Δ_i are disjoint since the A_i's are exclusive.

We illustrate case splittings and the corresponding partition of sets of conditionals Δ via the following example.

Example 4.1.1. *Let the signature* $\Sigma = \{a, b, c, d\}$ *and the set of conditionals* $\Delta = \{(c|ab),\ (d|a\overline{b}c), (d|a\overline{b}), (b|\overline{a}), (d|\overline{a}c)\}$ *be given. For* Δ*, it holds that* $\mathcal{P}_\Delta = \{a, \overline{a}\}$ *is a case splitting. The formulas* a *and* $\overline{a}$ *are exclusive and exhaustive, and it holds that for each conditional in* Δ*, its premise implies either* a *or* $\overline{a}$*. the case splitting* $\mathcal{P}_\Delta$ *induces the following partition of* $\Delta = \{(c|ab), (d|a\overline{b}c), (d|a\overline{b})\} \cup \{(b|\overline{a}), (d|\overline{a}c)\}$*.*

Note that the premises in the conditional belief base Δ have to imply one of the cases A_i and not be equivalent to it. Also, the subsets $\Delta_i \subseteq \Delta$ are not necessarily non-empty. For the special case $\Delta = \{(B|A)\}$, $\mathcal{P}_\Delta = \{A, \overline{A}\}$ is a case splitting, s.t. the case $\overline{A}$ induces the empty subset as part of the partition of Δ.

It is clear that for each set Δ, a case splitting exists, at least the trivial one $\mathcal{P}_\Delta = \{\top\}$, and thus, case splittings are not unique. In the following, we investigate different case splittings via a relation that compares case splittings.

4.2 Algorithm for the Finest Case Splitting

This section defines a *refinement* relation to delimit different case splittings. Employing this relation, we present an algorithm that computes the finest case splitting that is unique up to equivalences.

Example 4.1.1 from the previous section illustrates that various case splittings $\mathcal{P}_\Delta$ for the same set of conditionals Δ exist. However, it is evident that the more fine-grained $\mathcal{P}_\Delta$ is, the more fine-grained the partition of Δ is, which maximizes the effect of the conditionalization in our later presented Kinematics principle. We introduce the following refinement-relation to compare different case splitting on Δ:

Definition 4.2.1 (Refinement and specificity)**.** *Let Δ be a set of conditionals. For two case splittings $\mathcal{P}^1_\Delta = \{A_1, \dots, A_n\}$ and $\mathcal{P}^2_\Delta = \{B_1, \dots, B_{n'}\}$ of Δ, we say that $\mathcal{P}^1_\Delta$ is a* refinement *of $\mathcal{P}^2_\Delta$ iff every B_j is implied by some A_i:*

$$\mathcal{P}^1_\Delta \leqslant \mathcal{P}^2_\Delta \textit{ iff for all } B_j \in \mathcal{P}^2_\Delta, \textit{ there exists } A_i \in \mathcal{P}^1_\Delta \textit{ s.t. } A_i \models B_j.$$

We say that the A_i's are more specific *than the B_j's. Two case splittings $\mathcal{P}^1_\Delta$ and $\mathcal{P}^2_\Delta$ of Δ are equivalent iff $\mathcal{P}^1_\Delta \leqslant \mathcal{P}^2_\Delta$ and $\mathcal{P}^2_\Delta \leqslant \mathcal{P}^1_\Delta$*

For the Kinematics principle, the most interesting case splitting of Δ is the one that refines every other splitting. We call this splitting the *finest case splitting.* In the following theorem, we show that for every set Δ, there exists the finest case splitting, which is unique up to semantic equivalences.

Theorem 4.2.1. *For each set of conditional beliefs Δ there exists a unique finest case splitting (up to semantic equivalences and permutations).*

Proof. We start by defining a relation $\sim$ on $Prem(\Delta) = \{X \in \mathcal{L} | (B|X) \in \Delta\}$ the set of premises of Δ, s.t.

$$X \sim Y \text{ iff } XY \not\equiv \bot \text{ for } X, Y \in Prem(\Delta) \tag{4.1}$$

The relation $\sim$ is reflexive and symmetric, thus the transitive closure $\sim^*$ of $\sim$ is an equivalence relation on the elements of $Prem(\Delta)$ and it holds that $Prem(\Delta) = \bigcup_{i=1,\dots,n}[X_i]$, where $[X_i]$ are the equivalence classes of $\sim^*$. Let $A_i = \bigvee [X_i]$ for $i = 1, \dots, n$ and $A_0 = \neg(A_1 \vee \dots \vee A_n) \equiv \overline{A_1} \dots \overline{A_n}$, then $\mathcal{P}_\Delta = \{A_0, A_1, \dots, A_n\}$ defines a case splitting because for $i \neq j$ it holds that:

$$A_i A_j \equiv \big(\bigvee_{\tilde{X} \in [X_i], \tilde{Y} \in [X_j]} (\tilde{X}\tilde{Y}) \big) \equiv \bot,$$

and $A_0 \vee A_1 \vee \dots \vee A_n \equiv (\overline{A_1} \dots \overline{A_n}) \vee A_1 \dots \vee A_n \equiv \top$. It is clear that for every

Algorithm 1 Finest case splitting

Require: Finite set of conditionals Δ
Ensure: Unique finest case splitting of Δ

$Prem \leftarrow Prem(\Delta)$
$\mathcal{P}_\Delta = \emptyset$
while $Prem \neq \emptyset$ **do**
  Choose $X \in Prem$
  if there are $Y \in Prem$, $Y \neq X$, with $XY \not\equiv \bot$ **then**
    build $\chi = \{Y \in Prem | XY \not\equiv \bot\}$
    $A \leftarrow \bigvee \chi$
    $Prem \leftarrow (Prem \setminus \chi) \cup \{A\}$
  else
    $A \leftarrow X$
    $\mathcal{P}_\Delta \leftarrow \mathcal{P}_\Delta \cup \{A\}$
    $Prem \leftarrow Prem \setminus \{A\}$
  end if
end while
if $\bigvee \mathcal{P}_\Delta \not\equiv \top$ **then**
  $A_0 = \bigwedge_{A_i \in \mathcal{P}_\Delta} \overline{A_i}$
  $\mathcal{P}_\Delta = \mathcal{P}_\Delta \cup \{A_0\}$
end if
return $\mathcal{P}_\Delta$

premise $X \in Prem(\Delta) = \bigcup_{i=1,\dots,n}[X_i]$, there is one $A_i = \bigvee[X_i]$ which is implied by X, and that $\mathcal{P}_\Delta$ is unique up to permutation and semantic equivalences.

We still need to show that $\mathcal{P}_\Delta$ refines every other case splitting: Let $\mathcal{P}'_\Delta = \{B_1, \dots, B_{n'}\}$ be another case splitting of Δ. For $X, Y \in Prem(\Delta)$ it holds that, if $X \models B_i$ and $Y \models B_j$ with $i \neq j$, then $XY \equiv \bot$. This means that for $X, Y \in Prem(\Delta)$ with $XY \not\equiv \bot$, there is $i \in \{1, \dots, n'\}$ with $X \models B_i$ and $Y \models B_i$, which means $[X] \models B_i$. Hence, there exists $j \in 1, \dots, n$ with $\bigvee_{\tilde{X} \in [X_j]} \tilde{X} = A_j \models B_i$ and $\mathcal{P}_\Delta$ refines $\mathcal{P}'_\Delta$. □

The relation $\sim$ defined via (4.1) is used in the proof of Theorem 4.2.1 to define exclusive cases for $\mathcal{P}_\Delta$. We use this relation and present an algorithm that computes the finest case splitting for an arbitrary finite set of conditionals Δ: From the constructive proof of Theorem 4.2.1, we can conclude the following theorem:

Theorem 4.2.2. *Algorithm 1 terminates and is correct in the sense that it computes the unique finest case splitting for a finite set of conditionals Δ.*

It holds that the transitive closure of the relation $\sim$ is obtained by considering the disjunction over χ in line 7 and adding these to the set of premises $Prem$ in line 8. Then, for premises X_1, X_2, X_3 with $X_1X_2 \not\equiv \bot$ and $X_1X_3 \not\equiv \bot$, it might be the case that $X_2X_3 \equiv \bot$, but $(X_1 \vee X_2)X_3 \not\equiv \bot$ and therefore it holds that the A_i's in $\mathcal{P}_\Delta$ are exclusive. In this way, we ensure that the entire equivalence class is captured.

The running time of Algorithm 1 is determined by the SAT-Test in line 5 and 6. And in the worst case, it holds that the equivalence classes determined in the while-loop are singletons, and we obtain $\mathcal{O}(s^2)$, where s represents the runtime of the SAT-Test.

We continue Example 4.1.1 and use Algorithm 1 to determine the finest premise splitting for the conditional belief set Δ.

Example 4.2.1 (Continuing Example 4.1.1)**.** *We compute the finest premise splitting for $\Delta = \{(c|ab), (d|a\overline{b}c), (d|a\overline{b}), (b|\overline{a}), (d|\overline{a}c)\}$ from Example 4.1.1 using Algorithm 1.*
Initialize the set of premises $Prem$ of Δ and the case splitting $\mathcal{P}_\Delta$ as in line 1 and 2 from Algorithm 1

$$Prem = \{ab, a\overline{b}c, a\overline{b}, \overline{a}, \overline{a}c\} \text{ and } \mathcal{P}_\Delta = \emptyset.$$

For the first iteration, we choose $X = ab$ and line line 6 from Algorithm 1 yields $\chi = \{ab\}$. Thus, we add $A_1 = ab$ to $\mathcal{P}_\Delta = \{ab\}$. This leaves us with $Prem = \{a\overline{b}c, a\overline{b}, \overline{a}, \overline{a}c\}$ for the second iteration.
Now, we choose $X = a\overline{b}c$ with the corresponding set $\chi = \{a\overline{b}c, a\overline{b}\}$. Hence $A_2 = a\overline{b}c \vee a\overline{b} = a\overline{b}$ and therefore, we remove $a\overline{b}c$ from the set of premises and get $Prem = \{a\overline{b}, \overline{a}, \overline{a}e\}$.
In the next iteration, we take $X = a\overline{b}$. It holds that $XY \equiv \bot$ for all other $Y \in Prem$, therefore we jump to line 9 in Algorithm 1 and the case splitting does not change, s.t. $\mathcal{P}_\Delta = \{ab, a\overline{b}\}$ holds.
We continue with $X = \overline{a}$ and yield $\chi = \{\overline{a}, \overline{a}e\}$. So that, we add $A_3 = \overline{a}$ to the

case splitting and continue with $Prem = \{\overline{a}\}$*. For the last iteration, it holds that* $X = \overline{a} = A_3$*. Thus, we get the following case splitting* $\mathcal{P}_\Delta = \{ab, a\overline{b}, \overline{a}\}$ *and the set of premises* $Prem = \emptyset$*, i.e., we leave the while loop in Algorithm 1.*

It holds that $ab \vee a\overline{b} \vee \overline{a} \equiv \top$*, thus the algorithm terminates and returns the case splitting* $\mathcal{P}_\Delta = \{ab, a\overline{b}, \overline{a}\}$ *which determines a partitioning of* $\Delta = \{(c|ab)\} \cup \{(d|a\overline{b}c), (d|a\overline{b})\} \cup \{(b|\overline{a}), (d|\overline{a}c)\} = \Delta_1 \dot{\cup} \Delta_2 \dot{\cup} \Delta_3$*. This case splitting displays a unique finest case splitting of* Δ *according to Theorem 4.2.2. In particular, it holds that the splitting* $\mathcal{P}_\Delta = \{ab, a\overline{b}, \overline{a}\}$ *refines the splitting given in Example 4.1.1.*

Chapter 5

The Kinematics Principle for Ranking Functions

In the introduction of this part, we discussed the advanced belief revision problem (CondCS), which now leads us to the Kinematics principle that is our current primary subject of investigation. In this chapter, we present the Kinematics principle for ranking functions, which transfers the ideas of revision w.r.t. (disjoint) contextual information from probability theory to the framework of belief revision using OCFs. Based on the notion of locality introduced by Subset Independence via exclusive cases in [113], we define the *Kinematics Principle for ranking functions*, which guides the revision with respect to a case splitting of Δ. Thus, the conditional information decomposes naturally into disjoint contexts. We tackle the second set of research questions in the introduction during our investigation and start by presenting the Kinematics principle for OCFs, as a rational solution to the advanced belief revision problem (CondCS) in Section 5.1, and show that c-revisions provide a proof of concept for it. In the following Section 5.2, we investigate the implications of the Kinematics principle for propositional revision in cases where the exclusive cases A_i in the cases splitting of Δ fully capture the context of conditionals in Δ_i, i.e., when the premises of the conditionals in Δ_i are equivalent to A_i, instead of solely implying it. This restriction enables an elegant implementation of conditional revisions via conditionalization and propositional revision. In the last section 5.3, we show that, in the context of ranking functions, the Kinematics principle unfolds its total capacity and enables us to compose the full c-revised OCF $\kappa * \Delta$ from the local revisions $\kappa|A_i * \Delta_i$ reducing the complexity of the revision task dramatically depending on the size of the case splitting.

Bibliographic Remark. The contents of this part are based on joint work with Gabriele Kern-Isberner and Christoph Beierle [107, 111] (see Section 1.3).

5.1 Generalized Ranking Kinematics

In this section, we propose a definition of the Kinematics principle for OCFs and provide a proof of concept for this advanced belief revision axiom. Via the application of strategies, we can provide coherence across different revision scenarios considered in the Kinematics principle, thus leading to an elegant theorem that shows that strategic c-revisions are capable of dealing with (disjoint) contextual information as proposed by the Kinematics principle.

The *Kinematics principle for ranking functions* was initially introduced in [107] as *Generalized Ranking Kinematics* (GRK), therefore we sometimes abbreviate the Kinematics principle for OCFs with (GRK). The axiom (GRK) is defined for revision operators $*$ for OCFs taking conditional belief bases and additional sentences as input and yielding a revised OCF as output. Note that (GRK) makes use of conditionalization for OCFs as defined by Spohn in [119] (cf. (2.10) on page 37).

Definition 5.1.1 (Kinematics principle for OCFs, Generalized Ranking Kinematics (GRK))**.** *Let $A_1, \ldots, A_n$ be exhaustive and exclusive formulas. Let κ be a ranking function, and let $\Delta = \Delta_1 \cup \ldots \cup \Delta_n$ be a set of conditionals, with subsets Δ_i whose premises imply A_i, and $S = \bigvee_{j \in J} A_j$ with $\emptyset \neq J \subseteq \{1, \ldots, n\}$. The revision operator $*$ satisfies* Generalized Ranking Kinematics *iff*

$$\textbf{(GRK)} \qquad \kappa * (\Delta \cup \{S\})|A_i = (\kappa|A_i) * \Delta_i$$

The Kinematics principle offers a rational strategy for revising OCFs w.r.t. to information that is relevant solely in local (disjoint) contexts provided by cases A_i. Based on Subset Independence (3.4) two strong irrelevance assertions are introduced by (GRK): First, if the revised OCF $\kappa*(\Delta\cup\{S\})$ is conditionalized by a case A_i, then only the conditionals talking about this case, i.e., Δ_i, are relevant. The advantage of revision with respect to cases, i.e., local revision, is that it enables us to revise a more specific part of the OCF, i.e., $\kappa|A_i$, with more specific information, i.e., Δ_i, which are determined by the case A_i. The revision task tends to be less complex because we refrain from revising the full OCF with the full set Δ. Note that this notion of locality via contextual information is much closer to how we process information as humans, namely via concerning information w.r.t. to the relevant context [85, 58]. If we consider the particular case, where $S \equiv \top$ and presuppose that $*$ satisfies the vacuity postulate (cTV)[1], then it becomes apparent that the Kinematics principle for OCFs implements the idea that conditionalization and revision with contextual information, represented by subsets Δ_i, are commutable:

$$\kappa * \Delta|A_i \;=\; \kappa|A_i * \Delta_i. \tag{5.1}$$

Thus, it does not matter whether the agent first focuses on a specific case, defined by A_i, and then revises her beliefs or whether she first revises with the whole set of conditional information and then focuses on the case. Note that, in [107], (GRK) is introduced as *strong Generalized Ranking Kinematics*, which is accompanied by a weaker version called *weak Generalized Ranking Kinematics* ($\mathrm{GRK}^{\mathrm{weak}}$) that corresponds to the special case of $S \equiv \top$ stated above.

Second, in the context of revision, it is irrelevant for the local OCFs $\kappa|A_i$ whether a case A_i is more plausible than others. In the probabilistic case, we have $\mathcal{S}$, which represents the set of posterior probabilities for formulas A_i. Here, we consider a disjunction of the cases A_i, which indicates that one might be more plausible than the remaining cases. This second notion of irrelevance becomes evident for the special case of (GRK) with $\Delta = \emptyset$:

$$\kappa * \{S\}|A_i \;=\; \kappa|A_i\,.$$

As we can see, this corresponds to the notion of Probability Kinematics (3.3) in the context of semi-quantitative belief revision.

[1]Note that, in (3.2) on page 64 is defined for strategic c-revisions, to shorten the matter we do not rephrase it here with a general semi-quantitative revision operator $*$.

The Kinematics principle for ranking functions (GRK) is satisfied by c-revisions as the following theorem from [111] shows. Note that a similar theorem was presented in [107], but without employing strategies for c-revision, i.e., the theorem in [107] is much less clear and easy to understand.

Theorem 5.1.1. *Let $\Delta = \Delta_1 \cup \ldots \cup \Delta_n$ be a set of conditionals, with subsets $\Delta_i = \{(B_{i,j}|A_i C_{i,j})\}_{j=1,\ldots,n_i}$ with $n_i = |\Delta_i|$ for $i = 1, \ldots, n$ and a case splitting $\mathcal{P}_\Delta = \{A_1, \ldots, A_n\}$. Let $S = \bigvee_{j \in J} A_j$ and $\emptyset \neq J \subseteq \{1, \ldots, n\}$. If σ is a selection strategy that satisfies (IP-ESP$^\sigma$), then $*_\sigma$ is a strategic c-revision that satisfies (GRK).*

Proof. We first investigate the constraint satisfaction problems given in (GRK) for c-revisions, and then the general definition of the c-revision as in Definition 2.5.1. For all $\omega \in \Omega$, $\omega \models A_i$ holds for exactly one i. If $\omega \models A_i$ then all conditionals from Δ_k $(k \neq i)$, are not applicable and hence irrelevant. It holds for $\omega \models A_i$ that:

$$\kappa *_\sigma (\Delta \cup \{S\})(\omega) = \kappa_{0,\Delta} + \kappa(\omega) + \sum_{\substack{1 \leqslant j \leqslant n_i \\ \omega \models A_i C_{i,j} \overline{B}_{i,j}}} \eta_{i,j} + \begin{cases} \eta_S & \omega \not\models S \\ 0 & \text{otherwise} \end{cases} \tag{5.2}$$

with $\kappa_{0,\Delta}$ the corresponding normalization constant according to (2.15). $CR(\kappa, \Delta \cup \{S\})$ is defined by the following set of constraints for $\eta_{i,j}$, $i \in \{1, \ldots, n\}$ and $j \in \{1, \ldots, n_i\}$ and η_S:

$$\begin{aligned}
\eta_{i,j} > & \min_{\omega \models A_i C_{i,j} B_{i,j}} \{\kappa(\omega) + \sum_{1 \leqslant i \leqslant n} \sum_{\substack{l \neq j \\ \omega \models A_i C_{i,l} \overline{B}_{i,l}}} \eta_{i,l} + \begin{cases} \eta_S & \omega \not\models S \\ 0 & \text{otherwise} \end{cases} \} \\
& - \min_{\omega \models A_i C_{i,j} \overline{B}_{i,j}} \{\kappa(\omega) + \sum_{1 \leqslant i \leqslant n} \sum_{\substack{l \neq j \\ \omega \models A_i C_{i,l} \overline{B}_{i,l}}} \eta_{i,l} + \begin{cases} \eta_S & \omega \not\models S \\ 0 & \text{otherwise} \end{cases} \} \\
& = \min_{\omega \models A_i C_{i,j} B_{i,j}} \{\kappa(\omega) + \sum_{\substack{l \neq j \\ \omega \models A_i C_{i,l} \overline{B}_{i,l}}} \eta_{i,l}\} - \min_{\omega \models A_i C_{i,j} \overline{B}_{i,j}} \{\kappa(\omega) + \sum_{\substack{l \neq j \\ \omega \models A_i C_{i,l} \overline{B}_{i,l}}} \eta_{i,l}\}.
\end{aligned} \tag{5.3}$$

The impact factor for S η_S occurs only if $i \in J$ and either is cancelled out, or does not appear at all. In general, it holds for the impact factor η_S that:

$$\eta_S > \min_{\omega \models S}\{\kappa(\omega) + \sum_{j \in J} \sum_{\substack{1 \leqslant l \leqslant n_j \\ \omega \models A_j C_{j,l} \overline{B}_{j,l}}} \eta_{j,l}\} - \min_{\omega \models \overline{S}}\{\kappa(\omega) + \sum_{k \notin J} \sum_{\substack{1 \leqslant l \leqslant n_k \\ \omega \models A_k C_{k,l} \overline{B}_{k,l}}} \eta_{k,l}\}.$$

Conditioning on A_i yields for worlds $\omega \models A_i$

$$\begin{aligned}
&(\kappa *_\sigma (\Delta \cup \{S\}))|A_i\,(\omega) = \kappa *_\sigma (\Delta \cup \{S\})(\omega) - \kappa *_\sigma (\Delta \cup \{S\})(A_i) \\
&= \kappa_{0,\Delta} + \kappa(\omega) + \sum_{\substack{1 \leqslant j \leqslant n_i \\ \omega \models A_i C_{i,j} \overline{B}_{i,j}}} \eta_{i,j} + \begin{cases} \eta_S & \omega \not\models S \\ 0 & \text{otherwise} \end{cases} \\
&- \min_{\tilde{\omega} \models A_i}\{\kappa_{0,\Delta} + \kappa(\tilde{\omega}) + \sum_{\substack{1 \leqslant j \leqslant n_i \\ \tilde{\omega} \models A_i C_{i,j} \overline{B}_{i,j}}} \eta_{i,j} + \begin{cases} \eta_S & \omega \not\models S \\ 0 & \text{otherwise} \end{cases}\} \qquad (5.4) \\
&= \kappa(\omega) + \sum_{\substack{1 \leqslant j \leqslant n_i \\ \omega \models A_i C_{i,j} \overline{B}_{i,j}}} \eta_{i,j} \underbrace{- \min_{\tilde{\omega} \models A_i}\{\kappa(\tilde{\omega}) + \sum_{\substack{1 \leqslant j \leqslant n_i \\ \tilde{\omega} \models A_i C_{i,j} \overline{B}_{i,j}}} \eta_{i,j}\}}_{(*)}. \qquad (5.5)
\end{aligned}$$

Now, we turn to the case when we first conditionalize κ and then c-revise the conditionalized OCF by Δ_i. For $\omega \models A_i$, we obtain

$$\begin{aligned}
\kappa|A_i *_\sigma \Delta_i(\omega) &= \kappa_{0,i} + \kappa|A_i\,(\omega) + \sum_{\substack{1 \leqslant j \leqslant n_i \\ \omega \models A_i C_{i,j} \overline{B}_{i,j}}} \mu_{i,j} \\
&= \kappa_{0,i} + \kappa(\omega) - \kappa(A_i) + \sum_{\substack{1 \leqslant j \leqslant n_i \\ \omega \models A_i C_{i,j} \overline{B}_{i,j}}} \mu_{i,j} \\
&= \kappa(\omega) + \sum_{\substack{1 \leqslant j \leqslant n_i \\ \omega \models A_i C_{i,j} \overline{B}_{i,j}}} \mu_{i,j} + \underbrace{\kappa_{0,i} - \kappa(A_i)}_{(**)} \qquad (5.6)
\end{aligned}$$

with $\kappa_{0,i}$ the corresponding normalization constant for c-revisions of $\kappa|A_i$ with Δ_i. The constraint satisfaction problem $CR(\kappa|A_i, \Delta_i)$ is given by the following inequalities for $j = \{1, \ldots, n_i\}$:

$$\mu_{i,j} > \min_{\omega \models A_i C_{i,j} B_{i,j}} \{ \kappa|A_i\,(\omega) + \sum_{\substack{l \neq j \\ \omega \models A_i C_{i,l} \overline{B}_{l,i}}} \mu_{i,l} \} - \min_{\omega \models A_i C_{i,j} \overline{B}_{i,j}} \{ \kappa|A_i\,(\omega) + \sum_{\substack{l \neq j \\ \omega \models A_i C_{i,l} \overline{B}_{i,l}}} \mu_{i,l} \} \quad (5.7)$$

$$= \min_{\omega \models A_i C_{i,j} B_{i,j}} \{\kappa(\omega) + \sum_{\substack{l \neq j \\ \omega \models A_i C_{i,l} \overline{B}_{i,l}}} \mu_{i,l}\} - \kappa(A_i) - \min_{\omega \models A_i C_{i,j} \overline{B}_{i,j}} \{\kappa(\omega) + \sum_{\substack{l \neq j \\ \omega \models A_i C_{i,l} \overline{B}_{i,l}}} \mu_{i,l}\} + \kappa(A_i)$$

$$= \min_{\omega \models A_i C_{i,j} B_{i,j}} \{\kappa(\omega) + \sum_{\substack{l \neq j \\ \omega \models A_i C_{i,l} \overline{B}_{i,l}}} \mu_{l,i}\} - \min_{\omega \models A_i C_{i,j} \overline{B}_{i,j}} \{\kappa(\omega) + \sum_{\substack{l \neq j \\ \omega \models A_i C_{i,l} \overline{B}_{i,l}}} \mu_{i,l}\}. \quad (5.8)$$

In order to separate the two revisions from each other prima facie, we used impact factors μ for $\kappa|A_i *_\sigma \Delta_i$ instead of η as in $\kappa *_\sigma (\Delta \cup \{S\})$.

As we can see, for a fixed $i \in \{1, \ldots, n\}$, the projection of $CR(\kappa, \Delta \cup \{S\})$ to Δ_i, i.e., (5.3), and $CR(\kappa|A_i\,, \Delta_i)$, i.e., (5.8), are defined by the same inequalities. Thus, it holds that $CR(\kappa, \Delta \cup \{S\})_{\Delta_i} = CR(\kappa|A_i\,, \Delta_i)$. Therefore by (IP-ESP$^\sigma$), we get that $\vec{\eta}_i = \sigma(\kappa, \Delta \cup \{S\})_{\Delta_i} = \sigma(\kappa|A_i\,, \Delta_i) = \vec{\mu}_i$. To complete the proof, we still need to show that $(*) = (**)$ holds, i.e.,

$$- \min_{\omega \models A_i} \{\kappa(\omega) + \sum_{\substack{1 \leqslant j \leqslant n_i \\ \omega \models A_i C_{i,j} \overline{B}_{i,j}}} \eta_{i,j}\} \stackrel{!}{=} \kappa_{0,i} - \kappa(A_i). \quad (5.9)$$

We use (2.15) and $\vec{\eta}_i = \vec{\mu}_i$ to compute $\kappa_{0,i}$ and get

$$\kappa_{0,i} - \kappa(A_i) = - \min_{\omega \models A_i} \{\kappa(\omega|A_i) + \sum_{\substack{1 \leqslant j \leqslant n_i \\ \omega \models A_i C_{i,j} \overline{B}_{i,j}}} \mu_{i,j}\} - \kappa(A_i)$$

$$= - \min_{\omega \models A_i} \{\kappa(\omega) - \kappa(A_i) + \sum_{\substack{1 \leqslant j \leqslant n_i \\ \omega \models A_i C_{i,j} \overline{B}_{i,j}}} \eta_{i,j}\} - \kappa(A_i) = - \min_{\omega \models A_i} \{\kappa(\omega) + \sum_{\substack{1 \leqslant j \leqslant n_i \\ \omega \models A_i C_{i,j} \overline{B}_{i,j}}} \eta_{i,j}\}. \quad (5.10)$$

□

The proof of Theorem 5.1.1 provides some crucial insights about the CSP defining the corresponding c-revisions in (GRK) and also about the specific role of the normalization constants, which we summarize in the following proposition:

Proposition 5.1.2. *Presuppose the same prerequisites as in Theorem 5.1.1. We consider the strategic c-revisions* $(\kappa *_\sigma (\Delta \cup \{S\}))|A_i$ *and* $\kappa|A_i *_\sigma \Delta_i$ *with a selection strategy* σ *which satisfies (IP-ESP*$^\sigma$*), that are crucial for GRK. Then the following statements hold:*

1. *It holds that* $CR(\kappa, \Delta \cup \{S\})_{\Delta_i} = CR(\,\kappa|A_i\,, \Delta_i)$ *and thus, due to (IP-ESP*$^\sigma$*) we get that*

$$\sigma(\kappa, \Delta \cup \{S\})_{\Delta_i} = \sigma(\,\kappa|A_i\,, \Delta_i) \tag{5.11}$$

2. *It holds that*

$$\begin{aligned} &(\kappa *_\sigma (\Delta \cup \{S\}))|A_i\,(\omega) -\, \kappa|A_i\, *_\sigma \Delta_i(\omega) = \\ &-\min_{\omega \models A_i}\{\kappa(\omega) + \sum_{\substack{1 \leqslant j \leqslant n_i \\ \omega \models A_i C_{i,j} \overline{B}_{i,j}}} \eta_{i,j}\} - (\kappa_{0,i} - \kappa(A_i)) = 0, \end{aligned} \tag{5.12}$$

The first statement concerning the CSP follows immediately from the equality between (5.3) and (5.8) in the proof of Theorem 5.1.1, i.e., the equality of the inequalities defining the impact factors for the subsets Δ_i. However, the equality of the CSP is not enough to ensure the rank-wise equality of $(\kappa *_\sigma (\Delta \cup \{S\}))|A_i$ and $\kappa|A_i *_\sigma \Delta_i$, we also need a correspondence between the normalization constants as displayed in the second statement of the proposition. This statement follows from (5.5), (5.6) and the equality of $(*)$ from (5.5) and $(**)$ from (5.6) as was shown in (5.10). In Section 5.3, we use these statements to define the global revision $\kappa *_\sigma \Delta$ using the local ones in $\kappa|A_i *_\sigma \Delta_i$. We give an example of (GRK) with $S \equiv \top$ for c-revisions.

Example 5.1.1. *Let* $\Sigma = \{a, b, c, d\}$ *and* $\Delta_1 = \{(c|ab), (d|abc)\}$*,* $\Delta_2 = \{(d|a\overline{b})\}$*,* $\Delta_3 = \{(d|\overline{a}), (cd|\overline{a})\}$ *such that* $\Delta = \Delta_1 \cup \Delta_2 \cup \Delta_3$ *and* $\mathcal{P}_\Delta = \{ab, a\overline{b}, \overline{a}\}$ *as the finest premise splitting and* $S = \top$*. For the rest of the example* $i \in \{1, 2, 3\}$ *holds. The OCFs* κ *and* $\kappa|A_i$ *are depicted in Table 5.1, along with schematic c-revised* $\kappa *^{\mathrm{c}} \Delta$ *and the computed ranks for the special choice of impact factors. From Proposition 5.1.2, it follows that* $CR(\kappa, \Delta)_{\Delta_i} = CR(\,\kappa|A_i\,, \Delta_i)$*. Thus, applying (IP-ESP*$^\sigma$*) we*

$\omega \in \Omega$	κ	$\kappa\vert A_i$	$\kappa *^{\mathrm{c}} \Delta$	$\kappa^{\star}_{\Delta}$	$\kappa^{\star}_{\Delta}\vert A_i$	$(\,\kappa\vert A_i\,)^{\star}_{\Delta_i}$
$abcd$	5	3	-1 + 5	4	0	0
$abc\overline{d}$	2	0	-1 + 2 + $\eta_{1,2}$	5	1	1
$ab\overline{c}d$	3	1	-1 + 3 + $\eta_{1,1}$	5	1	1
$ab\overline{c}\overline{d}$	4	2	-1 + 4 + $\eta_{1,1}$	6	2	2
$a\overline{b}cd$	3	3	-1 + 3	2	2	2
$a\overline{b}c\overline{d}$	0	0	-1 + 0 + $\eta_{2,1}$	1	1	1
$a\overline{b}\overline{c}d$	1	1	-1 + 1	0	0	0
$a\overline{b}\overline{c}\overline{d}$	2	2	-1+ 2 + $\eta_{2,1}$	3	3	3
$\overline{a}bcd$	4	3	-1 + 4	3	0	0
$\overline{a}bc\overline{d}$	1	0	-1 + 1 + $\eta_{3,1}$ + $\eta_{3,2}$	4	1	1
$\overline{a}b\overline{c}d$	2	1	-1 + 2 + $\eta_{3,2}$	5	2	2
$\overline{a}b\overline{c}\overline{d}$	3	2	-1 + 3 +$\eta_{3,1}$ + $\eta_{3,2}$	6	3	3
$\overline{a}\overline{b}cd$	6	5	-1 + 6	5	2	2
$\overline{a}\overline{b}c\overline{d}$	3	2	-1 + 3 + $\eta_{3,1}$ + $\eta_{3,2}$	6	3	3
$\overline{a}\overline{b}\overline{c}d$	4	3	-1 + 4 + $\eta_{3,2}$	7	4	4
$\overline{a}\overline{b}\overline{c}\overline{d}$	5	4	-1 + 5 + $\eta_{3,1}$ + $\eta_{3,2}$	8	5	5
			$\kappa_{0,\Delta} = -1$			$\kappa_{0,1} = -3$, $\kappa_{0,2} = -1$, $\kappa_{0,3} = -3$

Table 5.1: OCF κ, $\kappa|A_i$ and c-revisions from Example 5.1.1. Worlds in $Mod(A_i)$ $(i = 1, 2, 3)$ are separated using different gray tones.

can choose the following Pareto-minimal impact factors

$$\sigma(\kappa, \Delta)_{\Delta_1} = (\eta_{1,1}, \eta_{1,2}) = (3, 4) = \sigma(\,\kappa|A_1\,, \Delta_1)$$
$$\sigma(\kappa, \Delta)_{\Delta_2} = (\eta_{2,1}) = (2) = \sigma(\,\kappa|A_2\,, \Delta_2)$$
$$\sigma(\kappa, \Delta)_{\Delta_3} = (\eta_{3,1}, \eta_{3,2}) = (0, 4) = \sigma(\,\kappa|A_3\,, \Delta_3).$$

In Table 5.1, we indicate the special choice of impact factors from above with an additional superscript and apply the following abbreviations for the corresponding c-revisions $\kappa *_\sigma \Delta = \kappa^{\star}_{\Delta}$, $(\kappa *_\sigma \Delta)|A_i = \kappa^{\star}_{\Delta}|A_i$ *and* $\kappa|A_i *_\sigma \Delta_i = (\,\kappa|A_i\,)^{\star}_{\Delta_i}$. *From Table 5.1, it is clear that* $\kappa|A_i *_\sigma \Delta_i(\omega) = (\kappa *_\sigma \Delta)|A_i\,(\omega)$, *i.e., (GRK) is satisfied.*

5.2 Reduction to Local Propositional Revision

In this section, we consider special cases in which for each subset Δ_i the case A_i is equivalent to all premises of conditionals in Δ_i, i.e., the case splitting fully captures the context of the corresponding new information. For these special cases, we show that conditional revision and propositional with solely the conclusions of the corresponding conditionals coincide, thus making the revision task in the Kinematics principle less complex. Also, we discuss the relationship between conditionalization and propositional revision in a broader context, leading to a conditionalized version of the well-known Ramsey Test.

In a broad sense, the Kinematics principle makes revision and conditionalization interchangeable in cases where the new information can be split into different cases. Local cases are set up via conditionalization, and the new information concerning only these cases can be processed independently of the remaining input. Thus, conditionalization is a powerful tool to introduce locality in the revision process and enables us to simplify it. The Kinematics principle uses general local cases introduced by case splittings $\mathcal{P}_\Delta = \{A_1, \ldots, A_n\}$ as defined in Definition 4.1.1, which can be further specified by the premises of the conditionals in the corresponding subsets Δ_i, which have to imply A_i. Hence, the conditionals in Δ_i might concern more specific situations within the case of A_i. Now, we focus on the special case, where the local cases from the case splitting $A_1, \ldots, A_n$ are not further refined by the conditionals' premises and fully represent the local contexts the Kinematics principle operates on. Thus, we assume that for each subset Δ_i, it holds that $\Delta_i = \{(A_i|B_{i,j})\}_{j=1}^n$, i.e., all conditionals in Δ_i have the same premise. In general, for a revision operator $*$ that satisfies (GRK), it holds that

$$\kappa * \Delta|A_i \;=\; \kappa|A_i * \Delta_i.$$

In the special case of subsets $\Delta_i = \{(A_i|B_j)\}_{j=1}^n$, where the OCFs $\kappa|A_i$ already condition on A_i and therefore set up the specific local cases, the premises in Δ_i seem superfluous, and the question arises whether revision by Δ_i can be implemented by multiple propositional revision. These cases are subsumed by the principle of *Local Propositional Revision.*

Definition 5.2.1. *Let $A_1, \ldots, A_n$ be a case splitting, κ be a ranking function, and let $\Delta = \Delta_1 \cup \ldots \cup \Delta_n$ be a set of conditionals with subsets $\Delta_i = \{(B_j|A_i)\}_{1 \leqslant j \leqslant n_i}$ and $\mathcal{B}_i = \{B_j\}_{1 \leqslant j \leqslant n_i}$ be the set of consequents of the conditionals in Δ_i. A revision operator $*$ for κ satisfies* Local Propositional Revision *iff*

$$\textbf{\textit{(LPR)}} \qquad \kappa|A_i * \Delta_i = \kappa|A_i * \mathcal{B}_i.$$

Note that, in this section we focus on the OCF version of (LPR) since we employ c-revision as a proof of concept. In Section 6.1, we also discuss (LPR) in the qualitative context.

The (LPR) principle connects conditional and propositional revision via conditionalization and, thus, brings the notion of locality introduced by conditionalization to light. It states that if the agent focuses on the specific case a conditional information is coming from, then it can be reduced to a piece of propositional information, namely just the consequent of $(B_j|A_i)$. Note that methods of conditional revision can be used for propositional revision by representing propositions as conditionals with tautological premises. Because conditional information differs fundamentally from propositional information (cf. Section 2.4), it holds that in general, methods of propositional revision cannot easily be adapted to conditional information, and the (LPR) principle is a simple yet elegant approach to fill this gap for sets of conditionals with identical premises.

Note that, in the special case of $\Delta = \{(B|A)\}$, i.e., if we revise with just a single conditional, we obtain a case splitting with $A_1 = A$ and $A_2 = \overline{A}$ and $\Delta_1 = \Delta$ and $\Delta_2 = \emptyset$.[2] Thus, (LPR) reduces to

$$\kappa|A * (B|A) = \kappa|A * B. \tag{5.13}$$

Conditionalization with A reduces the set of possible worlds to $\mathrm{Mod}(A) \subseteq \Omega$ and, thus, eliminates neutrality in the evaluation of $(B|A)$ w.r.t. a possible world ω (cf. formula (2.6) on page 34), s.t. each world either verifies or falsifies $(B|A)$. Therefore $\kappa|A$ shows the same response behavior for $\omega \in \mathrm{Mod}(A)$ w.r.t. a conditional $(B|A)$ or

[2]In Section 4.1, we already discussed that the subsets $\Delta_i \subseteq \Delta$ induced by a case splitting are not necessarily non-empty.

a proposition B. It holds that, in the context of conditionalized OCFs, the revision with a single conditional boils down to a propositional revision with the conclusion for revision operators satisfying (LPR).

Before we take a closer look on the relations between propositional and conditional revision introduced by (LPR) and provide a proof of concept, we show that $\mathcal{B}$, i.e., the set of all consequents for Δ_i, in (LPR) inherits consistency from the subsets Δ_i and therefore, all revisions in (LPR) are well-defined:

Lemma 5.2.1. *If $\Delta_i = \{(B_j|A_i)\}_{1\leqslant j\leqslant n_i}$ is consistent, then $\mathcal{B} = \{B_j\}_{1\leqslant j\leqslant n_i}$ is a consistent set of propositions.*

Proof. The following proof uses the notion of tolerance for sets of conditionals presented in [1]. According to [1],if Δ is consistent, then at least one conditional $(B_j|A_i) \in \Delta$ is tolerated by all remaining conditionals in Δ. Thus, there exists $\omega \in \Omega$ with

$$\omega \models A_iB_j \wedge \bigwedge_{\substack{1\leqslant \ell\leqslant n_i\\ \ell\neq i}} (A_i \Rightarrow B_\ell) \Leftrightarrow \omega \models A_iB_j \wedge \bigwedge_{\substack{1\leqslant \ell\leqslant n_i\\ \ell\neq i}} (\overline{A_i} \vee B_\ell) \Leftrightarrow \omega \models A_iB_1\cdots B_{n_i}.$$

So, $\mathrm{Mod}(B_1 \wedge \ldots \wedge B_{n_i}) \neq \emptyset$ and therefore $\mathcal{B} = \{B_1, \ldots, B_{n_i}\} = \mathcal{B}$ is consistent. □

As a proof of concept, we show that impact-preserving strategic c-revisions satisfy the (LPR) axiom.

Theorem 5.2.2. *Let κ be a ranking function. Let $\Delta = \Delta_1 \cup \ldots \Delta_n$ with $\Delta_i = \{(B_j|A_i)\}_{1\leqslant j\leqslant n_i}$ be a set of conditionals and $\mathcal{B}_i = \{B_j\}_{1\leqslant j\leqslant n_i}$ as specified in Definition 5.2.1. If σ satisfies (IP-ESP$^\sigma$) then $*_\sigma$ is a strategic c-revision operator that satisfies (LPR).*

Proof. The constraint satisfaction problem $CR(\kappa|A_i\,, \mathcal{B}_i)$ is given by the following set of the constraints for all $j = 1, \ldots, n_i$:

$$\begin{aligned}\eta_j &> \min_{\omega\models B_j}\{\,\kappa|A_i\,(\omega) + \sum_{\substack{j\neq k\\ \omega\models\overline{B}_k}} \eta_k\} - \min_{\omega\models\overline{B}_j}\{\,\kappa|A_i\,(\omega) + \sum_{\substack{j\neq k\\ \omega\models\overline{B}_k}} \eta_k\}\\ &= \min_{\omega\models A_iB_j}\{\kappa(\omega) - \kappa(A_i) + \sum_{\substack{j\neq k\\ \omega\models A\overline{B}_k}} \eta_k\} - \min_{\omega\models A_i\overline{B}_j}\{\kappa(\omega) - \kappa(A_i) + \sum_{\substack{j\neq k\\ \omega\models A\overline{B}_k}} \eta_k\} \quad (5.14)\end{aligned}$$

Equation (5.14) holds since we revise the conditionalized OCF $\kappa|A_i$ which is defined only on worlds in $\mathrm{Mod}(A_i)$. So, in this case, for all worlds $\omega \models B_j$ it holds $\omega \models A_iB_j$. The set of the constraints in (5.14) for all $j = 1, \ldots, n_i$ is the same as $CR(\kappa|A_i, \Delta)$ (compare (5.7) in the proof of Theorem 5.1.1). Thus, $CR(\kappa|A_i, \mathcal{B}_i) = CR(\kappa|A_i, \Delta_i)$ and therefore $\sigma(\kappa|A_i, \mathcal{B}_i) = \sigma(\kappa|A_i, \Delta_i)$ by (IP-ESP$^\sigma$).

We still need to show that $\kappa_1^{*_\sigma}(\omega) = \kappa|A_i *_\sigma \mathcal{B}_i(\omega) = \kappa|A_i *_\sigma \Delta_i(\omega) = \kappa_2^{*_\sigma}(\omega)$ for all $\omega \models A_i$. For $\omega \models A_i$, it holds that

$$\kappa_1^{*_\sigma}(\omega) = \kappa_{0,\Delta} + \kappa|A_i(\omega) + \sum_{\substack{j=1,\ldots,n_i \\ \omega \models \overline{B}_j}} \eta_j$$

$$= \kappa_{0,\Delta} + \kappa|A_i(\omega) + \sum_{\substack{j=1,\ldots,n_i \\ \omega \models A_i\overline{B}_j}} \eta_j \tag{5.15}$$

$$\text{and} \quad \kappa_2^{*_\sigma}(\omega) = \kappa_{0,i} + \kappa|A_i(\omega) + \sum_{\substack{j=1,\ldots,n_i \\ \omega \models A_i\overline{B}_j}} \eta_j. \tag{5.16}$$

Note that, due to $\sigma(\kappa|A_i, \mathcal{B}_i) = \sigma(\kappa|A_i, \Delta_i)$, the c-revisions use the same impact vector $\vec{\eta}$. Since $\kappa_{0,\Delta}, \kappa_{0,i}$ display normalization constants defined by (2.15), and both c-revisions in (5.15) and (5.16) use the same values, also the normalization constants must be the same, and we are done. □

Note that, it holds that $CR(\kappa|A_i, \mathcal{B}_i) = CR(\kappa|A_i, \Delta_i)$ and thus, (IP-ESP$^\sigma$) is crucial for the proof of the theorem above.

We illustrate (LPR) for strategic c-revisions in the following example.

Example 5.2.1. *In Table 5.2 a ranking function κ is given. Let $\Delta = \Delta_1 \cup \Delta_2 = \{(b|a), (\overline{c}|a)\} \cup \{(\overline{b}|\overline{a})\}$ with the finest case splitting $\mathcal{P}_\Delta = \{A_1, A_2\}$ with $A_1 = a$ and $A_2 = \overline{a}$ according to Algorithm 1, s.t. $\Delta = \Delta_1 \cup \Delta_2 = \{(b|a), (\overline{c}|a)\} \cup \{(\overline{b}|\overline{a})\}$ and therefore $\mathcal{B}_1 = \{c, \overline{b}\}$ and $\mathcal{B}_2 = \{\overline{b}\}$.*

We revise $\kappa|A_i$ with Δ_i resp. with $\mathcal{B}_i$ to illustrate the benefits of (LPR). We employ strategic c-revisions $*_\sigma$ with a strategy σ that satisfies (IP-ESP$^\sigma$).

We start with $i = 1$, i.e., $A_1 = a$, $\Delta_1 = \{(b|a), (\overline{c}|a)\}$, and $\mathcal{B}_1 = \{b, \overline{c}\}$ and investigate

$$\kappa|a *_\sigma \Delta_1 \ \textit{vs.} \ \kappa|a *_\sigma \mathcal{B}_1.$$

$\omega \in \Omega$	κ	$\kappa\|A_i$	$\kappa\|A_i *_\sigma \Delta_i$	$(\kappa\|A_i)^\star_{\Delta_i}$	$\kappa\|A_i *_\sigma \mathcal{B}_i$	$(\kappa\|A_i)^\star_{\mathcal{B}_i}$
abc	1	1	$\kappa_{0,1} + 1 + \eta_{1,2}$	$= 1$	$\kappa_{0,1} + 1 + \eta_2$	$= 1$
$ab\overline{c}$	0	0	$\kappa_{0,1} + 0$	$= 0$	$\kappa_{0,1} + 0$	$= 0$
$a\overline{b}c$	2	2	$\kappa_{0,1} + 2 + \eta_{1,1} + \eta_{1,2}$	$= 3$	$\kappa_{0,1} + 2 + \eta_1 + \eta_2$	$= 3$
$a\overline{b}\overline{c}$	0	0	$\kappa_{0,1} + 0 + \eta_{1,1}$	$= 1$	$\kappa_{0,1} + 0 + \eta_1$	$= 1$
$\overline{a}bc$	1	0	$\kappa_{0,2} + 0 + \eta_{2,1}$	$= 1$	$\kappa_{0,2} + 0 + \eta_1$	$= 1$
$\overline{a}b\overline{c}$	1	0	$\kappa_{0,2} + 0 + \eta_{2,1}$	$= 1$	$\kappa_{0,2} + 0 + \eta_1$	$= 1$
$\overline{a}\overline{b}c$	2	1	$\kappa_{0,2} + 1$	$= 0$	$\kappa_{0,2} + 1$	$= 0$
$\overline{a}\overline{b}\overline{c}$	2	1	$\kappa_{0,2} + 1$	$= 0$	$\kappa_{0,2} + 1$	$= 0$
	$\kappa(a) = 0$			$\kappa_{0,1} = 0$		$\kappa_{0,1} = 0$
	$\kappa(\overline{a}) = 0$			$\kappa_{0,2} = -1$		$\kappa_{0,2} = -1$

Table 5.2: OCF κ, $\kappa|A_i$ and c-revisions from Example 5.2.1. Worlds in $\mathrm{Mod}(A_i)$ $(i = 1, 2)$ are separated using different gray tones.

Via (3.1), *we get the following inequalities defining impact factors for conditionals in* Δ_1.

$$
\begin{aligned}
(b|a) \quad \eta_{1,1} &> \min\Big\{ \kappa|a\,(abc) + \eta_{1,2},\ \kappa|a\,(ab\overline{c}) \Big\} - \min\Big\{ \kappa|a\,(a\overline{b}c) + \eta_{1,2},\ \kappa|a\,(a\overline{b}\overline{c}) \Big\} \\
&= 0 - 0 = 0 \\
(\overline{c}|a) \quad \eta_{1,2} &> \min\Big\{ \kappa|a\,(ab\overline{c}),\ \kappa|a\,(a\overline{b}\overline{c}) + \eta_{1,1} \Big\} - \min\Big\{ \kappa|a\,(abc),\ \kappa|a\,(a\overline{b}c) + \eta_{1,1} \Big\} \\
&= 0 - 1 = -1
\end{aligned}
$$

Because σ *satisfies (IP-ESP*$^\sigma$*), we can choose* $\eta^\star_{1,1} = \eta^\star_1 = 1$ *and* $\eta^\star_{1,2} = \eta^\star_2 = 0$ *for* $\eta^\star_i$ *the corresponding impact factor for the propositional revision with* $\mathcal{B}_1$. *The superscript indicates the special choice of the impact factor. Table 5.2 shows the results of the revision as a schema as well as with the specific choice of impact factors indicated by a superscript. It holds that* $\kappa|a *_\sigma \Delta_1(\omega) = (\kappa|a)^\sigma_{\Delta_1} = (\kappa|a)^\sigma_{\mathcal{B}_1} = \kappa|a *_\sigma \mathcal{B}_1(\omega)$ *since* σ *satisfies (IP-ESP*$^\sigma$*), and therefore* $(\kappa|a)^\star_{\Delta_i} = (\kappa|a)^\star_{\mathcal{B}_i}$.
Now, we turn to $i = 2$, *i.e.,* $A_2 = \overline{a}$, $\Delta_2 = \{(\overline{b}|\overline{a})\}$, *and* $\mathcal{B}_2 = \{\overline{b}\}$. *As before, we investigate the conditional vs. the propositional revision*

$$\kappa|\overline{a} *_\sigma \Delta_2 \ \textit{vs.}\ \ \kappa|\overline{a} *_\sigma \mathcal{B}_2.$$

Via (3.1)*, we get the following inequality defining the impact factor for* $(\overline{b}|\overline{a})$ *in* Δ_2*.*

$$(\overline{b}|\overline{a}) \quad \eta_{2,1} > \min\Big\{ \kappa|\overline{a}\,(\overline{a}\overline{b}c),\, \kappa|\overline{a}\,(\overline{a}\overline{b}\overline{c})\Big\} - \min\Big\{ \kappa|\overline{a}\,(\overline{a}bc),\, \kappa|\overline{a}\,(\overline{a}b\overline{c})\Big\} = 1 - 0 = 1$$

Following the same argumentation as for $i = 1$*, we can choose* $\eta^\star_{2,1} = \eta^\star_1 = 2$ *for the corresponding impact factor* η_1 *for the propositional revision with* $\mathcal{B}_2 = \{\overline{b}\}$*. The results of the revision can be found in Table 5.2 and it holds that* $\kappa|\overline{a} *_\sigma \Delta_2(\omega) = (\kappa|\overline{a})^\sigma_{\Delta_2}(\kappa|\overline{a})^\sigma_{\mathcal{B}_2} = \kappa|\overline{a} *_\sigma \mathcal{B}_2(\omega)$ *since* σ *satisfies (IP-ESP*$^\sigma$*) and therefore* $(\kappa|\overline{a})^\star_{\Delta_2}(\kappa|\overline{a})^\star_{\mathcal{B}_2} =$*. All in all, it holds that strategic c-revisions* $*_\sigma$ *satisfy (LPR).*

The following proposition summarizes essential statements about the compatibility of conditionalization for OCFs [119] and revision operators $*$ for OCFs which satisfy the well-known Ramsey Test (cf. (2.2) on page 26).

Proposition 5.2.3. *Let* κ *be an OCF and* $*$ *a revision operator which satisfies the Ramsey Test, i.e.,* $\kappa \models (B|A)$ *iff* $\kappa * A \models B$*. Then the following statements hold:*

- *Compatibility of conditionalization with conditionals:*

$$\kappa|A \models B \text{ iff } \kappa \models (B|A)$$

- *Conditional Ramsey Test:*

$$\text{(CRT)} \quad \kappa|A \models B \text{ iff } \kappa * A \models B \tag{5.17}$$

Proof. We start by showing that OCF-conditionalization is compatible with conditional information. It holds that $\kappa|A \models B$ iff $\kappa|A\,(\overline{B}) > 0$ which is equivalent to $\kappa|A\,(B) < \kappa|A\,(\overline{B})$ due to the properties of OCFs (cf. page 36 in Section 2.4). Since $\kappa|A$ is defined on Mod(A), we can conclude that $\kappa|A\,(AB) < \kappa|A\,(A\overline{B})$ and therefore, also $\kappa(AB) < \kappa(A\overline{B})$, i.e., $\kappa \models (B|A)$.

Now, we show that (CRT) holds. From the first statement, we can conclude that $\kappa|A \models B$ iff $\kappa \models (B|A)$, which is equivalent to $\kappa * A \models B$ due to the classical Ramsey Test. □

The first statement compares the acceptance of a proposition in a specific context provided by conditionalization to the acceptance of a conditional. As we can see, conditionalization for OCFs and conditional information are compatible in a rational way, i.e., if we consider only worlds in which the premise of a conditional is true and therefore focus on a specific context provided by the premise, then we also must accept the consequents of the conditionals that deal with this case. The Conditional Ramsey Test builds up on this compatibility and connects the acceptance of propositional information in a conditionalized state with the acceptance of the same information for revised states. This corresponds to the notion that conditionalization operators implement special propositional revision operators, which was already discussed for probabilistic revision operators [88, 87].

5.3 Building a Global Revision From the Local Revision

The Kinematics principle for OCFs displays a powerful concept that deals with case-related information that allows us to reduce the computational complexity of the revision with a set of conditionals by employing the idea that only conditionals talking about a specific case are relevant in the corresponding case. Note that especially for c-revisions, the computational complexity is directly linked to the size of Δ resp. Δ_i, since for each conditional, a new constraint variable defined by (3.1) is added to the constraint satisfaction problem. Also, the size of Ω, i.e., the number of possible worlds, is relevant for the computational complexity. Here, the Kinematics principle for ranking functions offers a logically sound strategy how to reduce the complexity of the revision task by focusing on the particular case via conditionalization, i.e., reducing the number of possible worlds $\mathrm{Mod}(A_i) \subseteq \Omega$ and also the size of the input set $\Delta_i \subseteq \Delta$. Unfortunately, it is unclear how the plausibility ranks from different cases can be reassembled to recreate the globally revised OCF $\kappa^*_\Delta = \kappa * \Delta$. This is because conditionalization naturally leads to a loss of information on the relations between sub-OCFs from different cases, and the full global belief state κ^*_Δ is not uniquely determined by the local revisions from (GRK) $(\kappa * \Delta)|A_i = (\kappa^*_\Delta)|A_i$ resp. $\kappa|A_i * \Delta_i = (\kappa|A_i)^*_{\Delta_i}$.

So, in general, creating a globally revised OCF from the locally revised sub-OCFs is an open research question in general, which was tackled in [111] by employing c-revisions. In the following sections, we discuss a simplified form of ranking functions and a concatenation operator and show that with these technical tools at hand, we can define c-revision operators that recreate globally c-revised ranking functions from local ones in a straightforward way and thus providing the right-to-left direction of the following revision scenarios:

$$\kappa^*_\Delta = \kappa * \Delta \longleftrightarrow \kappa|A_i * \Delta_i = (\,\kappa|A_i\,)^*_{\Delta_i} \overset{\text{(GRK)}}{=} (\kappa * \Delta)|A_i \;=\; \kappa^*_\Delta|A_i$$

Note that the left-to-right connection is provided by (GRK) with $S = \top$ via conditionalization with the corresponding case A_i. Before we define the tools mentioned above, it is essential to acknowledge that the concepts presented in this section provide a way to reconstruct information that is usually lost during conditionalization and can only be applied to ranking functions. This is due to the arithmetic provided by OCFs, which is not provided in the qualitative framework. The numerical representation of plausibility ranks makes the mechanisms that come into play when conditionalizing a ranking function clearer and thus can be traced back. Qualitative representations of belief states as plausibilistic TPOs only provide relative information, i.e., the plausibility of a world can only be defined via its positioning to others, and therefore, they are less absolute.

5.3.1 Pre-OCFs and Concatenation of Ranking Functions

In Theorem 5.1.1, we have shown that c-revisions satisfy (GRK). To reconstruct the global c-revision $\kappa^{\mathrm{c}}_\Delta = \kappa *^{\mathrm{c}} \Delta$ from the local c-revisions $(\,\kappa|A_i\,)^{\mathrm{c}}_{\Delta_i} = \kappa|A_i *^{\mathrm{c}} \Delta_i$, we still need some technical tools, which we define in this section. From (5.11) in Proposition 5.1.2, it follows that the inequalities defining each impact factor η_i for conditionals in Δ are not affected by the switch from the local to the global scenarios, i.e., whether we consider subsets Δ_i or the global set of conditionals Δ. Yet, (5.12) from Proposition 5.1.2 shows that also the calculation of the normalization constants $\kappa_{0,i}$ for each local revision is crucial to ensure (GRK) for c-revisions since they play a distinctive role for $\kappa^{\mathrm{c}}_\Delta$ versus $(\,\kappa|A_i\,)^{\mathrm{c}}_{\Delta_i}$. We illustrate the role of normalization constants in the following example.

Example 5.3.1. *Let* $\Sigma = \{a, b, c\}$ *and* $\Delta = \{(b|a), (c|a), (b|\overline{a}), (\overline{c}|\overline{a})\}$ *with subsets*

$\Delta_1 = \{(b|a),(c|a)\}$ and $\Delta_2 = \{(b|\overline{a}),(\overline{c}|\overline{a})\}$ and $\mathcal{P}_\Delta = \{a,\overline{a}\}$ as finest case splitting. The prior ranking function κ as well as the conditionalized sub-OCFs $\kappa|A_i$ (with $A_1 = a$ and $A_2 = \overline{a}$) can be found in Table 5.3. For a better overview, we place the conditionalized OCFs beneath each other in one column. In Theorem 5.1.1, we showed that $CR(\kappa,\Delta)_{\Delta_i} = CR(\,\kappa|A_i\,,\Delta_i)$ holds, so using the ranking function κ from Table 5.3 we get the following system of inequalities:

$$
\begin{array}{ll}
(b|a) & \eta_{1,1} > \min\{3, 2+\eta_{1,2}\} - \min\{2, 1+\eta_{1,2}\} \\
(c|a) & \eta_{1,2} > \min\{3, 2+\eta_{1,1}\} - \min\{2, 1+\eta_{1,1}\} \\
(b|\overline{a}) & \eta_{2,1} > 4 - \min\{2, \eta_{2,2}\} \\
(\overline{c}|\overline{a}) & \eta_{2,2} > \min\{4, 2+\eta_{2,1}\} - \min\{4, \eta_{2,1}\}
\end{array}
$$

Employing (IP-ESP$^\sigma$), we choose a selection strategy σ, such that

$$
\begin{aligned}
\sigma(\kappa,\Delta)_{\Delta_1} &= (\eta_{1,1},\eta_{1,2}) = (2,2) = \sigma(\,\kappa|A_1\,,\Delta_1),\\
\sigma(\kappa,\Delta)_{\Delta_2} &= (\eta_{2,1}\eta_{2,2}) = (2,3) = \sigma(\,\kappa|A_2\,,\Delta_2)
\end{aligned}
$$

holds with Pareto-minimal impact factors. We indicate the special choice of impact factors and the corresponding strategic c-revisions with a superscript $\star$. The schematic revision results $\kappa *_\sigma \Delta$ resp. $\kappa|A_i *_\sigma \Delta_i$ $(i = 1,2)$ as well as their calculated ranks using the specific impact factors and the normalization constant κ_0 resp. $\kappa_{i,0}$ for $\kappa|A_i *_\sigma \Delta_i$ $(i = 1,2)$ are displayed in Table 5.3.

Example 5.3.1 shows, that for the local c-revision $\kappa|\overline{a} *_\sigma \Delta_2$, it holds that the global c-revision $\kappa *_\sigma \Delta$ uses the same plausibility ranks (since $\kappa(\overline{a}) = 0$) and the same impact factors $\eta_{2,j}$ $(j = 1,2)$. Only, the normalization constant for the global revision $\kappa_{0,\Delta}$ differs from the constants $\kappa_{0,1}$ and $\kappa_{0,2}$ for the local revisions in Table 5.3. This example, therefore, illustrates the crucial role of the normalization constants for our goal of expressing the global c-revision $\kappa^{\mathrm{c}}_{\Delta}$ by the local $(\,\kappa|A_i\,)^{\mathrm{c}}_{\Delta_i}$. Note that this vital role was already apparent in (5.12) from Proposition 5.1.2.

To exclude the effect of normalization, we need more general ranking functions, which we call *pre-ranking functions* or *pre-OCFs* that are not necessarily normalized. Note that normalization is an artifact that turns general rankings in terms of natural

$\omega \in \Omega$	κ	$\kappa\|A_i$	$\kappa *_\sigma \Delta$	$\kappa^\star_\Delta$	$\kappa\|A_i *_\sigma \Delta_i$	$\kappa^\star\|A_i{}_{\Delta_i}$
abc	3	2	$\kappa_0 + 3$	$= 0$	$\kappa_{0,1} + 2$	$= 0$
$ab\overline{c}$	2	1	$\kappa_0 + 2 + \eta_{1,2}$	$= 1$	$\kappa_{0,1} + 1 + \eta_{1,2}$	$= 1$
$a\overline{b}c$	2	1	$\kappa_0 + 2 + \eta_{1,1}$	$= 1$	$\kappa_{0,1} + 1 + \eta_{1,1}$	$= 1$
$a\overline{b}\overline{c}$	1	0	$\kappa_0 + 1 + \eta_{1,1} + \eta_{1,2}$	$= 2$	$\kappa_{0,1} + 0 + \eta_{1,1} + \eta_{1,2}$	$= 2$
$\overline{a}bc$	4	4	$\kappa_0 + 4 + \eta_{2,2}$	$= 4$	$\kappa_{0,2} + 4 + \eta_{2,2}$	$= 3$
$\overline{a}b\overline{c}$	4	4	$\kappa_0 + 4$	$= 1$	$\kappa_{0,2} + 4$	$= 0$
$\overline{a}\overline{b}c$	0	0	$\kappa_0 + 0 + \eta_{2,1} + \eta_{2,2}$	$= 2$	$\kappa_{0,2} + 0 + \eta_{2,1} + \eta_{2,2}$	$= 1$
$\overline{a}\overline{b}\overline{c}$	2	2	$\kappa_0 + 2 + \eta_{2,1}$	$= 1$	$\kappa_{0,2} + 2 + \eta_{2,1}$	$= 0$
	$\kappa(a) = 0$			$\kappa_0 = -3$		$\kappa_{0,1} = -2$
	$\kappa(\overline{a}) = 1$			$\kappa_0 = -3$		$\kappa_{0,2} = -4$

Table 5.3: OCF κ, $\kappa|A_i$ and c-revisions from Example 5.3.1. Worlds in $\mathrm{Mod}(A_i)$ $(i = 1, 2)$ are separated using different gray tones. The κ-ranks of a and $\overline{a}$ and the values of the normalization constants can be found in the last row.

numbers into OCFs. For c-revisions (2.14), normalization is used as a meta-concept that is applied to a structural schema, and normalization is not relevant for the constraint satisfaction problems 3.1.2.

Definition 5.3.1. *A* pre-OCF *or* pre-ranking function *is a function* $\kappa^{pre} : \Omega \to \mathbb{N} \cup \{\infty\}$.

Every pre-OCF κ^{pre} can be transformed to a classic ranking function via the *OCF-operator ocf* : $\kappa^{\mathrm{pre}} \mapsto \kappa$ with

$$\kappa(\omega) = ocf(\kappa^{\mathrm{pre}}) = \kappa^{\mathrm{pre}}(\omega) - \min_{\omega \in \Omega}\{\kappa^{\mathrm{pre}}(\omega)\}. \tag{5.18}$$

It holds that every standard OCF is also a pre-OCF, and the two concepts can be transferred into each other respectively in a straightforward way. However, the operator *ocf* is not injective since different pre-OCFs can yield the same ranking function applying the operator *ocf*. For a conditionalized ranking function $\kappa|A_i$ a natural candidate for an associated pre-OCF is the prior ranking function defined on the corresponding set of worlds $\mathrm{Mod}(A_i)$ only:

$$(\,\kappa|A_i\,)^{\mathrm{pre}}(\omega) = \kappa(\omega) \text{ for } \omega \models A_i.$$

The idea of neglecting normalization for ranking functions can also be utilized for (strategic) c-revisions. Thus, we define pre-c-revisions in full compliance with the standard strategic c-revisions by just leaving out the normalization constant:

Definition 5.3.2 (Strategic pre-c-revision)**.** *Let κ be a ranking function and let $\Delta = \{(B_1|X_1), \ldots, (B_m|X_m)\}$ be a set of conditionals. For a selection strategy σ for c-revisions, we define a* strategic pre-c-revision of κ by Δ *as a pre-OCF $(\kappa^{*_\sigma})^{pre}$ s.t. $(\kappa^{*_\sigma})^{pre}$ accepts Δ and is given by the impact vector $\vec{\eta} = \sigma(\kappa, \Delta)$:*

$$(\kappa^{*_\sigma})^{pre}(\omega) = \kappa *_\sigma^{pre} \Delta = \kappa(\omega) + \sum_{\substack{1 \leqslant i \leqslant m \\ \omega \models X_i \overline{B}_i}} \eta_i. \tag{5.19}$$

It is straightforward yet essential to acknowledge that the constraint satisfaction problem specifying η_i is not affected by the normalization constant and, thus, does not differ from the standard one for pre-c-revisions. As for standard c-revisions, the impact values η_i correspond to a single conditional $(B_i|X_i) \in \Delta_i$ each.

Another challenge we have to meet when setting up $\kappa^{\mathrm{c}}_{\Delta}$ from $(\,\kappa|A_i\,)^{\mathrm{c}}_{\Delta_i}$ is the reconstruction of the global OCF κ from conditionalized sub-OCFs $\kappa|A_i$. Conditionalization can be seen as a mapping from ranking functions defined on the whole set of worlds Ω to ranking functions defined on a subset $\mathrm{Mod}(A_i)$. Hence, conditionalization leads to irreversible loss of information concerning the ranks of worlds in $\Omega \setminus \mathrm{Mod}(A_i)$. In this general scenario, the advantages of employing numerical ranks to express the plausibility of worlds become particularly clear because numbers and the underlying arithmetic allow us to understand exactly to which extent the case implied by a world influences its global ranking in the OCF. The unique structure given by the case splitting in our Kinematics principle for OCFs is crucial for the following concatenation operator, which reconstructs a global pre-OCF from the local ones:

Definition 5.3.3 (Concatenation Operator)**.** *Let $\{A_1, \ldots, A_n\}$ be a case splitting and let $(\,\kappa|A_i\,)^{pre}$ be the corresponding pre-OCF defined on Mod(A_i). The* concatenation operator $\oplus$ *maps a set of conditional pre-OCFs $\{(\,\kappa|A_i\,)^{pre}\}_{i=1}^{n}$ to a single* *pre-OCF* $\oplus(\{(\kappa|A_i)^{pre}\}_{i=1}^n)$ *defined on the whole set* Ω *as follows:*

$$\begin{aligned}&\oplus(\{(\kappa|A_i)^{pre}\}_{i=1}^n):\Omega\to\mathbb{N}\cup\{\infty\},\\ &\qquad\qquad s.t.,\ \oplus(\{(\kappa|A_i)^{pre}\}_{i=1}^n)(\omega)=(\kappa|A_i)^{pre}(\omega)\ for\ \omega\models A_i.\end{aligned}\tag{5.20}$$

Note that, Definition 5.3.3 is defined for pre-OCFs, but since every standard OCF matches the definition of a pre-OCF, the concatenation operator is not solely restricted to pre-OCFs but is also defined for ranking functions. It holds that if one of the $(\kappa|A_i)^{\text{pre}}$ is a ranking function, i.e., there exists at least one world with rank zero, then $\oplus(\{(\kappa|A_i)^{\text{pre}}\}_{i=1}^n)$ is a ranking function, too, because concatenation does not change the ranks of $(\kappa|A_i)^{\text{pre}}$. Furthermore, for the operator $\oplus$, the guiding principle for combining the ranks from local cases to a global ranking function is simply to connect the ranks from disjoint local cases to a full pre-OCF on Ω. Thus, the operator $\oplus$ does not really combine information but rather glues different sub-OCFs together. This is possible due to the special structure of the local cases in the case splitting $\{A_1,\ldots,A_n\}$. Exclusivity guarantees that the plausibility rankings do not interfere. Thus, we can simply string them together, and exhaustiveness ensures that $\oplus(\{(\kappa|A_i)^{\text{pre}}\}_{i=1}^n)$ ranks each world in Ω. Concatenation, even in the case of exclusive cases, is not defined for more general epistemic states like TPOs. The plausibility of a world in a TPO depends on its relative positioning towards other worlds; this information is irretrievably lost during conditionalization as it displays a mapping from Ω to $\text{Mod}(A_i)$.

5.3.2 From Global to Local C-Revisions

Now, we have the toolbox to define the global c-revision $\kappa *^{\text{c}} \Delta$ with a set of conditionals Δ which satisfies the preamble of (GRK) from the local c-revisions $\kappa|A_i *^{\text{c}} \Delta_i$. By employing pre-OCFs from Definition 5.3.1 and the corresponding pre-c-revisions from Definition 5.3.2 together with the concatenation operation defined in (5.20) we get the following theorem:

Theorem 5.3.1. *Let* $\Delta=\Delta_1\cup\ldots\cup\Delta_n$ *be a set of conditionals and* $\mathcal{P}_\Delta=\{A_1,\ldots,A_n\}$ *be a case splitting of* Δ *specified as in the preamble of (GRK). Let* κ *be an OCF and let* σ *be a selection strategy that satisfies (IP-ESP$^\sigma$). It holds* *for strategic pre-c-revisions $*_\sigma^{pre}$ resp. strategic c-revisions $*_\sigma$ induced by the same strategy σ and the concatenation operator $\oplus(\cdot)$ that*

$$\kappa *_\sigma \Delta(\omega) = ocf(\oplus(\{\, \kappa|A_i \;*_\sigma^{pre} \Delta_i + \kappa(A_i)\}_{i=1}^n))(\omega) \tag{5.21}$$

for all $\omega \in \Omega$.

Proof. The theorem follows directly from (GRK) for ranking functions and the definition of strategic pre-c-revisions resp. c-revisions and the concatenation operator $\oplus(\cdot)$. For $\omega \in \mathrm{Mod}(A_i)$, it holds that

$$\kappa|A_i \;*_\sigma^{\mathrm{pre}} \Delta_i(\omega) + \kappa(A_i) = \;\kappa|A_i\,(\omega) + \sum_{\substack{1 \leqslant j \leqslant n_i \\ \omega \models A_i C_{i,j} \overline{B}_{i,j}}} \eta_{i,j} + \kappa(A_i) = \kappa *_\sigma^{\mathrm{pre}} \Delta(\omega) \tag{5.22}$$

Note that since, the A_i's are exclusive the (pre-)c-revision with Δ yields the same results on $\mathrm{Mod}(A_i)$ as the revisions with Δ_i (cf. Proposition 5.1.2). Since (5.22) holds for all $1 \leqslant i \leqslant n$, we get

$$\kappa *_\sigma^{\mathrm{pre}} \Delta(\omega) = \oplus(\{\, \kappa|A_i \;*_\sigma^{\mathrm{pre}} \Delta_i + \kappa(A_i)\}_{i=1}^n) \tag{5.23}$$

for all $\omega \in \Omega$ via employing the concatenation operator $\oplus$. Thus, (5.21) follows from (5.18) via normalization for all $\omega \in \Omega$. □

Theorem 5.3.1 shows that for sets of conditionals Δ, satisfying the prerequisites of (GRK), instead of revising the prior κ with the whole set Δ, we can revise with subsets Δ_i of Δ concerning just the local cases in $\kappa|A_i$ and then concatenate the results and normalize. The local revision $\kappa|A_i$ is sufficient to c-revise the prior ranking function κ with the set of conditionals Δ, i.e., we can reconstruct $\kappa_\Delta^{\mathrm{c}}$ from $(\,\kappa|A_i\,)_{\Delta_i}^{\mathrm{c}}$ for strategic c-revision satisfying (IP-ESP$^\sigma$) via normalizing the concatenated and revised pre-OCFs $\kappa|A_i *_\sigma^{\mathrm{pre}} \Delta_i$ after adding $\kappa(A_i)$ which is given by the prior κ. This way, we can compute the global c-revision $\kappa *^{\mathrm{c}} \Delta$ more efficiently since concatenation and normalization come at linear cost. Thus, splitting into local subcases reduces the exponential effort of the revision significantly.

We illustrate our results with an example.

$\omega \in \Omega$	κ	$\kappa\vert A_i$	$\kappa *_{\sigma}^{\text{pre}} \Delta$		$\oplus((\kappa\vert A_i)^{\text{pre}}_{\Delta_i}(\omega) + \kappa(A_i))$	
$abcd$	2	2	2	$= 2$	$2 + \kappa(ab)$	$= 2$
$abc\overline{d}$	2	2	$2 + \eta_{1,2}$	$= 3$	$2 + \eta_{1,2} + \kappa(ab)$	$= 3$
$ab\overline{c}d$	1	1	$1 + \eta_{1,1}$	$= 3$	$1 + \eta_{1,1} + \kappa(ab)$	$= 3$
$ab\overline{c}\overline{d}$	0	0	$0 + \eta_{1,1} + \eta_{1,2}$	$= 3$	$0 + \eta_{1,1} + \eta_{1,2} + \kappa(ab)$	$= 3$
$a\overline{b}cd$	1	0	$1 + \eta_{2,1}$	$= 2$	$1 + \eta_{2,1} + \kappa(a\overline{b})$	$= 2$
$a\overline{b}c\overline{d}$	3	2	$3 + \eta_{2,1} + \eta_{2,2}$	$= 4$	$3 + \eta_{2,1} + \eta_{2,2} + \kappa(a\overline{b})$	$= 4$
$a\overline{b}\overline{c}d$	1	0	1	$= 1$	$1 + \kappa(a\overline{b})$	$= 1$
$a\overline{b}\overline{c}\overline{d}$	2	1	$2 + \eta_{2,2}$	$= 2$	$2 + \eta_{2,2} + \kappa(a\overline{b})$	$= 2$
$\overline{a}bcd$	4	3	$4 + \eta_{3,2}$	$= 5$	$4 + \eta_{3,2} + \kappa(\overline{a})$	$= 5$
$\overline{a}bc\overline{d}$	3	2	3	$= 3$	$3 + \kappa(\overline{a})$	$= 3$
$\overline{a}b\overline{c}d$	1	0	$1 + \eta_{3,1} + \eta_{3,2}$	$= 4$	$1 + \eta_{3,1} + \eta_{3,2} + \kappa(\overline{a})$	$= 4$
$\overline{a}b\overline{c}\overline{d}$	4	3	$4 + \eta_{3,1}$	$= 6$	$4 + \eta_{3,1} + \kappa(\overline{a})$	$= 6$
$\overline{a}\overline{b}cd$	4	3	$4 + \eta_{3,2}$	$= 5$	$4 + \eta_{3,2} + \kappa(\overline{a})$	$= 5$
$\overline{a}\overline{b}c\overline{d}$	3	2	3	$= 3$	$3 + \kappa(\overline{a})$	$= 3$
$\overline{a}\overline{b}\overline{c}d$	2	1	$2 + \eta_{3,1} + \eta_{3,2}$	$= 5$	$2 + \eta_{3,1} + \eta_{3,2} + \kappa(\overline{a})$	$= 5$
$\overline{a}\overline{b}\overline{c}\overline{d}$	4	3	$4 + \eta_{3,1}$	$= 6$	$4 + \eta_{3,1} + \kappa(\overline{a})$	$= 6$

Table 5.4: OCF κ and $\kappa|A_i$ with $\kappa(ab) = 0, \kappa(a\overline{b}) = 1$ and $\kappa(\overline{a}) = 1$ and the strategic pre-c-revisions from Example 5.3.2 as schemas and with the corresponding calculated ranks $\kappa^{\star}_{\Delta}$ resp. $\oplus((\kappa|A_i)^{\star}_{\Delta_i}) + \kappa(A_i)$. For better readability, we employ different gray tones to distinguish models $\mathrm{Mod}(A_i)$.

Example 5.3.2. *Let* $\Sigma = \{a, b, c, d\}$ *and* $\Delta = \{(c|ab), (d|ab), (\overline{c}|a\overline{b}), (d|a\overline{b}), (c|\overline{a}), (\overline{d}|\overline{a})\}$ *a set of conditionals. It holds that* $\mathcal{P}_{\Delta} = \{ab, a\overline{b}, \overline{a}\}$ *displays the finest case splitting, which induces the following partition of* $\Delta = \Delta_1 \cup \Delta_2 \cup \Delta_3$ *with subsets* $\Delta_1 = \{(c|ab), (d|ab)\}$, $\Delta_2 = \{(\overline{c}|a\overline{b}), (d|a\overline{b})\}$ *and* $\Delta_3 = \{(c|\overline{a}), (\overline{d}|\overline{a})\}$. *The prior ranking function* κ *and the corresponding conditionalized sub-OCFs* $\kappa|A_i$ *for* $i = 1, 2, 3$ *can be found in Table 5.4. To compact the table, we place the conditionalized OCFs beneath each other in one column.*

Furthermore, we display the schematic pre-c-revised ranking functions $\kappa *_{\sigma}^{pre} \Delta$ *and the concatenation* $\oplus(\{\kappa|A_i *_{\sigma}^{pre} \Delta_i + \kappa(A_i)\}_{i=1}^{n})$ *in Table 5.4 with a selection strategy* σ *for c-revision which satisfies (IP-ESP$^{\sigma}$). From Theorem 5.3.1, we can conclude that the constraint satisfaction problem* $CR(\kappa|A_i, \Delta_i)$ *defining the c-revision* $\kappa|A_i *_{\sigma}^{pre} \Delta_i$

$\omega \in \Omega$	$\kappa *_\sigma \Delta$		$\kappa_0 + \oplus((\,\kappa\vert A_i\,)^{\text{pre}}_{\Delta_i}(\omega) + \kappa(A_i))$	
$abcd$	$\kappa_0 + 2$	$= 1$	$\kappa_0 + 2 + \kappa(ab)$	$= 1$
$abc\overline{d}$	$\kappa_0 + 2 + \eta_{1,2}$	$= 2$	$\kappa_0 + 2 + \eta_{1,2} + \kappa(ab)$	$= 2$
$ab\overline{c}d$	$\kappa_0 + 1 + \eta_{1,1}$	$= 2$	$\kappa_0 + 1 + \eta_{1,1} + \kappa(ab)$	$= 2$
$ab\overline{c}\overline{d}$	$\kappa_0 + 0 + \eta_{1,1} + \eta_{1,2}$	$= 2$	$\kappa_0 + 0 + \eta_{1,1} + \eta_{1,2} + \kappa(ab)$	$= 2$
$a\overline{b}cd$	$\kappa_0 + 1 + \eta_{2,1}$	$= 1$	$\kappa_0 + 1 + \eta_{2,1} + \kappa(a\overline{b})$	$= 1$
$a\overline{b}c\overline{d}$	$\kappa_0 + 3 + \eta_{2,1} + \eta_{2,2}$	$= 3$	$\kappa_0 + 3 + \eta_{2,1} + \eta_{2,2} + \kappa(a\overline{b})$	$= 3$
$a\overline{b}\overline{c}d$	$\kappa_0 + 1$	$= 0$	$\kappa_0 + 1 + \kappa(a\overline{b})$	$= 0$
$a\overline{b}\overline{c}\overline{d}$	$\kappa_0 + 2 + \eta_{2,2}$	$= 1$	$\kappa_0 + 2 + \eta_{2,2} + \kappa(a\overline{b})$	$= 1$
$\overline{a}bcd$	$\kappa_0 + 4 + \eta_{3,2}$	$= 4$	$\kappa_0 + 4 + \eta_{3,2} + \kappa(\overline{a})$	$= 4$
$\overline{a}bc\overline{d}$	$\kappa_0 + 3$	$= 2$	$\kappa_0 + 3 + \kappa(\overline{a})$	$= 2$
$\overline{a}b\overline{c}d$	$\kappa_0 + 1 + \eta_{3,1} + \eta_{3,2}$	$= 3$	$\kappa_0 + 1 + \eta_{3,1} + \eta_{3,2} + \kappa(\overline{a})$	$= 3$
$\overline{a}b\overline{c}\overline{d}$	$\kappa_0 + 4 + \eta_{3,1}$	$= 5$	$\kappa_0 + 4 + \eta_{3,1} + \kappa(\overline{a})$	$= 5$
$\overline{a}\overline{b}cd$	$\kappa_0 + 4 + \eta_{3,2}$	$= 4$	$\kappa_0 + 4 + \eta_{3,2} + \kappa(\overline{a})$	$= 4$
$\overline{a}\overline{b}c\overline{d}$	$\kappa_0 + 3$	$= 2$	$\kappa_0 + 3 + \kappa(\overline{a})$	$= 2$
$\overline{a}\overline{b}\overline{c}d$	$\kappa_0 + 2 + \eta_{3,1} + \eta_{3,2} + \kappa(\overline{a})$	$= 4$	$\kappa_0 + 2 + \eta_{3,1} + \eta_{3,2}$	$= 4$
$\overline{a}\overline{b}\overline{c}\overline{d}$	$4 + \eta_{3,1}$	$= 5$	$\kappa_0 + 4 + \eta_{3,1} + \kappa(\overline{a})$	$= 5$
	$\kappa_{0,\Delta}$ =-1		$\kappa_0 =$ -1	

Table 5.5: The c-revised OCF $\kappa *_\sigma \Delta$ and the pre-c-revised, concatenated and then normalized OCF $\kappa_0 + \oplus((\,\kappa\vert A_i\,)^{\text{pre}}_{\Delta_i}(\omega) + \kappa(A_i))$ as schemas and with the corresponding calculated ranks $\kappa^\star_\Delta$ resp. $\kappa_0 + \oplus((\,\kappa\vert A_i\,)^\star_{\Delta_i}) + \kappa(A_i)$.

and the projection of the constraint satisfaction problem $CR(\kappa, \Delta)_{\Delta_i}$ *for* $\kappa *^{pre}_\sigma \Delta$ *are the same. For the impact factors defining* $\kappa *_\sigma \Delta$ *we get the following inequalities.*

$$
\begin{aligned}
(c|ab) &\quad \eta_{1,1} > \min\{2, \eta_{1,2}\} - \min\{1, \eta_{1,2}\} \\
(d|ab) &\quad \eta_{1,2} > \min\{2, 1 + \eta_{1,1}\} - \min\{2, \eta_{1,1}\} \\
(\overline{c}|a\overline{b}) &\quad \eta_{2,1} > \min\{1, 2 + \eta_{2,2}\} - \min\{1, 3 + \eta_{2,2}\} \\
(d|a\overline{b}) &\quad \eta_{2,2} > \min\{1 + \eta_{2,1}, 1\} - \min\{3 + \eta_{2,1}, 2\} \\
(c|\overline{a}) &\quad \eta_{3,1} > \min\{4 + \eta_{3,2}, 3\} - \min\{1 + \eta_{3,2}, 4\} \\
(\overline{d}|\overline{a}) &\quad \eta_{3,2} > \min\{3, 4 + \eta_{3,1}\} - \min\{4, 1 + \eta_{3,1}\}
\end{aligned}
$$

Since σ satisfies (IP-ESP$^\sigma$), we can choose the same impact factors for both revision scenarios $CR(\kappa, \Delta)_{\Delta_i}$ and $CR(\,\kappa|A_i\,, \Delta_i)$.

$$\sigma(\kappa, \Delta)_{\Delta_1} = (\eta_{1,1}, \eta_{1,2}) = (2, 1) = \sigma(\,\kappa|A_1\,, \Delta_1)$$
$$\sigma(\kappa, \Delta)_{\Delta_2} = (\eta_{2,1}, \eta_{2,2}) = (1, 0) = \sigma(\,\kappa|A_2\,, \Delta_2)$$
$$\sigma(\kappa, \Delta)_{\Delta_3} = (\eta_{3,1}, \eta_{3,2}) = (2, 1) = \sigma(\,\kappa|A_3\,, \Delta_3)$$

The results of the revisions are displayed in Table 5.4. Note that we present the schematic revision as well as the calculated ranks corresponding to the special choice of impact factors, indicated by an additional superscript. It holds that $\kappa *^{pre}_{\sigma} \Delta(\omega) = \oplus(\{\,\kappa|A_i\, *^{pre}_{\sigma} \Delta_i(\omega) + \kappa(A_i)\}_{i=1}^{n})$ for all $\omega \in \Omega$.
Moreover, Table 5.5 depicts the schematic c-revision $\kappa *_\sigma \Delta$ and $ocf(\oplus(\{\,\kappa|A_i\, *^{pre}_{\sigma} \Delta_i + \kappa(A_i)\}_{i=1}^{n}))$ employing the same selection strategy σ as above and following the same steps for computing the corresponding c-revision. Again, we indicate the special choice of impact factors by an additional superscript. It holds $\kappa *_\sigma \Delta(\omega) = ocf(\oplus(\{\,\kappa|A_i\, *^{pre}_{\sigma} \Delta_i + \kappa(A_i)\}_{i=1}^{n}))$ for all $\omega \in \Omega$.

To sum up, we have shown that (GRK) offers a divide-and-conquer strategy for c-revisions since it is sufficient to reassemble the local revisions $\kappa|A_i\, * \Delta_i$, given by the case A_i. This yields the same result as revising with the complete set Δ. A concatenation operator executes the assembly of local revisions. Thus, c-revisions solve a merging problem in the context of local vs. global revisions, at least in the special case of (disjoint) cases.

Chapter 6

The Kinematics Principle in the Qualitative Framework

In this chapter, we investigate the qualitative Kinematics principle, which directly corresponds to the Kinematics principle for OCFs from Section 5.1. Note that, for the Kinematics principle for OCFs, the basic concepts needed were already present, i.e., there exists a concept of conditionalization for OCFs by Spohn [119], and c-revisions [63] provide a suitable revision operator for sets of conditionals. Hence, the main challenge in Section 5.1 was to prove that these concepts satisfy the Kinematics principle. Then, we continued with more high-level investigations, such as the relation to propositional revision (cf. Section 5.2) and how to build up the global c-revision via the local revisions (cf. Section 5.3).

However, we have much fewer technical resources in the qualitative framework since we lack a proper conditionalization mechanism and a suitable revision operator for sets of conditionals. Thus, the formulation of the qualitative Kinematics principle in the first section of this chapter represents a target and poses challenges, which we tackle in the following sections. First, we present a concept of qualitative conditionalization in Section 6.2. The second task, defining a suitable revision operator for sets of conditionals, is quite challenging because it goes far beyond established revision theories [2, 29]. However, we have seen that c-revisions comply with the Kinematics principle in the context of OCFs. In Section 6.4, we show that, via induction by c-revisions, we can revise plausibilistic TPOs by sets of conditionals. To define this qualitative version of c-revisions and show that it satisfies the Kinematics principle, an adequate transformation scheme between ranking functions and plausibilistic TPOs is crucial. This transformation scheme must meet high standards. It must adequately transfer the technique of c-revisions to total preorders while respecting crucial characteristics of conditionalization. Thus, the transformation scheme must comply with both mechanisms, c-revisions, and conditionalization, to the full extent. We present a transformation scheme between plausibilistic TPOs and OCFs in Section 6.3, which is intuitive and meets our requirements simultaneously. In Section 6.4, all our previously defined concepts and revision methods come together and enable us to present a proof of concept for the Kinematics principle at least for special cases.

The last section of this chapter is dedicated to investigating the Kinematics principle in the context of another conditional revision operator, namely the one presented by Chandler and Booth in [27]. Chandler and Booth's operator revises with only a single conditional. Thus, we can fall back on some of the results presented in Section 5.2 in the context of OCFs to compare their qualitative conditional operator with qualitative c-revisions, also w.r.t. the Kinematics principle.

Bibliographic Remark. The contents of this part are based on joint work with Gabriele Kern-Isberner and Christoph Beierle [70] (see Section 1.3).

6.1 The Qualitative Kinematics Principle

In this section, we define the qualitative Kinematics principle as a target that serves as guidance and sets the requirements for the investigations in the following sections of this chapter.

For the *qualitative Kinematics principle* Ψ, we assume that each epistemic state Ψ is equipped with a plausibilistic TPO $\preceq_\Psi$ as defined in Section 2.4. Using the direct correspondence between Ψ and $\preceq_\Psi$ we sometimes replace the abstract term Ψ for epistemic states by $\preceq_\Psi$. The revision operator $\bullet$ takes Ψ and a conditional belief base Δ resp. a set of propositional formulas $\mathcal{S}$ as input and delivers a revised epistemic state $\Psi^\bullet = \Psi \bullet (\Delta \cup \mathcal{S})$, i.e., a revised TPO $\preceq_{\Psi^\bullet} = \preceq_{\Psi \bullet (\Delta \cup \mathcal{S})}$ as output.

Definition 6.1.1 (Qualitative Kinematics (QK))**.** *Let* $A_1, \dots, A_n$ *be exhaustive and exclusive formulas. Let* $\Psi = (\Omega, \preceq_\Psi)$ *be an epistemic state, and let* $\Delta = \Delta_1 \cup \ldots \cup \Delta_n$ *be a set of conditionals, with subsets* Δ_i *whose premises imply* A_i*, and* $S = \bigvee_{j \in J} A_j$ *with* $\emptyset \neq J \subseteq \{1, \dots, n\}$*. A revision operator* $\bullet$ *satisfies* Qualitative Kinematics *iff*

$$\textbf{(QK)} \qquad \Psi \bullet (\Delta \cup \{S\})|A_i = (\, \Psi|A_i \,) \bullet \Delta_i$$

Actually, (QK) looks very similar to (GRK), except for replacing the OCF κ with a more general epistemic state Ψ. And the basic idea and the irrelevance assertions of (QK) and (GRK) are the same (cf. the explanation for Definition 5.1.1). It holds that (QK), like (GRK), has two crucial implications concerning the relevance of conditional information under revision. First, the information expressed in S, i.e., that one of the cases A_j ($j \in J$) might be more plausible does not affect the conditional beliefs for each case A_i. Second, for the posterior conditional beliefs given A_i only the respective new information Δ_i is relevant, i.e. conditionalization and revision are interchangeable.

Note that the subsets Δ_i can be empty in the partitioning of Δ. Also, we acknowledge that for the cases represented by the A_i's exclusivity is the crucial property since exhaustiveness can always be obtained by taking the negation of $\bigvee A_i$ as "the remaining case" and choosing the corresponding subset as $\Delta_i = \emptyset$ (cf. Algorithm 1 on page 72).

As for (GRK), it holds that if the plausibility of the cases is neglected, i.e., $S = \top$, a weaker version of (QK) follows for revision operators $\bullet$ that satisfy conditional Tautological Vacuity (cTV) from equation 3.2 on page 64[1]. Then (QK) together with (cTV) implies

$$\Psi \bullet \Delta|A_i = (\, \Psi|A_i \,) \bullet \Delta_i, \tag{6.1}$$

i.e., the revision w.r.t. to the specific case A_i, i.e., we conditionalize the agent's belief state and then revise with the corresponding set Δ_i, yields the same result as revising with the whole set of conditional information Δ and then focussing on the case A_i. We illustrate how (QK) can help us to significantly reduce the complexity of belief revision operators by an example.

[1]We obtain the qualitative version of (3.2) from page 64 by replacing κ by a more general epistemic state Ψ and strategic c-revisions $*_\sigma$ by the more general revision operator $\bullet$.

$\preceq_\Psi$:

$$
\begin{array}{ccccccc}
a\overline{b}\overline{c}\overline{d} & ab\overline{c}d & & abcd & abc\overline{d} & & a\overline{b}\overline{c}d \\
 & & \prec_\Psi & & & \prec_\Psi & \\
\overline{a}bcd & \overline{a}\overline{b}c\overline{d} & & \overline{a}b\overline{c}\overline{d} & \overline{a}\overline{b}\overline{c}\overline{d} & & \overline{a}\overline{b}cd
\end{array}
$$

$$\xrightarrow{\hspace{10cm}}$$

implausibility

Figure 6.1: Belief state Ψ given as TPO $\preceq_\Psi$.

Example 6.1.1. *An agent with the current epistemic state Ψ over the signature $\Sigma = \{a, b, c, d\}$ depicted in Figure 6.1, given as the TPO $\preceq_\Psi$, receives new conditional information concerning two disjoint scenarios $A_1 = a$ and $A_2 = \overline{a}$, s.t. $\Delta = \Delta_1 \cup \Delta_2$ with $\Delta_1 = \{(b|a), (c|a)\}$ and $\Delta_2 = \{(\overline{d}|\overline{a})\}$. We can derive from the minimal worlds of the TPO in Figure 6.1 that the agent is ignorant about the two cases a and $\overline{a}$ and that she deems $b \vee c \vee \overline{d}$ plausible. Note that in Figure 6.1, some worlds are excluded since these are considered to be impossible by the agent.*
This displays a situation where we can (hypothetically) apply the qualitative kinematics principle when incorporating the new conditional beliefs Δ given A_1 resp. A_2 and some additional propositional information about the plausibility of the cases A_1 resp. A_2 (e.g. that $A_1 = a$ is less plausible than $A_2 = \overline{a}$). Then (QK) would guide the revision process by ensuring that the conditional beliefs given A_1 resp. A_2 are influenced only by the respective new conditional information. We would get that

$$
\begin{aligned}
(\Psi \bullet (\Delta \cup \{\overline{a}\}))|a &= \Psi|a \bullet \{(b|a), (c|a)\} \\
(\Psi \bullet (\Delta \cup \{\overline{a}\}))|\overline{a} &= \Psi|\overline{a} \bullet \{(\overline{d}|\overline{a})\}
\end{aligned}
$$

Note that, new information S on A_1 vs. A_2 should be irrelevant for the conditional beliefs given A_i because these conditional beliefs always assume that A_i holds, independent of its plausibility in the revised state.

In Section 5.2, we introduced the postulate (LPR) in the context of OCFs, which relates conditional revision with sets $\Delta = \{(B_i|A_i)\}$ $(i = 1, \ldots, n)$ to propositional revision with the set of consequents $\mathcal{B} = \{B_i\}$. The following postulate displays a reformulation of (LPR) in the qualitative framework

$$\textbf{(LPR)} \quad \Psi|A_i * \Delta_i = \Psi|A_i * \mathcal{B}_i \tag{6.2}$$

If we neglect S and consider the special case of a single conditional $\Delta = \{(B|A)\}$, (LPR) reduces to

$$\Psi|A \bullet (B|A) = \Psi|A \bullet B. \tag{6.3}$$

Thus, making the conditional revision with $(B|A)$ superfluous in the context of conditionalization. Together with (QK) for a single conditional, we get

$$\Psi \bullet (B|A)|A = \Psi|A \bullet (B|A) = \Psi|A \bullet B, \tag{6.4}$$

i.e., we can express the conditional revision of Ψ with $(B|A)$ as a propositional revision with B, and it does not matter whether we first focus on the specific context provided by A and then revise accordingly or vice versa.

6.2 Qualitative Conditionalization

The qualitative Kinematics principle (QK) is a sophisticated axiom of conditional belief revision, which requires a concept of conditionalization for epistemic states. In general, conditionalization can be seen as a method to update an agent's conditional beliefs concerning specific evidence given as a propositional formula $A \in \mathcal{L}$. This concept is well-known in probability theory as conditional probabilities and has been transferred to ranking functions in [117] by Spohn. In this section, we define a qualitative conditionalization operator for epistemic states, represented as plausibilistic TPOs, and propositional formulas, which share important characteristics with Spohn's conditionalization for ranking functions.

There are several well-justified requirements that a conditionalization operation should meet. We illustrate them using ranking functions. First, conditionalization by A should not change the type of an epistemic state, and A should be believed after conditionalization. For any OCF κ, these properties hold for each conditionalized ranking functions $\kappa|A$, since $\kappa|A\,(A) = \kappa(A) - \kappa(A) = 0$, so $\kappa|A$ is an OCF on $\mathrm{Mod}(A)$. Moreover $\kappa|A \models A$ because $\omega \models A$ for all $\omega \in \mathrm{Mod}(A)$. And hence, trivially also for all ω such that $\kappa|A\,(\omega) = 0$. Furthermore, relations among the models of A are preserved: For $\omega_1, \omega_2 \in \mathrm{Mod}(A)$ it holds that

$$\kappa(\omega_1) \leqslant \kappa(\omega_2) \text{ iff } \kappa|A\,(\omega_1) \leqslant \kappa|A\,(\omega_2). \tag{6.5}$$

Finally, conditionalization is compatible with the acceptance of conditional beliefs, meaning that $\kappa \models (B|A)$ iff $\kappa|A \models B$ as was shown in Proposition 5.2.3. Summarizing, what we expect from a qualitative conditionalization $\Psi|A$ are the following properties:

(1) $\Psi|A \models A$;

(2) $\omega_1 \preceq_{\Psi|A} \omega_2$ iff $\omega_1 \preceq_\Psi \omega_2$ for all $\omega_1, \omega_2 \models A$;

(3) $\Psi|A \models B$ iff $\Psi \models (B|A)$.

We now turn to the definition of conditionalization of TPOs.

Definition 6.2.1 (Conditionalization of TPOs)**.** *Let $\preceq_\Psi$ be a TPO on Ω, let A be a propositional formula. The* conditionalization of Ψ on A*, denoted by $\Psi|A$, is defined as $\Psi|A = (Mod(A), \preceq_{\Psi|A})$ such that*

$$\omega_1 \preceq_{\Psi|A} \omega_2 \;\; \textit{iff} \;\; \omega_1 \preceq_\Psi \omega_2 \;\; \textit{for}\; \omega_1, \omega_2 \in Mod(A).$$

First, it is easy to see that conditionalization for TPOs yields a TPO $\preceq_{\Psi|A}$ on $\mathrm{Mod}(A)$. Moreover, this definition complies with all requirements (1) – (3) listed above. Indeed, we have

$$\omega_1 \preceq_\Psi \omega_2 \text{ iff } \omega_1 \preceq_{\Psi|A} \omega_2 \tag{6.6}$$

for $\omega_1, \omega_2 \in \mathrm{Mod}(A)$ by definition. Moreover, $\Psi|A \models A$ for trivial reasons. Last but not least, we compare the acceptance of conditionals with conditionalization for total preorders.

Proposition 6.2.1 (Compatibility with conditionals)**.** *For an epistemic state* $\Psi = (\Omega, \preceq_\Psi)$ *and a conditional* $(B|A)$ *it holds that:*

$$\Psi \models (B|A) \text{ iff } \Psi|A \models B. \tag{6.7}$$

Proof. $\Psi|A \models B$ iff $\min(\mathrm{Mod}(A), \preceq_{\Psi|A}) \models B$. Since $\Psi|A$ is defined on $\mathrm{Mod}(A)$, this is equivalent to $AB \prec_\Psi A\overline{B}$, i.e., to $\Psi \models (B|A)$. □

We continue with Example 3.1.1 to illustrate conditionalization for TPOs.

Example 6.2.1 (Continuing Example 3.1.1)**.** *Let* Ψ *be the TPO defined in Example 3.1.1. If we conditionalize* Ψ *by* $A = a$ *we get the TPO* $\Psi|A : a\bar{b}c \prec abc \prec ab\bar{c}, a\bar{b}\bar{c}$. *We illsustrate the compatibility with conditionals for the conditional* $(a|\overline{b})$*, it holds that* $\Psi \models (a|\overline{b})$ *because* $a\overline{b} \prec ab$ *and* $\Psi|a \models \overline{b}$.

To further elaborate on the relations between conditionalization for TPOs and conditionalization for OCFs, we first need to define operators that define transformations from ranking functions to TPOs and vice versa.

6.3 Transformation Operators for Epistemic States

The qualitative Kinematics principle (QK) makes conditionalization and revision interchangeable, i.e., a revision operator which satisfies this powerful axiom must provide coherence across revision scenarios with prior conditionalized states and the conditionalization of posterior revised state. In the framework of OCFs, c-revisions satisfy this powerful principle, thus, can serve as a blueprint for a qualitative revision operator. If one takes this idea seriously, the requirement for a transformation scheme between plausibilistic TPOs and OCFs arises, which is largely compatible with the conditionalization for ranking functions and TPOs. In this section, we introduce two intuitive and, at the same time, fully compliant transformation operators, at least for special cases. We illustrate their strengths by presenting a commutative diagram in Figure 6.2, which adequately connects the transformation and conditionalization.

Plausibilistic TPOs and OCFs are representations of epistemic states as total preorders indicating plausibility. Thus, both can be seen as qualitative representations of an agent's belief state. TPOs display plausibility relations between worlds by simply preordering them according to their plausibility. Ranking functions also display a preorder of worlds, but they associate worlds with a numerical rank of plausibility, making it easier to distinguish how much more (or less) plausible worlds are. In both formalisms, the most plausible worlds are in the lowest layer. In this section, we introduce two transformation operators between OCFs and TPOs that play a mediating role between the two frameworks and respect the given relations on worlds in each framework. These transformations are shown to respect conditionalization for both representations of epistemic states as far as possible and are useful to define a revision operator for TPOs. We start with the transformation from ranking functions to TPOs.

Definition 6.3.1. *Let κ be a ranking function on Ω. The* transformation operator *τ maps κ to a total preorder $\Psi_\kappa = (\Omega, \preceq_{\Psi_\kappa})$, $\tau : \kappa \mapsto \Psi_\kappa$, such that for all $\omega_1, \omega_2 \in \Omega$*

$$\omega_1 \preceq_{\Psi_\kappa} \omega_2 \text{ iff } \kappa(\omega_1) \leqslant \kappa(\omega_2) \tag{6.8}$$

holds.

τ is a well-defined operator, since we can define a TPO for every ranking function κ by simply transferring the ranking of worlds to a total preorder. Clearly, τ is surjective but not injective, as we see in the following example:

Example 6.3.1. *The ranking functions $\kappa_1, \kappa_2 : \{a, \bar{a}\} \to \mathbb{N} \cup \{\infty\}$ with $\kappa_1(a) = \kappa_2(a) = 0$ and $\kappa_1(\bar{a}) = 1$, $\kappa_2(\bar{a}) = 2$ are mapped to the same TPO by τ, namely $a \prec_{\Psi_{\kappa_i}} \bar{a}$ for $i = 1, 2$.*

In general, two OCFs $\kappa_1, \kappa_2 \in \tau^{-1}(\Psi)$ are equivalent with respect to the TPO on Ω that they induce which can be seen immediately from equation (6.8).[2]

Because τ is not injective, we cannot simply define a transformation from TPOs to ranking functions as the inverse function of τ. We introduce a new operator ρ that maps TPOs to selected ranking functions. Because of the empty layers, it is

[2]Equivalence for OCFs is defined on page 37

possible to match a TPO with an infinite number of ranking functions. We choose a minimal OCF in the sense that we do not allow empty layers.

Definition 6.3.2. *Let $\Psi = (\Omega, \preceq_\Psi)$ be an epistemic state. The* transformation operator *ρ maps Ψ to an OCF κ_Ψ by taking the minimal ranks of the inverse image of τ for Ψ, i.e., ρ is defined by $\rho : \Psi \mapsto \kappa_\Psi$ with*

$$\kappa_\Psi(\omega) = \min_{\kappa \in \tau^{-1}(\Psi)} \{\kappa(\omega)\}. \tag{6.9}$$

To show that ρ is a well-defined operator, we need to prove that κ_Ψ is an OCF. Moreover, we show that

$$\kappa_\Psi(\omega_1) \leqslant \kappa_\Psi(\omega_2) \text{ iff } \omega_1 \preceq_\Psi \omega_2 \tag{6.10}$$

holds, i.e. $\kappa_\Psi \in \tau^{-1}(\Psi)$ and the relations between worlds are maintained.

Proposition 6.3.1. *For all $\Psi = (\Omega, \preceq_\Psi)$ there is a unique minimal OCF $\kappa_\Psi \in \tau^{-1}(\Psi)$ as defined in* (6.9)*; furthermore, κ_Ψ satisfies* (6.10)*.*

Proof. First, we have to show that κ_Ψ is an OCF. Since τ is surjective, the set $\tau^{-1}(\Psi)$ is non-empty, so κ_Ψ is well-defined by (6.9). For $\omega \in \min(\Omega, \preceq_\Psi)$, it holds that $\kappa(\omega) = 0$ for all $\kappa \in \tau^{-1}(\Psi)$, thus $\kappa_\Psi(\omega) = 0$ and κ_Ψ displays an OCF.

Now, we turn to (6.10): For all $\kappa \in \tau^{-1}(\Psi)$ it holds that $\kappa(\omega_1) \leqslant \kappa(\omega_2)$ iff $\omega_1 \preceq_\Psi \omega_2$. This applies particularly if we take the minimal ranks on both sides, i.e.,

$$\min_{\kappa \in \tau^{-1}(\Psi)} \{\kappa(\omega_1)\} \leqslant \min_{\kappa \in \tau^{-1}(\Psi)} \{\kappa(\omega_2)\} \text{ iff } \omega_1 \preceq_\Psi \omega_2,$$

and therefore κ_Ψ satisfies (6.10) resp. (6.8), hence $\kappa_\Psi \in \tau^{-1}(\Psi)$. □

The following lemma collects useful statements about the transformation operators τ and ρ.

Lemma 6.3.2. *Let τ and ρ be as defined in Definitions 6.3.1 and 6.3.2. Then the following statements hold:*

1. *τ is surjective but not injective.*

2. ρ *is injective but not surjective.*

3. $\rho(\Psi) = \kappa_\Psi$ *is convex.*

4. $\tau \circ \rho = id$*, but* $\rho \circ \tau \neq id$ *in general. However,* κ *and* $\rho \circ \tau(\kappa)$ *are equivalent.*

Proof. (1) is clear from the above. Regarding (3), if κ_Ψ was not convex, it would have empty layers, which contradicts the minimum in (6.9). Due to $\kappa_\Psi \in \tau^{-1}(\Psi)$, we have $\tau \circ \rho(\Psi) = \tau(\rho(\Psi)) = \tau(\kappa_\Psi) = \Psi$. However, κ_2 from Example 6.3.1 shows that the converse does not hold because of $\rho \circ \tau(\kappa_2) = \rho(\tau(\kappa_2)) = \rho(\Psi_{\kappa_2}) = \kappa_1 \neq \kappa_2$. This proves (4). Nevertheless, for any two worlds ω_1, ω_2, we have $\kappa(\omega_1) \leqslant \kappa(\omega_2)$ iff $\omega_1 \preceq_{\Psi_\kappa} \omega_2$ (by (6.8)) iff $\kappa_{\Psi_\kappa}(\omega_1) \leqslant \kappa_{\Psi_\kappa}(\omega_2)$ (by (6.10)), and $\kappa_{\Psi_\kappa} = \rho \circ \tau(\kappa)$. This shows that κ and $\rho \circ \tau(\kappa)$ are equivalent. (2) is an easy consequence of (4). □

Since ρ maps each TPO to the unique minimal relation-preserving OCF, it is an injective function. It is clear that κ_Ψ as defined in (6.9) is convex because we choose minimal ranks for $\kappa_\Psi \in \tau^{-1}(\preceq_\Psi)$. While for the composition of our two transformation operators, we generally have $\rho \circ \tau \neq id$ because τ is not injective, for convex ranking functions, $\rho \circ \tau = id$ holds, as the next proposition shows.

Proposition 6.3.3. *Let* κ *be an OCF, and let* ρ *and* τ *be as defined in Definitions 6.3.1 and 6.3.2. If* κ *is convex, then* $\rho \circ \tau(\kappa) = \rho(\tau(\kappa)) = \kappa$ *holds.*

Proof. We set $\tau(\kappa) = \Psi_\kappa$ and $\rho(\Psi_\kappa) = \kappa_{\Psi_\kappa} = \kappa'$. We have to show that $\kappa' = \kappa$. By definition, $\kappa'(\omega) = \rho(\Psi_\kappa)(\omega) = \min_{\kappa'' \in \tau^{-1}(\Psi_\kappa)}\{\kappa''(\omega)\}$, and $\kappa \in \tau^{-1}(\Psi_\kappa)$ according to (6.8). Hence $\kappa'(\omega) \leqslant \kappa(\omega)$ for all $\omega \in \Omega$.

Assume there is ω_0 such that $\kappa'(\omega_0) < \kappa(\omega_0)$, and choose ω_0 with minimal $\kappa'(\omega_0) = r_0$, i.e. $\kappa'(\omega') = \kappa(\omega')$ for all ω' such that $\kappa'(\omega') < \kappa'(\omega_0)$. Since $r_0 \leqslant \max_{\omega \in \Omega} \kappa'(\omega) \leqslant \max_{\omega \in \Omega} \kappa(\omega)$ and κ is presupposed to be convex, it follows that there is ω_1 such that $\kappa(\omega_1) = r_0 < \kappa(\omega_0)$, and therefore $\omega_1 \prec_{\Psi_\kappa} \omega_0$. Because $\kappa' \in \tau^{-1}(\Psi_\kappa)$ holds via Proposition 6.3.1, we also have $\kappa'(\omega_1) < \kappa'(\omega_0)$. But on the other hand, since $\kappa'(\omega_0)$ was chosen minimally, we obtain $\kappa'(\omega_1) = \kappa(\omega_1) = r_0 = \kappa'(\omega_0)$, which yields a contradiction. □

Furthermore, if Ψ is convex with respect to a formula A, the transformation operator ρ preserves the property of convexity for the conditionalized OCF $\rho(\Psi)|A$:

$$\begin{tikzcd} \kappa \arrow[r, "\cdot|A"] \arrow[d, shift left, "\tau"] & \kappa|A \arrow[d, shift left, "\tau"] \\ \Psi \arrow[u, shift left, "\rho"] \arrow[r, "\cdot|A"'] & \Psi|A \arrow[u, shift left, "\rho"] \end{tikzcd}$$

Figure 6.2: Relations between the ranking functions κ resp. $(A|\kappa)$ and the TPOs Ψ resp. $(A|\Psi)$ using the transformation operators τ and ρ.

Proposition 6.3.4. *Let $\Psi = (\Omega, \preceq_\Psi)$ be an epistemic state that is convex with respect to A. Then $\rho(\Psi)|A = \kappa_\Psi|A$ is convex.*

Proof. We have to show that for all $r, 0 \leqslant r \leqslant \max_{\omega \models A}\{\kappa_\Psi|A\,(\omega)\}$, there is $\omega_0 \models A$ such that $\kappa_\Psi|A\,(\omega_0) = r$. For $r \geq 0$, we have $r \leqslant \max_{\omega \models A}\{\kappa_\Psi|A\,(\omega)\} = \max_{\omega \models A}\{\kappa_\Psi(\omega)\} - \kappa_\Psi(A)$ iff $r + \kappa_\Psi(A) \leqslant \max_{\omega \models A}\{\kappa_\Psi(\omega)\}$, which means that $r + \kappa_\Psi(A) \leqslant \max_{\omega \in \Omega}\{\kappa_\Psi(\omega)\}$. By Lemma 6.3.2, we can conclude that κ_Ψ is convex, so there is $\omega_0' \in \Omega$ such that $\kappa_\Psi(\omega_0') = r + \kappa_\Psi(A)$, i.e., $\kappa_\Psi(\omega_0') - \kappa_\Psi(A) = r$.

Choose $\omega_1, \omega_2 \models A$ such that $\kappa_\Psi|A\,(\omega_1) = 0$, and $\kappa_\Psi|A\,(\omega_2) = \max_{\omega \models A}\{\kappa_\Psi|A\,(\omega)\}$ holds, resp., if $r = 0$ or $r = \max_{\omega \models A}\{\kappa_\Psi|A\,(\omega)\}$, then one of ω_1, ω_2 can be chosen as $\omega_0 \models A$ with $\kappa_\Psi|A\,(\omega_0) = r$. So let $0 < r < \max_{\omega \models A}\{\kappa_\Psi|A\,(\omega)\}$. Then we have $\kappa_\Psi(\omega_1) - \kappa_\Psi(A) < \kappa_\Psi(\omega_0') - \kappa_\Psi(A) < \kappa_\Psi(\omega_2) - \kappa_\Psi(A)$, and therefore $\kappa_\Psi(\omega_1) < \kappa_\Psi(\omega_0') < \kappa_\Psi(\omega_2)$. Due to (6.10), this implies $\omega_1 \prec \omega_0' \prec \omega_2$, and because $\omega_1, \omega_2 \models A$ and $\mathrm{Mod}(A)$ is $\preceq$-convex, there is $\omega_0 \models A$ such that $\omega_0 \approx \omega_0'$. Again due to (6.10), we obtain $\kappa_\Psi(\omega_0) = \kappa_\Psi(\omega_0')$, and hence $\kappa_\Psi|A\,(\omega_0) = \kappa_\Psi(\omega_0) - \kappa_\Psi(A) = \kappa_\Psi(\omega_0') - \kappa_\Psi(A) = r$. $\square$

Propositions 6.3.3 and 6.3.4 allow us to study bidirectional translations, at least in special cases. We use this to prove the tight bounds between conditionalization for TPOs vs. conditionalization for OCFs by investigating the commutativity of the diagram displayed in Figure 6.2. We show that conditionalization for TPOs and conditionalization for OCFs can be mapped onto each other by using the transformation operators τ resp. ρ. First, we show that for ranking functions, the diagram commutes in general, i.e., that $\tau(\kappa)|A = \tau(\kappa|A)$ holds.

Theorem 6.3.5. *Let κ be an OCF, $A \in \mathcal{L}$ be a consistent formula, and τ the transformation operator as defined in Definition 6.3.1. Then the following holds:*

$$\tau(\kappa)|A = \Psi_\kappa|A = \Psi_{(\kappa|A)} = \tau(\kappa|A). \tag{6.11}$$

Proof. All $\Psi_\kappa|A, \Psi_{(\kappa|A)}, \kappa|A$ are defined on $\mathrm{Mod}(A)$. We need to show: For $\omega_1, \omega_2 \in \mathrm{Mod}(A)$, it holds that $\omega_1 \preceq_{(\Psi_\kappa)|A} \omega_2$ iff $\omega_1 \preceq_{\Psi_{(\kappa|A)}} \omega_2$.

Let $\omega_1, \omega_2 \in \mathrm{Mod}(A)$, then we obtain $\omega_1 \preceq_{(\Psi_\kappa)|A} \omega_2 \Leftrightarrow \omega_1 \preceq_{\Psi_\kappa} \omega_2 \overset{(6.8)}{\Leftrightarrow} \kappa(\omega_1) \leqslant \kappa(\omega_2) \Leftrightarrow \kappa|A(\omega_1) \leqslant \kappa|A(\omega_2) \overset{(6.8)}{\Leftrightarrow} \omega_1 \preceq_{\Psi_{(\kappa|A)}} \omega_2$. □

Theorem 6.3.5 shows that when starting with a ranking function κ, it does not change the result of the respective qualitative conditionalization whether we first transform κ to a TPO Ψ_κ and then conditionalize, or if we first conditionalize κ and then transform $\kappa|A$ to a TPO $\Psi_{\kappa|A}$. This is because τ and the conditionalization operation preserve relations among possible worlds. Therefore, our Definition 6.2.1 perfectly fits the requirements of OCF conditionalization.

Now, we turn to the transformation operator ρ. Here, however, ρ and conditionalization do not commute in general, i.e., we have $\rho(\Psi)|A = \kappa_\Psi|A \neq \kappa_{\Psi|A} = \rho(\Psi|A)$ generally. This is due to the fact that $\kappa_\Psi|A$ can have empty layers, whereas $\kappa_{\Psi|A}$ is always convex, cf. Lemma 6.3.2. We illustrate this with an example.

Example 6.3.2. *For $\Psi : \overline{a}b \prec ab \prec \overline{a}\overline{b} \prec a\overline{b}$, we have $\rho(\Psi) = \kappa_\Psi = \kappa$ with $\kappa(\overline{a}b) = 0$, $\kappa(ab) = 1$, $\kappa(\overline{a}\overline{b}) = 2$ and $\kappa(a\overline{b}) = 3$. Hence, $\kappa|a(ab) = 0$, and $\kappa|a(a\overline{b}) = 2$. As we can see $\kappa|a$ has an empty layer, i.e., there is no possible world with rank 1. If we first conditionalize Ψ by a, we obtain $\Psi|a : ab \prec a\overline{b}$. It holds that $\rho(\Psi|a) = \kappa_{\Psi|a}$ with $\kappa_{\Psi|a}(ab) = 0$, $\kappa_{\Psi|a}(a\overline{b}) = 1$ is convex, therefore $\kappa_\Psi|a \neq \kappa_{\Psi|a}$.*

Nevertheless, $\kappa_\Psi|A$ and $\kappa_{\Psi|A}$ yield the same result after application of τ:

Theorem 6.3.6. *Let Ψ be a total preorder, $A \in \mathcal{L}$ be a consistent formula, and ρ, τ the transformation operators as defined in Definitions 6.3.1 and 6.3.2. Then the following holds:*

$$\tau(\kappa_\Psi|A) = \tau(\kappa_{\Psi|A}),$$

i.e., $\tau(\rho(\Psi)|A) = \tau(\rho(\Psi|A)) = \Psi|A$.

Proof. All involved TPOs and OCFs are defined on Mod(A). We have to show that for all $\omega_1, \omega_2 \in \mathrm{Mod}(A)$, it holds that $\omega_1 \preceq_{\Psi_{(\kappa_\Psi|A)}} \omega_2$ iff $\omega_1 \preceq_{\Psi_{\kappa_{(\Psi|A)}}} \omega_2$. Let $\omega_1, \omega_2 \in \mathrm{Mod}(A)$. Then we obtain $\omega_1 \preceq_{\Psi_{(\kappa_\Psi|A)}} \omega_2 \overset{(6.8)}{\Leftrightarrow} \kappa_\Psi|A\,(\omega_1) \leqslant \kappa_\Psi|A\,(\omega_2) \overset{(6.5)}{\Leftrightarrow} \kappa_\Psi(\omega_1) \leqslant \kappa_\Psi(\omega_2) \overset{(6.10)}{\Leftrightarrow} \omega_1 \preceq_\Psi \omega_2 \overset{(6.6)}{\Leftrightarrow} \omega_1 \preceq_{\Psi|A} \omega_2 \overset{(6.10)}{\Leftrightarrow} \kappa_{\Psi|A}(\omega_1) \leqslant \kappa_{\Psi|A}(\omega_2) \overset{(6.8)}{\Leftrightarrow} \omega_1 \preceq_{\Psi_{\kappa_{(\Psi|A)}}} \omega_2$. Hence, we can conclude $\tau(\rho(\Psi|A)) = \Psi|A$ by Lemma 6.3.2. □

Theorem 6.3.6 shows that we get the same preorder $\Psi|A$, whether we first conditionalize the original TPO Ψ, then transform it to a ranking function $\kappa_{\Psi|A}$ and then use τ, or if we start with the transformation ρ, then conditionalize and then transform $\kappa_\Psi|A$ to a TPO $\Psi_{\kappa_\Psi|A}$. Although Example 6.3.2 shows that the diagram in Figure 6.2 does not commute with respect to ρ in general, we can heal this flaw by applying τ.

For total preorders which are convex with respect to the proposition A by which they are conditionalized, the diagram in Figure 6.2 commutes also in the ρ-direction:

Theorem 6.3.7. *Let Ψ be a total preorder which is convex with respect to A. Then it holds that*

$$\rho(\Psi)|A = \kappa_\Psi|A = \kappa_{\Psi|A} = \rho(\Psi|A).$$

Proof. Let Ψ be convex with respect to A. All involved TPOs and OCFs are defined on Mod(A). From Proposition 6.3.4, it follows that $\kappa_\Psi|A$ is convex. With Proposition 6.3.3, we obtain $\rho(\tau(\kappa_\Psi|A)) = \kappa_\Psi|A$. Therefore, $\kappa_\Psi|A = \rho(\tau(\kappa_\Psi|A)) \overset{Th.6.3.5}{=} \rho(\tau(\kappa_\Psi)|A) = \rho(\tau(\rho(\Psi))|A) \overset{Lem.6.3.2}{=} \rho(\Psi|A) = \kappa_{\Psi|A}$. □

In the following, we continue with Example 6.1.1 and illustrate the connections which we have worked out in Theorem 6.3.7.

Example 6.3.3 (Continuing Example 6.1.1)**.** *It holds that Ψ given by $\preceq_\Psi$ in Figure 6.1 on page 102 is convex w.r.t. the cases $A_1 = a$ and $A_2 = \overline{a}$ (cf. Definition 3.1.1 on page 61). Also, the agent is agnostic about the cases $A_1 = a$ and $A_2 = \overline{a}$ since she neither believes nor disapproves of one of them. So, to get to know more about what will plausibly happen in the two cases, she considers the conditionalized epistemic states $\Psi|a$ and $\Psi|\overline{a}$ substantiated by the TPOs in Figure 6.3. Both conditionalized epistemic states are derived by Ψ according to Definition 6.2.1.*

$\preceq_\Psi |a$:

$$a\overline{b}\,\overline{c}\,\overline{d} \quad ab\overline{c}d \quad \prec_{\Psi|a} \quad abcd \quad abc\overline{d} \quad \prec_{\Psi|a} \quad a\overline{b}\,\overline{c}d$$

$\preceq_\Psi |\overline{a}$:

$$\overline{a}bcd \quad \overline{a}\,\overline{b}c\overline{d} \quad \prec_{\Psi|\overline{a}} \quad \overline{a}b\overline{c}\,\overline{d} \quad \overline{a}\,\overline{b}\,\overline{c}\,\overline{d} \quad \prec_{\Psi|\overline{a}} \quad \overline{a}\,\overline{b}cd$$

$$\xrightarrow{\hspace{10cm}}$$

implausibility

Figure 6.3: Conditionalized belief states $\Psi|a$ resp. $\Psi|\overline{a}$ given as TPOs $\preceq_\Psi |a$ resp. $\preceq_\Psi |\overline{a}$.

Via the transformation operator ρ from Definition 6.3.2, we get the ranking function κ_Ψ which is depicted in the first column of Table 6.1 by assigning (minimal) ranks to the layers of $\preceq_\Psi$. The OCF κ_Ψ can be conditionalized via (2.10) *and we get $\kappa_\Psi|A_i$ ($i = 1, 2$) from Table 6.1. It holds that, if we transform the conditionalized states $\Psi|a$ and $\Psi|\overline{a}$ to OCFs $\kappa_{\Psi|a}$ and $\kappa_{\Psi|\overline{a}}$ via the operator ρ, we get exactly the same ranking function as $\kappa_\Psi|A_i$ for $i = 1, 2$, i.e.,*

$$\kappa_{\Psi|a}(\omega) = \kappa_\Psi|a\,(\omega) \text{ and } \kappa_{\Psi|\overline{a}}(\omega) = \kappa_\Psi|\overline{a}\,(\omega) \text{ for all } \omega \in \Omega.$$

Thus, as stated in Theorem 6.3.7 for convex TPOs conditionalization and the transformation with operator ρ are interchangeable.

To sum up, in this section, we have shown that conditionalization for TPOs from Definition 6.2.1 is compatible with the conditionalization for ranking functions to a large extent. Of course, we lose information and structure when we move from an OCF to a TPO. In particular, TPOs do not allow for the concept of difference, which is crucial for OCF conditionalization. Nevertheless, Theorems 6.3.5 – 6.3.7 show that we are able to maintain all relevant information for conditionalization at least for worlds in $\mathrm{Mod}(A)$. More precisely, we establish a one-to-one correspondence between TPOs and OCFs (normally, it is one-to-many) that is tight enough to transfer the

$\omega \in \Omega$	κ_Ψ	$\kappa_\Psi \vert A_i$
$a\overline{b}\overline{c}\overline{d}$	0	0
$ab\overline{c}d$	0	0
$abcd$	1	1
$abc\overline{d}$	1	1
$a\overline{b}\overline{c}d$	2	2
$\overline{a}bcd$	0	0
$\overline{a}\overline{b}c\overline{d}$	0	0
$\overline{a}b\overline{c}\overline{d}$	1	1
$\overline{a}\overline{b}\overline{c}\overline{d}$	1	1
$\overline{a}\overline{b}cd$	2	2

Table 6.1: Transformed OCF κ_Ψ and conditionalized OCFs $\kappa_\Psi|A_i$ with $A_1 = a$ and $A_2 = \overline{a}$. To separate worlds from $\mathrm{Mod}(A_i)$ we employ different gray tones.

concept of difference from OCFs to TPOs by making the transfer commutable with conditionalization, at least for special cases. This commutativity is an essential prerequisite for transferring formal revision approaches and properties from OCFs to TPOs, in particular concerning our qualitative Kinematics property.

6.4 Qualitative C-Revisions and the Kinematics Principle

In this section, we combine the results of our previous investigations and define a qualitative revision operator that uses (strategic) c-revisions for OCFs as a blueprint and is compatible with both concepts of conditionalization, i.e., the one presented in Section 6.2 and OCF conditionalization from Spohn [117]. After discussing some basic properties of qualitative c-revisions, we state the main result in this section, which shows that qualitative c-revisions which rely on an impact-preserving selection strategy satisfy (QK) for the case of convex TPOs.

With the help of the transformation operators τ and ρ and their special properties with respect to conditionalization proved in the previous section and illustrated via Figure 6.2, we can now easily transfer conditional revision operators for OCFs to

TPOs.

Definition 6.4.1 (Qualitative c-revision $\bullet^{\mathrm{c}}$, qualitative strategic c-revision $\bullet_\sigma$ for TPOs)**.** *Let $\Psi = (\Omega, \preceq)$ be an epistemic state with $\rho(\Psi) = \kappa_\Psi$ as the corresponding OCF, and let Δ be a set of conditionals. Let $*^{\mathrm{c}}$ be a c-revision operator according to Definition 2.5.1, and let σ be a selection strategy for c-revisions, inducing a strategic c-revision operator $*_\sigma$. A* qualitative c-revision $\bullet^{\mathrm{c}}$ *for Ψ is defined via κ_Ψ as follows:*

$$\Psi \bullet^{\mathrm{c}} \Delta = \tau(\kappa_\Psi *^{\mathrm{c}} \Delta). \tag{6.12}$$

Analogously, a qualitative strategic c-revision $\bullet_\sigma$ *based on strategy σ is defined by*

$$\Psi \bullet_\sigma \Delta = \tau(\kappa_\Psi *_\sigma \Delta). \tag{6.13}$$

Note that c-revisions fully comply with the AGM framework [2] and the Darwiche-Pearl postulates for iterated Belief Revision [29]. More precisely, they satisfy the principle of conditional preservation (PCP)* (cf. Subsection 2.5.2) as was shown in [63]. From this, we can derive that qualitative c-revisions satisfy the qualitative version of the (PCP)*, namely the *Qualitative Principle of Conditional Preservation* (PCP)$^\bullet$ from [65] (cf. Subsection 2.5.2). This was shown in [70] via the following theorem.

Theorem 6.4.1. *Any qualitative c-revision $\bullet^{\mathrm{c}}$ of a TPO $\Psi = (\Omega, \preceq_\Psi)$ by sets of conditionals Δ, i.e., $\Psi \bullet^{\mathrm{c}} \Delta = \tau(\kappa_\Psi *^{\mathrm{c}} \Delta)$, satisfies (PCP)$^\bullet$.*

Proof. Let $\Psi \bullet^{\mathrm{c}} \Delta = \tau(\kappa_\Psi *^{\mathrm{c}} \Delta)$ be a qualitative c-revision operation for TPOs and conditional belief sets $\Delta = \{(B_1|A_1), \ldots, (B_n|A_n)\}$ as defined by (6.12) with a qualitative c-revision operator $\bullet^{\mathrm{c}}$. Let two multisets of possible worlds $\tilde{\Omega}_1 = \{\omega_1, \ldots, \omega_k\}$ and $\tilde{\Omega}_2 = \{\omega'_1, \ldots, \omega'_k\}$ with the same cardinality be given such that $\# \mathit{Ver}_{(B_j|A_j)}(\tilde{\Omega}_1) = \# \mathit{Ver}_{(B_j|A_j)}(\tilde{\Omega}_2)$ and $\# \mathit{Fals}_{(B_j|A_j)}(\tilde{\Omega}_1) = \# \mathit{Fals}_{(B_j|A_j)}(\tilde{\Omega}_2)$ for conditionals $(B_j|A_j) \in \Delta$.

We prove (PCP)$^\bullet$ (1), (PCP)$^\bullet$ (2) is analogous. Presuppose that for all $i, 1 \leqslant i \leqslant k$, it holds that $\omega_i \preceq_\Psi \omega'_i$, and there is at least one $i, 1 \leqslant i \leqslant k$ such that $\omega_i \prec_\Psi \omega'_i$ holds. Then also for κ_Ψ, we have $\kappa_\Psi(\omega_i) \leqslant \kappa_\Psi(\omega'_i)$ for all $i, 1 \leqslant i \leqslant k$, and there is at least one $i, 1 \leqslant i \leqslant k$ such that $\kappa_\Psi(\omega_i) < \kappa_\Psi(\omega'_i)$ holds. This means, that for

$\omega \in \Omega$	κ	$\kappa_\Psi *_\sigma \Delta$	$\kappa^\star_\Delta$
$a\overline{b}\overline{c}\overline{d}$	0	$\kappa_0 + 0 + \eta_{1,1} + \eta_{1,2}$	$= 3$
$ab\overline{c}d$	0	$\kappa_0 + 0 + \eta_{1,2}$	$= 2$
$abcd$	1	$\kappa_0 + 1$	$= 1$
$abc\overline{d}$	1	$\kappa_0 + 1$	$= 1$
$a\overline{b}\overline{c}d$	2	$\kappa_0 + 2 + \eta_{1,1} + \eta_{1,2}$	$= 5$
$\overline{a}bcd$	0	$\kappa_0 + 0 + \eta_{2,1}$	$= 1$
$\overline{a}\overline{b}c\overline{d}$	0	$\kappa_0 + 0$	$= 0$
$\overline{a}b\overline{c}\overline{d}$	1	$\kappa_0 + 1$	$= 1$
$\overline{a}\overline{b}\overline{c}\overline{d}$	1	$\kappa_0 + 1$	$= 1$
$\overline{a}\overline{b}cd$	2	$\kappa_0 + 2 + \eta_{2,1}$	$= 3$
		$\kappa_0 = 0$	

Table 6.2: C-Revision $\kappa_\Psi *_\sigma \Delta$ with Δ of the transformed OCF κ_Ψ. We depict the schematic c-revision and $\kappa^\star_\Delta$ with the calculated ranks which corresponds to Example 6.1.1.

κ_Ψ, the left hand side of (6.12) is negative. Since $\kappa_\Psi *^{\mathrm{c}} \Delta = \kappa^{\mathrm{c}}_\Psi$ is a c-revision and satisfies (PCP)*, also the right hand side of (6.12) referring to κ^{c}_Ψ must be negative, i.e., there must be at least one $j, 1 \leqslant j \leqslant k$, such that $\kappa^{\mathrm{c}}_\Psi(\omega_j) < \kappa^{\mathrm{c}}_\Psi(\omega'_j)$ holds.

Now we have to relate this to $\Psi^{\mathrm{c}} = \Psi \bullet^{\mathrm{c}} \Delta = \tau(\kappa_\Psi *^{\mathrm{c}} \Delta)$. Note that $\rho \circ \tau(\kappa^{\mathrm{c}}_\Psi) = \rho \circ \tau(\kappa_\Psi *^{\mathrm{c}} \Delta) = \rho(\Psi \bullet^{\mathrm{c}} \Delta) = \rho(\Psi^{\mathrm{c}}) = \kappa_{\Psi^{\mathrm{c}}}$, hence $\kappa_{\Psi^{\mathrm{c}}}$ and κ^{c}_Ψ are equivalent according to Lemma 6.3.2. This means that $\kappa^{\mathrm{c}}_\Psi(\omega_j) < \kappa^{\mathrm{c}}_\Psi(\omega'_j)$ iff $\kappa_{\Psi^{\mathrm{c}}}(\omega_j) < \kappa_{\Psi^{\mathrm{c}}}(\omega'_j)$, which is equivalent to $\omega_j \prec_{\Psi^{\mathrm{c}}} \omega'_j$, which was to be shown. □

The Darwiche-Pearl postulates can be derived from (PCP)$^\bullet$, thus, are satisfied by qualitative c-revisions of TPOs.

Note that for this result, we do not have to use strategies because (PCP)$^\bullet$ deals with exactly one revision scenario $\Psi \bullet \Delta$. For (QK), however, results of different revision scenarios have to be compared, so strategies come into play to guide a coherent revision behavior across different scenarios. We apply qualitative strategic c-revisions to our running example 6.1.1.

Example 6.4.1 (Continuing Example 6.1.1)**.** *We qualitatively c-revise the agent's beliefs Ψ represented as a TPO $\preceq_\Psi$ in Figure 6.1 with the conditional beliefs $\Delta =$*

$\preceq_{\Psi^{\bullet^c}}$:

$$\overline{a}\overline{b}c\overline{d} \prec_{\Psi^{\bullet^c}} \begin{array}{cc} \overline{a}bcd & \\ abcd & abc\overline{d} \\ \overline{a}b\overline{c}\overline{d} & \overline{a}\overline{b}\overline{c}\overline{d} \end{array} \prec_{\Psi^{\bullet^c}} ab\overline{c}d \prec_{\Psi^{\bullet^c}} \begin{array}{c} a\overline{b}\overline{c}d \\ \overline{a}\overline{b}cd \end{array} \prec_{\Psi^{\bullet^c}} a\overline{b}\overline{c}d$$

$$\xrightarrow{\hspace{20em}}$$

implausibility

Figure 6.4: Qualitative c-revision $\Psi^{\bullet^c} = \Psi \bullet^c \Delta$ with Δ of belief state Ψ on the basis of the strategic c-revision $\kappa_\Psi *_\sigma \Delta$.

$\Delta_1 \cup \Delta_2$ with $\Delta_1 = \{(b|a), (c|a)\}$ and $\Delta_2 = \{(\overline{d}|\overline{a})\}$. Therefore, we make use of κ_Ψ from Table 6.1. The constraint satisfaction problem $CR(\kappa_\Psi, \Delta)$ yields the following inequalities defining the impact factors for the c-revision:

$$\begin{array}{ll} (b|a) & \eta_{1,1} > 0 - 0 = 0, \\ (c|a) & \eta_{1,2} > 1 - 0 = 1, \\ (\overline{d}|\overline{a}) & \eta_{2,1} > 0 - 0 = 0. \end{array}$$

A Pareto-minimal selection strategy σ chooses the impact vector $\sigma(\kappa_\Psi, \Delta) = (1, 2, 1)$. The result of the strategic c-revision $\kappa_\Psi *_\sigma \Delta$ is depicted in Table 6.2. Thus, we yield the qualitative strategic c-revised TPO $\Psi^{\bullet^c} = \Psi \bullet^c \Delta = \tau(\kappa_\Psi *_\sigma \Delta)$ from Figure 6.4 via the transformation operator τ and the strategic c-revision operator $*_\sigma$.

We show that qualitative c-revision from Definition 6.4.1 that employ a strategy satisfying (IP-ESP$^\sigma$) satisfy (QK) at least for special prior epistemic states Ψ.

Theorem 6.4.2. *Let $\Psi = (\Omega, \preceq_\Psi)$ be an epistemic state, and let σ be a selection strategy for c-revisions that satisfies (IP-ESP$^\sigma$). If Ψ is convex with respect to the cases $A_1, \ldots, A_n$ then $\Psi \bullet_\sigma (\Delta \cup \{S\})$ as defined in* (6.4.1) *satisfies* Qualitative Kinematics (QK)*, i.e.,*

$$(\Psi \bullet_\sigma (\Delta \cup \{S\}))|A_i = (\,\Psi|A_i\,) \bullet_\sigma \Delta_i,$$

where $\Delta = \Delta_1 \cup \ldots \cup \Delta_n$, the case splitting $\mathcal{P}_\Delta = \{A_1, \ldots, A_n\}$, and S are specified as in the preamble of (QK).

Proof. The proof is a straightforward consequence of what we have shown before:

$$\begin{aligned}
(\Psi \bullet_\sigma (\Delta \cup \{S\}))|A_i &= \tau(\kappa_\Psi *_\sigma (\Delta \cup \{S\}))|A_i \\
&\overset{Th. 6.3.5}{=} \tau(\,\kappa_\Psi *_\sigma (\Delta \cup \{S\})|A_i\,) \\
&\overset{Th. 5.1.1}{=} \tau(\,\kappa_\Psi|A_i \,*_\sigma \Delta_i) \\
&\overset{Th. 6.3.7}{=} \tau(\kappa_{\,\Psi|A_i} *_\sigma \Delta_i) = \Psi|A_i \bullet_\sigma \Delta_i.
\end{aligned}$$

□

In this section, all the definitions and results we presented in this chapter come into play. The transformation operator τ defined in Section 6.3 enables us to define qualitative strategic c-revisions from c-revisions for OCFs. Moreover, we have shown in Theorem 6.3.5 that the transformation with τ and conditionalization commute, thus we can employ the Kinematics principle for OCFS. Here, the impact-preserving strategy σ ensures that c-revisions and conditionalization are commutable. Since Ψ is convex w.r.t. to each case A_i, the transformation with ρ and the conditionalization with the corresponding case are compatible, i.e., $\kappa_\Psi|A_i = \kappa_{\,\Psi|A_i}$ holds and we can conclude that (QK) is satisfied. Note that the usage of strategies, i.e., the qualitative c-revision operator $\bullet_\sigma$ is crucial for (QK) since it is crucial for Theorem 5.1.1 which proves that c-revisions that employ a selection strategy which satisfies (IP-ESP$^\sigma$) satisfy (GRK), i.e., (QK) for OCFs.

Furthermore, note that the OCF-variant of the postulate for Local Propositional Revision (LPR) is satisfied by strategic c-revision $*_\sigma$ if σ satisfies (IP-ESP$^\sigma$). Using the transformation schemata (cf. Definition 6.3.2 on page 107 and Definition 6.3.1 on page 106), we can conclude that impact-preserving strategic c-revisions $\bullet_\sigma$ σ satisfy the qualitative version of (LPR) from (6.2) from page 103 at least in the special case of a convex TPO. We state a corollary from Theorem 5.2.2 and Theorem 6.4.2.

Corollary 6.4.3. *Let $\Psi = (\Omega, \preceq_\Psi)$ be an epistemic state and let $\Delta = \Delta_1 \cup \ldots \Delta_n$ with $\Delta_i = \{(B_j|A_i)\}_{1 \leqslant j \leqslant n_i}$ be a set of conditionals and $\mathcal{B}_i = \{B_j\}_{1 \leqslant j \leqslant n_i}$ as specified in Definition 5.2.1 on page 83. If σ satisfies (IP-ESP$^\sigma$) then $\bullet_\sigma$ is a strategic c-revision operator that satisfies (LPR).*

The Kinematics principle is also beneficial for the scenario in Example 6.1.1.

$\omega \in \Omega$	$\kappa_{\Psi\vert\overline{a}}$	$\kappa_{\Psi\vert\overline{a}} *_\sigma \Delta_2$	$\kappa^\star_\Delta$
$\overline{a}bcd$	0	$\kappa_0 + 0 + \eta_{2,1}$	$= 1$
$\overline{a}\overline{b}c\overline{d}$	0	$\kappa_0 + 0$	$= 0$
$\overline{a}b\overline{c}\overline{d}$	1	$\kappa_0 + 1$	$= 1$
$\overline{a}\overline{b}\overline{c}\overline{d}$	1	$\kappa_0 + 1$	$= 1$
$\overline{a}\overline{b}cd$	2	$\kappa_0 + 2 + \eta_{2,1}$	$= 3$
		$\kappa_0 = 0$	

Table 6.3: C-Revision $\kappa_{\Psi|\overline{a}} *_\sigma \Delta_2$ of the transformed OCF $\kappa_{\Psi|\overline{a}}$ with Δ_2. $\kappa^\star_\Delta$ corresponds to the c-revision with $\eta_{2,1} = 1$ of impact factors from Example 6.4.2.

$\preceq_{\Psi|\overline{a} \bullet_\sigma \Delta_2}$:

$$\overline{a}\overline{b}c\overline{d} \prec_{\Psi|\overline{a} \bullet_\sigma \Delta_2} \overline{a}bcd \quad \overline{a}b\overline{c}\overline{d} \quad \overline{a}\overline{b}\overline{c}\overline{d} \prec_{\Psi|\overline{a} \bullet_\sigma \Delta_2} \overline{a}\overline{b}cd$$

$$\xrightarrow[\text{implausibility}]{}$$

Figure 6.5: Qualitative c-revision $\Psi|\overline{a} \bullet_\sigma \Delta_2$ with Δ_2 of belief state $\Psi|\overline{a}$ on the basis of the strategic c-revision $\kappa_\Psi|\overline{a} *_\sigma \Delta_2$.

Example 6.4.2. *The conditional information in Δ concerns two disjoint scenarios, either $A_1 = a$ or $A_2 = \overline{a}$ holds. Now, the agent receives the new information that $\overline{a}$, i.e., the second case, is more plausible than the first one. Here, (QK) allows us to focus on the case of $A_2 = \overline{a}$ by revising solely $\Psi|\overline{a}$ with the new information concerning this case, i.e., Δ_2. For the qualitative strategic c-revision $\Psi|\overline{a} \bullet_\sigma \Delta_2 = \tau(\kappa_{\Psi|\overline{a}} *_\sigma \Delta_2)$, we employ $\Psi|\overline{a}$ from Figure 6.3 and the corresponding ranking function $\kappa_{\Psi|\overline{a}} = \kappa_\Psi|\overline{a}$ from Table 6.1. Let the selection strategy σ from the c-revision in Table 6.2 be impact preserving, i.e., it satisfies (IP-ESP$^\sigma$). Then, it holds for the projection of the constraint satisfaction problem $CR(\kappa_\Psi, \Delta)_{\Delta_2}$ and the constraint satisfaction problem $CR(\kappa_\Psi|\overline{a}, \Delta_2)$ that*

$$CR(\kappa_\Psi, \Delta)_{\Delta_2} = \eta_{2,1} = 1 = CR(\kappa_\Psi|\overline{a}, \Delta_2).$$

We yield the schematics strategic c-revisions $\kappa_{\Psi|A_2} *_\sigma \Delta_2$ resp. $\kappa^\star_\Delta$ with the calculated

rank from Table 6.3 via the conditionalized belief state $\Psi|A_2$ *in 6.3. The revised TPO* $\Psi|\overline{a} \bullet_\sigma \Delta_2$ *is depicted in Figure 6.5. Thus, in case* $A_2 = \overline{a}$ *holds, the agent would expect* $\overline{d}$ *to hold. Also, for* $\Psi \bullet \Delta$ *in Figure 6.4, the conditionalization with the specific case* $A_2 = \overline{a}$ *leads the agent to expect* $\overline{d}$ *to be true.*
Note that the same holds for $CR(\kappa_\Psi, \Delta)_{\Delta_1}$ *and* $CR(\kappa_\Psi|a, \Delta_1)$ *and it holds for the qualitative c-revision strategic c-revisions that* $(\Psi \bullet_\sigma \Delta|A_1 = (\Psi|A_1) \bullet_\sigma \Delta_1$.

6.5 Qualitative C-Revision in Comparison

In the previous section, we have seen that qualitative c-revisions provide a proof of concept for (QK). Now, we relate them to other revision scenarios resp. operators. First, in Subsection 6.5.1, we analyze c-revisions with a single conditional versus c-revisions with the material implication in the qualitative framework. Then, in Subsection 6.5.2, we investigate qualitative c-revision w.r.t. properties of the conditional revision provided by Chandler and Booth in [27]. And conversely, we verify the Kinematics principle for the revision operator from [27].

6.5.1 C-Revision With the Material Implication

In this subsection, we analyze qualitative c-revisions by a single conditional resp. by the material implication. We derive the c-revision with the material implication from our general schema of conditional revision by representing the propositional information $A \to B$ as a conditional with tautological premise $(A \to B|\top)$. Therefore, applying a top-down method for deriving c-revisions with the material implication, starting from the more sophisticated account of c-revision with a single conditional.

Compared to the general revision schema of c-revisions, a c-revision with a single conditional is pretty simple (cf. Section 2.5.3, page 47). From the c-revision with a single conditional in (2.20) and the minimal impact factor η_{m} from (2.22), we obtain the following minimal c-revision:

$$\kappa *^{\mathrm{c}}_{\min} (B|A)(\omega) = \\ -\kappa(\overline{A} \vee B) + \kappa(\omega) + \begin{cases} \kappa(AB) - \kappa(A\overline{B}) + 1, & \omega \models A\overline{B} \\ 0, & \omega \models \overline{A} \vee B \end{cases}, \tag{6.14}$$

s.t., it holds that

$$\eta_{\mathrm{m}} = \kappa(AB) - \kappa(A\overline{B}) + 1. \tag{6.15}$$

From Proposition 2.5.1, it follows that $\kappa_0 = \kappa(\overline{A} \vee B)$.

In Section 2.4, we noted that plausible formulas A could be embedded into conditionals via $(A|\top)$. This embedding enables us to identify the c-revision $\kappa *^{\mathrm{c}} (A \to B)$ as $\kappa *^{\mathrm{c}} (A \to B|\top)$, and we yield the following minimal c-revision $\kappa *^{\mathrm{c}}_{\mathrm{min}} (A \to B)$ via the same argumentation as for (6.14)

$$\kappa *^{\mathrm{c}}_{\mathrm{min}} (A \to B)(\omega) = \\ - \kappa(\overline{A} \vee B) + \kappa(\omega) + \begin{cases} \kappa(\overline{A} \vee B) - \kappa(A\overline{B}) + 1 & \omega \models A\overline{B} \\ 0 & \omega \models \overline{A} \vee B \end{cases}, \tag{6.16}$$

s.t., it holds for the minimal impact factor η'_{m} that

$$\eta'_{\mathrm{m}} = \kappa(\overline{A} \vee B) - \kappa(A\overline{B}) + 1. \tag{6.17}$$

It holds that $\kappa(AB) \geq \kappa(\overline{A} \vee B)$ and therefore $\eta_{\mathrm{m}} \geq \eta'_{\mathrm{m}}$, thus every c-revision by $(B|A)$ is also a c-revision by $A \to B$, but not the other way round. If $\kappa(AB) = \kappa(\overline{A} \vee B)$ holds, then the families of c-revisions by $(B|A)$ resp. $(A \to B)$ coincide.

We have seen that for c-revisions with a single conditional, like (6.14) or (6.16), it is straightforward to choose minimal impact factors according to (2.22). To maintain coherence with our previous results, we introduce a postulate from [70] for selection strategies σ, which chooses the impact factor η according to (2.22).

(Single$^{\sigma}_{min}$) A selection strategy σ is *single-minimal* if for any revision scenario with a single conditional $(B|A)$, we have $\sigma(\kappa, \{(B|A)\}) = \kappa(AB) - \kappa(A\overline{B}) + 1$.

Note that single-minimal c-revisions, i.e., $\kappa *_{\sigma} (B|A)$ resp. $\kappa *_{\sigma} (A \to B)$ with a σ that satisfies (Single$^{\sigma}_{min}$) correspond to the minimal c-revisions in (6.14) resp. (6.16).

Now, we transfer c-revision with a single conditional resp. the material implication to the qualitative framework employing the transformation schema between OCFs and TPOs defined in Section 6.3.

Definition 6.5.1. *Let* $\Psi = (\Omega, \preceq)$ *be an epistemic state with* $\rho(\Psi) = \kappa_\Psi$ *as the corresponding OCF. Let* $*^{\mathrm{c}}$ *be a c-revision operator for OCFs according to Definition 2.5.1, and let* σ *be a selection strategy for c-revisions, inducing a strategic c-revision operator* $*_\sigma$. *A* qualitative c-revision $\bullet^{\mathrm{c}}$ by a single conditional $(B|A)$ resp. by the material implication $(A \to B)$ *for* Ψ *is defined via* κ_Ψ *and* (6.14) *resp. by* (6.16) *as follows:*

$$\Psi \bullet^{\mathrm{c}} (B|A) = \tau(\kappa_\Psi *^{\mathrm{c}} (B|A)), \tag{6.18}$$

$$\Psi \bullet^{\mathrm{c}} (A \to B) = \tau(\kappa_\Psi *^{\mathrm{c}} (A \to B)). \tag{6.19}$$

Analogously, a qualitative strategic c-revision $\bullet_\sigma$ by $(B|A)$ resp. $(A \to B)$ *based on strategy* σ *is defined by*

$$\Psi \bullet_\sigma (B|A) = \tau(\kappa_\Psi *_\sigma (B|A)), \tag{6.20}$$

$$\Psi \bullet_\sigma (A \to B) = \tau(\kappa_\Psi *_\sigma (A \to B)). \tag{6.21}$$

Via the transformation operators τ and ρ, this follows immediately for qualitative c-revisions as defined above.

Example 6.5.1. *In Table 6.4 the single-minimal c-revision of a ranking function* κ *with the material implication* $(a \to b)$, $\kappa^{\mathrm{m}}_{A\to B}$, *resp. with conditional* $(b|a)$, $\kappa^{\mathrm{m}}_{(B|A)}$, *is depicted. For the c-revision with* $a \to b$, *we get the impact factor* $\eta'_{\mathrm{m}} = 2$, *and for the conditional c-revision with* $(b|a)$, *it holds that* $\eta_{\mathrm{m}} = 4$. *Note that, the different impact factors* $\eta_{\mathrm{m}} \neq \eta'_{\mathrm{m}}$ *are because while for* $\kappa *^{\mathrm{c}} (a \to b)$ *the acceptance of* $\overline{a}$ *in the posterior OCF suffices, the c-revision* $\kappa *^{\mathrm{c}} (b|a)$ *makes sure that the ab-worlds are strictly more plausible than the* $a\overline{b}$*-worlds. The normalization constant is* $\kappa_0 = -1$ *for both operations. As we can see, both c-revisions do not impact the ordering of worlds in* $Mod(\overline{a})$, *and the distances between the worlds in* $Mod(a\overline{b})$ *stay the same.*

6.5.2 Conditional Revision by Chandler and Booth

In this section, we investigate the conditional revision operator by Chandler and Booth [27] in terms of the Kinematics principle and qualitative c-revision in terms of properties of the conditional revision operator from [27].

$\omega \in \Omega$	κ	$\kappa *^{\mathrm{c}} (a \to b)$	$\kappa^{\mathrm{m}}_{A \to B}$	$\kappa *^{\mathrm{c}} (b\|a)$	$\kappa^{\mathrm{m}}_{(B\|A)}$
abc	4	$\kappa_0 + 4$	=3	$\kappa_0 + 4$	=3
$ab\overline{c}$	3	$\kappa_0 + 3$	=2	$\kappa_0 + 3$	=2
$a\overline{b}c$	4	$\kappa_0 + 4 + \eta'$	=5	$\kappa_0 + 4 + \eta$	=7
$a\overline{b}\overline{c}$	0	$\kappa_0 + 0 + \eta'$	=1	$\kappa_0 + 0 + \eta$	=3
$\overline{a}bc$	4	$\kappa_0 + 4$	=3	$\kappa_0 + 4$	=3
$\overline{a}b\overline{c}$	3	$\kappa_0 + 3$	=2	$\kappa_0 + 3$	=2
$\overline{a}\overline{b}c$	2	$\kappa_0 + 2$	=1	$\kappa_0 + 2$	=1
$\overline{a}\overline{b}\overline{c}$	1	$\kappa_0 + 1$	=0	$\kappa_0 + 1$	=0

Table 6.4: Single-minimal c-revision with the material implication $(a \to b)$, $\kappa^{\mathrm{m}}_{A \to B}$, resp. with the conditional $(b|a)$, $\kappa^{\mathrm{m}}_{(B|A)}$.

It holds that the conditional revision operator from [27] relies heavily on the revision with the material implication. However, identifying conditionals with the material implication leads to paradox and undesirable behavior [77], which is one of the main reasons why the account of conditionals as three-valued logical entities [30, 1] has become the main framework of conditional logic. However, the revision with the material implication provides some valuable insights about the revision with conditionals. Chandler and Booth's approach relies on fine-tuning the revision with the material implication to receive a conditional revision. Thus, they apply a bottom-up method for defining conditional revision.

To provide ground for our following investigations, we briefly recapitulate the main definitions and postulates from Chandler and Booth's work [27].

The authors in [27] start their investigation of a conditional revision operator by stating the following *Success*-condition for conditional revision:

(S$^\bullet$) $\min(\preceq_{\Psi \bullet (B|A)}, \mathrm{Mod}(A)) \subseteq \mathrm{Mod}(B)$.

In general, revision by a material implication is insufficient to accept the corresponding conditional in the resulting belief set. However, there are special cases where this relation between conditional revision and the revision with the material implication holds. In [27], these cases are subsumed via a *Vacuity* postulate.

(V$^\bullet$) If $\Psi \bullet (A \to B) \models (B|A)$, then $\Psi \bullet (B|A) = \Psi \bullet (A \to B)$

Revision with the material implication $(A \rightarrow B)$ displays the starting point of conditional revision defined by Chandler and Booth in [27]. The main idea is to modify this revised state (minimally) such that $(\mathrm{S}^\bullet)$ hold and some relevant features are retained. The retainment of relevant features of $\Psi \bullet (A \rightarrow B)$ is implemented by the following retainment postulate:

(Ret1$^\bullet$) If $\omega_1, \omega_2 \in \mathrm{Mod}(\overline{A} \vee B)$, then $\omega_1 \preceq_{\Psi \bullet (A \rightarrow B)} \omega_2$ iff $\omega_1 \preceq_{\Psi \bullet (B|A)} \omega_2$

The goal is to perform minimal modifications of $\Psi \bullet (A \rightarrow B)$, s.t. $(\mathrm{S}^\bullet)$ and $(\mathrm{Ret1}^\bullet)$ hold. This is called distance-minimization under constraints in [27]. The retainment principle in $(\mathrm{Ret1}^\bullet)$ claims that the internal ordering in $\mathrm{Mod}(\overline{A} \vee B)$ should be left untouched during the minimization process. Three additional retainment principles complement this principle:

(Ret2$^\bullet$) If $\omega_1, \omega_2 \in \mathrm{Mod}(A\overline{B})$, then $\omega_1 \preceq_{\Psi^\bullet_{(B|A)}} \omega_2$ iff $\omega_1 \preceq_{\Psi^\bullet_{(A \rightarrow B)}} \omega_2$

(Ret3$^\bullet$) If $\omega_1 \in \mathrm{Mod}(\overline{A} \vee B)$, $\omega_2 \in \mathrm{Mod}(A\overline{B})$, and $\omega_1 \prec_{\Psi^\bullet_{(B|A)}} \omega_2$ then $\omega_1 \prec_{\Psi^\bullet_{(A \rightarrow B)}} \omega_2$

(Ret4$^\bullet$) If $\omega_1 \in \mathrm{Mod}(\overline{A} \vee B)$, $\omega_2 \in \mathrm{Mod}(A\overline{B})$, and $\omega_1 \preceq_{\Psi^\bullet_{(B|A)}} \omega_2$ then $\omega_1 \preceq_{\Psi^\bullet_{(A \rightarrow B)}} \omega_2$

The retainment principles imply that the demotion of the plausibility of worlds in $\mathrm{Mod}(A\overline{B})$ in relation to worlds in $\mathrm{Mod}(\overline{A} \vee B)$ are the single admissible transformation when moving from $\Psi \bullet (A \rightarrow B)$ to $\Psi \bullet (B|A)$. Note that $(\mathrm{Ret1}^\bullet)$– $(\mathrm{Ret4}^\bullet)$ compare the conditional revision $\preceq_{\Psi^\bullet_{(B|A)}}$ to the revision $\preceq_{\Psi^\bullet_{(A \rightarrow B)}}$,w.r.t. to worlds that satisfy $A \rightarrow B$, i.e., $\omega \in \mathrm{Mod}(\overline{A} \vee B)$ and worlds that not satisfy $A \rightarrow B$, i.e., $\omega \in \mathrm{Mod}(A\overline{B})$. The first two principles, $(\mathrm{Ret1}^\bullet)$ resp. $(\mathrm{Ret2}^\bullet)$ ensure that the prior ordering of worlds in the set $\mathrm{Mod}(\overline{A} \vee B)$ resp. $\mathrm{Mod}(A\overline{B})$ are kept. The third and fourth postulates, $(\mathrm{Ret3}^\bullet)$ and $(\mathrm{Ret4}^\bullet)$, add the requirement that the upgrade in the plausibility of worlds in $\mathrm{Mod}(\overline{A} \vee B)$ in relation to worlds in $\mathrm{Mod}(A\overline{B})$ is retained when moving from $\Psi \bullet (A \rightarrow B)$ to $\Psi \bullet (B|A)$.

The revision by $(A \rightarrow B)$ is crucial for the conditional revision operator proposed by Chandler and Booth and is executed by one of the elementary revision operators $\bullet$ from Definition 2.3.2 on page 31). Taking the revision $\Psi \bullet (A \rightarrow B)$ as a basis the conditional revision operator $\bullet_{\mathcal{C}}$ from [27] uses lexicographic revision $\bullet^{\ell}$ [81], as in Definition 2.3.2, to define a conditional revision as follows:

Definition 6.5.2 (Conditional revision operator $\bullet_{\mathcal{C}}$, [27])**.** *Let $\bullet$ be an elementary revision operator and let $\bullet^{\ell}$ be the lexicographic revision operator as in Definition 2.3.2. The* conditional revision operator *$\bullet_{\mathcal{C}}$ maps an epistemic state $\Psi = (\preceq_\Psi, \Omega)$ and a conditional $(B|A)$ to an epistemic state $\Psi \bullet_{\mathcal{C}} (B|A) = (\preceq_{\Psi \bullet_{\mathcal{C}}(B|A)}, \Omega)$ via the lexicographic revision of $\Psi \bullet (A \to B)$ with a proposition corresponding to a set of worlds $D(\preceq_{\Psi\bullet(A\to B)}, AB) \cap Mod(\overline{A} \vee B)$:*

$$\Psi \bullet_{\mathcal{C}} (B|A) = (\Psi \bullet (A \to B)) \bullet^{\ell} (D(\preceq_{\Psi\bullet(A\to B)}, AB) \cap Mod(\overline{A} \vee B)). \tag{6.22}$$

where $D(\preceq_{\Psi\bullet(A\to B)}, AB)$ is called the down-set *of all models* $\min(\preceq_\Psi, AB)$ *and is defined by*

$$D(\preceq_{\Psi\bullet(A\to B)}, AB) = \{\omega \in \Omega \,|\, \omega \preceq \omega', \text{ for some } \omega' \in \min(\preceq_\Psi, AB)\}. \tag{6.23}$$

Note that, according to [27], the revision operator $\bullet_{\mathcal{C}}$ returns the minimal TPO that minimizes the distance d_K (Kemeny-distance [61]) to $\Psi \bullet (A \to B)$ given the constraints (S$^\bullet$) and (Ret1$^\bullet$), while satisfying the DP postulates and (Ret2$^\bullet$) – (Ret4$^\bullet$). The minimization is realized via the lexicographic revision of $\Psi \bullet (A \to B)$ by a proposition that corresponds to all models of $(A \to B)$ that are more or equally plausible as the minimal models of AB, i.e., all worlds in Mod$(A\overline{B})$ are shifted upwards. Still, the relations among the worlds in the down set $D(\preceq_{\Psi\bullet(A\to B)}, AB) \cap \mathrm{Mod}(\overline{A} \vee B)$ stay the same. In [27], the authors claim for elementary revision operators $\bullet$ from Definition 2.3.2 on page 31, that the conditional revision operator $\bullet_{\mathcal{C}}$ satisfies the following proposition.

Proposition 6.5.1 ([27])**.** *For $\Psi = (\preceq_\Psi, \Omega)$ an epistemic state and $\bullet$ an elementary revision operator from Definition 2.3.2, it holds that $\bullet$ and $\bullet_{\mathcal{C}}$ satisfy*

$$(\Psi \bullet_{\mathcal{C}} (B|A))|A = (\Psi|A) \bullet B. \tag{6.24}$$

The proposition states that if one excludes worlds where the antecedent of $(B|A)$ is false, the conditional revision results in a propositional revision with the consequent. This result is similar to the special case of (LPR) for revision with a single conditional in (6.4) we discussed on page 103. Note that the original proposition from [27] lacks a conditionalization operator, and the authors employed the intersection with $\mathrm{Mod}(A)$ to define a restriction of the TPO to $\mathrm{Mod}(A)$. Therefore, the postulate (LPR) for a general conditional revision operator, we discussed in Section 6.1 resp. in Section 5.2 for OCFs, provides a more high-level result on the relation between conditional revision and propositional revision going beyond the relation for conditional revision in the sense of [27] in the above-stated proposition. However, Proposition 6.5.1 is useful for investigating (QK) in the context of the conditional revision operator from Definition 6.5.2.

Now, we show that qualitative single-minimal c-revision satisfies the characteristic postulates for the conditional revision operator presented in Definition 6.5.2.

The following proposition shows that qualitative single-minimal c-revisions satisfy ($\mathrm{V}^\bullet$), i.e., if the revision with the material implication leads to the acceptance of the corresponding conditional, then both strategic c-revisions $\kappa \bullet_\sigma (B|A)$ and $\kappa \bullet_\sigma (A \to B)$ coincide.

Proposition 6.5.2. *Let* $\Psi = (\Omega, \preceq)$ *with* $\rho(\Psi) = \kappa_\Psi$ *from Definition 6.3.2. The postulate* ($\mathrm{V}^\bullet$) *holds for the qualitative strategic c-revision operator* $\bullet_\sigma$ *with a single-minimal strategy* σ *that satisfies (*$Single^\sigma_{min}$*).*

Proof. Presuppose that $\Psi^{\bullet_\sigma}_{(A\to B)} = \Psi \bullet_\sigma (A \to B) \models (B|A)$, i.e., $\Psi^{\bullet_\sigma}_{(A\to B)}(AB) < \Psi^{\bullet_\sigma}_{(A\to B)}(A\overline{B})$. Then, from (6.21), we can conclude for $\kappa_\Psi *_\sigma (A \to B)$ that

$$(*) \quad \kappa_\Psi *_\sigma (A \to B)(AB) < \kappa_\Psi *_\sigma (A \to B)(A\overline{B}).$$

Since σ is single-minimal, it holds that $\kappa_\Psi *_\sigma (A \to B)$ is defined as in (6.16) and we abbreviate $\kappa_\Psi *_\sigma (A \to B) = \kappa^{\mathrm{m}}_{A\to B}$.

Now, we have to show that $\kappa^{\mathrm{m}}_{A\to B}(\omega) = \kappa^{\mathrm{m}}_{(B|A)}(\omega)$, where $\kappa^{\mathrm{m}}_{(B|A)} = \kappa_\Psi *_\sigma (B|A)$ with single-minimal selection strategy σ is defined as in (6.14).

We distinguish the following cases:

1. For $\omega \models \overline{A} \vee B$, $\kappa^{\mathrm{m}}_{A\to B}(\omega) = \kappa^{\mathrm{m}}_{(B|A)}(\omega)$ follows from (6.14) and (6.16).

2.For $\omega \models A\overline{B}$, $\kappa^{\mathrm{m}}_{A\to B}(\omega) - \kappa^{\mathrm{m}}_{(B|A)}(\omega) = \kappa_\Psi(\overline{A} \vee B) - \kappa_\Psi(AB)$, i.e., in order to show $\kappa^{\mathrm{m}}_{A\to B}(\omega) = \kappa^{\mathrm{m}}_{(B|A)}(\omega)$, we have to show that $\kappa_\Psi(\overline{A} \vee B) = \kappa_\Psi(AB)$ holds.

"$\leqslant$": Due to (2.8) from page 36, it holds that

$$\kappa_\Psi(\overline{A} \vee B) = \min\{\kappa(AB), \kappa(A\overline{B}), \kappa(\overline{AB})\} \leqslant \kappa_\Psi(AB).$$

"$\geq$": Since $(*)$ holds for $\kappa^{\mathrm{m}}_{A\to B}$ and $\kappa^{\mathrm{m}}_{A\to B}(AB) = -\kappa_\Psi(\overline{A} \vee B) + \kappa_\Psi(AB)$ resp. $\kappa^{\mathrm{m}}_{A\to B}(A\overline{B}) = -\kappa_\Psi(\overline{A} \vee B) + \kappa_\Psi(A\overline{B}) + \kappa_\Psi(\overline{A} \vee B) - \kappa_\Psi(A\overline{B}) + 1 = 1$, we get that

$$\kappa^{\mathrm{m}}_{A\to B}(AB) < \kappa^{\mathrm{m}}_{A\to B}(A\overline{B}) \;\Leftrightarrow\; \kappa_\Psi(AB) < \kappa_\Psi(\overline{A} \vee B) + 1 \;\Rightarrow\; \kappa_\Psi(AB) \leqslant \kappa_\Psi(\overline{A} \vee B),$$

since OCF-ranks $\kappa_\Psi(AB), \kappa_\Psi(\overline{A} \vee B) \in \mathbb{N}$. □

The following proposition shows qualitative c-revisions with $(A \to B)$ resp. with $(B|A)$ satisfy the retainment principles and therefore connect revision with the material implication and the corresponding conditional in a rational manner according to the authors of [27].

Proposition 6.5.3. *Let $\Psi = (\Omega, \preceq)$ with $\rho(\Psi) = \kappa_\Psi$ as defined in* (6.10)*. (Ret1$^\bullet$), (Ret2$^\bullet$), (Ret3$^\bullet$) and (Ret4$^\bullet$) hold for the qualitative strategic c-revision operator $\bullet_\sigma$ with σ that satisfies (Single$^\sigma_{min}$).*

Proof. Since σ is single-minimal, it holds that $\kappa_\Psi *_\sigma (A \to B) = \kappa^{\mathrm{m}}_{A\to B}$ as in (6.16) and $\kappa^{\mathrm{m}}_{(B|A)} = \kappa_\Psi *_\sigma (B|A)$ as in (6.14).
(Ret1$^\bullet$): Let $\omega_1, \omega_2 \in \mathrm{Mod}(\overline{A} \vee B)$. It holds that

$$\begin{aligned}
\omega_1 \preceq_{\Psi\bullet_\sigma(A\to B)} \omega_2 &\overset{(6.8)}{\Leftrightarrow} \kappa^{\mathrm{m}}_{A\to B}(A \to B)(\omega_1) \leqslant \kappa^{\mathrm{m}}_{A\to B}(\omega_2)\\
&\overset{(6.16)}{\Leftrightarrow} -\kappa_\Psi(\overline{A} \vee B) + \kappa_\Psi(\omega_1) \leqslant -\kappa_\Psi(\overline{A} \vee B) + \kappa_\Psi(\omega_2)\\
&\overset{(6.14)}{\Leftrightarrow} \kappa^{\mathrm{m}}_{(B|A)}(\omega_1) \leqslant \kappa^{\mathrm{m}}_{(B|A)}(\omega_2) \overset{(6.8)}{\Leftrightarrow} \omega_1 \preceq_{\Psi\bullet_\sigma(B|A)} \omega_2.
\end{aligned}$$

(Ret2$^\bullet$): Let $\omega_1, \omega_2 \in \mathrm{Mod}(A\overline{B})$. It holds that

$$\begin{aligned}
\omega_1 \preceq_{\Psi\bullet_\sigma(A\to B)} \omega_2 &\overset{(6.8)}{\Leftrightarrow} \kappa^{\mathrm{m}}_{A\to B}(\omega_1) \leqslant \kappa^{\mathrm{m}}_{A\to B}(\omega_2)\\
&\overset{(6.16)}{\Leftrightarrow} -\kappa_\Psi(\overline{A} \vee B) + \kappa_\Psi(\omega_1) + \eta'_{\mathrm{m}} \leqslant -\kappa_\Psi(\overline{A} \vee B) + \kappa_\Psi(\omega_2) + \eta'_{\mathrm{m}}\\
&\Leftrightarrow -\kappa_\Psi(\overline{A} \vee B) + \kappa_\Psi(\omega_1) + \eta_{\mathrm{m}} \leqslant -\kappa_\Psi(\overline{A} \vee B) + \kappa_\Psi(\omega_2) + \eta_{\mathrm{m}}\\
&\overset{(6.14)}{\Leftrightarrow} \kappa_\Psi *_\sigma (B|A)(\omega_1) \leqslant \kappa_\Psi *_\sigma (B|A)(\omega_2) \overset{(6.8)}{\Leftrightarrow} \omega_1 \preceq_{\Psi\bullet_\sigma(B|A)} \omega_2.
\end{aligned}$$

Note that, to shorten the notation we used η_{m} as an abbreviation for the equation in (6.15) and η'_{m} as an abbreviation for (6.17).
(Ret3$^\bullet$): Let $\omega_1 \in \mathrm{Mod}(\overline{A} \vee B)$ and $\omega_2 \in \mathrm{Mod}(A\overline{B})$. It holds that

$$\begin{aligned}\omega_1 \prec_{\Psi \bullet_\sigma (A \to B)} \omega_2 &\overset{(6.8)}{\Leftrightarrow} \kappa_\Psi *_\sigma (A \to B)(\omega_1) < \kappa_\Psi *_\sigma (A \to B)(\omega_2) \\ &\overset{(6.16)}{\Leftrightarrow} -\kappa_\Psi(\overline{A} \vee B) + \kappa_\Psi(\omega_1) < -\kappa_\Psi(\overline{A} \vee B) + \kappa_\Psi(\omega_2) + \eta'_{\mathrm{m}}\end{aligned}$$

with single-minimal impact factor η'_{m} from (6.17), it holds that $\eta_{\mathrm{m}} \geq \eta'_{\mathrm{m}}$ for η_{m} from (6.15), since $\kappa(AB) \geq \kappa(\overline{A} \vee B)$ and we get that

$$\begin{aligned}&- \kappa_\Psi(\overline{A} \vee B) + \kappa_\Psi(\omega_1) < -\kappa_\Psi(\overline{A} \vee B) + \kappa_\Psi(\omega_2) + \eta'_{\mathrm{m}} \\ \Rightarrow\quad &-\kappa_\Psi(\overline{A} \vee B) + \kappa_\Psi(\omega_1) < -\kappa_\Psi(\overline{A} \vee B) + \kappa_\Psi(\omega_2) + \eta_{\mathrm{m}} \\ \overset{(6.14)}{\Leftrightarrow}\quad &\kappa_\Psi *_\sigma (B|A)(\omega_1) < \kappa_\Psi *_\sigma (B|A)(\omega_2) \overset{(6.8)}{\Leftrightarrow} \omega_1 \prec_{\Psi \bullet_\sigma (B|A)} \omega_2.\end{aligned}$$

(Ret4$^\bullet$) follows analogously to (Ret3$^\bullet$). □

The following example illustrates qualitative c-revision and conditional revision in the sense of [27].

Example 6.5.2. *In Figure 6.6 on page 129 a conditional revision of an epistemic state Ψ is depicted, this example was taken from [27]. On the right-hand side, the qualitative c-revision of Ψ with $(a \to b)$ and $(b|a)$ can be found, which is derived from the c-revision in Table 6.4. The numbers correspond to possible worlds; these are ordered from bottom to top, with the minimal worlds on the lowest level. As we can see, we get different results from $\Psi \bullet_{\mathcal{C}} (b|a)$ and $\Psi \bullet (b|a)$, because the $a\bar{b}\bar{c}$-world is inserted differently. But the relations among the worlds in Mod$(a \to b)$ are kept.*

Using Proposition (6.5.1), we can show that $\bullet_{\mathcal{C}}$ satisfies (QK).

Theorem 6.5.4. *Let $\Psi = (\preceq_\Psi, \Omega)$ be an epistemic state and $\bullet_{\mathcal{C}}$ be the conditional revision operator as defined in Definition 6.5.2. Then $\Psi \bullet_{\mathcal{C}} (B|A)$ satisfies (QK), i.e.*

$$(\Psi \bullet_{\mathcal{C}} (B|A))|A = (\Psi|A) \bullet_{\mathcal{C}} (B|A) \tag{6.25}$$

Proof. It is sufficient to show that $(\Psi|A) \bullet_{\mathcal{C}} (B|A) = (\Psi|A) \bullet B$, (QK) then follows immediately from (6.24). That $(\Psi|A) \bullet_{\mathcal{C}} (B|A) = (\Psi|A) \bullet B$ is shown in the proof of Proposition 10 in [27] via using (Ret1$^\bullet$), the DP postulates and the concept of *state restriction* that states that: If $\omega \in \min(\preceq_{\Psi|A}, \mathrm{Mod}(B))$ and $\omega' \in \mathrm{Mod}(\overline{B})$, then $\omega \preceq_{(\Psi|A)\bullet B} \omega'$, otherwise $\omega \preceq_{(\Psi|A)\bullet B} \omega'$ iff $\omega \preceq_{\Psi\bullet B} \omega'$. $\square$

Note that in (6.25), we omitted the term S in the Qualitative Kinematics principle, since for the revision with a single conditional, it holds that $S = A$, $S = \overline{A}$ or $S = \top$. In each of these cases, the revision with S is obsolete either due to the vacuity postulate ($S = \top$) or due to conditionalization with A resp. $\overline{A}$ for $S = A$ resp. $S = \overline{A}$. Note that, in Theorem 6.5.4, we neglected the prerequisite in (QK) concerning the convexity of Ψ w.r.t. the cases. This advantage arises from the fact that $\bullet_{\mathcal{C}}$ focuses on the revision with a single conditional in the qualitative context. However, this restriction to the special case of $\Delta = \{(B|A)\}$ in (QK) ultimately leads to the loss of one of the main strengths of (QK), namely, that (QK) enables us to revise with information from disjoint contexts simultaneously with maximal efficiency exploiting the disjoint cases $A_1, \ldots, A_n$.

The following example illustrates the result from Theorem 6.5.4, i.e., that the conditional revision operator from [27] satisfies (QK).

Example 6.5.3. *In Figure 6.7, the conditional revisions* $(\Psi \bullet_{\mathcal{C}} (b|a))|a$ *and* $(\Psi|a) \bullet_{\mathcal{C}} (b|a)$ *are depicted. For both operations we yield the same results and therefore (QK) holds. This example illustrates that for conditional revision with a single conditional, conditionalization with the premise and revision are interchangeable.*

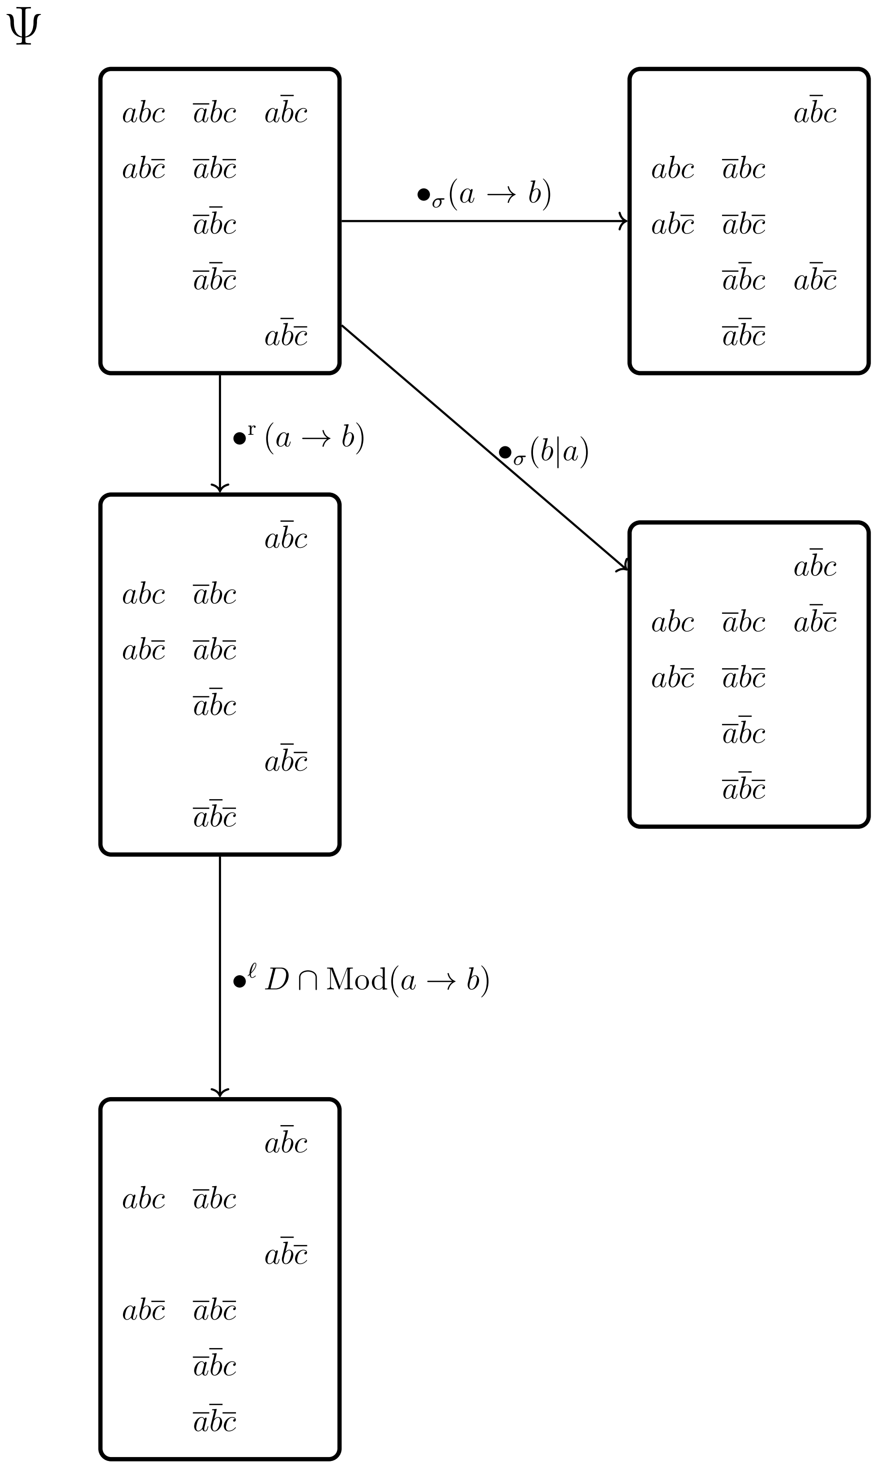

Figure 6.6: On the left side, we depict the conditional revision $\Psi \bullet_{\mathcal{C}} (b|a)$ with the restrained revision operator $\bullet^{\mathrm{r}}$ as elementary operator for the revision $\Psi \bullet^{\mathrm{r}} (a \to b)$. On the right side, we depict the qualitative c-revisions of $\Psi \bullet^{\mathrm{c}} (a \to b)$ and $\Psi \bullet^{\mathrm{c}} (b|a)$.

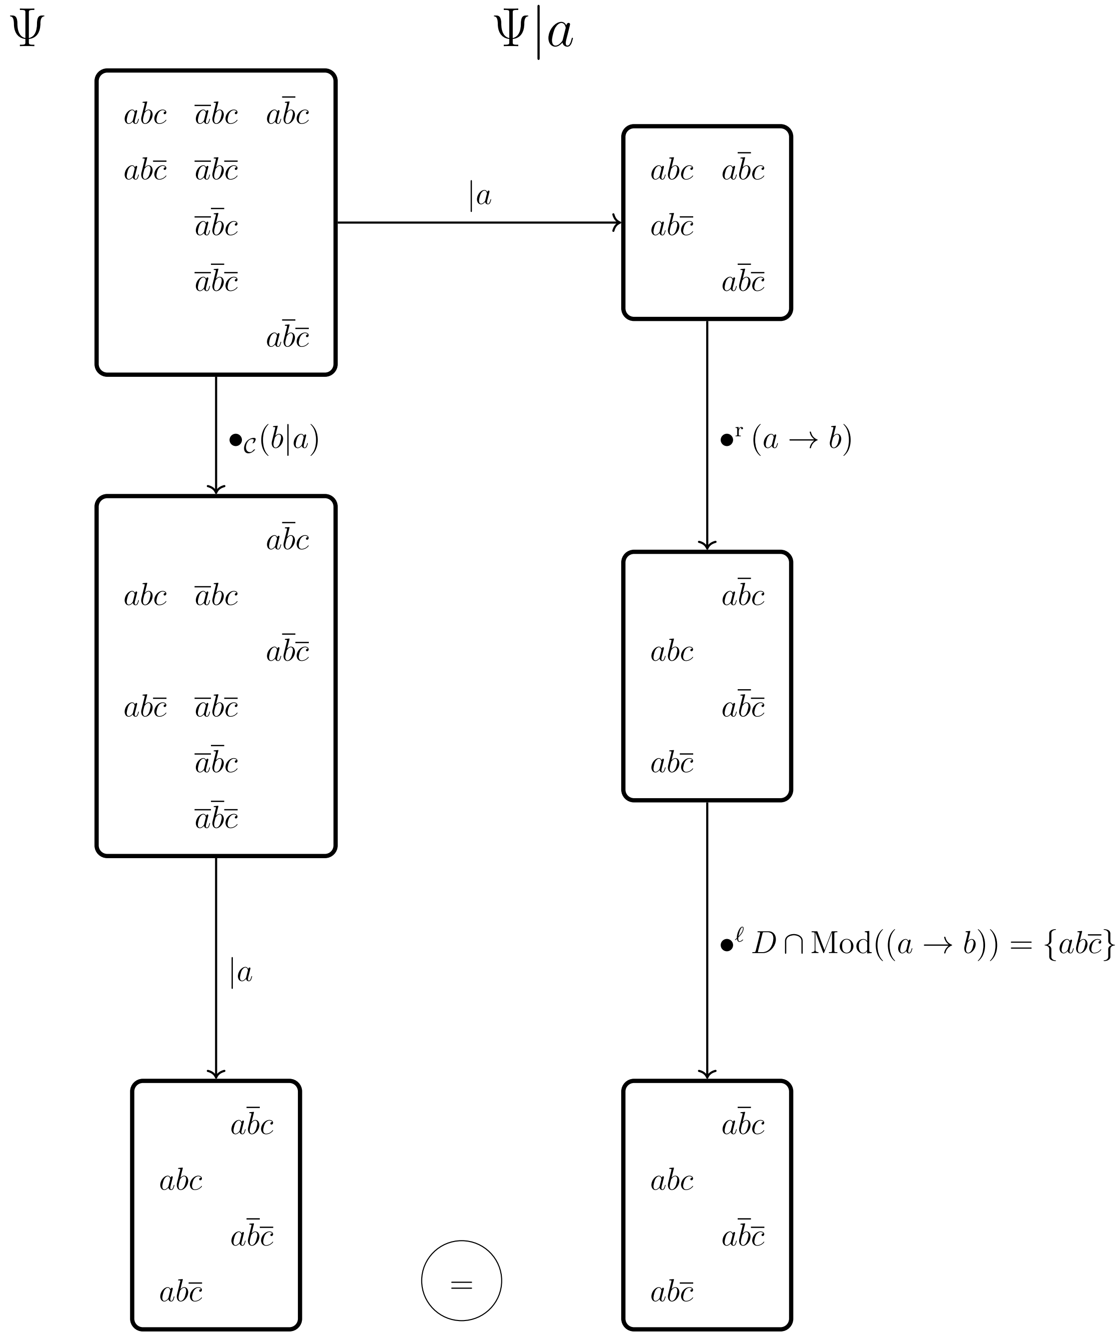

Figure 6.7: (QK) for conditional Revision Operator $\bullet_{\mathcal{C}}$ from [27]. It holds that $(\Psi \bullet_{\mathcal{C}} (b|a))|a = (\Psi|a) \bullet_{\mathcal{C}} (b|a)$.

Intermediate Summary for Part I

In this part, we thoroughly investigated a Kinematics principle for Belief Revision inspired by the locality notions implemented in Bayesian networks [88] and probabilistic reasoning [57, 113]. At the beginning of this part, we were confronted with the advanced belief revision problem (CondCS), which raises the question of how to revise an agent's epistemic state in the light of new information which decomposes naturally into disjoint contexts, called cases. Our Kinematics principle provides an answer to this question, which fits naturally into how human reasoners revise their beliefs and states that, when confronted with information from different cases, the new information should be relevant solely in parts of the belief state that correspond to the specific case. Conditionals provide us with the meta-structure needed to express this notion of locality in Belief Revision. They allow us to set up local context by connecting the information from the consequent to a premise, introducing a specific context in a natural way. On the other hand, conditionalization is crucial since it allows us to focus on specific parts of the belief state. Consequently, our Kinematics principle takes full advantage of conditionals in Belief Revision, and the thorough investigations in this part lead to meaningful insights, which we summarize in the following.

The Kinematics Principle for Ranking Functions. We started with a thorough investigation of the cases employed in (CondCS) and presented an algorithm that enables us to exploit the notion of locality on which our Kinematics principle is based to the full extent. Then we presented and analyzed the Kinematics principle in the context of OCFs in Section 5.1. We showed that c-revisions satisfy it. On the technical side, the Kinematics principle states that conditionalization and revision are interchangeable; this holds for c-revisions and the OCF-conditionalization presented by Spohn [119].

We elaborated the Kinematics principle and derived the principle of Local Propositional Revision (LPR), which rationally connects propositional and conditional revision via considering a particular case of the prerequisites of the Kinematics principle in Section 5.2. Ultimately, this led to a conditional variant of the well-known Ramsey Test. In Section 5.3, we solved the merging problem for c-revisions, at least in the unique setting of the Kinematics principle, via providing a concatenation operator that allows us to set up a global c-revision using local c-revision w.r.t. the specific cases.

The Kinematics Principle in the Qualitative Framework. While in the context of OCFs, sophisticated tools such as c-revisions and a conditionalization operator are available, the Kinematics principle in the qualitative framework remains a target. However, it provides a strong motivation for defining these concepts. We presented a qualitative conditionalization operator in Section 6.2, which is compatible with OCF-conditionalization. However, a qualitative revision operator for sets of conditionals goes far beyond the current state of the art in Belief Revision. Via powerful yet natural transformation operators for OCFs resp. TPOs presented in Section 6.3, we were able to transfer c-revisions to the qualitative framework in Section 6.4, which required a high level of compatibility down to the smallest level of detail. Lastly, in Section 6.5, we analyze qualitative c-revision with a single conditional resp. with the material implication, and confront them with the conditional revision operator of Chandler and Booth [27]. This leads to the insight that Chandler and Booth employ a bottom-up strategy for conditional revision starting from the revision of Ψ by the corresponding material implication $A \to B$ and subsequently fine-adjusting. In contrast, qualitative c-revisions follow a top-down approach that is self-contained.

Part II

Parameterized Belief Change

Chapter 7

Introduction to Part II

The vivid research area of Belief Revision [2, 29] investigates how an agent incorporates new information that may be inconsistent with their current epistemic state. During the revision, inconsistencies are cleared out, yet meta-information accompanying the new input, e.g., regarding reliability, is not considered. Extensive studies in psychology (see, e.g., [125, 116, 93]) have shown that the somewhat naive acceptance of any input information does not correspond to how individuals revise their beliefs in real-life settings. In fact, expertise, reliability, and other factors impact how we incorporate new information (or even if we incorporate it). Realistic revision models should provide the necessary means to process this meta-information. The idea of parameterized belief revision presented in this part is to take another (potential) belief as a point of reference that influences whether resp. to which extent the new belief should be incorporated into the agent's posterior belief state. We investigate the following parameterized belief revision problem:

(ParameterRev) Let Ψ be an epistemic state and β be new input information. We presuppose that β does not come isolated but is accompanied by some meta-information which can be expressed by a so-called reference sentence α. How should Ψ be rationally revised by β such that the parameter α is taken into account to yield a posterior state $\Psi \bullet_\alpha \beta$, and how is the belief change mechanism affected by the plausibility relations between α and β?

Our aim in this section is to exploit the internal strengths of conditionals to express the parameterized information in the setting of (ParamRev) as input for a conditional revision operator. Thereby, illustrating how conditionals enable us to implement meta-information accompanying propositional input as specific context.

Several approaches to belief revision incorporate parameters indicating trust or reliance towards the input information, see e.g., [56, 3]. The work in this part relies mainly on *Revision by Comparison* (RbC) [34] and *Bounded Revision* (BR) [100], firstly introduced by Eduardo Fermé and Hans Rott. Both mechanisms implement the idea that the input β comes together with a reference sentence α which guides the revision process. Despite its strengths, the formal implementation of both revision mechanisms is hard to understand, and the change mechanisms are not intuitively apparent. Together with the revision problem (ParameterRev), this leads to the following research questions that we aim to answer for both parameterized revision operators, Revision by Comparison, and Bounded Revision in the course of this part:

- *How can we (elegantly) reformulate each parameterized belief revision operator's change mechanism to clarify which worlds are affected by the change mechanism?*

- *Is it possible to simplify the parameterized revision rule while maintaining significant features and, therefore, the distinct character of Revision by Comparison resp. Bounded Revision?*

- *Can we define each revision mechanism as a revision with a (set of) conditionals so that the formerly parameterized information in the reference sentence and its corresponding role in the revision process is fully captured by a (directly accessible) input information?*

- *What are the similarities and differences between Revision by Comparison and Bounded Revision? How do these show in their corresponding conditional revision?*

In the following section, we present formal preliminaries for the results in this part. In Section 7.2, we summarize relevant related work, starting with a general overview of parameterized belief revision without claim to completeness. Then in Subsection 7.2.2, we recall basic definitions and notions of Revision by Comparison from [34], and in Subsection 7.2.3, we do the same for Bounded Revision from [100]. Our research and the results presented in this part are structured in consecutive organized chapters:

Chapter 8 We investigate the non-prioritized belief change mechanism of Revision by Comparison (RbC) in the framework of plausibilistic TPOs and make its underlying strategy more explicit in Section 8.1. Here, we introduce three simple yet elegant postulates which fully capture the change mechanism underlying RbC. Then we investigate the hybrid belief change character of RbC as an operation between revision with the input sentence and the contraction of the reference sentence in Section 8.2. In Section 8.3, we provide a realization of RbC in the ranking function framework and show that RbC corresponds to a special c-revision with a set of weak conditionals.

Chapter 9 In Section 9.1, we recapture the basics of Bounded Revision, firstly introduced by Rott in [100] as an iterated parameterized belief change mechanism, and present an elegant reformulation using TPOs and a representation theorem, therefore clarifying the strategy underlying BR. Following the line of Rott [100], we show that BR can be expressed as a lexicographic revision with a designated formula in Section 9.2 and investigate essential limiting cases. Then, in Section 9.3, we present a methodological implementation of BR in the framework of OCFs via c-revisions, i.e. implementing the parameterized change mechanism of BR as a conditional revision.

Chapter 10 After an intermediate summary of the previously presented results on RbC resp. BR, we compare the parameterized belief change mechanism RbC and BR for plausibilistic TPOs and OCFs. We closely examine similarities and differences concerning the change mechanism, the postulates satisfied by each revision operator, and the corresponding operators for OCFs. An example illustrates the main results of this comparison.

Bibliographic Remark. The contents of this part are based on joint work with Gabriele Kern-Isberner [109, 110] (see Section 1.3).

7.1 Formal Preliminaries For This Part

We introduce some notations and formal preliminaries relevant solely to this part of the thesis.

First, we briefly overview the well-known belief representation via a *system of spheres* in the style of Grove [42].

Definition 7.1.1 (System of spheres (SOS), [42]). *Let $\$$ be a finite, non-empty set of subsets $S \subseteq \Omega$, centered on some $\bigcap \$$. We call $\$$ a* system of spheres *(SOS) if it satisfies the following conditions:*

(S1) *$\$$ is totally ordered by set-inclusion $\subseteq$, i.e., $S, S' \in \$$, then $S \subseteq S'$ or $S' \subseteq S$*

(S2) *$\bigcap \$$ is the $\subseteq$-minimum if $\$$, i.e., $\bigcap \$ \in \$$ and if $S \in \$$ then $\bigcap \$ \subseteq S$*

(S3) *$\Omega \in \$$, i.e., Ω is the largest element in $\$$*

(S4) *If $A \in \mathcal{L}$ and there is a sphere in $\$$ intersecting $Mod(A)$, then there is a smallest sphere in $\$$ intersecting $Mod(A)$, i.e., there is a sphere S, s.t. $S \cap Mod(A) \neq \emptyset$ and $S' \cap Mod(A) \neq \emptyset$ implies $S \subseteq S'$ for all $S' \in \$$*

Grove's SOS model an agent's belief state via nested spheres (in terms of set inclusion) representing the intuition that the smallest sphere contains the most plausible worlds, the second innermost sphere (without the innermost one) contains the second most plausible worlds, and so on[1]. Thus, the belief set of an SOS $Bel(\$) = \bigcap \$$ is defined as the center of $\$$. Then an SOS is equivalent to a plausibilistic TPO [42, 91, 34] when we identify the layers of $\preceq$ with the newly added worlds compared to the previous spheres $S \in \$$, starting with the innermost one as the $\preceq$-minimal worlds.

Epistemic entrenchment (EE) was introduced by Gärdenfors in [43, 39] as binary relation $\leqslant_E$ ordering formulas in an agent's belief set. It holds that degrees of entrenchments are measured qualitatively, i.e., for two sentences $A, B \in \mathcal{L}$, the notation $A \leqslant_E B$ stands for "B is at least as epistemically entrenched as A" and $A <_E B$ is defined as $A \leqslant_E B$ but not $B \leqslant_E A$. In [39], Gärdenfors stated that

[1]Note that, Lewis [77] also defined an SOS to provide semantics for counterfactual logics. His definition differs from Grove's, e.g.; he defined an individual sphere for each world in Ω.

"the degree of epistemic entrenchment has a bearing on what is abandoned from a knowledge set and what is retained when a contraction or revision is carried out." So, in general, *standard EE relations* in the sense of [39] mirror an agent's attitude towards her current beliefs, i.e. represent an inner ordering of formulas in the belief set. They were used in many formal developments for belief change [39, 79, 97, 34]. Standard Epistemic Entrenchment plays a crucial role in AGM contractions, and Gärdenfors and Makinson proved in [39] a representation theorem that states that every AGM contraction can be generated from a standard EE relation and vice versa. The *Levi-Identity* [74, 2]

$$K{\star}A = (K{\div}\overline{A}){+}A$$

which defines AGM revision with A in terms of an AGM contraction with the contrary information $\overline{A}$, followed by an AGM expansion with A. Thus, EE relation can also be used for AGM revision, and in [39], it is proved that for standard EE relations generating an AGM contraction, we get an AGM revision operator via the Levi-Identity. Now that we clarified the role of EE relations in the sense of Gärdenfors and Makinson [39], we turn to a slightly generalized version firstly defined by Nayak in [79], which we call EE relations from here onwards.

Definition 7.1.2 ([79]). *A total preorder $\leqslant_E$ over $\mathcal{L}$ is called a relation of* epistemic entrenchment*, if it satisfies the following conditions:*

(E1)	*If $A \leqslant_E B$ and $B \leqslant_E C$, then $A \leqslant_E C$*	*(Transitivity)*
(E2)	*If $A \models B$, then $A \leqslant_E B$*	*(Dominance)*
(E3)	*$A \leqslant_E A \wedge B$ or $B \leqslant_E A \wedge B$*	*(Conjunctiveness)*
(E4)	*If $A \leqslant_E B$ for all $A \in \mathcal{L}$, then $B \equiv \top$*	*(Maximality)*

The original definition of Gärdenfors and Makinson in [39] consists of conditions (E1) – (E4) plus an additional minimality constraint that states that all non-beliefs are minimally entrenched. This definition is too restrictive, especially in the light of iterated belief change [79, 97, 49], and for our purposes, the Nayak version of EE relations is more useful[2]. Standard entrenchment relations from [39] are given

[2]In [46] Gärdenfors and Makinson themselves supported the idea of generalized EE relations

relative to a belief set. However, belief sets can be extracted from each epistemic entrenchment as they contain enough information by themselves (as noted in [79]) via

$$Bel(\leqslant_E) = \begin{cases} \{A \in \mathcal{L} \mid \bot <_E A\}, & \text{if } \bot <_E A \text{ for some } A \\ \mathcal{L} & \text{otherwise} \end{cases}.$$

Note that this belief set it deductively closed, and $Bel(\leqslant_E)$ collapses to $\mathcal{L}$ for absurd EE relations [79]. Also, EE relations $\leqslant_E$ from [79] unlike standard EE relations, do not assume maximally entrenched beliefs to be immutable logical truths[3].

So far, we have seen that EE-relations are TPOs that display a preference ordering of formulas in an agent's belief set. When changing an agent's beliefs, they take on the perspective of belief contraction since they guide which beliefs an agent should give up more easily than others. The qualitative framework most prominent in this thesis, the plausibilistic TPOs, on the other hand, are more focused on belief revision because they do not specify the preference relation of worlds within their corresponding belief set; the worlds constituting the belief set all are maximally plausible, but rather order worlds outside of the belief set. In this way, a plausibilistic TPO is more focused on guiding which worlds should become part of the posterior belief set during the revision process. Plausibilistic TPOs and epistemic entrenchment relations are dual approaches to belief representation and (via the Levi-Identity) both formalisms are fundamental to AGM Belief Revision. In the following, we present a proposition that summarizes the relationship between entrenchment relations and plausibility orderings, which was already discussed in [91, 42] and will prove to be helpful in the context of this paper:

Proposition 7.1.1 ([91])**.** *For each epistemic entrenchment relation $\leqslant_E$ with belief set $Bel(\leqslant_E)$ over $\mathcal{L}$,*

$$A \leqslant_E B \text{ iff } \overline{A} \preceq \overline{B} \tag{7.1}$$

defines a faithful TPO $\preceq$ with $Bel(\preceq)$ over $\mathcal{L}$ s.t. $Bel(\leqslant_E) = Bel(\preceq)$ and vice versa.

Intuitively, this means that if A is more entrenched than B, then the minimal models of $\overline{A}$ are more plausible than the minimal models of $\overline{B}$. Via Proposition

called *expectation orderings*, which are more valuable than the classical EE relations in the light of Belief Revision.

[3]For an intuitive explanation as to why this makes sense, see [79].

(7.1.1), we can lift the EE-relation on formulas to a plausibilistic TPO $\preceq$ over Ω. Therefore, we first have to note that, in general, each EE-relation is uniquely defined via the entrenchment classification of maximal disjunctions over the signature Σ, which constitutes $\mathcal{L}$. This holds because each formula in $A \in \mathcal{L}$ can be represented via a canonical conjunctive normal form

$$A \equiv A_1 \wedge \ldots \wedge A_n$$

with maximal disjunctions A_i, and from the properties defining an EE relation, it follows that

$$A \equiv A_1 \wedge \ldots \wedge A_n =_E \min_{i=1,\ldots,n} A_i, \tag{7.2}$$

for each EE relation $\leqslant_E$ (cf. [39, 79]), i.e., A is equally entrenched as the least entrenched maximal disjunction A_i. Applying (7.1) to the EE-relation over the maximal disjunctions leads to a plausibilistic TPO over complete conjunctions of all variables in the underlying signature. By a slight abuse of notation, we identify these complete conjunctions as possible worlds (cf. (2.1) on page 18). Since this applies especially to minimal worlds, we can conclude that $Bel(\leqslant_E) = Bel(\preceq)$ also holds for plausibilistic TPOs. This way, Proposition 7.1.1 allows us to lift the EE-relation over $\mathcal{L}$ to a plausibilistic TPO over Ω.

7.2 Related Work

In this section, we discuss the issue of parameterized belief revision in a more extensive research context and discuss related work. Since, in the context of this thesis, Revision by Comparison resp. Bounded Revision displays the most important works on parameterized belief revision; we discuss the original works [34] from Fermé and Rott for RbC resp. [100] from Rott for BR more thoroughly in Subsection 7.2.2 resp. Subsection 7.2.3.

7.2.1 Supplementary Information for Belief Change Operators

The basic idea of parameterized belief revision operators is to consider additional information when it comes to the revision of an input sentence. This rather general idea has already been discussed in various facets in the past. There have been several approaches to belief revision, with input information accompanied by some supplementary information.

For example, Spohn defines in [117] a revision operator on ranking functions where the input information β is accompanied by a degree of plausibility k with which β is to be accepted in the posterior state. Recently, there has been much research on revision operators that rely on some notion of trust as additional information [56, 3, 12, 124]. Booth and Hunter discuss in [12] how trust, represented as a collection of sets of possible worlds, in an agent's (domain-specific) expertise should influence the revision process. They handle this additional information as a precursor to belief revision, which relativizes the input formula. In [56], an approach of trust-influenced revision is presented for OCFs, where trust is defined in terms of a distance function between states.

The two main frameworks of parameterized revision we consider in this thesis, Revision by Comparison [34] and Bounded Revision [100], employ other beliefs as points of reference in order to compare the relative plausibility of the new information toward the existing belief state. Fermé and Rott motivated the usage of parameterized information consisting of an input and a reference sentence via the following example:

Example 7.2.1 ([34])**.** *Suppose a colleague tells us ´Lisa is negotiating with Candy company'. Should we accept this piece of information? Most belief revision mechanisms tell us yes, but also ask us to fix to which extent this new information should be accepted. What we can do now is ask our friend how sure she is of this piece of information. She might say that it is at least as well-confirmed as the claim that Lisa has got an offer from Healthy Company. Another way of obtaining the same sort of comparative information is by juxtaposing our assessments of the reliability of sources. If our colleague is at least as trustworthy and well-informed as another person who has testified to the truth about the offer from Healthy Company the other* *day, we will wish to accept the information about Lisa's negotiation at a level of certainty that is at least as high as our degree of belief in the offer from Healthy Company. If the latter belief, however, is not firm enough to overcome our doubts about Lisa's negotiations with Candy Company, we are likely to end up doubting both pieces of information.*

All in all, the choice of the reference might be determined by the input information's context (e.g., reliability or preferences) or, in a more technical or application-oriented environment, by the user's choice.

Before we present an overview of the results presented in [34] for Revision by Comparison in Subsection 7.2.2, we want to point briefly to some related work specific to Revision by Comparison. Furthermore, Revision by Comparison displays a non-prioritized revision mechanism, i.e., the acceptance of the input information is not guaranteed, and the kind of change depends on the interplay between input and reference sentence. This dynamic moves it in the vicinity of *credibility-limited revision operators* (CL revision operators) [54, 11] where a revision on a belief state is solely performed if the input is part of some set of credible formulas $\mathcal{C}$ otherwise the original belief state is kept. RbC is appealing as the basis for CL revision operators since it allows for a flexible and reasonable determination of $\mathcal{C}$ via choosing a suitable reference sentence.

7.2.2 Revision by Comparison by Fermé and Rott

In the following, we outline the fundamental approach and characteristics of Revision by Comparison (RbC), as described in [34]. Initially, we examine RbC as a revision method for Grovean SOS [42], and subsequently for EE-relations, as defined in [79]. We aim to review the key features and findings of RbC from [34], which serve as the foundation for our research in the subsequent chapter.

In [34], Fermé and Rott proposed a Revision by Comparison operator $\circledcirc_\alpha \beta$ as a new model for reasoning without numbers which takes as input two propositional sentences. The input sentence β displays the new information, and α represents additional meta-information. Fermé and Rott summarize the goal of RbC as follows:

"Accept β with a degree of plausibility that at least equals that of α"

In [34], each belief set is expected to be equipped with a representation of the corresponding belief state. Note that, in [34], the term plausibility is generally used to express rather a preference than a plausibility ordering over worlds.

In [34], the authors formally define Revision by Comparison in two qualitative frameworks of belief representation. To define the semantics of RbC in Section 2 of [34], the authors refer to a system of spheres in the style of Grove [42] (cf. Definition 7.1.1 on page 138).

Definition 7.2.1 (RbC for SOS, [34])**.** *Let $\$$ be an SOS, and $\$^{\circledcirc}_{\alpha,\beta}$ be the posterior SOS after RbC with α as a reference and β as the input sentence. Then $\$ \circledcirc_\alpha \beta = \$^{\circledcirc}_{\alpha,\beta}$ is an SOS is defined by*

$$\$^{\circledcirc}_{\alpha,\beta} = \{S \cup Mod(\beta) : S \in \$ \textit{ and } S \subseteq Mod(\alpha)\} \cup \{S : S \in \$ \textit{ and } S \not\subseteq Mod(\alpha)\}$$

Via this definition, the posterior SOS $\$^{\circledcirc}_{\alpha,\beta}$, generated by the RbC operator $\circledcirc_\alpha \beta$, is obtained by shifting the $\overline{\beta}$ worlds closer to the prior belief set than the closest $\overline{\alpha}$ worlds outwards up to the ring where the closest $\overline{\alpha}$-worlds reside. As we can see, this semantic recipe is not easily visible from the definition of $\$^{\circledcirc}_{\alpha,\beta}$. Therefore, the above-stated definition is accompanied by some observations for the possible worlds reading of SOS, which remain without proof and are solely stated as observations. Later, we clarify the semantics of RbC using plausibilistic TPOs in Section 8.1.

In [34], the authors consider special cases of SOS that do not contain the set of all possible worlds Ω and define RbC of those SOS as special cases. We exclude these particular class of SOSs and follow Grove's definition of SOS [42, 33].

While the authors used possible world representations of belief states, such as SOS, as motivation for RbC, they turned to epistemic entrenchment relations in the following investigations. The entrenchment relations in [34] satisfy Definition 7.1.2 [79], except that they omit the Maximality-condition (E4). This is because, in [34] revisions with *irrevocable sentences*, i.e., sentences that are no less entrenched than $\top$ are considered. We exclude such sentences, following Gärdenfors' intuition that no sentence $A \in \mathcal{L}$ shall be more entrenched than logical truths and assume that each entrenchment relation $\leqslant_E$ we consider satisfies Definition 7.1.2.[4] We extract

[4]Note that this does not affect the following results, since we exclude special cases that corre-

belief states from belief sets represented as epistemic entrenchment relations in the usual way and recall the definition of RbC of β w.r.t. α on epistemic entrenchment relations from [34]. Note that this definition is equivalent to the definition of RbC for SOS as was shown in [34] via a straightforward translation between EE-relations and SOSs presented in [34].

Definition 7.2.2 (Revision by Comparison for $\leqslant_E$, [34])**.** *Let $\leqslant_E$ be an epistemic entrenchment relation and $\leqslant_E^{\circledcirc_{\alpha,\beta}}$ be the posterior entrenchment relation after RbC with α as a reference and β as the input sentence. Then $\leqslant_E \circledcirc_\alpha \beta = \leqslant_E^{\circledcirc_{\alpha,\beta}}$ is defined by*

$$\gamma \leqslant_E^{\circledcirc_{\alpha,\beta}} \delta \text{ iff } \begin{cases} \alpha \wedge (\beta \Rightarrow \gamma) \leqslant_E (\beta \Rightarrow \delta), & \text{if } \gamma \leqslant_E \alpha \\ \gamma \leqslant_E \delta, & \text{otherwise} \end{cases} \tag{7.3}$$

for any arbitrary sentences $\gamma, \delta \in \mathcal{L}$.

Definition 7.2.2 is the "official definition" on which Revision by Comparison is based [34]. Fermé and Rott formulated in [34] properties that hold for $\leqslant_E^{\circledcirc_{\alpha,\beta}}$ to clarify what the operation changes in the prior epistemic entrenchment. For the posterior entrenchment relation after RbC of β w.r.t. α, $\leqslant_E^{\circledcirc_{\alpha,\beta}}$, it holds that $\leqslant_E^{\circledcirc_{\alpha,\beta}}$ satisfies the following properties

(RbC)$_E$ Input β is at least as entrenched as α, i.e., $\alpha \leqslant_E^{\circledcirc_{\alpha,\beta}} \beta$

(MinRbC)$_E$ Input β is not additionally lifted, i.e., $\alpha <_E \beta$ iff $\alpha <_E^{\circledcirc_{\alpha,\beta}} \beta$

(α-level)$_E$ For any $\gamma \in \mathcal{L}$, it holds that $\alpha <_E \gamma$ iff $\alpha <_E^{\circledcirc_{\alpha,\beta}} \gamma$

(α-relation)$_E$ For any $\gamma, \delta \in \mathcal{L}$, if $\alpha \leqslant_E \gamma$ and $\alpha \leqslant_E \delta$, then $\gamma \leqslant_E \delta$ iff $\gamma \leqslant_E^{\circledcirc_{\alpha,\beta}} \delta$ holds

(β-level)$_E$ For any $\gamma \in \mathcal{L}$, it holds that ($\gamma \leqslant_E \alpha$ or $\gamma \leqslant_E \beta$) iff $\gamma \leqslant_E^{\circledcirc_{\alpha,\beta}} \beta$

The property (RbC)$_E$ corresponds to the main idea of Revision by Comparison. And together with (MinRbC)$_E$, it states that β shall be at least as deeply entrenched

spond to the RbC operation where either the input α or the reference β is irrevocable. For more information on RbC with irrevocable inputs, see [34].

as α, but not further, except it was already more deeply entrenched in the prior relation $\leqslant_E$. Because of (α-level)$_E$, it holds for each sentence γ that is more deeply entrenched than α, that RbC does not change the prior EE-relation. From (α-relation)$_E$, it follows that RbC of β w.r.t. α does not change the entrenchment relations among worlds more or equally deep entrenched than α. Thus, properties (α-level)$_E$ and (α-relation)$_E$ state that for all worlds more or equally entrenched than α, RbC of β w.r.t. α does not change the prior ordering. From (β-level)$_E$, we can conclude that $\gamma \in \mathcal{L}$ is less or equally entrenched than the input β in the posterior ordering if and only if it was already less or equally entrenched than α or β in the prior ordering.

In [34], Fermé and Rott noted that the posterior entrenchment ordering $\leqslant_E^{\circledcirc_{\alpha,\beta}}$ crucially depends on the prior relation between input sentence β and reference sentence α. They distinguished four (exhaustive but not exclusive) cases that lead to different kinds of changes. We briefly summarize their results, which were formulated using epistemic entrenchment relations.

The intended case. If $\beta <_E \alpha$ and $\overline{\beta} <_E \alpha$, then (7.3) cannot be simplified, and $Bel(\leqslant_E \circledcirc_\alpha \beta) = Bel(\leqslant_E) \star \beta$, i.e., in this cases RbC coincides with an AGM revision with β [34]. Later, we call this case the β*-revision*.

The vacuous case. If $\alpha \leqslant_E \beta$ holds, i.e., (RbC)$_E$ is already satisfied in the prior EE-relation, then nothing changes, and it holds for the belief sets that $Bel(\leqslant_E \circledcirc_\alpha \beta) = Bel(\leqslant_E)$.

The unsuccessful case. If $\alpha \leqslant_E \overline{\beta}$, then (7.3) can be simplified to ($\gamma \leqslant_E^{\circledcirc_{\alpha,\beta}} \delta$ iff $\gamma \leqslant_E \alpha$ or $\gamma \leqslant_E \delta$) and, we obtain $Bel(\leqslant_E \circledcirc_\alpha \beta) = \{\gamma \in \mathcal{L} \mid \alpha <_E \gamma\}$. This does not correspond to an AGM contraction of α.

The epistemic collapse. This case occurs if α and $\overline{\beta}$ are irrevocable, i.e., $\top \leqslant_E \alpha, \overline{\beta}$ and leads to an inconsistent belief set. This case is not relevant in the context of this thesis since we excluded sentences that are more or equally entrenched than $\top$.

We discuss the mechanism behind the first three cases and the resulting properties more thoroughly in the context of plausibilistic TPOs in Section 8.2.

7.2.3 Bounded Revision by Rott

Now, we state the basic methodology and properties of Bounded Revision (BR) as it was presented in [100]. First, as an operation on Grovean SOS [42] and then via EE-relations in the sense of [79]. We recapture properties of BR and relevant results from [100] on which we base our investigations in Chapter 9.

In [100], Rott started the formal investigation of BR in the framework of system of spheres \$ in the style of Grove [42] (cf. Definition 7.1.1 on page 138) by stating a definition of the posterior SOS $\$ \circ_\alpha \beta = \$^\circ_{\alpha,\beta}$ resulting from a BR of β w.r.t. α.

Definition 7.2.3 (BR for SOS, [34]). *Let \$ be an SOS and $\$^\circ_{\alpha,\beta}$ be the posterior SOS after BR with α as reference and β as input sentence. Then $\$ \circ_\alpha \beta = \$^\circ_{\alpha,\beta}$ is an SOS is defined by*

$$\begin{aligned}\$^\circ_{\alpha,\beta} =& \{S \cap Mod(\beta) : S \in \$, S \cap Mod(\beta) \neq \emptyset \text{ and } S \subseteq S_{\alpha,\beta}\} \\ &\cup \{S \cup (S_{\alpha,\beta} \cap Mod(\alpha)) : S \in \$\}\end{aligned}$$

with $S_{\alpha,\beta}$ the smallest sphere S in \$, s.t. $S \cap Mod(\beta) \not\subseteq Mod(\alpha)$; if there is no such sphere, take $S_{\alpha,\beta}$ to be the largest sphere.

Rott explains the intuition behind $\$^\circ_{\alpha,\beta}$ that the best β-worlds are moved to the center, as long as α holds and even a little longer. The main idea is expressed in [100] as follows:

> "Accept β as long as α holds along with β, and just a little further."

Even though the possible worlds semantic in SOS is appealing, the recipe for BR in Defintion 7.2.3 is fairly complicated and most of the formal developments of BR in [100] are based on the following definition of RbC for epistemic entrenchement relations in the style of [79].

Definition 7.2.4 ([100]). *Let $\leqslant_E$ be an entrenchment relation and $\alpha, \beta \in \mathcal{L}$. The* Bounded Revision by β w.r.t. α of $\leqslant_E$ is an entrenchment relation $\leqslant_E \circ_\alpha \beta = \leqslant_E^{\circ_{\alpha,\beta}}$ *defined as follows for any arbitrary sentences $\gamma, \delta \in \mathcal{L}$:*

$$\gamma \leqslant_E^{\circ_{\alpha,\beta}} \delta \textit{ iff } \begin{cases} (\beta \Rightarrow \gamma) \leqslant_E (\beta \Rightarrow \delta), & \textit{if } \beta \Rightarrow (\gamma \wedge \delta) \leqslant_E (\beta \Rightarrow \alpha) \\ \gamma \leqslant_E \delta, & \textit{otherwise} \end{cases} \tag{7.4}$$

It holds that $\leqslant_E^{\circ_{\alpha,\beta}}$ is a well-defined entrenchment relation and $\circ_\alpha \beta$ satisfies the DP postulates [100][5]. Rott already mentioned in [100] that the meaning of this recipe for BR is not easy to understand. In general, it holds that BR is motivated by the same concerns as Revision by Comparison from [34], combined with the desire to preserve conditional beliefs in a DP manner presented in Section 2.3. Accordingly, the reference sentence α functions for BR as a measure of how firmly entrenched β should be in the agent's posterior belief state. Yet, it is not clear what "just a little further" means. In a way, α serves as a bound for the acceptance of β. But how exactly does the entrenchment of β depend on the entrenchment of α? And how are the posterior belief set and the reference sentence related to each other? These intuitions do not become clear from Definitions 7.2.3 and 7.2.4. In the course of the following sections, we present more simple yet elegant postulates defining BR in analogy to (7.4) in the context of TPOs, which provide grounds for further investigations.

We continue with general results about BR stated in [100] that are relevant in the context of this thesis. From Definition 7.2.4 Rott observed that for the reference sentence $\alpha \equiv \bot$, we get the same result as for the natural revision [15] with input β and for $\alpha \equiv \top$, we get a lexicographic revision [79] with β. Thus, for the limiting cases in which α is either never or always true, we can reconstruct two well-known iterated belief revision operators. A definition of these operators for plausibilistic TPOs is given in Definition 2.3.2 on page 31. In Section 9.2, following the line of Rott [100], we discuss these limiting cases more thoroughly and prove that it is also possible to reconstruct natural resp. lexicographic revision from our definition of BR for plausibilistic TPO. For now, we continue by stating two properties that are crucial for BR, which we recapture from [100] employing entrenchment relations.

The informal goal of BR is partly formalized by the success condition for BR:

(BR)$_E$ β is strictly more entrenched than α: $\alpha <_E^{\circ_{\alpha,\beta}} \beta$

Note that, from Definition 7.2.4 and (BR)$_E$ it remains unclear how much more plausible the input β shall be in the posterior ordering, i.e., the question of how the entrenchment of α and β are related to one another is not fully answered.

[5] In [100] a reformulation of the DP postulates for parameterized revision is given. More on this in Chapter 10.

We have already mentioned that the BR operator $\circ_\alpha \beta$ satisfies the DP postulates for iterated Belief Revision with β. This holds regardless of the choice of α [100]. Furthermore, it holds that the resulting belief set $Bel(\leqslant_E^{\circ_{\alpha,\beta}})$ is insensitive to the choice of the reference sentence α. Rott calls this property the *Same Beliefs Condition (SBC)$_E$*.

(SBC)$_E$ $\quad Bel(\leqslant_E \circ_\alpha \beta) = Bel(\leqslant_E \circ_\gamma \beta)$ for any $\alpha, \gamma \in \mathcal{L}$

Note that, since BR displays above all an iterated belief revision with β, the standard success condition for belief revision operators $\beta \in Bel(\leqslant_E^{\circ_{\alpha,\beta}})$ holds. BR implements the idea that β is accepted in the revised state independent from α. The role of the reference sentence is limited to determining how firmly β is entrenched in the posterior state. Clarifying the role of α and the interplay between reference and input sentence in BR for TPOs on possible worlds is the goal of the investigations in Chapter 9.

Chapter 8

Revision by Comparison

Revision by Comparison (RbC), firstly presented in [34] by Fermé and Rott, is motivated by the idea to take propositional sentences α as points of references when revising with β which influence the acceptance of β in the posterior state. So, RbC $\circledcirc_\alpha \beta$ displays a belief revision mechanism for epistemic states Ψ that takes as input two propositional pieces of information, a reference sentence α and an input sentence β, and maps them onto a revised epistemic state $\Psi \circledcirc_\alpha \beta$. The intuitive idea behind RbC now suggests that the level of preference of input sentence β in the posterior belief state is constrained via a designated reference sentence α.

Even though the basic idea might sound simple, it is not a priori clear how α influences the acceptance of β. RbC, as presented in [34], actually heavily depends on the interplay between α and β in the initial belief state. So much that not even the acceptance of the input is guaranteed in general, and there frequently occur cases where instead of revising with β the agent ends up giving up his former belief in α. RbC is a non-prioritized belief change mechanism with a versatile character and interesting dynamics. This chapter aims to transfer the parameterized input of RbC to the object level employing conditionals and combine this reformulation with c-revisions, leading to a simple yet elegant realization of parameterized belief change that fully captures the mechanism of RbC. Realizing RbC as a c-revision with weak conditionals provides us with new insights into the change mechanism and demonstrate the underlying characterization via an iterated contraction operator. On the way to achieving this goal, among other things, we transfer RbC to the framework of plausibilistic TPOs via a representation theorem, clarify the role of α, and further investigate the hybrid belief change character of RbC.

The following sections of this chapter are organized as follows: In Section 8.1, we make the strategy underlying RbC more explicit using plausibilistic TPOs and present a representation theorem that elegantly captures the change mechanism of RbC via three postulates. The hybrid belief change character implemented by RbC, illustrating the operations versatility, is discussed in the context of plausibilistic TPOs in Section 8.2. The mechanism and its intuitive strengths are transferred to the semi-quantitative framework of ranking functions in Section 8.3, leading to an elegant (methodological) implementation. Finally, we present an implementation of RbC as a c-revision with weak conditionals, allowing us to transfer the parameterized input information from the meta-level to the directly usable object level.

Bibliographic Remark. The contents of this part are based on joint work with Gabriele Kern-Isberner [109] (see Section 1.3).

8.1 Mechanism of Revision by Comparison for TPOs

In this section, we investigate of the methodology of Revision by Comparison $\circledcirc_\alpha \beta$ as a parameterized revision method, using plausibilistic TPOs over possible worlds $\preceq$ as representation of belief states. Via more comprehensible constraints for TPOs, we are able to specify which worlds are affected by the revision via a single formula, thus leading to clarification of the role of α. Also, we provide semantic postulates and a representation theorem for RbC in the context of TPOs.

In Section 7.1, we have investigated the relationship between plausibility orderings on possible worlds and epistemic entrenchment relations. Via equation (7.1) from Proposition 7.1.1 on page 140, we transfer the constraints from Definition 7.2.2 to qualitative constraints in the framework of plausibilistic TPOs $\preceq$ resp. $\preceq \circledcirc_\alpha \beta = \preceq^{\circledcirc}_{\alpha,\beta}$, which correspond to the entrenchment relations $\leqslant_E$ resp. $\leqslant_E^{\circledcirc_{\alpha,\beta}}$.

Definition 8.1.1 (Revision by Comparison for $\preceq$ [109])**.** *Let $\preceq$ be a plausibilistic TPO. The posterior plausibilistic TPO after RbC with α as reference and β as input* *sentence* $\preceq \circledcirc_\alpha \beta = \preceq^{\circledcirc}_{\alpha,\beta}$ *is defined as*

$$\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega' \ \textit{iff} \ \begin{cases} \alpha \Rightarrow (\beta\omega) \preceq \beta\omega', \ \textit{if} \ \omega \preceq \overline{\alpha} & \text{(I)} \\ \omega \preceq \omega', \ \textit{otherwise} & \text{(II)} \end{cases} \tag{8.1}$$

for any arbitrary sentences $\gamma, \delta \in \mathcal{L}$.

These constraints follow immediately from (7.3) via (7.1) for γ, δ as maximal disjunctions. It holds that the negation of γ, δ then display maximal conjunctions, which correspond to possible worlds (cf. (2.1) on page 18), yielding (8.1) for possible worlds ω and ω' (cf. Section 7.1). Employing (7.1) again, we transfer the properties of Revision by Comparison to the framework of plausibilistic TPOs.

(RbC)$_\preceq$ $\overline{\beta}$ is at least as plausible as $\overline{\beta}$, i.e., $\overline{\alpha} \preceq^{\circledcirc}_{\alpha,\beta} \overline{\beta}$

(MinRbC)$_\preceq$ $\overline{\alpha}$ is not additionally lifted, i.e., $\overline{\alpha} \prec \overline{\beta}$ iff $\overline{\alpha} \prec^{\circledcirc}_{\alpha,\beta} \overline{\beta}$

(α**-level)**$_\preceq$ For any $\gamma \in \mathcal{L}$, it holds that $\overline{\alpha} \prec \gamma$ iff $\overline{\alpha} \prec^{\circledcirc}_{\alpha,\beta} \gamma$

(α**-relation)**$_\preceq$ For any $\gamma, \delta \in \mathcal{L}$, if $\overline{\alpha} \preceq \gamma$ and $\overline{\alpha} \preceq \delta$, then $\gamma \preceq \delta$ iff $\gamma \preceq^{\circledcirc}_{\alpha,\beta} \delta$ holds

(β**-level)**$_\preceq$ For any $\gamma \in \mathcal{L}$, it holds that ($\gamma \preceq \overline{\alpha}$ or $\gamma \preceq \overline{\beta}$) iff $\gamma \preceq^{\circledcirc}_{\alpha,\beta} \overline{\beta}$

From (7.3) and thus, also from (8.1), it remains unclear prima facie which worlds exactly are affected by the change mechanism implemented by RbC of β w.r.t. α, also the crucial role of the relative positioning of the input β to the reference α is not clearly recognizable. The following proposition states an equivalent reformulation of (8.1), which fully integrates RbC of β w.r.t. α in the possible worlds reading via exclusive and exhaustive cases.

Proposition 8.1.1 ([109]). *For RbC of* β *w.r.t.* α *on TPOs, it holds that* (8.1) *is equivalent to the following constraints:*

$$\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega' \ \textit{iff} \ \begin{cases} \omega \preceq \omega', \ \textit{if} \ (\omega, \omega' \models \beta \ \textit{and} \ \omega \preceq \overline{\alpha}) & \text{(I)} \\ \qquad\quad \textit{or} \ \overline{\alpha} \prec \omega & \text{(II)} \\ \top, \ \textit{if} \ (\omega' \models \overline{\beta} \ \textit{and} \ \omega \preceq \overline{\alpha}) & \text{(III)} \\ \qquad\quad \textit{or} \quad (\omega \models \overline{\beta}, \omega' \models \beta, \ \textit{and} \ \omega \preceq \overline{\alpha} \preceq \omega') & \text{(IV)} \end{cases} \tag{8.2}$$

Proof. To prove the equivalence between (8.1) and (8.2), we expand and equivalently reorder the constraints in (8.1) until we receive (8.2).

Before we turn to the case $\bar{\alpha} \prec \omega$, we start with (I) on the right hand side of (8.1) and presuppose that $\omega \preceq \bar{\alpha}$ holds throughout the following case distinctions. Via the first case distinction, we consider either $\omega \models \overline{\beta}$ or $\omega \models \beta$. For each of these cases, we open two more cases concerning scenarios in which $\omega' \models \overline{\beta}$ or $\omega' \models \beta$. Thus, we get four disjunct cases, leading to different constraints defining the RbC change mechanism.

1. Case $\omega \models \overline{\beta}$: If $\omega \models \overline{\beta}$, we get for (I) in (8.1) that $\alpha \Rightarrow (\beta\omega) \equiv \bar{\alpha} \vee \bot \equiv \bar{\alpha} \preceq \beta\omega'$.

   Presuppose that $\omega' \models \overline{\beta}$**,** then it holds that

   $$\bar{\alpha} \preceq \beta\omega' \equiv \bot, \tag{8.3}$$

   which is always satisfied.
   Presuppose that $\omega' \models \beta$**,** then we get for (I) in (8.1) the following constraint

   $$\bar{\alpha} \preceq \beta\omega' \equiv \omega'. \tag{8.4}$$

2. Case $\omega \models \beta$: For $\omega \models \beta$ we can conclude that $\alpha \Rightarrow (\beta\omega) \equiv \bar{\alpha} \vee \omega \preceq \beta\omega'$

   Presuppose that $\omega' \models \overline{\beta}$**,** then it holds that

   $$\bar{\alpha} \vee \omega \preceq \beta\omega' \equiv \bot, \tag{8.5}$$

   which is always satisfied.
   Presuppose that $\omega' \models \beta$**.** It holds that $\bar{\alpha} \vee \omega \preceq \omega$, since $\omega \preceq \bar{\alpha}$ holds in (I) from (8.1). We get the following constraint

   $$\bar{\alpha} \vee \omega \approx \omega \preceq \beta\omega' \equiv \omega'. \tag{8.6}$$

For the cases ($\omega \models \overline{\beta}$ and $\omega' \models \overline{\beta}$) and ($\omega \models \beta$ and $\omega' \models \beta$) the constraint in (I) from (8.1) is satisfied trivially. Hence, we can summarize these cases via a single constraint. Together with (II) from (8.1), we get the following equivalent reformulation for (8.1) via (8.3), (8.4), (8.5) and (8.6).

$$\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega' \text{ iff } \begin{cases} \omega \preceq \omega' & \text{if } \underbrace{(\omega \preceq \overline{\alpha}, \omega, \omega' \models \beta)}_{(8.6)} \text{ or } \underbrace{\overline{\alpha} \prec \omega}_{\text{(II) in (8.1)}} \\ \top & \text{if } \underbrace{(\omega \preceq \overline{\alpha}, \omega' \models \overline{\beta})}_{(8.3)+(8.5)} \text{ or } \\ & \underbrace{(\omega \preceq \overline{\alpha}, \omega \models \overline{\beta} \text{ and } \omega' \models \beta, \overline{\alpha} \preceq \omega')}_{(8.4)} \end{cases}$$

□

The constraints given in (8.2) implement RbC of β w.r.t. α via exclusive and exhaustive cases. The following proposition summarizes the characteristics of the change induced by RbC.

Proposition 8.1.2. *Let $\omega, \omega' \in \Omega$ be possible worlds. For a plausibilistic TPO $\preceq$ and an RbC of β w.r.t. α, $\preceq \circledcirc_\alpha \beta = \preceq^{\circledcirc}_{\alpha,\beta}$, which satisfies the constraints in* (8.2)*, the following statements are true*

1. *If $\overline{\alpha} \preceq \omega, \omega'$, then it holds that $\omega \preceq \omega'$ iff $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$*
2. *If $\omega, \omega' \models \beta$, then it holds that $\omega \preceq \omega'$ iff $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$*
3. *If $\omega \models \overline{\beta}$ and $\omega \prec \overline{\alpha}$, then it holds that $\omega \approx^{\circledcirc}_{\alpha,\beta} \overline{\alpha}$*

Proof. We prove the constraints.

1. Let $\overline{\alpha} \preceq \omega, \omega'$.
   Case 1: For $\overline{\alpha} \prec \omega$, $\omega \preceq \omega'$ iff $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ follows from case (II) in (8.2).
   Case 2: Let $\overline{\alpha} \approx \omega$. Due to $\overline{\alpha} \preceq \omega, \omega'$, it holds that $\omega \approx \overline{\alpha} \preceq \omega'$ and we need to show that $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ holds.
   We consider the following (exclusive and exhaustive) subcases:
   (a) $\underline{\omega, \omega' \models \beta\text{:}}$ $\omega \preceq \omega'$ iff $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ follows from case (I) in (8.2).
   (b) $\underline{\omega \models \beta \text{ and } \omega' \models \overline{\beta}\text{:}}$ $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ follows from case (III) in (8.2).

(c) $\underline{\omega \models \overline{\beta} \text{ and } \omega' \models \beta\text{:}}$ $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ follows from case (IV) in (8.2) and $\omega \approx \overline{\alpha} \preceq \omega'$.

(d) $\underline{\omega, \omega' \models \overline{\beta}\text{:}}$ $\omega \preceq \omega'$ iff $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ follows from case (III) in (8.2).

2. Let $\omega, \omega' \models \beta$.
Case 1: Let $\omega \preceq \overline{\alpha}$. $\omega \preceq \omega'$ iff $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ follows from case (I) in (8.2).
Case 2: Let $\overline{\alpha} \prec \omega$. $\omega \preceq \omega'$ iff $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ follows from case (II) in (8.2).

3. Let $\omega \models \overline{\beta}$ and $\omega \prec \overline{\alpha}$.
Let $\omega_{\overline{\alpha}} \in \min(\overline{\alpha}, \preceq)$ s.t. $\omega_{\overline{\alpha}} \approx \overline{\alpha}$ and thus, in particular $\omega_{\overline{\alpha}} \preceq \overline{\alpha}$ and $\overline{\alpha} \preceq \omega_{\overline{\alpha}}$.

(a) We show $\overline{\alpha} \preceq^{\circledcirc}_{\alpha,\beta} \omega$.
It holds that $\omega \models \overline{\beta}$, thus $\omega_{\overline{\alpha}} \preceq^{\circledcirc}_{\alpha,\beta} \omega$ follows from (III) in (8.2). And we get that $\overline{\alpha} \preceq^{\circledcirc}_{\alpha,\beta} \omega$ holds.

(b) We show $\omega \preceq^{\circledcirc}_{\alpha,\beta} \overline{\alpha}$.
Case 1: Let $\omega_{\overline{\alpha}} \models \beta$, then $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega_{\overline{\alpha}}$ follows from (IV) in (8.2).
Case 2: Let $\omega_{\overline{\alpha}} \models \overline{\beta}$, then $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega_{\overline{\alpha}}$ follows from (III) in (8.2).

Thus, it holds that $\omega \preceq^{\circledcirc}_{\alpha,\beta} \overline{\alpha}$. All in all, we get that $\omega \approx^{\circledcirc}_{\alpha,\beta} \overline{\alpha}$ holds.

□

The first two statements from Proposition 8.1.2 deal with cases where the plausibility relations are not affected by RbC. All worlds less or equally plausible than $\overline{\alpha}$ and all β-worlds keep their relative positioning towards each other in $\preceq$. Note here that the β-worlds at most as plausible as $\overline{\alpha}$ are included in the first and the second statement to avoid lengthy case distinctions. The third statement of Proposition 8.1.2 is of particular importance. While the first two statements refer to cases in which the plausibility relations do not change, i.e., $\omega \preceq \omega'$ iff $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$, the third one deals with the actual plausibility shift implemented by the RbC-operator. Thus it summarizes that which worlds affected by RbC, namely all $\overline{\beta}$-worlds that are more or equally plausible than $\overline{\alpha}$ which follows from (III) and (IV) in (8.2). These worlds are shifted up to the level of plausibility where the $\overline{\alpha}$-worlds reside. This displays the main shift implemented by RbC, and hence, these worlds are crucial for RbC $\circledcirc_\alpha \beta$. Note that only the $\overline{\beta}$-worlds strictly more plausible than $\overline{\alpha}$ are actually shifted, i.e., their plausibility in fact decreases, since the $\overline{\beta}$-worlds as plausible as $\overline{\alpha}$ in the prior ordering remain on the same level after RbC. Moreover, the shift of $\overline{\beta}$-worlds strictly more plausible than $\overline{\alpha}$ has an indirect effect on β-worlds strictly more plausible than $\overline{\alpha}$. Because these worlds remain on their plausibility level, they get promoted de facto compared to $\overline{\beta}$-worlds strictly more plausible than $\overline{\alpha}$. We summarize these important previous sets of worlds via the following formulas in order to ease the notation for the following results.

Definition 8.1.2 (Penalty and indirect reward formula of RbC)**.** *Let $\preceq$ be a plausibilistic TPO and $\circledcirc_\alpha \beta$ be an RbC operator of β w.r.t. α.*

- *The* penalty formula *of RbC of β w.r.t. α is defined as*

$$\psi^{\circledcirc}_{\alpha,\beta} = \overline{\beta} \wedge (\bigvee_{\omega \prec \overline{\alpha}} \omega).$$

  We call the set of possible worlds $\Psi^{\circledcirc}_{\alpha,\beta} = Mod(\psi^{\circledcirc}_{\alpha,\beta}) = \{\omega \in \Omega \mid \omega \in Mod(\overline{\beta}), \omega \prec \overline{\alpha}\}$ the penalty set *of RbC.*

- *The* indirect reward formula *of RbC of β w.r.t. α is defined as*

$$\theta^{\circledcirc}_{\alpha,\beta} = \beta \wedge (\bigvee_{\omega \prec \overline{\alpha}} \omega).$$

  We call the set of possible worlds $\Theta^{\circledcirc}_{\alpha,\beta} = Mod(\theta^{\circledcirc}_{\alpha,\beta}) = \{\omega \in \Omega \mid \omega \in Mod(\beta), \omega \prec \overline{\alpha}\}$ the reward set *of RbC.*

It holds that $\psi^{\circledcirc}_{\alpha,\beta}$ and $\theta^{\circledcirc}_{\alpha,\beta}$ are exclusive formulas, i.e., $\psi^{\circledcirc}_{\alpha,\beta} \wedge \theta^{\circledcirc}_{\alpha,\beta} \equiv \perp$ holds. Hence, they implement a (disjoint) partition of all worlds strictly more plausible than $\overline{\alpha}$ and it holds that $\psi^{\circledcirc}_{\alpha,\beta} \vee \theta^{\circledcirc}_{\alpha,\beta} \equiv \bigvee_{\omega \prec \overline{\alpha}} \omega$.

We illustrate the mechanism of RbC via the following example.

Example 8.1.1. *In Figure 8.1a a plausibilistic TPO $\preceq$ over the signature $\Sigma = \{a, b, c, d\}$ is given. Note that, $\bar{\Omega}$ subsumes all worlds, whose plausibility ranking is not explicitly given and which reside on the same plausibility level above the given, more plausible, worlds. We perform an RbC of b w.r.t. a following the constraints given in Proposition 8.1.1 and get the TPO $\preceq \circledcirc_a b = \preceq^{\circledcirc}_{a,b}$. From* (I) *and* (II) *in* (8.2)*,*

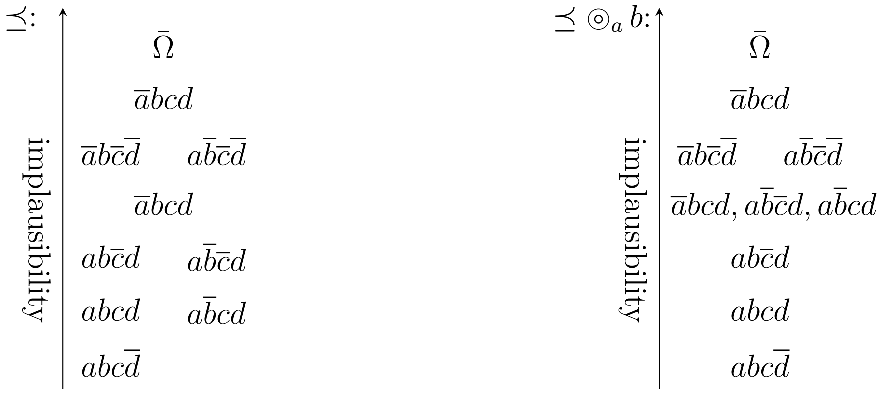

a) Plausibilistic TPO $\preceq$. b) RbC-revised TPO $\preceq \circledcirc_a b$.

Figure 8.1: Revision by Comparison by a w.r.t. b.

we can conclude that the relations among the worlds $\overline{a}bcd, \overline{a}b\overline{c}\overline{d}, a\overline{b}\overline{c}\overline{d}, ab\overline{c}d, abcd, abc\overline{d}$ *and* $\overline{a}bcd$ *do not change, this corresponds to statements 1 and 2 in Proposition 8.1.2, since* $abc\overline{d}, abcd, ab\overline{c}d \models \theta^{\circledcirc}_{a,b}$.
For $\preceq$*, it holds that* $a\overline{b}cd, a\overline{b}\overline{c}d \models \psi^{\circledcirc}_{a,b}$ *and* $\min(\bar{a}, \preceq) = \{\bar{a}bcd\}$*, which we notate in the following as* $\omega_{\bar{a}}$*, since it determines the plausibility rank of* $\bar{a}$ *in* $\preceq$*. For all* $\omega' \models \psi^{\circledcirc}_{a,b}$*, we get from* (III) *in* (8.2)*, that* $\omega_{\bar{a}} \preceq^{\circledcirc}_{a,b} \omega'$ *holds and from* (IV) *in* (8.2) *that* $\omega' \preceq^{\circledcirc}_{a,b} \omega_{\bar{a}}$ *holds. Thus,* $a\overline{b}cd, a\overline{b}\overline{c}d \approx^{\circledcirc}_{a,b} \overline{a}bcd \approx^{\circledcirc}_{a,b} \bar{a}$ *holds.*

The following lemma summarizes useful statements about $\psi^{\circledcirc}_{\alpha,\beta}$ and $\theta^{\circledcirc}_{\alpha,\beta}$ and is useful for the proofs of following theorems.

Lemma 8.1.3. *For* $\psi^{\circledcirc}_{\alpha,\beta}$ *and* $\theta^{\circledcirc}_{\alpha,\beta}$ *be as defined in Definition 8.1.2, the following statements are true:*

- *For* $\overline{\alpha} \preceq \omega$*, it holds that* $\omega \not\models \psi^{\circledcirc}_{\alpha,\beta} \vee \theta^{\circledcirc}_{\alpha,\beta}$
- *For* $\omega \prec \overline{\alpha}$*, it holds that* $\omega \models \psi^{\circledcirc}_{\alpha,\beta} \vee \theta^{\circledcirc}_{\alpha,\beta}$ *and* $\omega \models \alpha$

The statements from Lemma 8.1.3 follow directly from the disjunction of worlds used to define $\psi^{\circledcirc}_{\alpha,\beta}$ resp. $\theta^{\circledcirc}_{\alpha,\beta}$ in Definition 8.1.2 and the minimality of ranks (cf. (2.7) on page 36).

We use the penalty resp. indirect reward formula of RbC to present a representation theorem for Revision by Comparison. In contrast to the constraints from

(8.2), we focus on worlds strictly less plausible than $\overline{\alpha}$, since only these $\overline{\beta}$-worlds are shifted and pressed onto the plausibility level where the $\overline{\alpha}$-worlds reside. The representation theorem provides semantic constraints that make the revision mechanism of RbC of β w.r.t. α more explicit via three simple yet elegant postulates.

Theorem 8.1.4 (Representation Theorem for RbC). *Let $\circledcirc_\alpha \beta$ be an RbC operator of β w.r.t. α. Let $\preceq$ be a plausibilistic TPO and $\preceq \circledcirc_\alpha \beta = \preceq^{\circledcirc}_{\alpha,\beta}$ be the corresponding RbC-revised plausibilistic TPO. Then $\preceq$ and $\preceq^{\circledcirc}_{\alpha,\beta}$ satisfy* (8.2) *iff $\preceq$ and $\preceq^{\circledcirc}_{\alpha,\beta}$ satisfy:*

(RbC1) *If $\omega, \omega' \not\models \psi^{\circledcirc}_{\alpha,\beta}$, then $\omega \preceq \omega'$ iff $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$*

(RbC2) *If $\omega_{\overline{\alpha}} \in \min(\bar{\alpha}, \preceq)$ and $\omega' \models \psi^{\circledcirc}_{\alpha,\beta}$, then $\omega' \approx^{\circledcirc}_{\alpha,\beta} \omega_{\overline{\alpha}} \approx^{\circledcirc}_{\alpha,\beta} \overline{\alpha}$*

(RbC3) *If $\omega \models \theta^{\circledcirc}_{\alpha,\beta}$ and $\omega' \models \psi^{\circledcirc}_{\alpha,\beta}$ then $\omega \prec^{\circledcirc}_{\alpha,\beta} \omega'$*

Proof. "$\Leftarrow$"

Presuppose that $\preceq, \preceq^{\circledcirc}_{\alpha,\beta}$ satisfy (RbC1) – (RbC3). We show that $\preceq, \preceq^{\circledcirc}_{\alpha,\beta}$ satisfy (8.2).

(I) in (8.2). Let $\omega \preceq \bar{\alpha}$ and $\omega, \omega' \models \beta$. Then it holds that $\omega, \omega' \not\models \psi^{\circledcirc}_{\alpha,\beta}$ (cf. Lemma 8.1.3) and we can conclude from (RbC1) that $\omega \preceq \omega'$ iff $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ holds.

(II) in (8.2). Let $\bar{\alpha} \prec \omega$.
Case 1: Let $\omega' \not\models \psi^{\circledcirc}_{\alpha,\beta}$. Due to $\bar{\alpha} \prec \omega$, it holds that $\omega \not\models \psi^{\circledcirc}_{\alpha,\beta}$ (cf. Lemma 8.1.3) and thus we conclude that $\omega \preceq \omega'$ iff $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ holds via (RbC1).
Case 2: For $\omega' \models \psi^{\circledcirc}_{\alpha,\beta}$, it holds that $\omega' \prec \bar{\alpha} \prec \omega$. Thus, $\omega \preceq \omega'$ does not hold and therefore the equivalence (II) in (8.2) is satisfied trivially.

(III) in (8.2). Let $\omega \preceq \bar{\alpha}$ and $\omega' \models \bar{\beta}$. We need to show that $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ holds. We distinguish four exclusive cases for ω, on the top level, we need to distinguish whether $\omega \models \beta$ or $\omega \models \overline{\beta}$. From there, we take a closer look on the plausibility relation towards $\overline{\alpha}$, if $\omega \prec \overline{\alpha}$ then it holds that $\omega \models \theta^{\circledcirc}_{\alpha,\beta}$ or $\omega \models \psi^{\circledcirc}_{\alpha,\beta}$ (cf. Lemma 8.1.3). Otherwise it holds that $\omega \in \min(\bar{\alpha}, \preceq)$.

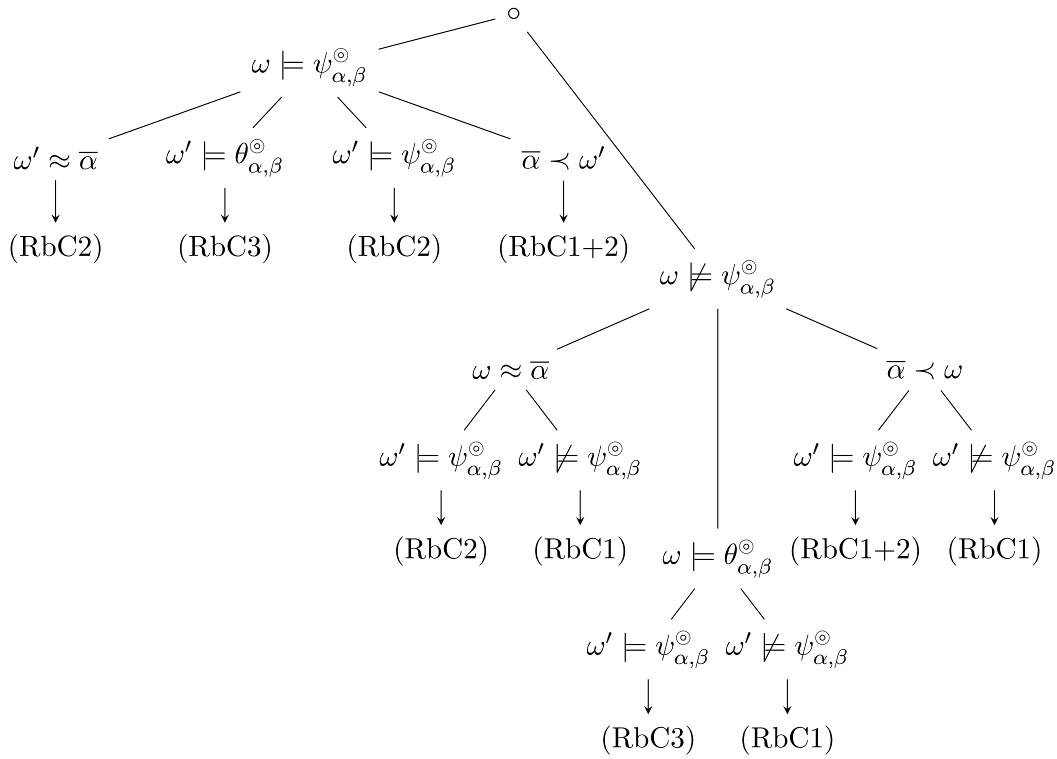

Figure 8.2: Schematic sketch of all use cases of the postulates (RbC1) – (RbC3) from Theorem 8.1.4 for Revision by Comparison of β w.r.t. α for a plausibilistic TPO $\preceq$. The case (RbC1+2) follows immediately from Corollary 8.1.5.

(a) Let $\omega \models \beta$ and $\omega \prec \overline{\alpha}$, i.e., $\omega \models \theta^{\circledcirc}_{\alpha,\beta}$.
We presuppose that $\omega' \models \overline{\beta}$ and make the following case distinction:
For $\omega' \models \psi^{\circledcirc}_{\alpha,\beta}$, $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ follows from (RbC3).
For $\omega' \not\models \psi^{\circledcirc}_{\alpha,\beta}$, $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ follows from (RbC1), since $\omega \prec \overline{\alpha} \preceq \omega'$ (cf. Lemma 8.1.3).

(b) Let $\omega \models \beta$ and $\omega \in \min(\overline{\alpha}, \preceq)$, i.e., $\omega \not\models \theta^{\circledcirc}_{\alpha,\beta}$.
For $\omega' \models \psi^{\circledcirc}_{\alpha,\beta}$, $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ follows from (RbC2).
For $\omega' \not\models \psi^{\circledcirc}_{\alpha,\beta}$, $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ follows from (RbC1), since $\omega \preceq \overline{\alpha} \preceq \omega'$ (cf. Lemma 8.1.3).

(c) Let $\omega \models \overline{\beta}$ and $\omega \prec \overline{\alpha}$, i.e., $\omega \models \psi^{\circledcirc}_{\alpha,\beta}$.
From (RbC2), we can conclude that $\omega \approx^{\circledcirc}_{\alpha,\beta} \overline{\alpha}$.
Case 1: For $\omega' \models \psi^{\circledcirc}_{\alpha,\beta}$, we can conclude from (RbC2) that $\omega' \approx^{\circledcirc}_{\alpha,\beta} \overline{\alpha}$.

Thus, $\omega \approx^{\circledcirc}_{\alpha,\beta} \omega'$, i.e., especially $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$.
Case 2: For $\omega' \not\models \psi^{\circledcirc}_{\alpha,\beta}$, it holds that $\overline{\alpha} \preceq \omega'$ (cf. Lemma 8.1.3). Since for all $\omega_{\overline{\alpha}} \in \min(\overline{\alpha}, \preceq)$, $\omega_{\overline{\alpha}} \not\models \psi^{\circledcirc}_{\alpha,\beta}$ holds, we can conclude from $\overline{\alpha} \preceq \omega'$ that $\overline{\alpha} \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ holds via (RbC1). Together with $\omega \approx^{\circledcirc}_{\alpha,\beta} \overline{\alpha}$ via (RbC2), we conclude that $\omega \preceq^{\circledcirc}_{\alpha,\beta} \overline{\alpha} \preceq^{\circledcirc}_{\alpha,\beta} \omega'$.

(d) Let $\omega \models \overline{\beta}$ and $\omega \in \min(\overline{\alpha}, \preceq)$, i.e., $\omega \not\models \psi^{\circledcirc}_{\alpha,\beta}$.
First, we show that $\omega \approx^{\circledcirc}_{\alpha,\beta} \overline{\alpha}$ holds.
From $\omega \in \min(\bar{\alpha}, \preceq)$, it follows that $\omega \models \overline{\alpha}$. Presuppose for contradiction that $\omega \not\approx^{\circledcirc}_{\alpha,\beta} \overline{\alpha}$, i.e., there exists $\omega' \models \overline{\alpha}$, s.t. $\omega' \prec^{\circledcirc}_{\alpha,\beta} \omega$. Since $\omega' \models \overline{\alpha}$, it holds that $\overline{\alpha} \preceq \omega'$ and therefore $\omega' \not\models \psi^{\circledcirc}_{\alpha,\beta}$ (cf. Lemma 8.1.3). From (RbC1), it follows from $\omega' \prec^{\circledcirc}_{\alpha,\beta} \omega$ that $\omega' \prec \omega$ holds, which contradicts $\omega \in \min(\bar{\alpha}, \preceq)$. Thus, for all worlds $\omega' \models \overline{\alpha}$, it holds that $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ and therefore $\omega \in \min(\overline{\alpha}, \preceq^{\circledcirc}_{\alpha,\beta})$, s.t. $\omega \approx^{\circledcirc}_{\alpha,\beta} \overline{\alpha}$ holds.
Case 1: For $\omega' \models \psi^{\circledcirc}_{\alpha,\beta}$, we can conclude from (RbC2) that $\omega' \approx^{\circledcirc}_{\alpha,\beta} \overline{\alpha}$, i.e., $\omega \approx^{\circledcirc}_{\alpha,\beta} \omega'$ holds.
Case 2: For $\omega' \not\models \psi^{\circledcirc}_{\alpha,\beta}$, it holds that $\overline{\alpha} \preceq \omega'$ (cf. Lemma 8.1.3). Therefore, it holds that $\omega \approx \overline{\alpha} \preceq \omega'$ and thus we can conclude $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ from (RbC1).

1. (IV) in (8.2). Let $\omega \preceq \bar{\alpha}, \omega \models \overline{\beta}$ and $\omega' \models \beta, \bar{\alpha} \preceq \omega'$. It holds that $\omega' \not\models \psi^{\circledcirc}_{\alpha,\beta}$. We need to show that $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ holds.

   (a) For $\omega \not\models \psi^{\circledcirc}_{\alpha,\beta}$, we can conclude that $\omega \approx \overline{\alpha}$ holds and thus $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ follows from $\omega \approx \overline{\alpha} \preceq \omega'$ via (RbC1).

   (b) For $\omega \models \psi^{\circledcirc}_{\alpha,\beta}$, we can conclude from (RbC2) that $\omega \approx^{\circledcirc}_{\alpha,\beta} \overline{\alpha}$ holds.
   Since $\omega' \not\models \psi^{\circledcirc}_{\alpha,\beta}$ and for all $\omega_{\overline{\alpha}} \in \min(\overline{\alpha}, \preceq)$, it holds that $\omega_{\overline{\alpha}} \not\models \psi^{\circledcirc}_{\alpha,\beta}$ (cf. Lemma 8.1.3), we can conclude from $\overline{\alpha} \preceq \omega'$ that $\overline{\alpha} \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ via (RbC1). Thus, we get that $\omega \approx^{\circledcirc}_{\alpha,\beta} \overline{\alpha} \preceq^{\circledcirc}_{\alpha,\beta} \omega'$.

"$\Rightarrow$"

Presuppose that $\preceq, \preceq^{\circledcirc}_{\alpha,\beta}$ satisfy (8.2). We show that $\preceq, \preceq^{\circledcirc}_{\alpha,\beta}$ satisfy (RbC1) – (RbC3).

(RbC1): Let $\omega, \omega' \not\models \psi^{\circledcirc}_{\alpha,\beta}$.

Thus, for $\tilde{\omega} \in \{\omega, \omega'\}$, it holds that

$$\tilde{\omega} \models \overline{\psi^{\circledcirc}_{\alpha,\beta}} \equiv \beta \vee \bigvee_{\overline{\alpha} \preceq \omega} \omega.$$

We employ the following case distinction to show that $\omega \preceq \omega'$ iff $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ holds.

(a) Let $\overline{\alpha} \preceq \omega, \omega'$. Then $\omega \preceq \omega'$ iff $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ follows immediately from the first statement in Proposition 8.1.2.

(b) Let $\omega, \omega' \models \beta$. Then $\omega \preceq \omega'$ iff $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ follows immediately from the second statement in Proposition 8.1.2.

(c) Let $\overline{\alpha} \preceq \omega$ and $\omega' \models \beta$.
For $\overline{\alpha} \prec \omega$, the statement $\omega \preceq \omega'$ iff $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ follows from (II) in (8.2).
Let $\overline{\alpha} \approx \omega$.
Case 1: Let $\omega \models \beta$. Then, $\omega \preceq \omega'$ iff $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ follows from (b).
Case 2: Let $\omega \models \overline{\beta}$.
For $\overline{\alpha} \preceq \omega'$, $\omega \preceq \omega'$ iff $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ follows from (a).
For $\omega' \prec \overline{\alpha}$, it holds that $\omega' \prec \overline{\alpha} \approx \omega$. Thus, $\omega \preceq \omega'$ does not hold and therefore the equivalence $\omega \preceq \omega'$ iff $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ is satisfied trivially.

Note that, if we swap ω and ω', (c) also covers the case $\overline{\alpha} \preceq \omega'$ and $\omega \models \beta$.

(RbC2): Let $\omega_{\overline{\alpha}} \in \min(\overline{\alpha}, \preceq)$ and $\omega' \models \psi^{\circledcirc}_{\alpha,\beta}$. Since $\omega' \models \psi^{\circledcirc}_{\alpha,\beta}$, it holds that $\omega' \models \overline{\beta}$ and $\omega' \prec \overline{\alpha}$, thus (RbC2) follows immediately from statement 3 in Proposition 8.1.2.

(RbC3): Let $\omega \models \theta^{\circledcirc}_{\alpha,\beta}$ and $\omega' \models \psi^{\circledcirc}_{\alpha,\beta}$. It holds that $\omega, \omega' \prec \overline{\alpha}$.
The inequality $\omega \preceq^{\circledcirc}_{\alpha,\beta} \omega'$ follows from (III) in (8.2).
We show that $\omega' \preceq^{\circledcirc}_{\alpha,\beta} \omega$ does not follow from the constraints in (8.2). Due to the equivalence in (8.2), we can then conclude that the strict inequality $\omega \prec^{\circledcirc}_{\alpha,\beta} \omega'$ holds.
Since $\omega' \models \psi^{\circledcirc}_{\alpha,\beta}$, it holds that $\omega' \not\models \beta$, i.e., (I) from (8.2) does not apply, and due to $\omega, \omega' \prec \overline{\alpha}$ part (II) from (8.2) does not apply.
Since $\omega \models \theta^{\circledcirc}_{\alpha,\beta}$, it holds that $\omega \not\models \overline{\beta}$, i.e., (III) in (8.2) does not apply, and due

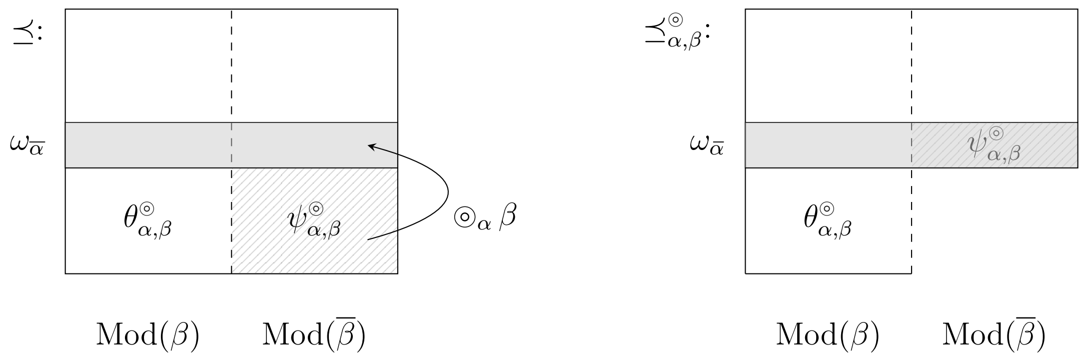

Figure 8.3: Schematic representation of Revision by Comparison of β w.r.t. α for a plausibilistic TPO $\preceq$.

to $\omega \prec \overline{\alpha}$ (IV) does not apply. Hence, we can conclude that $\omega' \preceq^{\circledcirc}_{\alpha,\beta} \omega$ does not hold otherwise it would contradict the equivalency in (8.2).

□

The following Corollary is a direct consequence of the (RbC1) – (RbC3).

Corollary 8.1.5. *Let $\circledcirc_\alpha \beta$ be an RbC operator of β w.r.t. α. Let $\preceq$ be a plausibilistic TPO and $\preceq \circledcirc_\alpha \beta = \preceq^{\circledcirc}_{\alpha,\beta}$ be the corresponding RbC-revised plausibilistic TPO. For $\omega_1, \omega_2 \models \psi^{\circledcirc}_{\alpha,\beta}$ and ω' with $\overline{\alpha} \prec \omega'$ it holds that*

$$\omega_1 \approx^{\circledcirc}_{\alpha,\beta} \omega_2 \approx^{\circledcirc}_{\alpha,\beta} \overline{\alpha} \prec^{\circledcirc}_{\alpha,\beta} \omega' \tag{8.7}$$

For $\omega_1, \omega_2 \models \psi^{\circledcirc}_{\alpha,\beta}$ in (8.7), we can follow from (RbC2) that $\omega_2 \approx^{\circledcirc}_{\alpha,\beta} \overline{\alpha}$ holds. Then $\omega_1 \approx^{\circledcirc}_{\alpha,\beta} \omega_2 \approx^{\circledcirc}_{\alpha,\beta} \overline{\alpha}$ is a direct consequence of (RbC2) since $\omega_1 \approx^{\circledcirc}_{\alpha,\beta} \overline{\alpha}$ and $\omega_2 \approx^{\circledcirc}_{\alpha,\beta} \overline{\alpha}$. The second part $\overline{\alpha} \prec^{\circledcirc}_{\alpha,\beta} \omega'$ follows from (RbC1), since $\omega_{\overline{\alpha}}, \omega' \not\models \psi^{\circledcirc}_{\alpha,\beta}$ holds for all minimal worlds $\omega_{\overline{\alpha}}$ in $\preceq$. Moreover, it holds that the postulates (RbC1) – (RbC3) are exclusive and exhaustive. The exclusivity follows immediately from the exclusivity of $\psi^{\circledcirc}_{\alpha,\beta}$ and $\theta^{\circledcirc}_{\alpha,\beta}$. For exhaustiveness, it is crucial that the roles of ω and ω' in (RbC1) – (RbC3) can be swapped. A schematic sketch of all use cases of ω and ω' when determining the posterior RbC-revised TPO $\preceq^{\circledcirc}_{\alpha,\beta}$ is presented in Figure 8.2.

The core shift of worlds performed by RbC is depicted in the schematic representation of RbC in Figure 8.3. Also, in Figure 8.3, the (indirect) plausibility increase

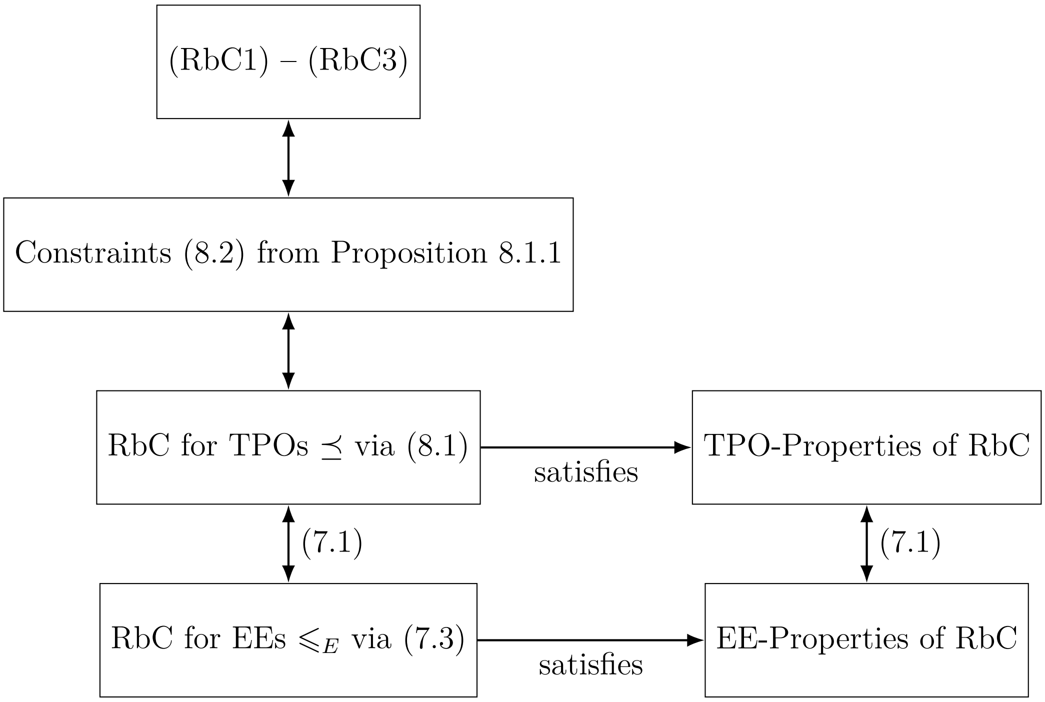

Figure 8.4: Overview of constraints and postulates defining RbC.

for worlds satisfying $\theta^{\circledcirc}_{\alpha,\beta}$ towards worlds in $\Psi^{\circledcirc}_{\alpha,\beta}$ becomes apparent, which is substantiated in (RbC3). We state a corollary following from Theorem 8.1.4 concerning the properties of RbC.

Corollary 8.1.6. *Let $\preceq$ be a plausibilistic TPO and $\circledcirc_\alpha \beta$ be an RbC operator of β w.r.t. α, s.t. $\preceq$ and $\preceq^{\circledcirc}_{\alpha,\beta}$ satisfy (RbC1) – (RbC3). Then the properties (RbC)$_\preceq$, (MinRbC)$_\preceq$, (α-level)$_\preceq$, (α-relation)$_\preceq$ and (β-level)$_\preceq$ hold for $\preceq$ and $\preceq^{\circledcirc}_{\alpha,\beta}$.*

The proof of the corollary summarizes the results of this section and a schematic sketch of it is given in Figure 8.4. From Theorem 8.1.4, we have seen that (RbC1) – (RbC3) are equivalent to (8.2), which is an equivalent reformulation from (8.1) which defines RbC for TPOs $\preceq$. (8.1) can be transformed into RbC for epistemic entrenchment relations $\leqslant_E$ in (7.3) from Definition 7.2.2. RbC for $\leqslant_E$ satisfies the epistemic entrenchment version of the RbC properties, which correspond to the TPO-version (RbC)$_\preceq$, (MinRbC)$_\preceq$, (α-level)$_\preceq$, (α-relation)$_\preceq$ and (β-level)$_\preceq$. Due to the direct correspondence between all of the above mentioned definitions of RbC, we can conclude that Corollary 8.1.6 holds and $\preceq^{\circledcirc}_{\alpha,\beta}$ satisfies all properties of RbC.

8.2 Hybrid Belief Change Character of Revision by Comparison

In this section, we discuss the hybrid belief change character of Revision by Comparison in the framework of plausibilistic TPOs employing our previous definitions and results.

Following the line of Fermé and Rott [34], we show that RbC results in a hybrid belief change operator for plausibilistic TPOs between revision with the input information and contraction of the reference sentence. In order to investigate the flexible approach of RbC, we make use of the penalty formula $\psi^{\circledcirc}_{\alpha,\beta}$ and the indirect reward formula $\theta^{\circledcirc}_{\alpha,\beta}$ of RbC for input β and reference α defined in the previous section. The compact formulation of $\psi^{\circledcirc}_{\alpha,\beta}$ resp. $\theta^{\circledcirc}_{\alpha,\beta}$ allows us to easily verify which case of RbC applies and makes the resulting change more comprehensible. In the following, we distinguish and discuss three basic cases which exhaust the space of logical possibilities. Note that they are not disjoint, but they yield identical results where more than one applies. To investigate and explain each case, we make use of the semantic constraints from (8.2) on page 153 and the subsequent postulates (RbC1) – (RbC3), because they break down the link between the penalty resp. indirect reward formula and the mechanism of RbC. The type of change, then depends on whether the corresponding sets of models for $\psi^{\circledcirc}_{\alpha,\beta}$ resp. $\theta^{\circledcirc}_{\alpha,\beta}$ are empty or not.

The β-Revision. As a starting point of the first basic case, we consider plausibilistic TPOs in which the plausibility of the reference sentence α is sufficiently high, where 'sufficiently' means that $\overline{\alpha}$ is strictly less plausible than α, also w.r.t. to the input sentence β and its negation $\overline{\beta}$. In terms of formulas and a prior TPO $\preceq$, this means that

$$\alpha\overline{\beta} \prec \overline{\alpha} \text{ and } \alpha\beta \prec \overline{\alpha} \tag{8.8}$$

holds. Thus, $\omega_{\alpha\overline{\beta}} \in \min(\alpha\overline{\beta}, \preceq)$ is a $\overline{\beta}$-world strictly less plausible than $\overline{\alpha}$, and $\omega_{\alpha\overline{\beta}} \models \psi^{\circledcirc}_{\alpha,\beta}$. Analog, $\omega_{\alpha\beta} \models \theta^{\circledcirc}_{\alpha,\beta}$ and we can conclude that $\Psi^{\circledcirc}_{\alpha,\beta} \neq \emptyset$ and $\Theta^{\circledcirc}_{\alpha,\beta} \neq \emptyset$. Fermé and Rott call this case the *intended case*, since it is the generic case for Revision by Comparison where the operation yields a revision with β, s.t. β is believed in the posterior belief state.

Theorem 8.2.1. *Let $\preceq$ be a plausibilistic TPO and $\circledcirc_\alpha \beta$ be an RbC operator of β w.r.t. α, s.t. $\preceq$ and $\preceq^{\circledcirc}_{\alpha,\beta}$ satisfy (RbC1) – (RbC3). It holds that $\psi^{\circledcirc}_{\alpha,\beta} \not\equiv \perp$ if and only if $\alpha\overline{\beta} \prec \overline{\alpha}$, and, $\theta^{\circledcirc}_{\alpha,\beta} \not\equiv \perp$ if and only if $\alpha\beta \prec \overline{\alpha}$. If $\psi^{\circledcirc}_{\alpha,\beta} \not\equiv \perp$ and $\theta^{\circledcirc}_{\alpha,\beta} \not\equiv \perp$, then $\preceq^{\circledcirc}_{\alpha,\beta} \models \beta$.*

Proof. We show that $\psi^{\circledcirc}_{\alpha,\beta} \not\equiv \perp$ iff $\alpha\overline{\beta} \prec \overline{\alpha}$ holds.
"$\Rightarrow$": Let $\psi^{\circledcirc}_{\alpha,\beta} \not\equiv \perp$, then there exists a minimal model $\omega_{\psi^{\circledcirc}_{\alpha,\beta}}$ of $\psi^{\circledcirc}_{\alpha,\beta}$, s.t. $\omega_{\psi^{\circledcirc}_{\alpha,\beta}} \prec \overline{\alpha}$ and $\omega_{\psi^{\circledcirc}_{\alpha,\beta}} \models \overline{\beta}$ minimally. From Lemma 8.1.3, we can conclude that $\omega_{\psi^{\circledcirc}_{\alpha,\beta}} \models \alpha$, i.e., $\omega_{\psi^{\circledcirc}_{\alpha,\beta}}$ is also a minimal model of $\alpha\overline{\beta}$ and therefore $\alpha\overline{\beta} \approx \omega_{\psi^{\circledcirc}_{\alpha,\beta}} \prec \overline{\alpha}$.
"$\Leftarrow$": Let $\alpha\overline{\beta} \prec \overline{\alpha}$, then there exists a $\alpha\overline{\beta}$-world $\omega_{\alpha\overline{\beta}}$, s.t. $\omega_{\alpha\overline{\beta}} \prec \overline{\alpha}$ and $\omega_{\alpha\overline{\beta}} \models \overline{\beta}$. Thus, $\omega_{\alpha\overline{\beta}} \models \psi^{\circledcirc}_{\alpha,\beta}$ and $\psi^{\circledcirc}_{\alpha,\beta} \not\equiv \perp$ holds.
The analog argumentation applies for $\theta^{\circledcirc}_{\alpha,\beta} \not\equiv \perp$ iff $\alpha\beta \prec \overline{\alpha}$ if we replace $\overline{\beta}$ by β.

Since $\psi^{\circledcirc}_{\alpha,\beta} \not\equiv \perp$, it holds that there exists $\omega_{\psi^{\circledcirc}_{\alpha,\beta}}$, s.t. $\omega_{\psi^{\circledcirc}_{\alpha,\beta}} \in \min(\psi^{\circledcirc}_{\alpha,\beta}, \preceq)$. Thus, $\omega_{\psi^{\circledcirc}_{\alpha,\beta}}$ satisfies $\omega_{\psi^{\circledcirc}_{\alpha,\beta}} \models \overline{\beta}$ and $\omega_{\psi^{\circledcirc}_{\alpha,\beta}} \prec \overline{\alpha}$ minimally, which means that there exists no $\omega' \models \overline{\beta}$ with $\omega' \prec \omega$ and we can conclude that $\omega_{\psi^{\circledcirc}_{\alpha,\beta}} \in \min(\overline{\beta}, \preceq)$.
For the RbC-revised TPO $\preceq^{\circledcirc}_{\alpha,\beta}$, due to (RbC2), it holds that $\omega_{\psi^{\circledcirc}_{\alpha,\beta}} \approx^{\circledcirc}_{\alpha,\beta} \omega' \approx^{\circledcirc}_{\alpha,\beta} \overline{\alpha}$ for all $\omega' \models \psi^{\circledcirc}_{\alpha,\beta}$. And for $\overline{\beta}$-worlds which do not satisfy $\psi^{\circledcirc}_{\alpha,\beta}$, i.e., $\tilde{\omega} \not\models \psi^{\circledcirc}_{\alpha,\beta}$, it holds that $\overline{\alpha} \preceq \tilde{\omega}$. From (RbC1), we can conclude that their plausibility relations do not change and $\overline{\alpha} \preceq^{\circledcirc}_{\alpha,\beta} \tilde{\omega}$ holds. So, all in all, we get $\omega_{\psi^{\circledcirc}_{\alpha,\beta}} \preceq^{\circledcirc}_{\alpha,\beta} \omega$ for all $\omega \models \overline{\beta}$, s.t., $\omega_{\psi^{\circledcirc}_{\alpha,\beta}} \in \min(\overline{\beta}, \preceq^{\circledcirc}_{\alpha,\beta})$ holds.
Following the same argumentation, as above we get that $\omega_{\theta^{\circledcirc}_{\alpha,\beta}} \in \min(\theta^{\circledcirc}_{\alpha,\beta}, \preceq)$ satisfies β minimally, i.e., $\omega_{\theta^{\circledcirc}_{\alpha,\beta}} \in \min(\beta, \preceq)$ holds.
For the RbC-revised TPO $\preceq^{\circledcirc}_{\alpha,\beta}$, it holds for all $\omega' \models \beta$ that they do not satisfy $\psi^{\circledcirc}_{\alpha,\beta}$ and therefore, we can conclude from (RbC1) that the relations among all β-worlds are kept and that $\omega_{\theta^{\circledcirc}_{\alpha,\beta}} \in \min(\beta, \preceq^{\circledcirc}_{\alpha,\beta})$ holds.

To sum up, a world $\omega_{\psi^{\circledcirc}_{\alpha,\beta}}$ resp. $\omega_{\theta^{\circledcirc}_{\alpha,\beta}}$ that satisfies $\psi^{\circledcirc}_{\alpha,\beta}$ resp. $\theta^{\circledcirc}_{\alpha,\beta}$ minimally, also satisfies $\overline{\beta}$ resp. β minimally for the prior $\preceq$ and the posterior $\preceq^{\circledcirc}_{\alpha,\beta}$ TPO. We get

$$\omega_{\psi^{\circledcirc}_{\alpha,\beta}} \approx \overline{\beta} \text{ resp. } \omega_{\theta^{\circledcirc}_{\alpha,\beta}} \approx \beta \text{ and } \omega_{\psi^{\circledcirc}_{\alpha,\beta}} \approx^{\circledcirc}_{\alpha,\beta} \overline{\beta} \text{ resp. } \omega_{\theta^{\circledcirc}_{\alpha,\beta}} \approx^{\circledcirc}_{\alpha,\beta} \beta \tag{8.9}$$

Hence, via (RbC3) it follows that $\omega_{\theta^{\circledcirc}_{\alpha,\beta}} \prec^{\circledcirc}_{\alpha,\beta} \omega_{\psi^{\circledcirc}_{\alpha,\beta}}$, i.e. $\beta \approx^{\circledcirc}_{\alpha,\beta} \omega_{\theta^{\circledcirc}_{\alpha,\beta}} \prec^{\circledcirc}_{\alpha,\beta} \omega_{\psi^{\circledcirc}_{\alpha,\beta}} \approx^{\circledcirc}_{\alpha,\beta} \overline{\beta}$

holds due to (8.9) and therefore $\preceq^{\circledcirc}_{\alpha,\beta} \models \beta$. □

In the intended case, where the reference α is selected as a comparatively plausible belief, s.t. $\bar{\alpha}$ is less plausible than both β and $\overline{\beta}$, the resulting belief set of RbC $Bel(\preceq \circledcirc_\alpha \beta)$ coincides with an AGM revision on $Bel(\preceq)$ with β as far as one-step revision is considered [34], i.e., $Bel(\preceq \circledcirc_\alpha \beta)$ satisfies the postulates (AGM$\star$1) –(AGM$\star$8) from page 21. Here, the worlds satisfying $\theta^{\circledcirc}_{\alpha,\beta}$ get promoted indirectly since their relative positioning towards worlds satisfying $\psi^{\circledcirc}_{\alpha,\beta}$ is promoted without actually lifting their plausibility level but rather by decreasing the plausibility of worlds satisfying $\psi^{\circledcirc}_{\alpha,\beta}$.

We present a schematic representation of the β-revision case in Figure 8.5a. In the following Example 8.1.1 is recaptured, which illustrates the β-revision case.

Example 8.2.1 (Continuing Example 8.1.1). *In Figure 8.1a a plausibilistic TPO $\preceq$ over the signature $\Sigma = \{a, b, c, d\}$ is given. Note that, $\bar{\Omega}$ subsumes all worlds, whose plausibility ranking is not explicitly given and which reside on the same plausibility level above the given, more plausible, worlds. We perform an RbC of b w.r.t. a, $\preceq \circledcirc_a b = \preceq^{\circledcirc}_{a,b}$, with $\Psi^{\circledcirc}_{a,b} = \{a\bar{b}\bar{c}d, a\bar{b}cd\}$ and $\Theta^{\circledcirc}_{a,b} = \{ab\bar{c}d, abcd, abc\bar{d}\}$. The posterior TPO $\preceq^{\circledcirc}_{a,b}$ is depicted in Figure 8.1b and it holds that $\preceq^{\circledcirc}_{a,b} \models b$.*

The Vacuous Case. This case applies for plausibilistic TPOs in which (RbC)$_{\preceq}$ holds in the prior epistemic state, i.e., $\overline{\alpha}$ is more or equally plausible as $\overline{\beta}$. So, there are no worlds that satisfy the penalty formula $\psi^{\circledcirc}_{\alpha,\beta}$, and thus RbC does not change anything, and the prior ordering is preserved. Figure 8.5b shows a schematic representation of the vacuous case.

Theorem 8.2.2. *Let $\preceq$ be a plausibilistic TPO and $\circledcirc_\alpha \beta$ be an RbC operator of β w.r.t. α, s.t. $\preceq$ and $\preceq^{\circledcirc}_{\alpha,\beta}$ satisfy (RbC1) – (RbC3). It holds that $\psi^{\circledcirc}_{\alpha,\beta} \equiv \bot$ if and only if $\overline{\alpha} \preceq \overline{\beta}$. Hence, if $\psi^{\circledcirc}_{\alpha,\beta} \equiv \bot$, then $\preceq = \preceq^{\circledcirc}_{\alpha,\beta}$.*

Proof. We show that $\psi^{\circledcirc}_{\alpha,\beta} \equiv \bot$ iff $\overline{\alpha} \preceq \overline{\beta}$ holds.
<u>"⇒":</u> Let $\psi^{\circledcirc}_{\alpha,\beta} \equiv \bot$, then it holds for all $\overline{\beta}$-worlds ω that $\overline{\alpha} \preceq \omega$. This applies in particular for minimal $\overline{\beta}$-worlds in $\preceq$, i.e., $\overline{\alpha} \preceq \overline{\beta}$ holds.
<u>"⇐":</u> Let $\overline{\alpha} \preceq \overline{\beta}$, then it holds for a minimal $\overline{\beta}$-world $\omega_{\overline{\beta}}$ in $\preceq$ that $\overline{\alpha} \preceq \omega_{\overline{\beta}} \preceq \omega'$ for

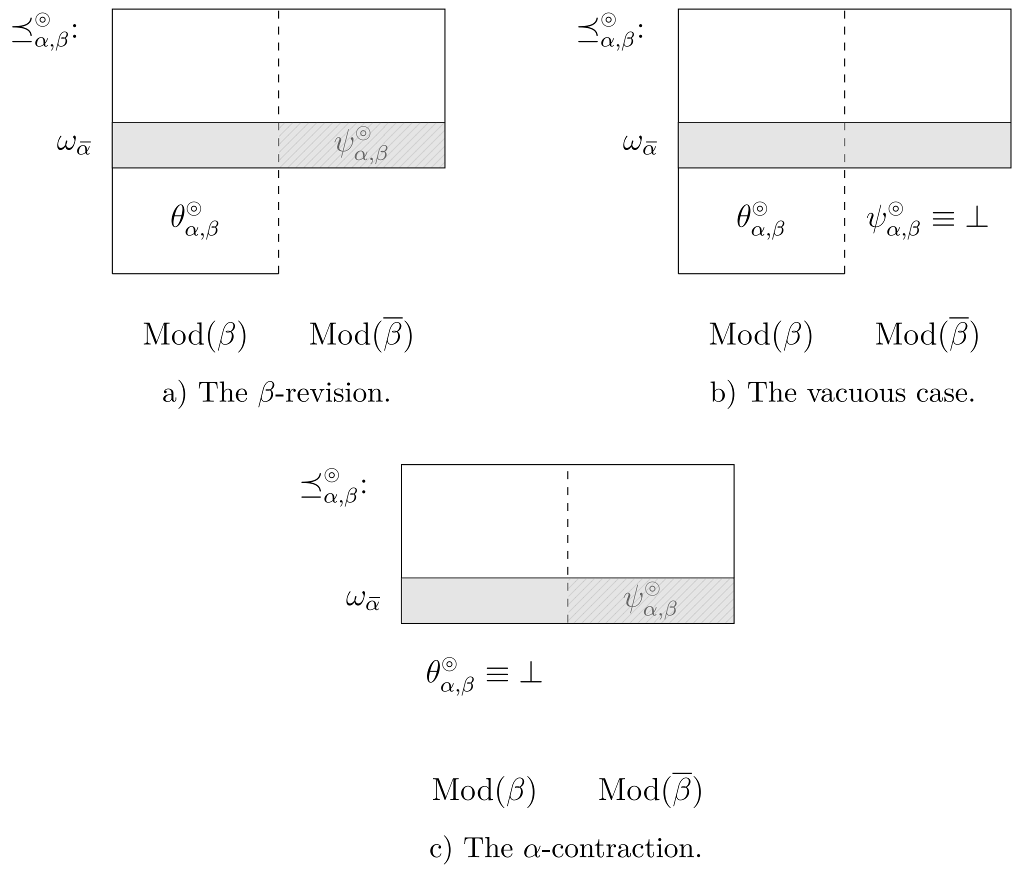

Figure 8.5: Schematic illustration of the hybrid belief change character implemented by Revision by Comparison of β w.r.t. α for a plausibilistic TPO $\preceq$.

all $\omega' \models \overline{\beta}$. Thus, there exists no $\omega' \models \overline{\beta}$ s.t. $\omega' \prec \overline{\alpha}$ and therefore $\psi^{\circledcirc}_{\alpha,\beta} \equiv \bot$. Because $\psi^{\circledcirc}_{\alpha,\beta} \equiv \bot$, it holds for all worlds in Ω that $\omega \not\models \psi^{\circledcirc}_{\alpha,\beta}$, and we can conclude that $\preceq = \preceq^{\circledcirc}_{\alpha,\beta}$ holds via (RbC1). $\square$

The vacuous case also applies for the special case of tautological revision, i.e., if $\beta \equiv \top$, s.t. RbC satisfies the tautological vacuity principle (TV) (3.2) from page 64 regardless of the choice of the reference sentence α. This follows immediately from the form of $\psi^{\circledcirc}_{\alpha,\beta}$ for $\beta \equiv \top$, which is

$$\psi^{\circledcirc}_{\alpha,\top} = \bot \wedge \bigvee_{\omega \prec \alpha} \equiv \bot, \tag{8.10}$$

and via Theorem 8.2.2, it holds that RbC for $\beta \equiv \top$ does not change the ordering.

Example 8.2.2. *We consider the plausibilistic TPO $\preceq$ from Figure 8.1a and perform an RbC of c w.r.t. b, $\preceq \circledcirc_b c = \preceq^{\circledcirc}_{b,c}$. It holds that $\Psi^{\circledcirc}_{b,c} = \emptyset$ and $\Theta^{\circledcirc}_{b,c} = \{abc\overline{d}\}$, i.e., the vacuous case of RbC applies. The posterior TPO $\preceq^{\circledcirc}_{b,c}$ is depicted in Figure 8.6a and it holds that $\preceq^{\circledcirc}_{b,c} = \preceq$.*

The α-Contraction. Revision by Comparison does not prioritize the new input information β for general prior plausibilistic TPO. In fact, if the negation of the reference sentence $\overline{\alpha}$ is more plausible than the input β, then the shifting of worlds satisfying $\psi^{\circledcirc}_{\alpha,\beta}$ up to the plausibility level of $\overline{\alpha}$ results in a contraction of the former belief α, s.t. $\preceq^{\circledcirc}_{\alpha,\beta} \not\models \alpha$ and the acceptance of the new input β is not guaranteed. Thus, RbC displays a non-prioritized belief revision mechanism. This case is called the *unsuccessful case* in [34], and it corresponds to the idea that the reference sentence displays the source β is coming from so that the plausibility of α can be interpreted as the reliability of this source, which decreases for a sufficiently implausible input.

Theorem 8.2.3. *Let $\preceq$ be a plausibilistic TPO and $\circledcirc_\alpha \beta$ be an RbC operator of β w.r.t. α, s.t. $\preceq$ and $\preceq^{\circledcirc}_{\alpha,\beta}$ satisfy (RbC1) – (RbC3). It holds that $\theta^{\circledcirc}_{\alpha,\beta} \equiv \bot$ if and only if $\overline{\alpha} \preceq \beta$. If $\theta^{\circledcirc}_{\alpha,\beta} \equiv \bot$, then $\preceq^{\circledcirc}_{\alpha,\beta} \not\models \alpha$.*

Proof. We show that $\theta^{\circledcirc}_{\alpha,\beta} \equiv \bot$ iff $\overline{\alpha} \preceq \beta$ holds.
<u>"⇒":</u> Let $\theta^{\circledcirc}_{\alpha,\beta} \equiv \bot$, then it holds for all β-worlds ω that $\overline{\alpha} \preceq \omega$. This applies in particular for minimal β-worlds in $\preceq$, i.e., $\overline{\alpha} \preceq \beta$ holds.
<u>"⇐":</u> Let $\overline{\alpha} \preceq \beta$, then it holds for a minimal β-world ω_β in $\preceq$ that $\overline{\alpha} \preceq \omega_\beta \preceq \omega'$ for all $\omega' \models \beta$. Thus, there exists no $\omega' \models \beta$ s.t. $\omega' \prec \overline{\alpha}$ and therefore $\theta^{\circledcirc}_{\alpha,\beta} \equiv \bot$.
<u>1. case:</u> Presuppose that $\Psi^{\circledcirc}_{\alpha,\beta} \neq \emptyset$.
From $\Psi^{\circledcirc}_{\alpha,\beta} \neq \emptyset$, we can conclude that there exists a minimal world $\omega_{\psi^{\circledcirc}_{\alpha,\beta}} \in \min(\psi^{\circledcirc}_{\alpha,\beta}, \preceq)$ s.t. $\omega_{\psi^{\circledcirc}_{\alpha,\beta}} \prec \overline{\alpha}$ and therefore $\omega_{\psi^{\circledcirc}_{\alpha,\beta}} \models \alpha$. Due to $\omega_{\psi^{\circledcirc}_{\alpha,\beta}} \prec \overline{\alpha}$, we can conclude that $\omega_{\psi^{\circledcirc}_{\alpha,\beta}} \in \min(\alpha, \preceq)$. And via (RbC2), we get that $\omega_{\psi^{\circledcirc}_{\alpha,\beta}} \approx^{\circledcirc}_{\alpha,\beta} \overline{\alpha}$, s.t. there is no α-world in $\preceq^{\circledcirc}_{\alpha,\beta}$ that is more plausible than $\overline{\alpha}$. Hence, $\preceq^{\circledcirc}_{\alpha,\beta} \not\models \alpha$.

<u>2. case:</u> Presuppose that $\Psi^{\circledcirc}_{\alpha,\beta} = \emptyset$.
Since $\Psi^{\circledcirc}_{\alpha,\beta}, \Theta^{\circledcirc}_{\alpha,\beta} = \emptyset$ and the second statement in Lemma 8.1.3 hold, there does not exist a world ω which is strictly more plausible than $\overline{\alpha}$, Thus, $\overline{\alpha} \preceq \alpha$ and due to Theorem 8.2.2, we can conclude that $\overline{\alpha} \preceq^{\circledcirc}_{\alpha,\beta} \alpha$ holds. $\square$

Note that, if $\psi^{\circledcirc}_{\alpha,\beta}, \theta^{\circledcirc}_{\alpha,\beta} \equiv \bot$ then the vacuous case and α-contraction coincide. In this case RbC does not change the prior ordering $\preceq$ and $\alpha \notin Bel(\preceq^{\circledcirc}_{\alpha,\beta})$ follows from $\alpha \notin Bel(\preceq)$.

Moreover, we show that in the α-contraction case the RbC does not satisfy the KM style contraction postulates (KM−1) – (KM−7) for a contraction of $\preceq$ by α. Therefore, we employ Theorem 2.5.2 on page 52 from [71] and show that

$$Bel(\preceq -\alpha) = Bel(\preceq) \cup \min(\overline{\alpha}, \preceq) \tag{8.11}$$

does not hold for the α-contraction case of RbC via the following counterexample:

Example 8.2.3. *We consider the following TPO over the signature* $\Sigma = \{a, b, c\}$

$$a\bar{b}\bar{c} \prec a\bar{b}c \prec abc, \bar{a}bc, \bar{a}\bar{b}c \prec \bar{\Omega}.$$

Note that, $\bar{\Omega}$ *denotes all remaining worlds which are not shown explicitly, on the same level of plausibility. We consider the RbC with input sentence* $\alpha = a$ *and reference sentence* $\beta = b$*. For this TPO, it holds that* $\Theta^{\circledcirc}_{\alpha,\beta} = \emptyset$ *and* $\Psi^{\circledcirc}_{\alpha,\beta} = \{a\bar{b}\bar{c}, a\bar{b}c\}$ *with the belief set* $Bel(\preceq) = \{a\bar{b}\bar{c}\}$*. RbC* $\preceq \circledcirc_a b$ *yields the following posterior TPO:*

$$a\bar{b}\bar{c}, a\bar{b}c, abc, \bar{a}bc, \bar{a}\bar{b}c \prec^{\circledcirc_{\alpha,\beta}}_{a,b} \bar{\Omega}.$$

with the belief set $Bel(\preceq \circledcirc_a b) = \{a\bar{b}\bar{c}, a\bar{b}c, abc, \bar{a}bc, \bar{a}\bar{b}c\}$*. It holds that*

$$\begin{aligned} Bel(\preceq) \cup \min(\overline{\alpha}, \preceq) &= \{a\bar{b}\bar{c}\} \cup \{\bar{a}bc, \bar{a}\bar{b}c\} \\ &= \{a\bar{b}\bar{c}, \bar{a}bc, \bar{a}\bar{b}c\} \subset Bel(\preceq \circledcirc_a b). \end{aligned}$$

Thus, RbC does not satisfy (8.11)*, i.e. the unsuccessful case does not correspond to a contraction of the reference sentence satisfying (KM−1) – (KM−7).*

The example illustrates that for RbC, all worlds in $\Psi^{\circledcirc}_{\alpha,\beta}$ are shifted up to the level of plausibility where $\overline{\alpha}$ resides, unlike, for KM style contractions, where only minimal $\overline{\alpha}$-worlds are relevant for the contraction. Due to $\Theta^{\circledcirc}_{\alpha,\beta} = \emptyset$, this leads to a posterior belief set which consists of all worlds on the same plausibility level as $\overline{\alpha}$ together with worlds in $\Psi^{\circledcirc}_{\alpha,\beta}$. Thus, the belief set of $\preceq^{\circledcirc}_{\alpha,\beta}$ is, in general, larger than

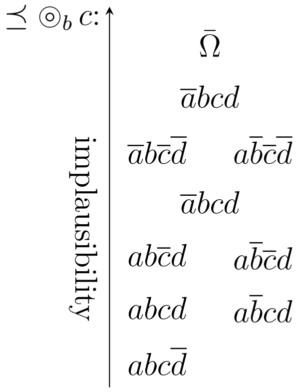

a) RbC-revised TPO $\preceq \circledcirc_b c$.

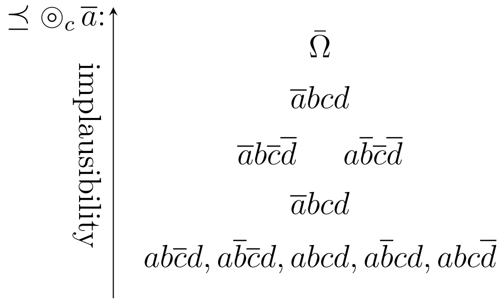

b) RbC-revised TPO $\preceq \circledcirc_c \overline{a}$.

Figure 8.6: Vacuous and α-contraction case of Revision by Comparison.

$Bel(\preceq) \cup \min(\overline{\alpha}, \preceq)$ in terms of set-inclusion. We present a schematic representation of the α-contraction case in Figure 8.5c and illustrate the α-contraction case via the following example.

Example 8.2.4. *We consider the plausibilistic TPO $\preceq$ from Figure 8.1a and perform an RbC of $\overline{a}$ w.r.t. c, $\preceq \circledcirc_c \overline{a} = \preceq^{\circledcirc}_{c,\overline{a}}$. It holds that $\Psi^{\circledcirc}_{c,\overline{a}} = \{abcd, a\overline{b}cd, abc\overline{d}\}$ and $\Theta^{\circledcirc}_{c,\overline{a}} = \emptyset$, i.e., the α-contraction case of RbC applies. The posterior TPO $\preceq^{\circledcirc}_{c,\overline{a}}$ is depicted in Figure 8.6b and it holds that $\preceq^{\circledcirc}_{c,\overline{a}} \not\models c$.*

RbC adapts its belief change to the prior belief state of an agent, more explicitly to the plausibility of the reference sentence α, which in the paradigm case should be believed with sufficiently high plausibility – where 'sufficiently high' means that α itself should be continued to be believed after RbC – and thus links reliability resp. priority of the new information to a reference sentence which is either specified via the input or can be selected freely by an agent. Except for a weaker form of belief change operators recently proposed in [105], this dynamic form of belief change is unique to the methodology of RbC.

To sum up, the three cases of RbC show that for RbC, the reference sentence can act as a marker of reliability and allows us to revise with a seemingly implausible new information only to a certain degree or even the devaluation of the reference.

8.3 Revision by Comparison for Ranking Functions

In this section, we turn to the semi-quantitative framework of ranking functions and present two types of methodological implementations for RbC, starting with a straightforward one presented in Subsection 8.3.1. Taking this formulation of RbC for OCFs as a basis, we define and investigate a set of weak conditionals which characterizes the change mechanism of RbC in Subsection 8.3.2. This set of weak conditionals finally characterizes the meta-information from the parameterized revision mechanism in RbC on the object level so that it can be used as input for revision operators capable of revising with sets of weak conditionals. Thus, the results from Subsection 8.3.2 enable us to define a c-revision with weak conditionals that implements RbC for OCFs and fully captures its versatile belief change character as a parameterized belief change operator.

Throughout this section, we use the penalty formula $\psi^{\circledcirc}_{\alpha,\beta}$ and the indirect reward formula $\theta^{\circledcirc}_{\alpha,\beta}$ for RbC, as defined in the previous section. Originally, these two formulas were defined for TPOs employing the condition $\omega \prec \overline{\alpha}$ to define a disjunction of possible worlds in $\psi^{\circledcirc}_{\alpha,\beta}$ resp. $\theta^{\circledcirc}_{\alpha,\beta}$. This condition can be easily transferred to the framework of OCFs via the translation (2.9), and we get for an OCF κ the following corresponding penalty resp. indirect reward formulas:

$$\psi^{\circledcirc}_{\alpha,\beta} = \overline{\beta} \wedge (\bigvee_{\kappa(\omega)<\kappa(\overline{\alpha})} \omega) \quad \text{and} \quad \theta^{\circledcirc}_{\alpha,\beta} = \beta \wedge (\bigvee_{\kappa(\omega)<\kappa(\overline{\alpha})} \omega)$$

We stick to the previous notation to avoid re-definition of these formulas.

8.3.1 Realization of Revision by Comparison for Ranking Functions

In the following, we present a realization of Revision by Comparison in the framework of ranking functions via a straightforward implementation of the postulates (RbC1) – (RbC3) from Theorem 8.1.4.

Definition 8.3.1 (RbC for OCFs, [109])**.** *Let κ be a ranking function. The* Revision by Comparison $\circledcirc_\alpha \beta$ with input sentence β and reference sentence α for OCFs *is defined as follows*

$$\kappa \circledcirc_\alpha \beta(\omega) = \kappa^{\circledcirc}_{\alpha,\beta}(\omega) = \kappa_0 + \begin{cases} \kappa(\bar{\alpha}), & \omega \models \psi^{\circledcirc}_{\alpha,\beta} \\ \kappa(\omega), & \textit{otherwise} \end{cases} \tag{8.12}$$

where $\kappa_0 = -\min\{\kappa(\bar{\alpha}), \kappa(\beta)\}$ is a normalization constant.

From (2.15), it follows for κ_0 that $\kappa_0 = -\min_{\omega \not\models \psi^{\circledcirc}_{\alpha,\beta}}\{\kappa(\bar{\alpha}), \kappa(\omega)\}$ holds. For $\omega \notin \Psi^{\circledcirc}_{\alpha,\beta}$, it holds that $\omega \models \beta$ or $\kappa(\omega) \geq \kappa(\bar{\alpha})$. Hence, only $\omega \models \beta$ are relevant for the minimum defining κ_0. Thus, we can conclude $\kappa_0 = -\min\{\kappa(\bar{\alpha}), \kappa(\beta)\}$ due to the minimality of ranks.

The definition of $\kappa^{\circledcirc}_{\alpha,\beta}$ in (8.12) displays a semi-quantitative version of RbC that implements the semantical recipe of RbC given in (8.2) for ranking functions κ in a simple, yet elegant way. The use of ranking functions makes the change in the prior belief state and the dependence on the relation between input and reference information in RbC more explicit and presents a direct translation of the change mechanism depicted in Figure 8.3. Thus, the methodology of RbC can be seen directly from Definition 8.3.1. Now, it is easy to see that the worlds satisfying $\psi^{\circledcirc}_{\alpha,\beta}$, among them the minimal $\overline{\beta}$-worlds, are shifted to the plausibility level of the minimal $\bar{\alpha}$-worlds, s.t. the following proposition holds.

Proposition 8.3.1. *Let $\kappa^{\circledcirc}_{\alpha,\beta} = \kappa \circledcirc_\alpha \beta(\omega)$ be the RbC-revised OCF from Definition 8.3.1. It holds that $\kappa^{\circledcirc}_{\alpha,\beta}(\overline{\alpha}) \leqslant \kappa^{\circledcirc}_{\alpha,\beta}(\overline{\beta})$.*

Proof. It holds that $\kappa^{\circledcirc}_{\alpha,\beta}(\overline{\alpha}) = \min_{\omega \models \overline{\alpha}}\{\kappa_0 + \kappa(\omega)\} = \kappa_0 + \kappa(\overline{\alpha})$. For $\kappa^{\circledcirc}_{\alpha,\beta}(\overline{\beta})$ we distinguish the following cases:
Case 1: Let $\psi^{\circledcirc}_{\alpha,\beta} \equiv \bot$. From Theorem 8.2.2, we can conclude that $\kappa(\omega) = \kappa^{\circledcirc}_{\alpha,\beta}(\omega)$ for all $\omega \in \Omega$ and $\kappa(\overline{\alpha}) \leqslant \kappa(\overline{\beta})$. Thus, $\kappa^{\circledcirc}_{\alpha,\beta}(\overline{\beta}) = \kappa_0 + \kappa(\overline{\beta}) \geq \kappa_0 + \kappa(\overline{\alpha}) = \kappa^{\circledcirc}_{\alpha,\beta}(\overline{\alpha})$.
Case 2: Let $\psi^{\circledcirc}_{\alpha,\beta} \not\equiv \bot$. Then there exists a minimal $\psi^{\circledcirc}_{\alpha,\beta}$-world $\omega_{\psi^{\circledcirc}_{\alpha,\beta}}$, s.t. $\omega_{\psi^{\circledcirc}_{\alpha,\beta}} \models \overline{\beta}$ minimally and we can conclude that $\kappa(\omega_{\psi^{\circledcirc}_{\alpha,\beta}}) = \kappa(\overline{\beta})$ and therefore $\kappa^{\circledcirc}_{\alpha,\beta}(\overline{\beta}) = \kappa_0 + \kappa(\overline{\alpha}) = \kappa^{\circledcirc}_{\alpha,\beta}(\overline{\alpha})$. $\square$

For worlds $\omega' \models \theta^{\circledcirc}_{\alpha,\beta}$, it holds in the RbC-revised OCF that they are strictly more plausible than worlds $\omega \models \psi^{\circledcirc}_{\alpha,\beta}$, i.e., they are indirectly rewarded by RbC.

Furthermore, it is clear that the plausibility relations among worlds not satisfying $\psi^{\circledcirc}_{\alpha,\beta}$ are kept. The corresponding ranks only change according to the normalization of the posterior ranking function. Whereas worlds in $\Psi^{\circledcirc}_{\alpha,\beta}$ are compressed onto a single plausibility level, s.t. their plausibility relations are lost. Via the following theorem, we show that (8.12) displays a proper translation of the change mechanism in (RbC1) – (RbC3) to the framework of OCFs via using the direct correspondence between OCFs and plausibilistic TPOS presented in (2.9) on page 37.

Theorem 8.3.2. *Let κ be a ranking function and $\kappa^{\circledcirc}_{\alpha,\beta} = \kappa \circledcirc_\alpha \beta$ be the RbC-revised ranking function from* (8.12). *Then the associated TPO $\preceq_{\kappa^{\circledcirc}_{\alpha,\beta}}$ via* (2.9) *satisfies (RbC1) – (RbC3).*

Proof. Via (2.9), we get the TPO $\preceq$ for the prior ranking function κ and the TPO $\preceq_{\kappa^{\circledcirc}_{\alpha,\beta}}$ from the RbC-revised OCF $\kappa^{\circledcirc}_{\alpha,\beta}$.

(RbC1): Let $\omega, \omega' \not\models \psi^{\circledcirc}_{\alpha,\beta}$, then it holds that $\kappa^{\circledcirc}_{\alpha,\beta}(\omega) = -\min\{\kappa(\overline{\alpha}), \kappa(\beta)\} + \kappa(\omega) \leqslant -\min\{\kappa(\overline{\alpha}), \kappa(\beta)\} + \kappa(\omega') = \kappa^{\circledcirc}_{\alpha,\beta}(\omega')$ holds if, and only if $\kappa(\omega) \leqslant \kappa(\omega')$, since $-\min\{\kappa(\overline{\alpha}), \kappa(\beta)\}$ is a constant factor. And therefore, via (2.9), we can conclude that (RbC1) is satisfied for $\preceq$ and $\preceq_{\kappa^{\circledcirc}_{\alpha,\beta}}$.

(RbC2): Let $\omega_{\overline{\alpha}} \in \min(\overline{\alpha}, \kappa)$ and $\omega' \models \psi^{\circledcirc}_{\alpha,\beta}$. For $\omega_{\overline{\alpha}}$, it holds that $\omega_{\overline{\alpha}} \not\models \psi^{\circledcirc}_{\alpha,\beta}$ and $\kappa(\omega_{\overline{\alpha}}) = \kappa(\overline{\alpha})$, thus $\kappa^{\circledcirc}_{\alpha,\beta}(\omega_{\overline{\alpha}}) = -\min\{\kappa(\overline{\alpha}), \kappa(\beta)\} + \kappa(\omega_{\overline{\alpha}}) = -\min\{\kappa(\overline{\alpha}), \kappa(\beta)\} + \kappa(\overline{\alpha}) = \kappa(\omega')$. Therefore, we can conclude via (2.9) that (RbC2) is satisfied and it holds that $\omega_{\overline{\alpha}} \approx_{\kappa^{\circledcirc}_{\alpha,\beta}} \omega'$.

(RbC3): Let $\omega \models \theta^{\circledcirc}_{\alpha,\beta}$ and let $\omega' \models \psi^{\circledcirc}_{\alpha,\beta}$. Since $\omega \models \theta^{\circledcirc}_{\alpha,\beta}$, it holds that $\kappa(\omega) < \kappa(\overline{\alpha})$. Thus, it holds that $\kappa^{\circledcirc}_{\alpha,\beta}(\omega) = -\min\{\kappa(\overline{\alpha}), \kappa(\beta)\} + \kappa(\omega) < -\min\{\kappa(\overline{\alpha}), \kappa(\beta)\} + \kappa(\overline{\alpha}) = \kappa^{\circledcirc}_{\alpha,\beta}(\omega')$. And via (2.9), we can conclude that (RbC3) holds for $\preceq_{\kappa^{\circledcirc}_{\alpha,\beta}}$. □

As a corollary of Theorem 8.3.2, it follows that RbC for OCFs satisfies the properties of RbC via the transformation from equation (2.9).

Theorem 8.3.3 ([109]). *Let κ be a ranking function and $\kappa^{\circledcirc}_{\alpha,\beta} = \kappa \circledcirc_\alpha \beta$. The associated TPO $\preceq_{\kappa^{\circledcirc}_{\alpha,\beta}}$ satisfies (RbC)$_\preceq$, (MinRbC)$_\preceq$, (α-level)$_\preceq$, (α-relation)$_\preceq$ and (β-level)$_\preceq$.*

Theorem 8.3.2 and 8.3.3 together show that Definition 8.3.1 is a suitable definition of RbC for ranking functions.

8.3.2 Revision by Comparison as Conditional Revision

We present and investigate a designated set of weak conditionals that characterizes Revision by Comparison's change mechanism, making the ensuing application of RbC more explicit. We show that via this set of weak conditionals, we can transfer the parameterized input for RbC to the directly usable object level and, therefore, illustrate the versatility and expressiveness of conditionals.

Consider the following set of weak conditionals.

Definition 8.3.2 ([109])**.** *Let κ be an OCF and $\psi^{\circledcirc}_{\alpha,\beta}$ as defined in Definition 8.1.2. We call the following set of weak conditionals* RbC base *and each weak conditional it contains an* RbC base conditional

$$\Delta^{\circledcirc}_{\alpha,\beta} = \{(|\overline{\alpha}|\overline{\alpha} \vee \omega|) \,|\, \omega \models \psi^{\circledcirc}_{\alpha,\beta}\}.$$

It holds that $\Delta^{\circledcirc}_{\alpha,\beta}$ is consistent since, each set of weak conditionals is consistent (cf. Proposition 2.4.1). In order to clarify the connection between $\Delta^{\circledcirc}_{\alpha,\beta}$ and RbC by an OCF κ, we examine the verification resp. falsification of each base conditional $(|\overline{\alpha}|\overline{\alpha} \vee \omega|)$. For the verification we get

$$(\overline{\alpha} \vee \omega) \wedge \overline{\alpha} \equiv \overline{\alpha} \tag{8.13}$$

and for the falsification

$$(\overline{\alpha} \vee \omega) \wedge \alpha \equiv \omega \wedge \alpha \equiv \omega, \tag{8.14}$$

since $\omega \models \psi^{\circledcirc}_{\alpha,\beta}$, i.e., $\kappa(\omega) < \kappa(\overline{\alpha})$, and therefore $\omega \models \alpha$ due to the minimaliy of ranks (cf. equation (2.7)).

Now, we investigate the usage of $\Delta^{\circledcirc}_{\alpha,\beta}$ for revisions of OCFs. Presuppose that $*$ displays a conditional revision operator for OCFs, which guarantees the acceptance of the input information, s.t. $\kappa * \Delta^{\circledcirc}_{\alpha,\beta} = \kappa^{*}_{\Delta^{\circledcirc}_{\alpha,\beta}} \models \Delta^{\circledcirc}_{\alpha,\beta}$ as the only condition for a successful revision.

The acceptance condition defined by each weak conditional in $\Delta^{\circledcirc}_{\alpha,\beta}$ provides guidance towards the necessary transformations of the prior OCF in order to achieve the success condition $(\text{RbC})_{\preceq}$ of Revision by Comparison. Each conditional $(|\overline{\alpha}|\overline{\alpha} \vee$

$\omega|) \in \Delta^{\circledcirc}_{\alpha,\beta}$ encodes a loss in plausibility for worlds falsifying the conditional. Due to the usage of numerical ranks, OCFs specify this loss in plausibility as the distance between verification and falsification. Here, the minimal loss in plausibility for each $\omega \in \Psi^{\circledcirc}_{\alpha,\beta}$, i.e., the distance between the prior and posterior ranks, is given as the absolute value norm between the corresponding verification and falsification of $(|\overline{\alpha}|\overline{\alpha} \vee \omega|)$, i.e. $|\kappa(\overline{\alpha}) - \kappa(\omega)|$. Note that this notion of distance is applicable due to the usage of integers to express plausibility, and it holds that the result of $\kappa(\overline{\alpha}) - \kappa(\omega)$ is always positive since $\kappa(\omega) < \kappa(\overline{\alpha})$ for $\omega \models \psi^{\circledcirc}_{\alpha,\beta}$.

It is clear from the definition of the weak conditionals in $\Delta^{\circledcirc}_{\alpha,\beta}$ that the loss in plausibility may differ for each world in the penalty set. Furthermore, it holds that $|\kappa(\overline{\alpha}) - \kappa(\omega)|$ displays only the minimal loss that must be covered to accept $\Delta^{\circledcirc}_{\alpha,\beta}$ since each impact factor is defined by an inequality, rather than an exact distance. Revision operators for OCFs can generally introduce empty layers or specify the distance between verification and falsification by adding a constant rank to falsifying worlds. We deal with the impact of empty layers in a later section of the following chapter. For now, we state that, generally, the acceptance of $\Delta^{\circledcirc}_{\alpha,\beta}$ is not enough to display the RbC mechanism fully. From (RbC2), we know that after an RbC of β w.r.t. α, all worlds $\omega \in \Psi^{\circledcirc}_{\alpha,\beta}$ are on the same plausibility level as minimal worlds satisfying $\overline{\alpha}$. Expressed via OCFs, this implies that they share the same rank, $\kappa^*_{\Delta^{\circledcirc}_{\alpha,\beta}}(\overline{\alpha}) = \kappa^*_{\Delta^{\circledcirc}_{\alpha,\beta}}(\omega)$ for each $\omega \models \psi^{\circledcirc}_{\alpha,\beta}$, as it is the case for $\kappa^{\circledcirc}_{\alpha,\beta}$ from Definition 8.3.1. The acceptance of $\Delta^{\circledcirc}_{\alpha,\beta}$ guarantees only one part of this equality, namely $\kappa^*_{\Delta^{\circledcirc}_{\alpha,\beta}}(\overline{\alpha}) \leqslant \kappa^*_{\Delta^{\circledcirc}_{\alpha,\beta}}(\omega)$. For the other direction of the inequality, we need a different set of weak conditionals which has a similar structure as $\Delta^{\circledcirc}_{\alpha,\beta}$, however, all conditionals $(|\overline{\alpha}|\overline{\alpha} \vee \omega|)$ are negated.

Definition 8.3.3. *Let κ be an OCF and $\psi^{\circledcirc}_{\alpha,\beta}$ as defined in Definition 8.1.2. We call the following set of weak conditionals* inverse RbC base *and each weak conditional it contains an* inverse RbC base conditional

$$(\Delta^{\circledcirc}_{\alpha,\beta})^{-1} = \{(|\alpha|\overline{\alpha} \vee \omega|) \,|\, \omega \models \psi^{\circledcirc}_{\alpha,\beta}\}.$$

Each conditional $(|\alpha|\overline{\alpha} \vee \omega|)$ in $(\Delta^{\circledcirc}_{\alpha,\beta})^{-1}$ displays the negated version of the corresponding conditional $(|\overline{\alpha}|\overline{\alpha} \vee \omega|) \in \Delta^{\circledcirc}_{\alpha,\beta}$. Thus, verification and falsification of

$(|\alpha|\overline{\alpha} \vee \omega|)$ are interchanged, s.t. (8.14) displays the verification of $(|\alpha|\overline{\alpha} \vee \omega|)$ and (8.13) its falsification. We can conclude for each OCF κ which accepts $(\Delta^{\circledcirc}_{\alpha,\beta})^{-1}$ that $\kappa(\omega) \leqslant \kappa(\overline{\alpha})$ for all $\omega \models \psi^{\circledcirc}_{\alpha,\beta}$. So, for each OCF that satisfies $\Delta^{\circledcirc}_{\alpha,\beta}$ and $(\Delta^{\circledcirc}_{\alpha,\beta})^{-1}$, i.e., $\kappa \models \Delta^{\circledcirc}_{\alpha,\beta} \cup (\Delta^{\circledcirc}_{\alpha,\beta})^{-1}$, it holds that $\kappa(\overline{\alpha}) = \kappa(\omega)$ for all $\omega \models \psi^{\circledcirc}_{\alpha,\beta}$. Note that, as for $\Delta^{\circledcirc}_{\alpha,\beta}$, it holds that $(\Delta^{\circledcirc}_{\alpha,\beta})^{-1}$ is consistent, and also, their union is consistent since all these sets solely consist of weak conditionals. Regarding RbC and its success condition, the inverse RbC base seems superfluous at first since for each $\omega \models \psi^{\circledcirc}_{\alpha,\beta}$, it holds that $\kappa(\omega) < \kappa(\overline{\alpha})$ and therefore either $(\Delta^{\circledcirc}_{\alpha,\beta})^{-1} = \emptyset$ because $\Psi^{\circledcirc}_{\alpha,\beta}$ is empty or κ already accepts $(\Delta^{\circledcirc}_{\alpha,\beta})^{-1}$. Yet, it adds to the understanding of the mechanism of RbC since the two sets $\Delta^{\circledcirc}_{\alpha,\beta}$ and $(\Delta^{\circledcirc}_{\alpha,\beta})^{-1}$ together fixate the exact transformation for all worlds $\omega \models \psi^{\circledcirc}_{\alpha,\beta}$. So, again, if we consider a general revision operator for OCFs $*$ which satisfies the (|Success|) condition from page 49, s.t. $\kappa * \Delta^{\circledcirc}_{\alpha,\beta} \cup (\Delta^{\circledcirc}_{\alpha,\beta})^{-1} = \kappa^{*}_{\Delta^{\circledcirc}_{\alpha,\beta} \cup (\Delta^{\circledcirc}_{\alpha,\beta})^{-1}} \models \Delta^{\circledcirc}_{\alpha,\beta} \cup (\Delta^{\circledcirc}_{\alpha,\beta})^{-1}$, then it holds that

$$\kappa^{*}_{\Delta^{\circledcirc}_{\alpha,\beta} \cup (\Delta^{\circledcirc}_{\alpha,\beta})^{-1}}(\omega) = \kappa(\omega) + |\kappa(\overline{\alpha}) - \kappa(\omega)| = \kappa(\overline{\alpha})$$

for $\omega \models \psi^{\circledcirc}_{\alpha,\beta}$, where $|\kappa(\overline{\alpha}) - \kappa(\omega)|$ displays the change of ranks which we need to implement to achieve $\kappa^{*}_{\Delta^{\circledcirc}_{\alpha,\beta} \cup (\Delta^{\circledcirc}_{\alpha,\beta})^{-1}} \models (|\overline{\alpha}|\overline{\alpha} \vee \omega|)$ and $\kappa^{*}_{\Delta^{\circledcirc}_{\alpha,\beta} \cup (\Delta^{\circledcirc}_{\alpha,\beta})^{-1}} \models (|\alpha|\overline{\alpha} \vee \omega|)$ at the same time. As we can see here, this corresponds to the definition of ranks for worlds $\omega \models \psi^{\circledcirc}_{\alpha,\beta}$ from (8.12), i.e., RbC for OCFs. So, for $\kappa^{*}_{\Delta^{\circledcirc}_{\alpha,\beta} \cup (\Delta^{\circledcirc}_{\alpha,\beta})^{-1}}$, we can conclude for each conditional revision operator, which ensures the acceptance of the input information, that the postulate (RbC2) is satisfied. We sum this up in the following theorem.

Theorem 8.3.4. *For an OCF κ and a conditional revision operator $*$ which takes a weak conditional belief base Δ^{w} as input, s.t. $\kappa * \Delta^{\mathrm{w}} \models \Delta^{\mathrm{w}}$, it holds that $\kappa * \Delta^{\circledcirc}_{\alpha,\beta} \cup (\Delta^{\circledcirc}_{\alpha,\beta})^{-1}$ satisfies (RbC2), i.e., the acceptance of $\Delta^{\circledcirc}_{\alpha,\beta} \cup (\Delta^{\circledcirc}_{\alpha,\beta})^{-1}$ corresponds to the constraint in (RbC2).*

Note that, in this section, the quantitative revision operator $*$ remains vague apart from some success condition that allows us to accept the new input. This is because, in this section, we focused on the core aspect of the RbC mechanism also depicted in our schematic illustration in Figure 8.3; namely, that worlds satisfying the penalty formula are shifted up to the level of plausibility where the $\overline{\alpha}$-worlds remain. To show the remaining postulates (RbC1) and (RbC3) and thus capture the full RbC mechanism, we need some additional minimal change paradigm implemented in the revision.

The advantage of capturing RbC as a revision with sets of weak conditionals lies in transforming the operation from the meta-level to the object level. The sets $\Delta^{\circledcirc}_{\alpha,\beta}$ and $(\Delta^{\circledcirc}_{\alpha,\beta})^{-1}$ substantiate the supplementary information given in the reference sentence α. And hence, provide a more clearly defined and directly usable tool for Revsion by Comparison. Yet, the revision task is quite a challenging one. We need a revision operator that is capable of dealing with sets of weak conditionals simultaneously. In the next section, we show that c-revisions with sets of weak conditionals serve as a proof of concept for RbC of β w.r.t. α.

8.3.3 Realization of Revision by Comparison as C-Revision

Previously, we discussed how the change mechanism underlying RbC can be realized via a set of weak conditionals. Now, taking these conceptual results, we realize RbC for OCFs from Definition 8.3.1 as a c-revision with sets of weak conditionals. This reveals the real character of RbC as an iterated contraction operation and makes RbC directly usable for existing frameworks of belief revision which are capable of revising with conditional information, such as the one presented in [47].

So far, we have seen that RbC of β w.r.t. α yields a posterior ordering of worlds by shifting worlds from the penalty set $\Psi^{\circledcirc}_{\alpha,\beta}$ up to the level of plausibility where $\overline{\alpha}$-worlds reside. This core shift of RbC is subsumed by the second postulate for Revision by Comparison (RbC2) from the Theorem 8.1.4. In the previous section, we have shown that we can capture this shift via a revision with sets of weak conditionals $\Delta^{\circledcirc}_{\alpha,\beta}$ and $(\Delta^{\circledcirc}_{\alpha,\beta})^{-1}$. Now, we focus on the revision with $\Delta^{\circledcirc}_{\alpha,\beta}$ and show that it is sufficient to c-revise solely with this set of weak conditionals to capture the mechanism of RbC.

C-Revisions provide a highly general framework for revising OCFs with sets of standard and weak conditionals. They, therefore can be used as a proof of concept for RbC as a conditional revision for OCFs. In Definition 2.5.2 from page 49 we presented c-revisions with sets of weak conditionals. Consider the set of weak conditionals $\Delta^{\circledcirc}_{\alpha,\beta}$ from Definition 8.3.2. For the c-revision $\kappa *^{c} \Delta^{\circledcirc}_{\alpha,\beta}$, we get via (2.24) from page 50

$$\kappa *^{\mathrm{c}} \Delta^{\circledcirc}_{\alpha,\beta}(\omega) = \kappa^{\mathrm{c}}_{\alpha,\beta}(\omega) = \kappa_0 + \kappa(\omega) + \sum_{\substack{\omega \models ((\overline{\alpha} \vee \omega') \wedge \alpha), \\ \omega' \models \psi^{\circledcirc}_{\alpha,\beta}}} \eta_{\omega'} \tag{8.15}$$

$$= \kappa_0 + \kappa(\omega) + \sum_{\substack{\omega \models \omega', \\ \omega' \models \psi^{\circledcirc}_{\alpha,\beta}}} \eta_{\omega'}. \tag{8.16}$$

Note that, (8.16) follows from (8.15) due to the equivalence given in (8.14) for $\omega' \models \psi^{\circledcirc}_{\alpha,\beta}$. It holds that each $\omega \models \psi^{\circledcirc}_{\alpha,\beta}$ falsifies its corresponding condition $(|\overline{\alpha}|\overline{\alpha} \vee \omega|)$ from $\Delta^{\circledcirc}_{\alpha,\beta}$. Thus, for each conditional $(|\overline{\alpha}|\overline{\alpha} \vee \omega|) \in \Delta^{\circledcirc}_{\alpha,\beta}$ only a single conditional is falsified and therefore the definition of c-revision with weak conditionals in (2.24) reduces to the following:

$$\kappa^{\mathrm{c}}_{\alpha,\beta}(\omega) = \kappa_0 + \kappa(\omega) + \begin{cases} \eta_\omega, & \omega \models \psi^{\circledcirc}_{\alpha,\beta} \\ 0, & \text{othw.} \end{cases} \tag{8.17}$$

The normalization constant κ_0, can be further substantiated via (2.25) from Definition 2.5.2

$$\kappa_0 = -\min_{\omega \in \Omega} \left\{ \kappa(\omega) + \begin{cases} \eta_\omega, & \omega \models \psi^{\circledcirc}_{\alpha,\beta} \\ 0, & \text{othw.} \end{cases} \right\}. \tag{8.18}$$

So, via the inequalities in (2.26), which constrain the impact factors for a general c-revision with weak conditionals, and the verification (8.13) resp. falsification (8.14) of $(|\overline{\alpha}|\overline{\alpha} \vee \omega|) \in \Delta^{\circledcirc}_{\alpha,\beta}$, we get the following inequality defining the impact factor for the corresponding conditional from $(|\overline{\alpha}|\overline{\alpha} \vee \omega|) \in \Delta^{\circledcirc}_{\alpha,\beta}$:

$$\eta_\omega \geq \min_{\omega' \models \overline{\alpha}} \{\kappa(\omega') + \sum_{\substack{\omega' \models \tilde{\omega}, \\ \tilde{\omega} \models \psi^{\circledcirc}_{\alpha,\beta}, \tilde{\omega} \neq \omega}} \eta_{\tilde{\omega}}\} - \min_{\omega' \models \omega} \{\kappa(\omega') + \sum_{\substack{\omega' \models \tilde{\omega}, \\ \tilde{\omega} \models \psi^{\circledcirc}_{\alpha,\beta}, \tilde{\omega} \neq \omega}} \eta_{\tilde{\omega}}\} = \kappa(\overline{\alpha}) - \kappa(\omega). \tag{8.19}$$

Note that, in both minima the sums equal zero, since for possible worlds $\tilde{\omega} \neq \omega$

$\omega' \models \tilde{\omega}$ is never satisfied. Then, via employing the definition of OCF ranks, we obtain this compact inequality. In general, the impact factors defining the c-revisions are not uniquely defined since the solution of the system of inequalities defining each c-revision is not unique. This is because verification and falsification from different conditionals may interact with each other, introducing dependencies among the choices of impact factors for c-revisions. For the c-revision with $\Delta^{\circledcirc}_{\alpha,\beta}$, it holds that each world from $\Psi^{\circledcirc}_{\alpha,\beta}$ falsifies a single conditional, and thus these conditionals do not interact. This enables us to define η_ω unambiguously by integers that satisfy the inequality in (8.19). A straightforward choice for η_ω is to choose minimal impact factors, s.t.

$$\eta^{\min}_{\omega} = \kappa(\overline{\alpha}) - \kappa(\omega) \tag{8.20}$$

holds. Note that, $\eta^{\min}_{\omega}$ is always non-negative since $\kappa(\omega) < \kappa(\overline{\alpha})$ for all $\omega \models \psi^{\circledcirc}_{\alpha,\beta}$. So that, we obtain the following minimal c-revision w.r.t. the impact factors.

$$\kappa^{\mathrm{c,min}}_{\alpha,\beta}(\omega) = \kappa_0 + \kappa(\omega) + \begin{cases} \kappa(\overline{\alpha}) - \kappa(\omega), & \omega \models \psi^{\circledcirc}_{\alpha,\beta} \\ 0, & \text{othw.} \end{cases} \tag{8.21}$$

$$= \kappa_0 + \begin{cases} \kappa(\overline{\alpha}), & \omega \models \psi^{\circledcirc}_{\alpha,\beta} \\ \kappa(\omega), & \text{othw.} \end{cases} \tag{8.22}$$

We can further substantiate the normalization constant κ_0 from (8.22) for the minimal c-revision via (8.18) as follows

$$\begin{aligned} \kappa_0 &= -\min\{\kappa(\overline{\alpha}), \min_{\omega \not\models \psi^{\circledcirc}_{\alpha,\beta}} \kappa(\omega)\} = -\min\{\kappa(\overline{\alpha}), \kappa(\overline{\psi^{\circledcirc}_{\alpha,\beta}})\} \\ &= -\min\{\kappa(\overline{\alpha}), \kappa(\beta \vee \bigvee_{\kappa(\overline{\alpha}) \leqslant \kappa(\omega)} \omega)\} = -\min\{\kappa(\overline{\alpha}), \min\{\kappa(\beta), \kappa(\overline{\alpha})\}\} \\ &= -\min\{\kappa(\overline{\alpha}), \kappa(\beta)\} \end{aligned} \tag{8.23}$$

The normalization constant from (8.23) is the same as the normalization constant for RbC for OCFs from (8.12). Now, it is easy to see that the minimal c-revision in (8.22) is the same as RbC for OCFs in (8.12) from Definition 8.3.1. We can state the following theorem:

Theorem 8.3.5. *Let κ be a ranking function. For the minimal c-revision $\kappa_{\alpha,\beta}^{\mathrm{c,min}}$ from* (8.22) *and $\kappa_{\alpha,\beta}^{\circledcirc} = \kappa \circledcirc_\alpha \beta$ from* (8.12)*, it holds that*

$$\kappa_{\alpha,\beta}^{\mathrm{c,min}}(\omega) = \kappa_{\alpha,\beta}^{\circledcirc}(\omega)$$

for all $\omega \in \Omega$.

Proof. From (8.12) and (8.22), it is immediately apparent that, for all $\omega \in \Omega$ it holds that $\kappa_{\alpha,\beta}^{\mathrm{c,min}}(\omega) = \kappa_{\alpha,\beta}^{\circledcirc}(\omega)$. □

Choosing the minimal impact factor $\eta_\omega^{\min}$ in (8.22) is crucial to obtain the equality $\kappa_{\alpha,\beta}^{\mathrm{c,min}} = \kappa_{\alpha,\beta}^{\circledcirc}$. Theorem 8.3.5 reveals a significant new insight into the change mechanism behind RbC. Since we can characterize it as a c-revision with weak conditionals, it holds that RbC corresponds to an iterated contraction operator with a designated set of conditionals in the context of OCFs as discussed in Section 2.5.4. For each world $\omega \models \psi_{\alpha,\beta}^{\circledcirc}$, we can define a weak conditional corresponding to the negated information we want to revise with resp. the corresponding standard conditional we want to contract. Then it holds that RbC satisfies the principle of conditional preservation, at least concerning this special set of conditionals. Since c-revisions are capable of revising with sets of weak conditionals simultaneously, we can take the whole set $\Delta_{\alpha,\beta}^{\circledcirc}$ as input leading to an RbC-revised OCF in a single revision step.

The following corollary follows directly from the transformation (2.9).

Corollary 8.3.6 ([109])**.** *Let κ be a ranking function. For the minimal c-revision $\kappa_{\alpha,\beta}^{\mathrm{c,min}}$ with $\Delta_{\alpha,\beta}^{\circledcirc}$ from* (8.22) *and $\kappa_{\alpha,\beta}^{\circledcirc} = \kappa \circledcirc_\alpha \beta$ from* (8.12)*, it holds that the corresponding plausibilistic TPOs $\preceq_{\kappa_{\alpha,\beta}^{\mathrm{c}}}$ and $\preceq_{\kappa_{\alpha,\beta}^{\circledcirc}}$ are the same, i.e.,*

$$\omega \preceq_{\kappa_{\alpha,\beta}^{\mathrm{c,min}}} \omega' \textit{ iff } \omega \preceq_{\kappa_{\alpha,\beta}^{\circledcirc}} \omega'.$$

Proof. From (8.12) and (8.22), it is immediately apparent that, for all $\omega \in \Omega$ it holds that $\kappa_{\alpha,\beta}^{\mathrm{c,min}}(\omega) = \kappa_{\alpha,\beta}^{\circledcirc}(\omega)$. And therefore, we can conclude via (2.9) that $\preceq_{\kappa_{\alpha,\beta}^{\mathrm{c,min}}} = \preceq_{\kappa_{\alpha,\beta}^{\circledcirc}}$ holds. □

$\omega \in \Omega$	κ	$\kappa^{\circledcirc}_{a,b}$	$\kappa^{\mathrm{c,min}}_{a,b}$	$\kappa^{\circledcirc}_{c,\bar{a}}$	$\kappa^{\mathrm{c,min}}_{c,\bar{a}}$
$abc\overline{d}$	0	0	0	0	0
$abcd$	1	1	1	0	0
$a\bar{b}cd$	1	3	3	0	0
$ab\bar{c}d$	2	2	2	0	0
$a\bar{b}\bar{c}d$	2	3	3	0	0
$\bar{a}bcd$	3	3	3	1	1
$\bar{a}b\bar{c}\overline{d}$	4	4	4	2	2
$a\bar{b}\bar{c}\overline{d}$	4	4	4	2	2
$\bar{a}bcd$	5	5	5	3	3
$\bar{\Omega}$	6	6	6	4	4

Table 8.1: Prior κ and the RbC-revised $\kappa^{\circ}_{\alpha,\beta}$ resp. c-revised $\kappa^{\mathrm{c}}_{\alpha,\beta}$.

Since $\preceq_{\kappa^{\circledcirc}_{\alpha,\beta}}$ satisfies (RbC1) – (RbC3), it follows immediately from Theorem 8.3.5 that $\preceq_{\kappa^{\mathrm{c,min}}_{\alpha,\beta}}$ also satisfies (RbC1) – (RbC3). Note that, for $\preceq_{\kappa^{\mathrm{c,min}}_{\alpha,\beta}}$ to satisfy (RbC2), i.e., $\omega \approx_{\kappa^{\mathrm{c,min}}_{\alpha,\beta}} \overline{\alpha}$ for $\omega \models \psi^{\circledcirc}_{\alpha,\beta}$, it is crucial to choose the minimal impact factor $\eta^{\min}_{\omega}$ in (8.22).

In general, because we chose minimal impact factors that satisfy (8.19), it suffices to revise with $\Delta^{\circledcirc}_{\alpha,\beta}$ to obtain (RbC2) and we can omit the revision with the additional set $(\Delta^{\circledcirc}_{\alpha,\beta})^{-1}$, leading to a leaner revision mechanism than in Theorem 8.3.4.

We illustrate Theorem 8.3.5 via the following example.

Example 8.3.1. *In Table 8.1 the convex prior ranking function κ which corresponds to $\preceq_\kappa$ from Example 8.1.1 is depicted, alongside with the two ranking functions $\kappa^{\circledcirc}_{\alpha,\beta} = \kappa \circledcirc_\alpha \beta$ resp. $\kappa^{\mathrm{c,min}}_{\alpha,\beta} = \kappa \circ \delta_{\alpha,\beta}$ for each corresponding RbC.*

For the first RbC by $\beta = b$ w.r.t. $\alpha = a$, we get the following penalty set $\Phi^{\circledcirc}_{a,b} = Mod(\bar{b} \wedge (\bigvee_{\kappa(\omega)<\bar{a}} \omega)) = \{a\bar{b}\bar{c}d, a\bar{b}cd\}$. RbC for OCFs $\kappa^{\circledcirc}_{a,b}(\omega)$ as defined in Definition 8.3.1 lifts these worlds onto the plausibility level of minimal worlds satisfying $\bar{a}$, here $\kappa(\bar{a}bcd) = 3 = \kappa(\bar{a})$. For the set of weak conditionals $\Delta^{\circledcirc}_{a,b} = \{(|\bar{a} \vee a\bar{b}\bar{c}d|\bar{a}|), (|\bar{a} \vee a\bar{b}cd|\bar{a}|)\}$, we get the minimal c-revision $\kappa^{\mathrm{c,min}}_{a,b}$ from (8.22). It holds that $\kappa^{\circledcirc}_{a,b}(\omega) = \kappa^{\mathrm{c,min}}_{a,b}(\omega)$ yield the same posterior OCF. Note that, for both mechanisms $\kappa^{\circledcirc}_{a,b}$ and $\kappa^{\mathrm{c,min}}_{a,b}$ we get the same normalization constant $\kappa_0 = -\min\{\kappa(\overline{a}), \kappa(b)\} = 0$. Both ranking functions are depicted in Table 8.1.

For the next RbC with reference $\alpha = \bar{a}$ *and input sentence* $\beta = c$*, it holds that* $\Psi^{\circledcirc}_{c,\overline{a}} = \{abcd, a\bar{b}cd, abc\bar{d}\}$ *and* $\Theta^{\circledcirc}_{c,\overline{a}} = \emptyset$*, i.e., the* α*-contraction case of RbC applies and it holds that* $\kappa^{\circledcirc}_{c,\overline{a}}, \kappa^{\mathrm{c,min}}_{c,\overline{a}} \not\models c$*. For the normalization constants of* $\kappa^{\circledcirc}_{c,\overline{a}}$*, we get* $\kappa_0 = -\min\{\kappa(\overline{c}), \kappa(\overline{a})\} = -2$*. For* $\kappa^{\mathrm{c,min}}_{c,\overline{a}}$*, we c-revise* κ *with the RbC base* $\Delta^{\circledcirc}_{c,\overline{a}} = \{(|\bar{a} \vee abc\bar{d}|\bar{c}|), (|\bar{a} \vee abcd|\bar{c}|), (|\bar{a} \vee a\bar{b}cd|\bar{c}|)\}$*. The normalization constant of* $\kappa^{\circledcirc}_{c,\overline{a}}$ *and* $\kappa^{\mathrm{c,min}}_{c,\overline{a}}$ *are the same, i.e.,* $\kappa_0 = \min\{\kappa(\overline{c}), \kappa(\overline{a})\} = -2$ *also for the c-revision. The posterior OCFs* $\kappa^{\circledcirc}_{c,\overline{a}}$ *and* $\kappa^{\mathrm{c,min}}_{c,\overline{a}}$ *are equal, both are depicted in Table 8.1.*

Note that, it holds for all RbC-revised and c-revised ranking functions in the above stated example, that their corresponding plausibilistic TPOs coincide with the RbC-revised TPOs from Figures 8.1b and 8.5c.

Chapter 9

Bounded Revision

Bounded Revision (BR) was firstly defined by Rott in [100] and takes as input a piece of new information $\beta \in \mathcal{L}$ which is accompanied by a designated reference sentence $\alpha \in \mathcal{L}$, representing some kind of meta-information about the input. Thus, it displays a belief revision operator that offers a strategy for solving the advanced belief revision problem (ParameterRev) stated in the introduction of this part.

Similar to RbC, the general idea of BR is that α acts as a guiding parameter for the depth with which β shall be accepted in the posterior state. In this respect, it shares the same underlying motivation with Revision by Comparison, the parameterized belief change operator we discussed in the previous chapter. While RbC and BR may appear similar on a basic level, essential divergences distinguish these two operations from one another. To elucidate these distinctions, we briefly examine the fundamental differences between RbC and BR, which helps to motivate our forthcoming investigations on BR. First, RbC tends to coarsen the agent's belief state by lifting worlds from different plausibility levels onto the same level. This makes it less suitable for iteration since the number of levels within the plausibilistic TPO shrinks. RbC does not satisfy the DP postulates for iterated belief revision (C1) – (C4). On the other hand, Bounded Revision is an iterated revision operator in the sense of Darwiche and Pearl [100], which leads us to the next fundamental difference between RbC and BR. While RbC implements a hybrid belief change operator between revision with the input and contraction of the reference sentence (cf. Section 8.2), BR displays a "real" revision operator, in the sense that the input is always accepted independent from the reference sentence, this provides coherence across revision scenarios with different reference sentences. The distinctive character of BR as an operator for iterated revision of beliefs makes it an attractive parameterized revision operator that we consider and study independently from RbC. And it holds that, in contrast to RbC, for BR, the influence of the parameter α is less apparent on the belief set because α is always accepted. Note that a thorough comparison of RbC vs. BR concludes this part. For now, our investigation of BR is mainly motivated by the fact that, to the best of our knowledge, it displays the only iterated parameterized revision operator in the sense of the belief revision problem (ParameterRev).

Rott investigated BR in [100] in terms of qualitative approaches and showed that BR satisfies the DP postulates for iterated revision and established connections with two well-known iterated revision operators. Starting from the investigations in [100], our primary goal is to incorporate the parameterized information in a single, more easily accessible input information while keeping relevant features of BR. Ultimately we present a conditional that subsumes the iterated change of BR providing grounds for (at least prototypical) applications in the framework of c-revisions (or other conditional revision operators). This implementation of BR underpins the diversity and flexibility that conditionals offer us as input for belief revision operators.

The following sections of this chapter are organized as follows: In Section 9.1, we investigate BR in the framework of plausibilistic TPOs, making its underlying change strategy more explicit via a unique formula and a representation theorem. After considering some limiting special cases in Section 9.2, we turn to the framework of OCFs in Section 9.3, where we first present a straightforward implementation of BR in Subsection 9.3.1. Eventually, by employing our previous results, we are able to present a designated conditional that elegantly subsumes the mechanism of BR and allows us to (methodologically) implement BR as a conditional c-revision.

Bibliographic Remark. The contents of this part are based on joint work with Gabriele Kern-Isberner [110] (see Section 1.3).

9.1 Mechanism of Bounded Revision for TPOs

In this section, we investigate the mechanism of BR as a parameterized iterated belief revision with an input sentence β and a reference sentence α for plausibilistic TPOs $\preceq$ as a representation of belief states. Via more comprehensible constraints for TPOs, we can specify crucial worlds characterizing the change mechanism of BR, leading to a representation theorem characterizing BR and, thus, clarifying the role of parameter α.

The constraints in (7.4) on page 147, defining BR for epistemic entrenchment relations, can be transferred to constraints for plausibilistic TPOs via (7.1) on page 140. Thus, the constraints from Definition 7.2.4 for $\leqslant_E$ resp. $\leqslant_E^{\circ_{\alpha,\beta}}$ yield the following qualitative constraints for TPOs $\preceq$ resp. $\preceq^{\circ}_{\alpha,\beta}$

Definition 9.1.1 ([110])**.** *Let $\preceq$ be a plausibilistic TPO and $\alpha, \beta \in \mathcal{L}$. The* Bounded Revision by β w.r.t. α of the plausibilistic TPO $\preceq$, $\preceq^{\circ}_{\alpha,\beta} = \preceq \circ_\alpha \beta$*, is defined as follows*

$$\omega \preceq^{\circ}_{\alpha,\beta} \omega' \;\; \textit{iff} \; \begin{cases} \beta\omega \preceq \beta\omega', \; \textit{if } \beta \wedge (\omega \vee \omega') \preceq \bar{\alpha}\beta & \text{(I)} \\ \quad \omega \preceq \omega', \; \textit{otherwise} & \text{(II)} \end{cases} \tag{9.1}$$

As for RbC, these constraints follow immediately from (7.4) via (2.9) from page 37 if we take maximal disjunctions for γ and δ. Then their negations $\bar{\gamma}$ and $\bar{\delta}$ are maximal conjunctions, which correspond to possible worlds ω and ω' in the set Ω.

Again, via applying (7.1), we transfer the success condition for BR $(\mathrm{BR})_E$ and the Same Beliefs Condition $(\mathrm{SBC})_E$ to the framework of plausibilistic TPOs.

(BR)$_\preceq$ $\bar{\alpha}$ is strictly more plausible than $\bar{\beta}$: $\bar{\alpha} \prec^{\circ}_{\alpha,\beta} \bar{\beta}$

(SBC)$_\preceq$ $Bel(\preceq \circ_\alpha \beta) = Bel(\preceq \circ_\gamma \beta)$ for any $\alpha, \gamma \in \mathcal{L}$

Since (9.1) is an equivalent reformulation of Definition 7.2.4 in the context of TPOs, it is obvious that $(\mathrm{BR})_\preceq$ and $(\mathrm{SBC})_\preceq$ hold for each BR revised TPO $\preceq^{\circ}_{\alpha,\beta}$ defined by (9.1), and BR for plausibilistic TPOs satisfies the DP postulates since BR for entrenchment relations does [100]. Also, $\preceq^{\circ}_{\alpha,\beta} \models \beta$ holds, i.e., $\preceq^{\circ}_{\alpha,\beta}$ displays a revision with β.

From (7.4) and thus, also from (9.1), it remains unclear prima facie which worlds exactly are affected by the change mechanism implemented by BR of β w.r.t. α. Also

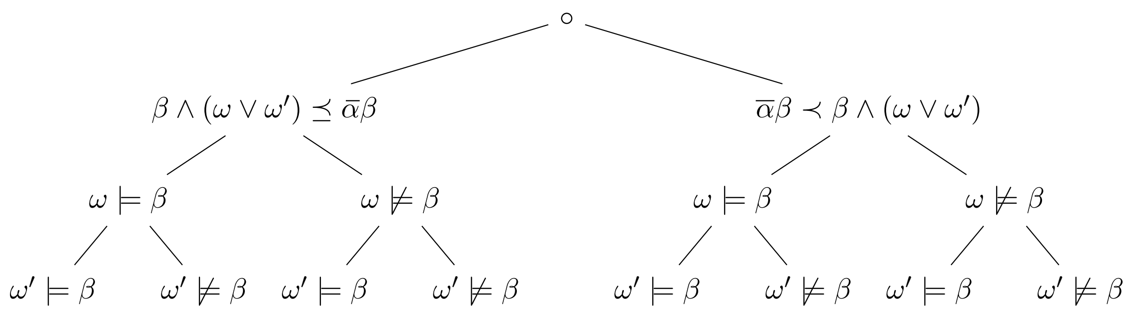

Figure 9.1: Overview of the case distinction from the proof of Proposition 9.1.1.

the crucial role of the relative positioning of the input β to the reference α is not clearly recognizable. The following proposition fully integrates BR by β w.r.t. α in the possible worlds reading via constraints. These constraints are helpful to clarify the semantical recipe behind BR.

Proposition 9.1.1. *For BR $\preceq^{\circ}_{\alpha,\beta}$ by β w.r.t. α, it holds that* (9.1) *is equivalent to the following constraints:*

$$\omega \preceq^{\circ}_{\alpha,\beta} \omega' \textit{ iff } \begin{cases} \omega \preceq \omega', \textit{ if } \big(\omega, \omega' \models \beta \textit{ and } (\omega, \omega' \preceq \overline{\alpha}\beta \textit{ or } \overline{\alpha}\beta \prec \omega, \omega')\big) & \text{(I) } \textit{or} \text{ (II)} \\ \qquad \textit{or } \big(\omega, \omega' \models \overline{\beta}\big) & \text{(III)} \\ \qquad \textit{or } \big(\omega \models \beta,\, \omega' \models \overline{\beta} \textit{ and } \overline{\alpha}\beta \prec \omega\big) & \text{(IV)} \\ \qquad \textit{or } \big(\omega' \models \beta,\, \omega \models \overline{\beta} \textit{ and } \overline{\alpha}\beta \prec \omega'\big) & \text{(V)} \\ \\ \top, \textit{ if } \big(\omega \models \beta, \omega \preceq \overline{\alpha}\beta \textit{ and} & \\ \quad (\omega' \models \overline{\beta} \textit{ or } \omega' \models \beta, \overline{\alpha}\beta \prec \omega')\big) & \text{(VI) } \textit{or} \text{ (VII)} \end{cases} \tag{9.2}$$

Proof. In order to prove equivalence between (9.1) and (9.2), we expand and equivalently reorder the constraints in (9.1) until we receive (9.2).
We start with (I) from (9.1) and presuppose that $\beta \wedge (\omega \vee \omega') \preceq \overline{\alpha}\beta$ holds. Then, we distinguish four exclusive and exhaustive cases with regard to whether $\omega, \omega' \models \beta$ or not and investigate the effects on $\beta\omega \preceq \beta\omega'$. Afterwards, we continue with (II) from (9.1) and presuppose that $\overline{\alpha}\beta \prec \beta \wedge (\omega \vee \omega')$ holds and split up four cases again, with regard to whether $\omega, \omega' \models \beta$ or not. An overview of the exhaustive and exclusive case distinction considered in this proof can be found in Figure 9.1.

Presuppose that $\beta \wedge (\omega \vee \omega') \preceq \bar{\alpha}\beta$.

1. $\underline{\omega \models \beta:}$

   (a) **Presuppose that $\omega' \models \beta$.** It holds that

   $$\beta \wedge (\omega \vee \omega') \equiv (\omega \vee \omega') \preceq \overline{\alpha}\beta \text{ and } \beta\omega \preceq \beta\omega' \text{ iff } \omega \preceq \omega'.$$

   From the first condition, $(\omega \vee \omega') \preceq \overline{\alpha}\beta$, we get another case distinction:

   i. If $\omega, \omega' \preceq \overline{\alpha}\beta$, then (I) from (9.1) is equivalent to $\omega \preceq^{\circ}_{\alpha,\beta} \omega'$ iff $\omega \preceq \omega'$.

   ii. If $\omega \preceq \overline{\alpha}\beta$ and $\overline{\alpha}\beta \prec \omega'$, then $\omega \preceq \omega'$ holds trivially and we get that (I) from (9.1) is equivalent to $\omega \preceq^{\circ}_{\alpha,\beta} \omega'$ iff $\top$.

   iii. If $\omega' \preceq \overline{\alpha}\beta$ and $\overline{\alpha}\beta \prec \omega$, then $\omega \preceq \omega'$ is never satisfied and (9.1) is not applicable.

   (b) **Presuppose that $\omega' \not\models \beta$.** It holds that

   $$\beta \wedge (\omega \vee \omega') \equiv \omega \preceq \overline{\alpha}\beta \text{ and } \beta\omega \preceq \beta\omega' \text{ iff } \omega \preceq \bot \text{ iff } \top.$$

   s.t. (I) from (9.1) is equivalent to ($\top$, if $\omega \models \beta$ and $\omega' \not\models \beta$ and $\omega \preceq \overline{\alpha}\beta$)

2. $\underline{\omega \not\models \beta:}$

   (a) **Presuppose that $\omega' \models \beta$.** It holds that $\beta \wedge (\omega \vee \omega') \equiv \omega' \preceq \overline{\alpha}\beta$ and $\beta\omega \preceq \beta\omega'$ iff $\bot \preceq \omega'$. s.t. (I) from (9.1) is not applicable.

   (b) **Presuppose that $\omega' \not\models \beta$.** It holds that $\beta \wedge (\omega \vee \omega') \equiv \bot \preceq \overline{\alpha}\beta$ s.t. (I) from (9.1) is not applicable.

From these case distinctions, we get that (9.1) is equivalent to

$$\omega \preceq^{\circ}_{\alpha,\beta} \omega' \text{ iff } \begin{cases} \omega \preceq \omega', \text{ if } \omega, \omega' \models \beta \text{ and } \omega, \omega' \preceq \overline{\alpha}\beta & \text{via case 1.a.i)} \\ \quad \top, \text{ if } \omega \models \beta, \omega \preceq \bar{\alpha}\beta \text{ and} & \\ \quad\quad \left(\omega' \models \overline{\beta} \text{ or } \omega' \models \beta, \bar{\alpha}\beta \prec \omega'\right) & \text{via case 1.b) or 1.a.ii)} \\ \omega \preceq \omega', \text{ otherwise} & \text{via (II) from (9.1)} \end{cases} \tag{9.3}$$

Note that, since case 1.a.iii) and case 2 are not applicable, they are not included in (9.3). The constraint from case 1.a.i) resp. case 1.b) resp. case 1.a.ii) corresponds to (I) resp. (VI) resp. (VII) from (9.2).

The last case in (9.3) corresponds to the case, when $\beta \wedge (\omega \vee \omega') \preceq \bar{\alpha}\beta$ does not hold. **Presuppose that $\overline{\alpha}\beta \prec \beta \wedge (\omega \vee \omega')$.**

1. <u>$\omega \models \beta$.</u>

   (a) **Presuppose that** $\omega' \models \beta$. It holds that $\overline{\alpha}\beta \prec \beta \wedge (\omega \vee \omega') \equiv (\omega \vee \omega')$. s.t. (II) from (9.2) is equivalent to $\omega \preceq^{\circ}_{\alpha,\beta} \omega'$ iff $\omega \preceq \omega'$, if $\omega, \omega' \models \beta$ and $\overline{\alpha}\beta \prec \omega, \omega'$.

   Note that, we do not need any more fine-grained case distinction for $\overline{\alpha}\beta \prec (\omega \vee \omega')$, since the plausibility of $(\omega \vee \omega')$ is determined by minimal worlds satisfying the formula.

   (b) **Presuppose that** $\omega' \not\models \beta$**.** It holds that $\overline{\alpha}\beta \prec \beta \wedge (\omega \vee \omega') \equiv \omega$. s.t. (IV) from (9.2) is equivalent to $\omega \preceq^{\circ}_{\alpha,\beta} \omega'$ iff $\big(\omega \preceq \omega'$, if $\omega \models \beta$, $\overline{\alpha}\beta \prec \omega$ and $\omega' \not\models \beta\big)$.

2. <u>$\omega \not\models \beta$:</u>

   (a) **Presuppose that** $\omega' \models \beta$**.** It holds that $\overline{\alpha}\beta \prec \beta \wedge (\omega \vee \omega') \equiv \omega'$. s.t. (V) from (9.2) is equivalent to $\omega \preceq^{\circ}_{\alpha,\beta} \omega'$ iff $\omega \preceq \omega'$, if $\omega \not\models \beta$, $\omega' \models \beta$ and $\overline{\alpha}\beta \prec \omega'$.

   (b) **Presuppose that** $\omega' \not\models \beta$**.** It holds that $\overline{\alpha}\beta \prec \beta \wedge (\omega \vee \omega') \equiv \perp$, s.t. we are always in case (II) from (9.1), i.e., $\omega \preceq^{\circ}_{\alpha,\beta} \omega'$ iff $\omega \preceq \omega'$, and (III) from (9.2) holds.

From these case distinctions, we get that (9.3) is equivalent to

$$\omega \preceq^{\circ}_{\alpha,\beta} \omega' \text{ iff } \begin{cases} \omega \preceq \omega', \text{ if } \omega,\omega' \models \beta \text{ and } \omega,\omega' \preceq \overline{\alpha}\beta & \\ \qquad \top, \text{ if } \omega \models \beta, \omega \preceq \overline{\alpha}\beta \text{ and} & \\ \qquad\quad \big(\omega' \models \overline{\beta} \text{ or } \omega' \models \beta, \overline{\alpha}\beta \prec \omega'\big) & \\ \omega \preceq \omega', \text{ if } \omega,\omega' \models \beta \text{ and } \overline{\alpha}\beta \prec \omega,\omega' & \text{via case 1.a)} \\ \omega \preceq \omega', \text{ if } \omega \models \beta,\ \omega' \models \overline{\beta} \text{ and } \overline{\alpha}\beta \prec \omega & \text{via case 1.b)} \\ \omega \preceq \omega', \text{ if } \omega \models \overline{\beta},\ \omega' \models \beta \text{ and } \overline{\alpha}\beta \prec \omega' & \text{via case 2.a)} \\ \omega \preceq \omega', \text{ if } \omega,\omega' \models \overline{\beta} & \text{via case 2.b)} \end{cases} \tag{9.4}$$

Equation (9.4) displays all constraints defining (9.1) resulting from the previous case distinctions and equivalent reformulations. The constraint from case 1.a) resp. case 1.b) resp. case 2.a) resp. case 2.b) corresponds to (II) resp. (IV) resp. (V) resp. (III) from (9.2). Eventually, we get the constraints from (9.2):

$$\omega \preceq^{\circ}_{\alpha,\beta} \omega' \text{ iff } \begin{cases} \omega \preceq \omega', \text{ if } \big(\omega,\omega' \models \beta \text{ and } (\omega,\omega' \preceq \overline{\alpha}\beta \text{ or } \overline{\alpha}\beta \prec \omega,\omega')\big) & \text{(I) or (II)} \\ \qquad \text{or } \big(\omega,\omega' \models \overline{\beta}\big) & \text{(III)} \\ \qquad \text{or } \big(\omega \models \beta,\ \omega' \models \overline{\beta} \text{ and } \overline{\alpha}\beta \prec \omega\big) & \text{(IV)} \\ \qquad \text{or } \big(\omega' \models \beta,\ \omega \models \overline{\beta} \text{ and } \overline{\alpha}\beta \prec \omega'\big) & \text{(V)} \\ \quad \top, \text{ if } \big(\omega \models \beta, \omega \preceq \overline{\alpha}\beta \text{ and} & \\ \qquad (\omega' \models \overline{\beta} \text{ or } \omega' \models \beta, \overline{\alpha}\beta \prec \omega')\big) & \text{(VI) or (VII)} \end{cases}$$

□

In (9.2), we employ exclusive cases leading to more comprehensible constraints for BR. In Figure 9.2, we present an overview of the exclusive and exhaustive use cases of the constraints in (9.2). Note that, exhaustiveness of the cases follows immediately from the equivalence between 9.1 and (9.2).

First, we explain all cases in which BR does not change the prior plausibilistic relation: Case (I) and (II) deal with β-worlds and imply that either if both worlds are more or equally plausible than $\overline{\alpha}\beta$ or if both worlds are less plausible than $\overline{\alpha}\beta$ their relations are kept, i.e., these cases clearly discriminate β-worlds along the plausibility limit specified by $\overline{\alpha}\beta$ while keeping the prior plausibility relation between worlds either more or equally resp. less plausible than $\overline{\alpha}\beta$. Case (III) deals with

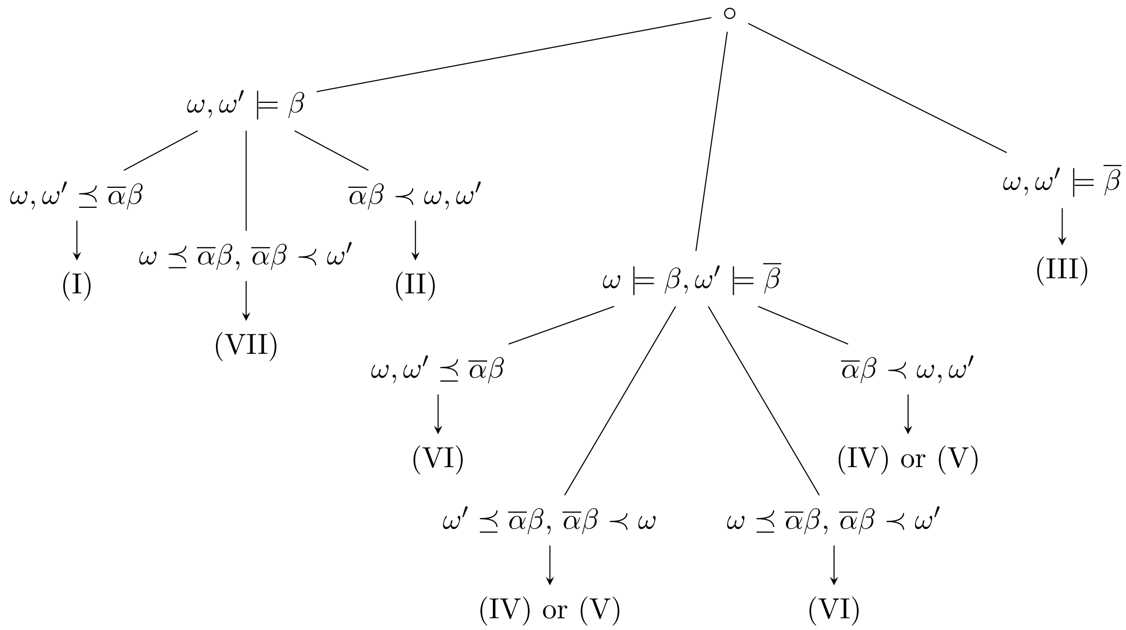

Figure 9.2: Schematic sketch of all use cases of the constraints (I) – (VII) in (9.2) defining BR from Proposition 9.1.1. Note that, for (IV) and (V) the roles of ω and ω' can be swapped.

all $\overline{\beta}$-worlds, implying conservation of the corresponding plausibility relations. Case (IV) and (V) from (9.2) together close the gap to the β-worlds less plausible than $\overline{\alpha}\beta$ and show that also here the prior plausibility relations between those worlds and $\overline{\beta}$-worlds are not affected by BR. BR promotes β-worlds that are at least as plausible as $\bar{\alpha}\beta$, s.t. they are strictly more plausible than the rest of Ω in the posterior state. The strictness of the inequality follows from the fact that we cannot swap the roles of ω and ω' in the cases (VI) and (VII) in (9.2). All in all, we can derive from (9.2) that worlds $\omega \models \beta$, s.t. $\omega \preceq \bar{\alpha}\beta$ are at the center of the iterated belief change implemented by BR by β w.r.t. α. We summarize them by a unique formula.

Definition 9.1.2 (Core of Bounded Revision, [110])**.** *For a plausibilistic TPO $\preceq$ and a BR operator $\circ_\alpha \beta$, we define the* core of Bounded Revision by β w.r.t. α *as the formula*

$$\varphi^{\circ}_{\alpha,\beta} = \beta \wedge (\bigvee_{\omega \preceq \bar{\alpha}\beta} \omega).$$

We call the set of possible worlds satisfying $\varphi^\circ_{\alpha,\beta}$ *the* core set of Bounded Revision *and notate* $\Phi^\circ_{\alpha,\beta} = Mod(\varphi^\circ_{\alpha,\beta}) = \{\omega \in \Omega \,|\, \omega \in Mod(\beta), \omega \preceq \overline{\alpha}\beta\}$.

The formula $\varphi^\circ_{\alpha,\beta}$ specifies via the reference sentence α, which β-worlds are *sufficiently* plausible to be promoted by BR by β w.r.t. α, i.e., to which extent the plausibility of β shall be increased relative to the remaining worlds depending on the plausibility of $\overline{\alpha}$. Thus, the disjunction in the core formula $\varphi^\circ_{\alpha,\beta}$ substantiates the idea to "accept the input sentence as far as the reference sentence and *just a little further*" from [100] by linking the choice of β-worlds that BR promotes to the plausibility of $\overline{\alpha}$. This places $\varphi^\circ_{\alpha,\beta}$ at the core of BR. If we look at the final posterior belief state, we cannot distinguish whether a BR decreases the plausibility of worlds outside the core or increases the plausibility of worlds within the core set. Because plausibilities are given only qualitatively, we cannot determine the shift's direction from the final result. Before we illustrate BR implemented via the constraints in (9.2) in an example, we reformulate and compress the constraints using the formula $\varphi^\circ_{\alpha,\beta}$. This leads to an equivalent formulation of (9.2) which we state in the following proposition. We employ a proposition for the reformulation, which is useful in the proof of our following representation theorem for BR.

Proposition 9.1.2. *Let* $\varphi^\circ_{\alpha,\beta}$ *be the core formula of BR by* β *w.r.t.* α *as in Definition 9.1.2. It holds that the constraints defining BR* $\preceq^\circ_{\alpha,\beta} = \preceq \circ_\alpha \beta$ *from Proposition 9.1.1 are equivalent to the following reformulation of* (9.2)

$$\omega \preceq^\circ_{\alpha,\beta} \omega' \textit{ iff } \begin{cases} \omega \preceq \omega', \textit{ if } \omega, \omega' \models \varphi^\circ_{\alpha,\beta} & \text{(i)} \\ \qquad \textit{or } \omega, \omega' \not\models \varphi^\circ_{\alpha,\beta} & \text{(ii)} \\ \top, \textit{ if } \omega \models \varphi^\circ_{\alpha,\beta} \textit{ and } \omega' \not\models \varphi^\circ_{\alpha,\beta} & \text{(iii)} \end{cases} \tag{9.5}$$

Proof. For each case in (9.2), we check whether ω resp. ω' is a model of $\varphi^\circ_{\alpha,\beta}$.

(I) For each ω, ω' it holds that $\omega, \omega \models \beta$ and $\omega, \omega' \preceq \overline{\alpha}\beta$, thus $\omega, \omega' \models \varphi^\circ_{\alpha,\beta}$. And for each ω, ω' that satisfies $\varphi^\circ_{\alpha,\beta}$, it holds that case (i) applies.

(II) For each ω, ω' it holds that $\overline{\alpha}\beta \prec \omega, \omega'$, thus $\omega, \omega' \not\models \varphi^\circ_{\alpha,\beta}$.

(III) For each ω, ω' it holds that $\omega, \omega \not\models \beta$, thus $\omega, \omega' \not\models \varphi^\circ_{\alpha,\beta}$.

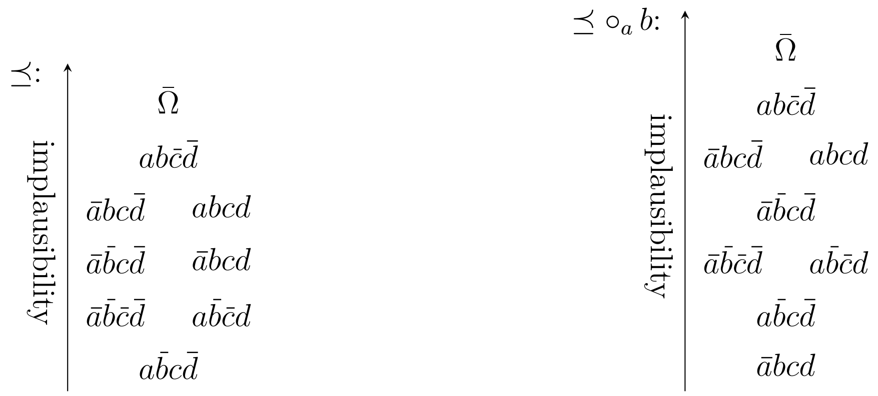

a) Plausibilistic TPO $\preceq$. b) BR revised TPO $\preceq \circ_a b$.

Figure 9.3: BR of the TPO $\preceq$ by a w.r.t. b.

(IV) For each ω it holds that $\overline{\alpha}\beta \prec \omega$, thus $\omega \not\models \varphi^{\circ}_{\alpha,\beta}$. And for each ω' it holds that $\omega' \not\models \beta$, thus $\omega' \not\models \varphi^{\circ}_{\alpha,\beta}$.

(V) For each ω' it holds that $\overline{\alpha}\beta \prec \omega'$, thus $\omega' \not\models \varphi^{\circ}_{\alpha,\beta}$. And for each ω it holds that $\omega \not\models \beta$, thus $\omega \not\models \varphi^{\circ}_{\alpha,\beta}$.

(VI) For each ω it holds that $\omega \models \beta$ and $\omega \preceq \overline{\alpha}\beta$, thus $\omega \models \varphi^{\circ}_{\alpha,\beta}$. And for each ω' it holds that $\omega' \not\models \beta$, thus $\omega' \not\models \varphi^{\circ}_{\alpha,\beta}$.

(VII) For each ω it holds that $\omega \models \beta$ and $\omega \preceq \overline{\alpha}\beta$, thus $\omega \models \varphi^{\circ}_{\alpha,\beta}$. And for each ω' it holds that $\omega' \models \beta$ and $\overline{\alpha}\beta \prec \omega'$, thus $\omega' \not\models \varphi^{\circ}_{\alpha,\beta}$

Due to the exhaustiveness of the cases (I) – (VII) of (9.2), we can subsume the cases (II) to (V) and get case (ii) of (9.5). And it holds that if ω satisfies $\varphi^{\circ}_{\alpha,\beta}$ and $\omega' \not\models \varphi^{\circ}_{\alpha,\beta}$, i.e., $\omega' \models \overline{\beta}$ or $\omega' \models \beta, \overline{\alpha}\beta \prec \omega'$, then case (VI) resp. (VII) apply and we get case (iii) of (9.5). □

We present the following example for BR to illustrate the mechanism of BR via our constraints (9.2) resp. (9.5):

Example 9.1.1. *In Figure 9.3a a plausibilistic TPO $\preceq$ with signature $\Sigma = \{a, b, c, d\}$ is given, where $\bar{\Omega}$ denotes all remaining worlds on the same level of plausibility, which are not shown explicitly. In Figure 9.3b the BR $\preceq \circ_a b$ with input $\beta = b$ and reference* *sentence $\alpha = a$ is depicted. It holds that $\Phi_{a,b} = \{\bar{a}bcd\}$ and thus, it follows from constraint (iii) in* (9.5)*, that $\bar{a}bcd \prec \circ_a b\, \omega'$ for all $\omega' \in (\Omega \setminus \Phi_{a,b})$, i.e., $\omega' \not\models \varphi^{\circ}_{\alpha,\beta}$. And from (ii) in* (9.5)*, it follows that the plausibility relations among all worlds outside of $\Phi_{a,b}$ remain the same.*

We state the following representation theorem, which proves that $\varphi^{\circ}_{\alpha,\beta}$ displays the right choice to characterize the change mechanism of BR by β w.r.t. α and provides us with simple yet elegant postulates axiomatizing the change mechanism of BR.

Theorem 9.1.3 (Representation Theorem for BR, [110])**.** *Let $\circ_\alpha \beta$ be a BR operator by β w.r.t. α. Let $\preceq$ be a plausibilistic TPO and $\preceq \circ_\alpha \beta = \preceq^{\circ}_{\alpha,\beta}$ be the corresponding BR revised plausibilistic TPO. Then $\preceq$ and $\preceq^{\circ}_{\alpha,\beta}$ satisfy* (9.5) *iff $\preceq$ and $\preceq^{\circ}_{\alpha,\beta}$ satisfy:*

(BR1) *If $\omega, \omega' \models \varphi^{\circ}_{\alpha,\beta}$, then $\omega \preceq \omega'$ iff $\omega \preceq^{\circ}_{\alpha,\beta} \omega'$*

(BR2) *If $\omega, \omega' \not\models \varphi^{\circ}_{\alpha,\beta}$, then $\omega \preceq \omega'$ iff $\omega \preceq^{\circ}_{\alpha,\beta} \omega'$*

(BR3) *If $\omega \models \varphi^{\circ}_{\alpha,\beta}$ and $\omega' \not\models \varphi^{\circ}_{\alpha,\beta}$ then $\omega \prec^{\circ}_{\alpha,\beta} \omega'$*

Proof. <u>$\Rightarrow$:</u> Presuppose that $\preceq$ and $\preceq^{\circ}_{\alpha,\beta}$ satisfy (9.5).

(BR1): For $\omega, \omega' \models \varphi^{\circ}_{\alpha,\beta}$, (BR1), i.e., $\omega \preceq \omega'$ iff $\omega \preceq^{\circ}_{\alpha,\beta} \omega'$ follows from case (i) in (9.5).

(BR2): For $\omega, \omega' \not\models \varphi^{\circ}_{\alpha,\beta}$, (BR2), i.e., $\omega \preceq \omega'$ iff $\omega \preceq^{\circ}_{\alpha,\beta} \omega'$ follows from case (ii) in (9.5).

(BR3): If $\omega \models \varphi^{\circ}_{\alpha,\beta}$ and $\omega' \not\models \varphi^{\circ}_{\alpha,\beta}$, then case (iii) in (9.5) applies and therefore $\omega \preceq^{\circ}_{\alpha,\beta} \omega'$ and (BR3) holds.
In general, it holds that if $\omega \models \varphi^{\circ}_{\alpha,\beta}$ then $\omega \models \beta$ and $\omega \preceq \bar{\alpha}\beta$, and for ω', it holds that either $\omega' \models \overline{\beta}$ or $\bar{\alpha}\beta \prec \omega'$. We exclude the doubly named worlds in $\mathrm{Mod}(\overline{\beta})$ and presuppose that $\omega' \models \overline{\beta}$ or ($\omega' \models \beta$ and $\bar{\alpha}\beta \prec \omega'$). Then the conditions of case (VI) or (VII) in (9.2) are satisfied and, we can conclude that $\omega \preceq^{\circ}_{\alpha,\beta} \omega'$ but not $\omega' \preceq^{\circ}_{\alpha,\beta} \omega$, i.e. (BR3) holds.

<u>$\Leftarrow$:</u> Presuppose that $\preceq$ and $\preceq^{\circ}_{\alpha,\beta}$ satisfy (BR1) – (BR3).

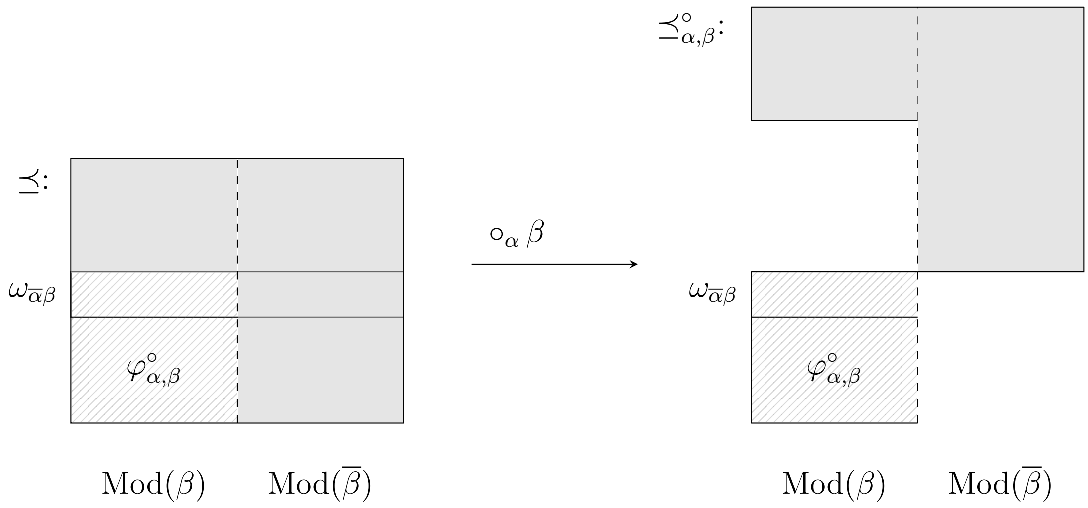

Figure 9.4: Schematic representation of BR of β w.r.t. α for a TPO $\preceq$.

(i) For $\omega, \omega' \models \varphi^{\circ}_{\alpha,\beta}$, we can conclude from (BR1) that $\omega \preceq \omega'$ iff $\omega \preceq^{\circ}_{\alpha,\beta} \omega'$ holds, i.e., case (i) in (9.5) is satisfied.

(ii) For $\omega, \omega' \not\models \varphi^{\circ}_{\alpha,\beta}$, we can conclude from (BR2) that $\omega \preceq \omega'$ iff $\omega \preceq^{\circ}_{\alpha,\beta} \omega'$ holds, i.e., case (ii) in (9.5) is satisfied.

(iii) For $\omega \models \varphi^{\circ}_{\alpha,\beta}$ and $\omega' \not\models \varphi^{\circ}_{\alpha,\beta}$, we can conclude from (BR3) that $\omega \prec^{\circ}_{\alpha,\beta} \omega'$, i.e., case (iii) in (9.5) is satisfied.

For each case (i) – (iii), (9.5) holds if $\preceq$ and $\preceq^{\circ}_{\alpha,\beta}$ satisfy (BR1) – (BR3). $\square$

The shift of worlds that BR performs is depicted in the schematic representation of BR in Figure 9.4. Note that, for plausibilistic TPOs, we cannot specify whether BR increases the plausibility of worlds in the core set $\Phi^{\circ}_{\alpha,\beta}$ by shifting them downwards to lower implausibility levels than worlds outside the core set or whether the posterior BR-revised TPO is yielded by shifting worlds outside the core set upwards to levels with higher implausibility. That means we cannot specify the direction of the shift for TPOs illustrated in Figure 9.4.

Next, we present a proposition which summarizes the characteristics and properties of BR defined via (BR1) – (BR3).

Proposition 9.1.4 ([110])**.** *For a plausibilistic TPO $\preceq$ and BR operator $\circ_{\alpha}\,\beta$ by β w.r.t. α, s.t. (BR1) – (BR3) hold for $\preceq \circ_{\alpha} \beta = \preceq^{\circ}_{\alpha,\beta}$, the following statements hold:*

1. *$\beta \in Bel(\preceq^{\circ}_{\alpha,\beta}) = Th(\min(Mod(\beta), \preceq))$*

2. *$\bar{\alpha} \prec^{\circ}_{\alpha,\beta} \overline{\beta}$, i.e., $(BR)_{\preceq}$ holds*

3. *$Bel(\preceq \circ_{\alpha} \beta) = Bel(\preceq \circ_{\gamma} \beta)$ for $\gamma \in \mathcal{L}$, so, $(SBC)_{\preceq}$ holds*

Proof. We show the statements employing (BR1) – (BR3) for BR of β w.r.t. α.

1. It holds that

$$\begin{aligned} Bel(\preceq \circ_{\alpha} \beta) &= Th(\{\omega \in \Omega |\, \omega \preceq^{\circ}_{\alpha,\beta} \omega' \text{ for all } \omega' \in \Omega\}) \\ &\overset{(\mathrm{BR3})}{=} Th(\{\omega \in \Phi^{\circ}_{\alpha,\beta} |\, \omega \preceq^{\circ}_{\alpha,\beta} \omega' \text{ for all } \omega' \in \Omega\}) \\ &\overset{(\mathrm{BR1})}{=} Th(\{\omega \in \Phi^{\circ}_{\alpha,\beta} |\, \omega \in \min(\Phi^{\circ}_{\alpha,\beta}, \preceq)\}) \\ &= Th(\min(\mathrm{Mod}(\beta), \preceq)) \end{aligned}$$

   The last equation holds, since for $\omega \in \min(\Phi^{\circ}_{\alpha,\beta}, \preceq)$, it holds that $\omega \models \beta$ and the restriction $\omega \preceq \bar{\alpha}\beta$ is obsolete. We can conclude that β holds in $Bel(\preceq \circ_{\alpha} \beta)$.

2. It follows from (BR3), that $\bar{\alpha}\beta \preceq^{\circ}_{\alpha,\beta} \bar{\alpha}\overline{\beta}$, because for each minimal $\bar{\alpha}\beta$-world, represented by $\omega_{\bar{\alpha}\beta}$, it holds that $\omega_{\bar{\alpha}\beta} \models \varphi^{\circ}_{\alpha,\beta}$ and for each minimal $\bar{\alpha}\overline{\beta}$-world, represented by $\omega_{\bar{\alpha}\overline{\beta}}$, it holds that $\omega_{\bar{\alpha}\overline{\beta}} \not\models \varphi^{\circ}_{\alpha,\beta}$. Since this holds for all minimal $\bar{\alpha}\beta$-worlds resp. $\bar{\alpha}\overline{\beta}$-worlds, we can conclude that $\bar{\alpha} \approx^{\circ}_{\alpha,\beta} \omega_{\bar{\alpha}\beta}$.
   Moreover, it holds for each minimal $\overline{\beta}$-world, represented by $\omega_{\overline{\beta}}$, that $\omega_{\overline{\beta}} \not\models \varphi^{\circ}_{\alpha,\beta}$. Now, via (BR3), we can conclude that $\bar{\alpha} \approx^{\circ}_{\alpha,\beta} \omega_{\bar{\alpha}\beta} \prec^{\circ}_{\alpha,\beta} \omega_{\overline{\beta}} \approx^{\circ}_{\alpha,\beta} \overline{\beta}$ holds, thus $\preceq^{\circ}_{\alpha,\beta}$ satisfies $(\mathrm{BR})_{\preceq}$.

3. $(\mathrm{SBC})_{\preceq}$ follows directly from the first statement, since the belief set of the BR revised plausibilistic TPO does not depend on the reference sentence, but solely on the input. Thus, $Bel(\preceq \circ_{\alpha} \beta) = Bel(\preceq \circ_{\gamma} \beta)$ holds for each $\alpha, \gamma \in \mathcal{L}$.

□

A schematic sketch of the results in this section is given in Figure 9.5 In the following example, we illustrate the versatile character of BR which depends on the interplay of input and reference sentence.

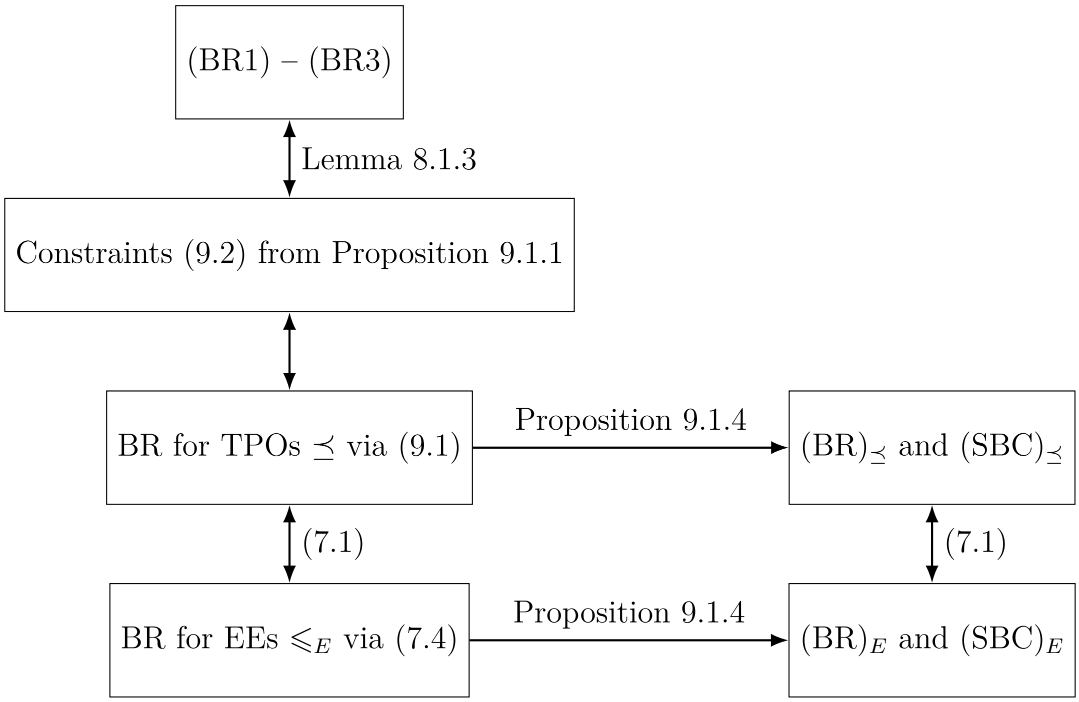

Figure 9.5: Overview of constraints and postulates defining BR.

Example 9.1.2. *In Figure 9.3a a plausibilistic TPO $\preceq$ with signature $\Sigma = \{a, b, c, d\}$ is given, where $\bar{\Omega}$ denotes all remaining worlds on the same level of plausibility, which are not shown explicitly. We perform three different BR operations in Figures 9.3b, 9.6a and 9.6b to illustrate the strength and special features of BR. In Figure 9.3b the BR $\preceq \circ_a b$ with input $\beta = b$ and reference sentence $\alpha = a$ is depicted. The BR $\preceq \circ_a b$ was already discussed in the context of our constraints from* (9.2) *in Example 9.1.1. As we can see, it holds that $\Phi_{a,b} = \{\bar{a}bcd\}$ and obviously $\preceq^*_{a,b} \models b$ and (BR)$_\preceq$ hold. Note that BR yields the same belief set of the posterior TPO for a different reference sentence, like, e.g., $\alpha = \bar{a}$ due to (SBC)$_\preceq$. In Figure 9.6a, the outcome of $\preceq \circ_b c$ with reference sentence $\alpha = b$ and input $\beta = c$ is illustrated. BR does not change the prior ordering since (BR1) – (BR3) are already satisfied. This example shows that the change implemented by BR is vacuous under the condition that all worlds in $\Phi^\circ_{\alpha,\beta}$ are already more plausible than worlds outside of it in the prior ordering. Otherwise, (BR1) – (BR3) imply the strengthening of the β-belief as can be seen in Figure 9.6b, where BR is performed with input $\beta = \bar{d}$ and reference sentence $\alpha = a$. Here, the input sentence $\bar{d}$ gets promoted via BR, even though it is already believed*

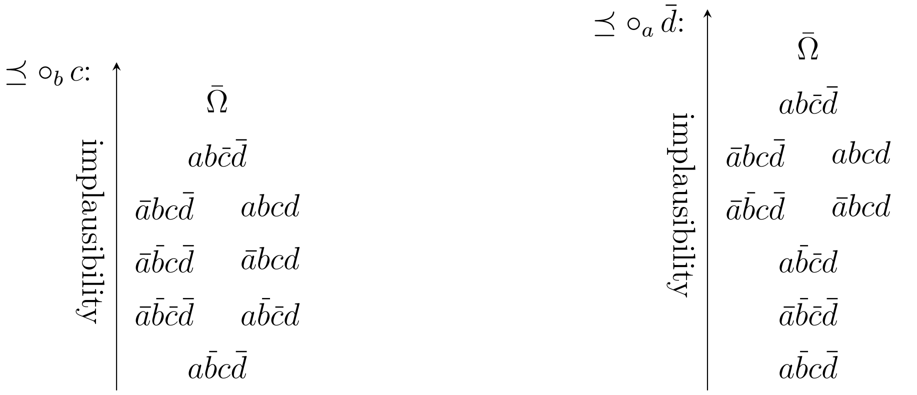

a) BR revised TPO $\preceq \circ_b c$. b) BR revised TPO $\preceq \circ_a \bar{d}$.

Figure 9.6: BR by c resp. $\bar{d}$ w.r.t. b resp. a.

in the prior ordering. Note that this promotion of the input belief depends on the choice of the reference sentence because the reference sentence specifies how much more plausible the input shall be in the posterior ordering.

Note that for BR by β w.r.t. α, the reference sentence α is used to specify the distance between β- and $\overline{\beta}$-worlds in the posterior state. This is because the reference sentence determines to which extent β-worlds are part of the core set $\Phi^{\circ}_{\alpha,\beta}$. In general, they have to be more or equally plausible than $\overline{\alpha}\beta$. Also, it holds that $\omega \models \beta$ for each $\omega \in \Phi^{\circ}_{\alpha,\beta}$ and if $\omega' \models \overline{\beta}$ it holds that $\omega' \notin \Phi^{\circ}_{\alpha,\beta}$. Thus, we get from (BR3) that each world satisfying $\varphi^{\circ}_{\alpha,\beta}$ is strictly more plausible than worlds that do not satisfy $\varphi^{\circ}_{\alpha,\beta}$ and $\overline{\alpha}\beta$ marks the plausibility level to which worlds are part of the core set.

Apart from cases where $\overline{\alpha} \prec \overline{\beta}$ is already present in the prior ordering, the shift implemented by BR satisfies a minimal change principle. We summarize some characteristics of the BR shift via the following proposition, which summarize the idea of minimality implemented in BR.

Proposition 9.1.5. *For a plausibilistic TPO $\preceq$ and BR operator $\circ_\alpha \beta$ by β w.r.t. α, s.t. (BR1) – (BR3) hold for $\preceq \circ_\alpha \beta = \preceq^{\circ}_{\alpha,\beta}$, the following statements hold:*

1. $\overline{\alpha}\beta \approx^{\circ}_{\alpha,\beta} \overline{\alpha}$
2. *If* $\{\omega \in Mod(\overline{\beta}) \mid \omega \preceq \overline{\alpha}\beta\} = \emptyset$, *then* $\preceq = \preceq^{\circ}_{\alpha,\beta}$

3. *If $\{\omega \in Mod(\overline{\beta}) \,|\, \omega \preceq \overline{\alpha}\beta\} \neq \emptyset$, then there exists no $\omega \in \Omega$, s.t. $\overline{\alpha} \prec^{\circ}_{\alpha,\beta} \omega \prec^{\circ}_{\alpha,\beta} \overline{\beta}$, i.e., $(BR)_{\preceq}$ is satisfied in the slightest possible way*

Proof. 1. For all minimal worlds $\omega_{\overline{\alpha}\beta} \models \overline{\alpha}\beta$, it holds that $\omega_{\overline{\alpha}\beta} \models \varphi^{\circ}_{\alpha,\beta}$, and, for all minimal worlds $\omega_{\overline{\alpha}\overline{\beta}} \models \overline{\alpha}\overline{\beta}$ holds that $\omega_{\overline{\alpha}\overline{\beta}} \not\models \varphi^{\circ}_{\alpha,\beta}$. Thus, we can conclude from (BR3) that $\omega_{\overline{\alpha}\beta} \prec^{\circ}_{\alpha,\beta} \omega_{\overline{\alpha}\overline{\beta}}$ and therefore $\overline{\alpha} \approx^{\circ}_{\alpha,\beta} \min\{\overline{\alpha}\beta, \overline{\alpha}\overline{\beta}\} \approx^{\circ}_{\alpha,\beta} \overline{\alpha}\beta$.

2. Since $\{\omega \in \mathrm{Mod}(\overline{\beta}) \,|\, \omega \preceq \overline{\alpha}\beta\} = \emptyset$, it holds for all worlds $\omega \in \Omega$ that either $\omega \models \varphi^{\circ}_{\alpha,\beta}$ or $\overline{\alpha}\beta \prec \omega$.

   Hence, for all worlds $\omega \not\models \varphi^{\circ}_{\alpha,\beta}$, it holds that $\overline{\alpha}\beta \prec \omega$. $(\star)$

   We distinguish the following three cases, which correspond to (BR1) – (BR3) and show that in each case the previous ordering does not change after BR.
   1. case: If $\omega_1, \omega_2 \models \varphi^{\circ}_{\alpha,\beta}$, then $\omega_1 \preceq^{\circ}_{\alpha,\beta} \omega_2$ iff $\omega_1, \preceq \omega_2$ follows from (BR1).
   2. case: If $\omega_1, \omega_2 \not\models \varphi^{\circ}_{\alpha,\beta}$, then $\omega_1 \preceq^{\circ}_{\alpha,\beta} \omega_2$ iff $\omega_1, \preceq \omega_2$ follows from (BR1).
   3. case: Let $\omega_1 \models \varphi^{\circ}_{\alpha,\beta}$ and $\omega_2 \not\models \varphi^{\circ}_{\alpha,\beta}$. We can conclude from $(\star)$ that $\overline{\alpha}\beta \prec \omega_2$ and therefore it holds that $\omega_1 \prec \omega_2$. And $\omega_1 \prec^{\circ}_{\alpha,\beta} \omega_2$ follows from (BR3).

3. For all $\omega \models \varphi^{\circ}_{\alpha,\beta}$, it holds that $\omega \preceq^{\circ}_{\alpha,\beta} \overline{\alpha}\beta \approx^{\circ}_{\alpha,\beta} \overline{\alpha}$.
   Since $\{\omega \in \mathrm{Mod}(\overline{\beta}) \,|\, \omega \preceq \overline{\alpha}\beta\} \neq \emptyset$ there exists a minimal model $\omega_{\overline{\beta}}$, s.t. $\omega_{\overline{\beta}} \models \overline{\beta}$ and $\omega_{\overline{\beta}} \preceq \overline{\alpha}\beta$. So, $\omega_{\overline{\beta}}$ minimally satisfies $\overline{\beta}$ in $\preceq$.
   Moreover, we can conclude that $\omega_{\overline{\beta}}$ also minimally satisfies $\overline{\varphi^{\circ}_{\alpha,\beta}}$, since for all β-worlds ω' in $\mathrm{Mod}(\overline{\varphi^{\circ}_{\alpha,\beta}})$, it holds that $\overline{\alpha}\beta \prec \tilde{\omega}$. Thus, it holds that $\omega_{\overline{\beta}} \approx \overline{\beta} \preceq \omega'$ for all $\omega' \models \overline{\varphi^{\circ}_{\alpha,\beta}}$. From (BR2), we can conclude that $\omega_{\overline{\beta}} \approx^{\circ}_{\alpha,\beta} \overline{\beta} \preceq^{\circ}_{\alpha,\beta} \omega'$ holds since all worlds involved in this inequality are not models of $\varphi^{\circ}_{\alpha,\beta}$.
   All in all, we get for $\omega \models \varphi^{\circ}_{\alpha,\beta}$ and $\omega' \not\models \varphi^{\circ}_{\alpha,\beta}$ that

$$\omega \overset{(BR1)}{\preceq^{\circ}_{\alpha,\beta}} \overline{\alpha}\beta \overset{1.}{\preceq^{\circ}_{\alpha,\beta}} \overline{\alpha} \overset{(BR3)}{\prec^{\circ}_{\alpha,\beta}} \overline{\beta} \approx^{\circ}_{\alpha,\beta} \omega_{\overline{\beta}} \overset{(BR2)}{\preceq^{\circ}_{\alpha,\beta}} \omega',$$

   i.e., there exists no world in between $\overline{\alpha}$ and $\overline{\beta}$.

□

From the Proposition above, we can conclude that $\bar{\alpha}\beta \approx^{\circ}_{\alpha,\beta} \bar{\alpha} \prec^{\circ}_{\alpha,\beta} \bar{\beta}$ holds due to the definition of $\varphi^{\circ}_{\alpha,\beta}$ and $(\mathrm{BR})_{\preceq}$. Yet, in general, BR does not guarantee that

$(BR)_{\preceq}$ is satisfied in the slightest possible way, i.e., it is possible that there exists a world ω s.t. $\overline{\alpha} \prec^{\circ}_{\alpha,\beta} \omega \prec^{\circ}_{\alpha,\beta} \overline{\beta}$. In the special case, if $\overline{\alpha}\beta \approx \overline{\alpha} \prec \omega \prec \overline{\beta}$ holds, the precondition of the first statement in Proposition 9.1.5 is satisfied, thus, $\preceq = \preceq^{\circ}_{\alpha,\beta}$ holds and BR with β w.r.t. α does not change the prior ordering, so that the world ω remains on a plausibility level in between $\overline{\alpha}$ and $\overline{\beta}$. We illustrate this special case via the following example.

Example 9.1.3. *Let $\bar{a}b \prec ab \prec a\bar{b} \prec \bar{a}\bar{b}$ be a plausibilistic TPO over $\Sigma = \{a, b\}$. For BR with reference $\alpha = a$ and input sentence $\beta = b$, it holds that $\Phi_{a,b} = \{\bar{a}b\}$ and thus, the constraints (BR1) – (BR3) hold and BR does not change the prior ordering. Note that, $(BR)_{\preceq}$ is satisfied, but for the posterior ordering it holds that $\bar{a} \prec^{\circ}_{\alpha,\beta} ab \prec^{\circ}_{\alpha,\beta} \bar{b}$*

In [100], Rott described the goal of BR to “accept β as long as α holds along with β, and just a little further.” The third statement in Proposition 9.1.5 substantiates what is meant by “just a little further”, namely that in terms of plausibility, $\overline{\alpha}$ is strictly more plausible than $\overline{\beta}$, but there exists no level of plausibility in between these two formulas, at least in cases where BR performs a real change.

9.2 Lexicographic and Natural Revision via Bounded Revision

This section shows that limiting cases of BR w.r.t. the reference sentence α correspond to lexicographic [82] resp. natural revision [17].

Rott observed this unique feature of BR in [100] and stated that BR is a parameterized belief change operation in between these two well-known revision mechanisms[1]. We make this connection more clear by applying the postulates (BR1) – (BR3) from our Representation Theorem 9.1.3. But first, we show that BR, in general, corresponds to a particular lexicographic revision and thus implements a more fine-grained revision mechanism via supplementary information in the form of a reference sentence α.

[1] In [100], natural revision is called *conservative revision* and lexicographic revision is used under the name *moderate revision*. These names can be traced back to the paper [98] by Rott.

Lexicographic revision displays an iterated belief revision operator that takes as input a sentence γ and is semantically defined on plausibilistic TPOs (cf. Definition 2.3.2 on page 31). We expand the constraints defining the lexicographic revision operator $\bullet^\ell$ [82] from Definition 2.3.2 a bit in order to compare them to (BR1) – (BR3) more easily, which characterize BR by β w.r.t. α, and get the following proposition:

Proposition 9.2.1. *Let $\preceq$ be a plausibilistic TPO and $\gamma \in \mathcal{L}$. The lexicographic revision $\preceq \bullet^\ell \gamma = \preceq^\ell_\gamma$ of $\preceq$ by γ is characterized by the following constraints*

(L1) *If $\omega, \omega' \models \gamma$, then $\omega \preceq^\ell_\gamma \omega'$ iff $\omega \preceq \omega'$*

(L2) *If $\omega, \omega' \not\models \gamma$, then $\omega \preceq^\ell_\gamma \omega'$ iff $\omega \preceq \omega'$*

(L3) *If $\omega \models \gamma$ and $\omega' \not\models \gamma$, then $\omega \prec^\ell_\gamma \omega'$*

We notice that lexicographic revision and Bounded Revision are related to each other. (BR1) – (BR3) correspond to (L1) – (L3), if we replace γ with the core formula $\varphi^\circ_{\alpha,\beta}$ for BR. The following theorem proves that BR by β w.r.t. α displays a lexicographic revision with the corresponding core of BR.

Theorem 9.2.2. *Let $\preceq$ be a plausibilistic TPO, $\circ_\alpha \beta$ be the BR operator by β w.r.t. α and $\bullet^\ell$ be the lexicographic revision operator. For $\preceq \circ_\alpha \beta$ and $\preceq \bullet^\ell \varphi^\circ_{\alpha,\beta}$ with $\varphi^\circ_{\alpha,\beta}$ being the core of BR by β w.r.t. α, it holds that*

$$\preceq \circ_\alpha \beta = \preceq \bullet^\ell \varphi^\circ_{\alpha,\beta}$$

s.t. $\omega\ (\preceq \circ_\alpha \beta)\ \omega'$ iff $\omega\ (\preceq \bullet^\ell \varphi^\circ_{\alpha,\beta})\ \omega'$ holds for all $\omega, \omega' \in \Omega$.

Proof. We compare the lexicographic revision operator $\bullet^\ell$ defined by (L1) – (L3) and the BR operator $\circ_\alpha \beta$ defined by (BR1) – (BR3).

It holds that that (L1) corresponds to (BR1) with $\gamma = \varphi^\circ_{\alpha,\beta}$, (L2) corresponds to (BR2) with $\gamma = \varphi^\circ_{\alpha,\beta}$ and (L3) corresponds to (BR3) with $\gamma = \varphi^\circ_{\alpha,\beta}$.

The postulates (L1) – (L3) and (BR1) – (BR3) are analog if we choose $\gamma = \varphi^\circ_{\alpha,\beta}$ as input for the lexicographic revision operator and both yield unique plausibilistic TPOs $\preceq^\ell_{\alpha,\beta}$ resp. $\preceq^\circ_{\alpha,\beta}$. Thus, we can conclude that $\omega \preceq^\circ_{\alpha,\beta} \omega'$ iff $\omega \preceq^\ell_{\alpha,\beta} \omega'$ holds for all ω, ω' and therefore, the plausibilistic TPOs $\preceq^\ell_{\alpha,\beta}$ and $\preceq^\circ_{\alpha,\beta}$ are the same. $\square$

Theorem 9.2.2 states that BR by β w.r.t. α corresponds to a lexicographic revision by $\varphi^{\circ}_{\alpha,\beta}$. Thus, we encode the supplementary information in the reference sentence α from the meta-level to the more clearly defined and directly usable object level via $\varphi^{\circ}_{\alpha,\beta}$. Hence, lexicographic revision with $\varphi^{\circ}_{\alpha,\beta}$ can be seen as a reduction of BR by β w.r.t. α to the framework of lexicographic revisions, which makes it directly usable for lexicographic revision solvers (see for e.g., [4]). Moreover, in contrast to a standard lexicographic revision with β, where (L3) forces quite rough changes on the prior ordering by making all β-worlds more plausible than $\overline{\beta}$-worlds, BR implements a more fine-grained revision with input β. The incorporation of α leads to the corresponding revision with $\varphi^{\circ}_{\alpha,\beta}$, where α marks to which plausibility level worlds in $\mathrm{Mod}(\beta)$ are promoted in the posterior TPO. This supports the idea that α serves as an indicator for the reliability of β and provides ground for further applications of supplementary information.

Now, we consider the limiting cases for BR of β w.r.t. α. That means cases where the reference sentence represents a logical truth $\alpha \equiv \top$ or a fallacy $\alpha \equiv \bot$. Later, we also deal with tautologies as input information.

For the limiting cases concerning the reference sentence, we suppose that either α never or always holds along with input sentence β. For the first case, we consider $\alpha \equiv \top$ and investigate the corresponding core formula. It holds that

$$\overline{\alpha} \wedge \beta \equiv \bot \wedge \beta \equiv \bot \tag{9.6}$$

and thus, we get for $\varphi_{\top,\beta}$

$$\varphi_{\top,\beta} \equiv \beta \wedge \bigvee_{\omega \preceq \bot} \omega \equiv \beta \wedge \top \equiv \beta \tag{9.7}$$

since $\omega \prec \bot$ holds for all worlds ω in $\preceq$ (cf. (2.4) on page 28), s.t. $\bigvee_{\omega \preceq \bot} \omega$ exhausts the space of logical possibilities. We can conclude from (9.7) that $\Phi_{\top,\beta} = \mathrm{Mod}(\beta)$. Applying this, we get the following reformulation of our postulates (BR1) – (BR3) for $\alpha \equiv \top$:

(BR1) If $\omega, \omega' \models \beta$, then $\omega \preceq \omega'$ iff $\omega \preceq^{\circ}_{\top,\beta} \omega'$

(BR2) If $\omega, \omega' \not\models \beta$, then $\omega \preceq \omega'$ iff $\omega \preceq^{\circ}_{\top,\beta} \omega'$

(BR3) If $\omega \models \beta$ and $\omega' \not\models \beta$ then $\omega \prec^{\circ}_{\top,\beta} \omega'$

Now, if we compare these postulates to (L1) – (L3), it is clear to see that (BR1) – (BR3) for $\alpha \equiv \top$ corresponds to a lexicographic revision with input β. We summarize this via the following theorem.

Theorem 9.2.3. *Let $\preceq$ be a plausibilistic TPO, $\circ_\alpha \beta$ be the BR operator by β w.r.t. α and $\bullet^\ell$ be the lexicographic revision operator. For $\preceq \circ_\top \beta$ with tautological reference sentence α, it holds that*

$$\preceq \circ_\top \beta = \preceq \bullet^\ell \beta$$

s.t. ω $(\preceq \circ_\top \beta)$ ω' iff ω $(\preceq \bullet^\ell \beta)$ ω' holds for all $\omega, \omega' \in \Omega$.

Proof. We compare the lexicographic revision operator $\bullet^\ell$ defined by (L1) – (L3) and the BR operator $\circ_\top \beta$ defined by (BR1) – (BR3), which follow immediately from (9.7) and (9.6). The postulates (L1) – (L3) and (BR1) – (BR3) are analog if we choose $\gamma = \beta$ as input for the lexicographic revision operator and both yield unique plausibilistic TPOs $\preceq \bullet^\ell \beta$ resp. $\preceq \circ_\top \beta$. Thus, we can conclude that $\omega(\circ_\top \beta)\omega'$ iff $\omega(\preceq \bullet^\ell \beta)\omega'$ holds for all ω, ω' and therefore, the plausibilistic TPOs are the same. □

So far, we have shown that lexicographic revision can be seen as a limiting case of BR with a tautological reference sentence. In the following, we investigate the relation between natural revision and Bounded Revision and consider the opposite limiting case where the reference sentence is represented by a logical contradiction.

We suppose that $\alpha \equiv \bot$. It holds that

$$\overline{\alpha} \wedge \beta \equiv \top \wedge \beta \equiv \beta \tag{9.8}$$

and thus, we get for $\varphi_{\bot,\beta}$

$$\varphi_{\bot,\beta} \equiv \beta \wedge \bigvee_{\omega \preceq \beta} \omega \equiv \bigvee_{\omega \in \min(\beta,\preceq)} \omega \tag{9.9}$$

Because the plausibility of β in $\preceq$ is defined via minimal worlds satisfying β, it holds that only minimal β-worlds satisfy $\omega \preceq \beta$ and β at the same time. Hence, we can conclude from (9.9) that $\Phi_{\bot,\beta} = \min(\beta, \preceq)$. We adjust the postulates (BR1) – (BR3) accordingly:

(BR1) If $\omega, \omega' \in \min(\beta, \preceq)$, then $\omega \preceq \omega'$ iff $\omega \preceq^{\circ}_{\perp,\beta} \omega'$

(BR2) If $\omega, \omega' \notin \min(\beta, \preceq)$, then $\omega \preceq \omega'$ iff $\omega \preceq^{\circ}_{\perp,\beta} \omega'$

(BR3) If $\omega \in \min(\beta, \preceq)$ and $\omega' \notin \min(\beta, \preceq)$ then $\omega \prec^{\circ}_{\perp,\beta} \omega'$

For $\alpha \equiv \perp$, BR of β solely affects the most plausible worlds satisfying the input sentence and pushes them to the lowermost level of the posterior TPO, thus satisfying the success condition for the revision with β in the slightest possible way. This rather conservative approach to belief revision corresponds to natural revision $\bullet^{\mathrm{n}}$ as defined in Definition 2.3.2 on page 31. We expand the constraints defining the natural revision operator $\bullet^{\mathrm{n}}$ [17] from Definition 2.3.2 a bit in order to compare them to (BR1) – (BR3) more easily in the case of $\perp$ as reference sentence.

Proposition 9.2.4. *Let $\preceq$ be a plausibilistic TPO and $\gamma \in \mathcal{L}$. The natural revision $\preceq \bullet^{\mathrm{n}} \gamma = \preceq^{\mathrm{n}}_{\gamma}$ of $\preceq$ by γ is characterized by the following constraints*

(N1) *If $\omega, \omega' \in \min(\gamma, \preceq)$, then $\omega \preceq^{\mathrm{n}}_{\gamma} \omega'$ iff $\omega \preceq \omega'$*

(N2) *If $\omega, \omega' \notin \min(\gamma, \preceq)$, then $\omega \preceq^{\mathrm{n}}_{\gamma} \omega'$ iff $\omega \preceq \omega'$*

(N3) *If $\omega \in \min(\gamma, \preceq)$ and $\omega' \notin \min(\beta, \preceq)$, then $\omega \prec^{\mathrm{n}}_{\gamma} \omega'$*

Now, it is clear to see, that BR with reference sentence $\alpha \equiv \perp$ implemented via (BR1) – (BR3) corresponds to a natural revision with input sentence β and we get the following theorem:

Theorem 9.2.5. *Let $\preceq$ be a plausibilistic TPO, $\circ_{\alpha} \beta$ be a BR operator by β w.r.t. α and $\bullet^{\mathrm{n}}$ be a natural revision operator. For $\preceq \circ_{\perp} \beta$ with contradictory reference sentence α, it holds that*

$$\preceq \circ_{\perp} \beta = \preceq \bullet^{\mathrm{n}} \beta$$

s.t. ω $(\preceq \circ_{\perp} \beta)$ ω' iff ω $(\preceq \bullet^{\mathrm{n}} \beta)$ ω' holds for all $\omega, \omega' \in \Omega$.

Proof. We compare the natural revision operator $\bullet^{\ell}$ defined by (N1) – (N3) and the BR operator $\circ_{\perp} \beta$ defined by (BR1) – (BR3), which follow immediately from (9.9) and (9.8). The postulates (N1) – (N3) and (BR1) – (BR3) are analog if we choose $\gamma = \beta$ as input for the natural revision operator and both yield unique plausibilistic

TPOs $\preceq \bullet^{\mathrm{n}} \beta$ resp. $\preceq \circ_{\perp} \beta$. Thus, we can conclude that $\omega (\circ_{\perp} \beta)\, \omega'$ iff $\omega (\preceq \bullet^{\mathrm{n}} \beta)\, \omega'$ holds for all ω, ω' and therefore, the plausibilistic TPOs are the same. □

For the limiting cases $\alpha \equiv \top$ resp. $\alpha \equiv \perp$, it holds that BR displays a non-parameterized belief revision mechanism essentially because it either corresponds to natural resp. to lexicographic revision, as was already stated in [100]. This shows that Bounded Revision naturally fills the space between these two revision operators. This is also reflected in the corresponding core sets of BR w.r.t. a standard reference sentences α resp. with $\alpha \equiv \perp$ or $\alpha \equiv \top$ as the following proposition shows.

Proposition 9.2.6. *Let $\preceq$ be a plausibilistic TPO, $\circ_\alpha \beta$ be the BR operator by β w.r.t. α and $\Phi^\circ_{\alpha,\beta}$ be the corresponding core set. It holds that $\Phi^\circ_{\perp,\beta} \subseteq \Phi^\circ_{\alpha,\beta} \subseteq \Phi^\circ_{\top,\beta}$.*

Proof. In general, it holds for the core set of BR of β w.r.t. α that $\Phi^\circ_{\alpha,\beta} = \{\omega \in \Omega \,|\, \omega \in \mathrm{Mod}(\beta), \omega \preceq \bar{\alpha}\beta\} = \mathrm{Mod}(\beta \wedge (\bigvee_{\omega \preceq \bar{\alpha}\beta} \omega))$ according to Definition 9.1.2.
Via (9.8), we get for $\alpha \equiv \perp$ the following core set

$$\Phi^\circ_{\perp,\beta} = \{\omega \in \Omega \,|\, \omega \in \mathrm{Mod}(\beta), \omega \preceq \beta\} = \min(\beta, \preceq) \tag{9.10}$$

and via (9.6), we get for $\alpha \equiv \top$ the core set

$$\Phi^\circ_{\top,\beta} = \{\omega \in \Omega \,|\, \omega \in \mathrm{Mod}(\beta)\} = \mathrm{Mod}(\beta) \tag{9.11}$$

Thus, we can conclude that

$$\underbrace{\min(\beta, \preceq)}_{\Phi^\circ_{\perp,\beta}} \subseteq \underbrace{\mathrm{Mod}(\beta) \cap \mathrm{Mod}(\bigvee_{\omega \preceq \bar{\alpha}\beta} \omega)}_{\Phi^\circ_{\alpha,\beta}} \subseteq \underbrace{\mathrm{Mod}(\beta)}_{\Phi^\circ_{\top,\beta}}$$

holds. □

We illustrate the relation between natural and lexicographic revision versus Bounded Revision via the following example.

Example 9.2.1. *In Figure 9.3a a plausibilistic TPO $\preceq$ with signature $\Sigma = \{a, b, c, d\}$ is given, where $\bar{\Omega}$ denotes all remaining worlds on the same level of plausibility, which are not shown explicitly. We compare three revision operators via the revision with*

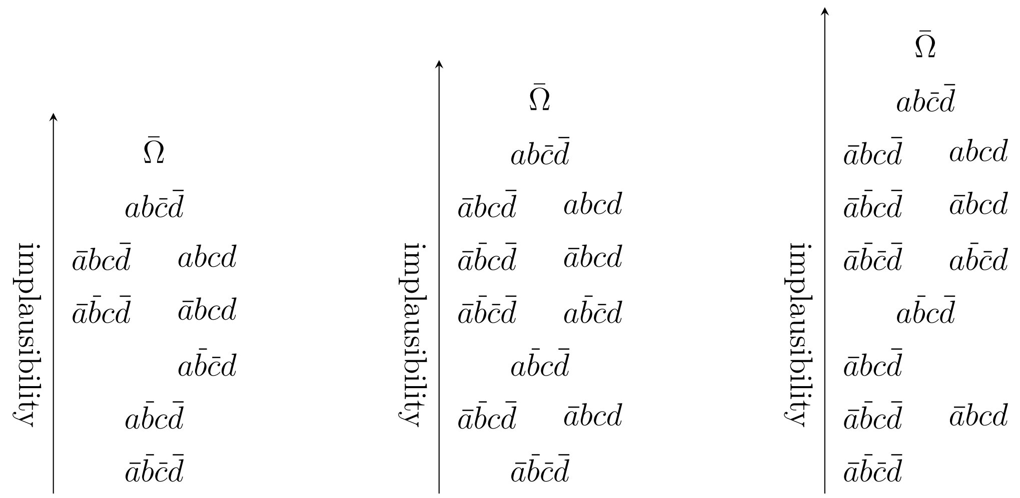

a) Posterior TPO $\preceq \bullet^{\mathrm{n}}\, \bar{a}$. b) Posterior TPO $\preceq \circ_{\bar{c}}\, \bar{a}$. c) Posterior TPO $\preceq \bullet^{\ell}\, \bar{a}$.

Figure 9.7: Natural and lexicographic revision as limiting cases of BR.

$\bar{a}$. The natural revision $\preceq \bullet^{\mathrm{n}}\, \bar{a}$ is depicted in Figure 9.7a. The BR with $\bar{a}$ w.r.t. $\bar{c}$, $\preceq \circ_{\bar{c}}\, \bar{a}$ is depicted in Figure 9.7b. In Figure 9.7c, we illustrate the lexicographic revision $\preceq \bullet^{\ell}\, \bar{a}$.

It holds for the natural revision $\preceq \bullet^{\mathrm{n}}\, \bar{a}$ in Figure 9.7a that solely the world $\min(\bar{a}, \preceq) = \{\bar{a}\bar{b}\bar{c}\bar{d}\}$ *is shifted to the lowermost levels. This shift corresponds to the shift implemented by BR of $\bar{a}$ w.r.t. to the contradictory reference sentence $\beta \equiv \bot$. Here it holds for the core set that* $\Phi^{\circ}_{\bot,\bar{a}} = \min(\bar{a}, \preceq)$.

BR with $\bar{a}$ w.r.t. $\bar{c}$, depicted in Figure 9.7b, implements a less conservative change, here the worlds in $\Phi^{\circ}_{\bar{c},\bar{a}} = \{\bar{a}\bar{b}\bar{c}\bar{d}, \bar{a}\bar{b}c\bar{d}, \bar{a}bcd\}$ are shifted to the lowermost levels while maintaining their inner plausibility orderings.

Figure 9.7c illustrates the lexicographic revision $\preceq \bullet^{\ell}\, \bar{a}$, and therefore, the most rigorous change on the prior TPO $\preceq$. Here, all worlds in Mod($\bar{a}$) = $\{\bar{a}\bar{b}\bar{c}\bar{d}, \bar{a}\bar{b}c\bar{d}, \bar{a}bcd, \bar{a}bc\bar{d}\}$ are made more plausible than the remaining a-worlds. For a BR of $\bar{a}$ w.r.t. to a tautological reference sentence, s.t. $\Phi^{\circ}_{\top,\bar{a}}$ = Mod($\bar{a}$) we get the same posterior TPO as for the lexicographic revision.

As we can see, it holds that $\Phi^{\circ}_{\bot,\bar{a}} \subseteq \Phi^{\circ}_{\bar{c},\bar{a}} \subseteq \Phi^{\circ}_{\top,\bar{a}}$. Thus, BR with $\bar{a}$ w.r.t. $\bar{c}$ fills the space between natural resp. lexicographic revision with $\bar{a}$.

We have recovered the two well-known belief revision mechanisms, natural and lexicographic revision, as limiting cases of BR. However, both natural and lexicographic revision seem to be defective. Natural revision accepts the new piece of evidence, but it accords it only to the lowest possible degree so that new evidence gets immediately lost if any contradiction with the next piece of evidence arises. On the other hand, lexicographic revision suffers from the opposite defect by accepting the input information very firmly and arguably too firmly. For a lexicographically TPO, it holds that all β-worlds are more plausible than all $\overline{\beta}$-worlds. So, while natural revision is too conservative, lexicographic revision seems too radical. BR has the advantage of systematically covering the whole range between these two extremes.

In the following, we deal with the special case of tautological revision, i.e., cases when the input $\beta \equiv \top$.

Proposition 9.2.7. *Let $\preceq$ be a plausibilistic TPO, $\circ_\alpha \beta$ be a BR operator by β w.r.t. α. For $\preceq \circ_\alpha \top$ with tautological input sentence β, it holds that $\preceq \circ_\alpha \top = \preceq$.*

Proof. For $\beta \equiv \top$, it holds that $\overline{\beta} \equiv \bot$, i.e.,

$$\{\omega \in \mathrm{Mod}(\overline{\beta}) \,|\, \omega \preceq \overline{\alpha}\beta\} = \emptyset$$

and hence, $\preceq \circ_\alpha \top = \preceq$ follows from the second statement in Proposition 9.1.5. $\square$

The proposition shows that BR satisfies a tautological vacuity principle for plausibilistic TPOs since the revision with $\top$ does not change the prior ordering regardless of the choice of the reference sentence α.

9.3 Bounded Revision for Ranking Functions

Now, we investigate Bounded Revision in the semi-quantitative framework of ranking functions and present two types of methodological implementations of BR. We start with a straightforward realization of the BR mechanism for OCFs in Subsection 9.3.1. This realization for ranking functions provides the foundation for the subsequent investigation of BR as a conditional c-revision in Subsection 9.3.2. We show that the change mechanism implemented in BR is fully captured by a single conditional, which relies on the interplay between the input sentence β and the reference sentence

α. Via this conditional, the parameterized belief revision via BR can be reduced to pleasantly easy c-revision by a single conditional.

Throughout this section, we use the core formula $\varphi^{\circ}_{\alpha,\beta}$ for BR as defined in Definition 9.1.2. The original definition uses TPOs to define the disjunction over all worlds more or equally plausible than $\overline{\alpha}\beta$. This condition can be easily transferred to the framework of OCFs via translation (2.9). Hence, we get for a ranking function κ the following corresponding core formula for BR:

$$\varphi^{\circ}_{\alpha,\beta} = \beta \wedge (\bigvee_{\kappa(\omega) \leqslant \kappa(\overline{\alpha}\beta)} \omega)$$

We stick to the previous notation to avoid a lengthy but uninformative re-definition of $\varphi^{\circ}_{\alpha,\beta}$.

9.3.1 Realization of Bounded Revision for Ranking Functions

In the following, we present a straightforward realization of BR for ranking functions, making the BR mechanism more explicit. In this subsection, we show that it satisfies the postulates (BR1) – (BR3).

Definition 9.3.1 ([110])**.** *Let κ be a ranking function.* Bounded Revision $\kappa \circ_\alpha \beta = \kappa^{\circ}_{\alpha,\beta}$ with input β and reference α for OCFs *is defined as follows:*

$$\kappa^{\circ}_{\alpha,\beta}(\omega) = \kappa_0 + \kappa(\omega) + \begin{cases} \kappa(\overline{\alpha}\beta) + 1, & \omega \not\models \varphi^{\circ}_{\alpha,\beta} \\ 0, & \mathit{otherwise} \end{cases} \tag{9.12}$$

with $\kappa_0 = -\kappa(\beta)$ as a normalization constant.

In general, for the normalization constant κ_0 in (9.12) it holds that

$$\kappa_0 = -\min\{\min_{\omega \models \varphi^{\circ}_{\alpha,\beta}} \{\kappa(\omega)\}, \min_{\omega' \not\models \varphi^{\circ}_{\alpha,\beta}} \{\kappa(\omega') + \kappa(\overline{\alpha}\beta) + 1\}\} = -\min\{\kappa(\varphi^{\circ}_{\alpha,\beta})\} = -\kappa(\beta), \tag{9.13}$$

because for worlds $\omega \models \varphi^{\circ}_{\alpha,\beta}$ it holds that $\omega \models \beta$ and $\kappa(\omega) \leqslant \kappa(\overline{\alpha}\beta)$. Thus,

$\kappa(\omega) < \kappa(\overline{\alpha}\beta)+1$ and therefore all worlds $\omega' \not\models \varphi^{\circ}_{\alpha,\beta}$ are irrelevant for the minimum. Employing the condition $\kappa(\omega) \leqslant \kappa(\overline{\alpha}\beta)$ from $\varphi^{\circ}_{\alpha,\beta}$ again, we conclude that each minimal β world satisfies $\varphi^{\circ}_{\alpha,\beta}$ minimally and we get that $\kappa(\varphi^{\circ}_{\alpha,\beta}) = \kappa(\beta)$. Hence $\kappa_0 = -\kappa(\beta)$ is the right choice for the normalization constant.

The semi-quantitative recipe for BR displayed in (9.12) reveals the basic mechanism of BR in a simple yet elegant way, and it displays a direct translation of the schematic BR operation from Figure 9.4. The plausibility of all worlds outside $\Phi^{\circ}_{\alpha,\beta}$ is reduced by increasing their implausibility ranks, leaving worlds satisfying $\varphi^{\circ}_{\alpha,\beta}$ among the most plausible ones. Thus, equation (9.12) implements BR by shifting worlds outside the core set in Figure 9.4 upwards. Relations within and outside the set $\Phi^{\circ}_{\alpha,\beta}$ are kept, because we add a constant factor, namely $\kappa(\overline{\alpha}\beta)+1$, to each world outside $\Phi^{\circ}_{\alpha,\beta}$. However, (9.12) is not minimal, in the sense that it does not change the prior ranking function when it is not necessary, i.e., when (BR1) – (BR3) already hold. In these cases, BR for OCFs strengthens the (already accepted) input information, which can also be advantageous, at least in special cases. In the following section, we discuss a minimal implementation of (BR1) – (BR3) for c-revisions more thoroughly and illustrate the different approaches via an example.

Note that BR by β w.r.t. α for ranking functions as defined in (9.12) is a variant of a realization of BR for OCFs. Due to the use of numerical ranks, other realizations are possible. For example, variants that implement a greater stretch between worlds $\omega \models \varphi^{\circ}_{\alpha,\beta}$ and $\omega' \not\models \varphi^{\circ}_{\alpha,\beta}$ by adding a constant number higher than 1 to worlds $\omega \not\models \varphi^{\circ}_{\alpha,\beta}$ in (9.12). Yet, all variants correspond to the same TPO since adding a constant number greater than one solely adds empty layers to $\kappa^{\circ}_{\alpha,\beta}$. The following theorem shows that BR for OCFs $\kappa^{\circ}_{\alpha,\beta}$ satisfies (BR1) – (BR3).

Theorem 9.3.1 ([110])**.** *Let κ be a ranking function and $\kappa \circ_{\alpha} \beta = \kappa^{\circ}_{\alpha,\beta}$ be the BR-revised OCF as defined in* (9.12)*. Then (BR1) – (BR3) hold for the corresponding plausibilistic TPOs $\preceq_{\kappa}$ and $\preceq_{\kappa^{\circ}_{\alpha,\beta}}$.*

Proof. Via (2.9), we get the TPO $\preceq_{\kappa}$ for the prior ranking function κ and the TPO $\preceq_{\kappa^{\circ}_{\alpha,\beta}}$ from the BR-revised OCF $\kappa^{\circ}_{\alpha,\beta}$.

(BR1): Let $\omega, \omega' \in \Phi^{\circ}_{\alpha,\beta}$, then $\kappa^{\circ}_{\alpha,\beta}(\omega) = -\kappa(\beta)+\kappa(\omega) \leqslant -\kappa(\beta)+\kappa(\omega') = \kappa^{\circ}_{\alpha,\beta}(\omega')$ follows immediately from $\kappa(\omega) \leqslant \kappa(\omega')$ and vice versa, since $\kappa_0 = -\kappa(\beta)$ is a constant factor. And therefore, via (2.9), we can conclude that (BR1) is satisfied for $\preceq_\kappa$ and $\preceq_{\kappa^\circ_{\alpha,\beta}}$.

(BR2): Let $\omega, \omega' \notin \Phi^\circ_{\alpha,\beta}$, then $\kappa^\circ_{\alpha,\beta}(\omega) = -\kappa(\beta) + \kappa(\omega) + \kappa(\overline{\alpha}\beta) + 1 \leqslant -\kappa(\beta) + \kappa(\omega') + \kappa(\overline{\alpha}\beta) + 1 = \kappa^\circ_{\alpha,\beta}(\omega')$ follows immediately from $\kappa(\omega) \leqslant \kappa(\omega')$ and vice versa, since $\kappa_0 = -\kappa(\beta)$ and $\kappa(\overline{\alpha}\beta)$ are constant factors. Therefore, we can conclude via (2.9) that (BR2) is satisfied.

(BR3): Let $\omega \in \Phi^\circ_{\alpha,\beta}$ and $\omega' \notin \Phi^\circ_{\alpha,\beta}$. For ω it holds that $\omega \models \beta$ and $\kappa(\omega) \leqslant \kappa(\overline{\alpha}\beta)$. Now, we consider the BR revised OCF with $\kappa^\circ_{\alpha,\beta}(\omega)$ and $\kappa^\circ_{\alpha,\beta}(\omega')$. It holds that

$$\begin{aligned}\kappa^\circ_{\alpha,\beta}(\omega) &= -\,\kappa(\beta) + \kappa(\omega) < -\kappa(\beta) + \kappa(\overline{\alpha}\beta) + 1 \\ &\leqslant -\kappa(\beta) + \kappa(\omega') + \kappa(\overline{\alpha}\beta) + 1 = \kappa^\circ_{\alpha,\beta}(\omega')\end{aligned}$$

since $\omega \models \varphi^\circ_{\alpha,\beta}$ and $\kappa(\omega') \geq 0$. Thus, via (2.9), we can conclude that (BR3) is satisfied for $\preceq_\kappa$.

□

From Proposition 9.1.4 and Theorem 9.3.1 it follows that, $\kappa^\circ_{\alpha,\beta}(\overline{\alpha}) < \kappa^\circ_{\alpha,\beta}(\overline{\beta})$ and $Bel(\kappa \circ_\alpha \beta) = Bel(\kappa \circ_\gamma \beta)$ for $\alpha, \gamma \in \mathcal{L}$ hold, i.e., $(\mathrm{BR})_\preceq$ and $(\mathrm{SBC})_\preceq$ reformulated via (2.9) for ranking functions are satisfied.

9.3.2 Realization of Bounded Revision as C-Revision

In this section, we realize BR for OCFs from (9.12) via a conditional revision of ranking functions with a single conditional. This makes the ensuing application of BR more explicit, and it makes BR directly usable for existing frameworks of belief revision, such as the one presented in [47], which are capable of revising with conditional information. Defining BR as a conditional revision illustrates the versatility of conditionals in Belief Revision, since conditionals enable us to express information from a meta-level, i.e., the reference sentence α, via a single logical entity. The revision with the conditional presented in this section manipulates the agent's belief state according to the rationale behind the parameterized revision operation BR.

We have seen, so far, that BR yields a posterior ordering by making worlds outside the core set of BR strictly less plausible than worlds from the core set. Now, we show that this shift can be characterized by a designated conditional which uses the core formula $\varphi^{\circ}_{\alpha,\beta}$ corresponding to a BR by β w.r.t. α:

$$\delta_{\alpha,\beta} = (\varphi^{\circ}_{\alpha,\beta} | \overline{\varphi^{\circ}_{\alpha,\beta}} \vee (\overline{\alpha}\beta))$$

We call $\delta_{\alpha,\beta}$ the *core conditional of BR by β w.r.t. α* and yield the corresponding verification resp. falsification formulas as follows

$$(\overline{\varphi^{\circ}_{\alpha,\beta}} \vee (\overline{\alpha}\beta)) \wedge (\varphi^{\circ}_{\alpha,\beta}) \equiv \varphi^{\circ}_{\alpha,\beta} \wedge \overline{\alpha}\beta \equiv \overline{\alpha}\beta \wedge (\bigvee_{\omega \preceq \overline{\alpha}\beta} \omega) \tag{9.14}$$

$$(\overline{\varphi^{\circ}_{\alpha,\beta}} \vee (\overline{\alpha}\beta)) \wedge (\overline{\varphi^{\circ}_{\alpha,\beta}}) \equiv \overline{\varphi^{\circ}_{\alpha,\beta}} \tag{9.15}$$

Equation (9.14) corresponds to the verification of $\delta_{\alpha,\beta}$ and (9.15) to its falsification. For the c-revision with our designated core conditional, we consider the definition of c-revisions with a single conditional from (2.20) on page 47 and get the following c-revision $\kappa *^{\mathrm{c}} \delta_{\alpha,\beta} = \kappa^{\mathrm{c}}_{\alpha,\beta}$ via the falsification of $\delta_{\alpha,\beta}$ in (9.15), s.t.

$$\kappa *^{\mathrm{c}} \delta_{\alpha,\beta}(\omega) = \kappa^{\mathrm{c}}_{\alpha,\beta}(\omega) = \kappa_0 + \kappa(\omega) + \begin{cases} \eta_\delta, & \text{for } \omega \models \overline{\varphi^{\circ}_{\alpha,\beta}} \\ 0, & \text{otherwise} \end{cases}, \tag{9.16}$$

holds. For c-revisions with a single conditional, the normalization constant κ_0 is defined in (2.15) on page 44. Hence, for the c-revision $\kappa^{\mathrm{c}}_{\alpha,\beta}$ we get for κ_0 that

$$\kappa_0 = -\kappa((\varphi^{\circ}_{\alpha,\beta} \wedge (\beta \Rightarrow \alpha)) \vee \varphi^{\circ}_{\alpha,\beta}) = -\kappa(\varphi^{\circ}_{\alpha,\beta}) = -\kappa(\beta) \tag{9.17}$$

holds, because $(\varphi^{\circ}_{\alpha,\beta} \wedge (\beta \Rightarrow \alpha)) \vee \varphi^{\circ}_{\alpha,\beta} \equiv \varphi^{\circ}_{\alpha,\beta}$. And, as before for (9.12), we can conclude for all minimal worlds satisfying $\varphi^{\circ}_{\alpha,\beta}$ that they also satisfy β minimally. Thus, we get that $-\kappa(\varphi^{\circ}_{\alpha,\beta}) = -\kappa(\beta)$ for the normalization constant in (9.16) and $\kappa *^{\mathrm{c}} \delta_{\alpha,\beta}$ is well-defined. For the non-negative impact factor η_δ, we get via (2.21) on page 48 and the corresponding verification (9.14) resp. falsification (9.15) of $\delta_{\alpha,\beta}$

the following inequality

$$
\begin{aligned}
\eta_\delta &> \kappa(\overline{\alpha}\beta \wedge (\bigvee_{\omega \preceq \overline{\alpha}\beta} \omega)) - \kappa(\overline{\varphi^\circ_{\alpha,\beta}}) \\
&= \kappa(\overline{\alpha}\beta) - \kappa(\overline{\varphi^\circ_{\alpha,\beta}})
\end{aligned}
\tag{9.18}
$$

which ensures that $\kappa^{\mathrm{c}}_{\alpha,\beta} \models \delta_{\alpha,\beta}$. In (9.18), $\kappa(\overline{\alpha}\beta \wedge (\bigvee_{\omega \preceq \overline{\alpha}\beta} \omega)) = \kappa(\overline{\alpha}\beta)$ follows immediately from the fact that only minimal models of $\overline{\alpha}\beta$ satisfy $\overline{\alpha}\beta$ and $\omega \preceq \overline{\alpha}\beta$ at the same time. In general, the plausibility of formulas in OCFs is defined via minimal models and thus, the OCF ranking of (9.14) equals $\kappa(\overline{\alpha}\beta)$.

The impact factor η_δ specifies the plausibility stretch between the input and the reference sentence in the revised state. This is crucial to reduce the meta-level revision with supplementary information in (9.12) to the object level encoded as the c-revision with a single conditional, which captures all relevant features of BR, as the following theorem shows.

Theorem 9.3.2 ([110]). *Let κ be a ranking function. For the c-revision $\kappa *^{\mathrm{c}} \delta_{\alpha,\beta} = \kappa^{\mathrm{c}}_{\alpha,\beta}$ as defined in* (9.22) *and $\kappa \circ_\alpha \beta = \kappa^\circ_{\alpha,\beta}$ from* (9.12), *it holds that the corresponding plausibilistic TPOs $\preceq_{\kappa^{\mathrm{c}}_{\alpha,\beta}}$ and $\preceq_{\kappa^\circ_{\alpha,\beta}}$ are the same, i.e.,*

$$\omega \preceq_{\kappa^{\mathrm{c}}_{\alpha,\beta}} \omega' \textit{ iff } \omega \preceq_{\kappa^\circ_{\alpha,\beta}} \omega'.$$

Proof. We show that TPO $\preceq_{\kappa^{\mathrm{c}}_{\alpha,\beta}}$ corresponding to the c-revision $\kappa^{\mathrm{c}}_{\alpha,\beta}$ satisfies (BR1) – (BR3) from Theorem 9.1.3. Since (BR1) – (BR3) defines a unique TPO and $\preceq_{\kappa^\circ_{\alpha,\beta}}$, i.e., the TPO corresponding to $\kappa^\circ_{\alpha,\beta}$, also satisfies (BR1) – (BR3) as we have shown in Theorem 9.3.1, we can immediately conclude that $\preceq_{\kappa^{\mathrm{c}}_{\alpha,\beta}} = \preceq_{\kappa^\circ_{\alpha,\beta}}$ holds.

(BR1): Let $\omega, \omega' \models \varphi^\circ_{\alpha,\beta}$. Then both worlds do not falsify $\delta_{\alpha,\beta}$, i.e., the second case in (9.16) applies. We get that $\kappa^{\mathrm{c}}_{\alpha,\beta}(\omega) = -\kappa(\beta) + \kappa(\omega) \leqslant -\kappa(\beta) + \kappa(\omega') = \kappa^{\mathrm{c}}_{\alpha,\beta}(\omega')$ iff $\kappa(\omega) \leqslant \kappa(\omega')$ and therefore, (BR1) follows for $\preceq_\kappa$ resp. $\preceq_{\kappa^\circ_{\alpha,\beta}}$ via (2.9).

(BR2): Let $\omega, \omega' \not\models \varphi^\circ_{\alpha,\beta}$. Then $\omega, \omega' \models \overline{\varphi^\circ_{\alpha,\beta}}$ holds and both worlds falsify $\delta_{\alpha,\beta}$. For the c-revision with $\delta_{\alpha,\beta}$ and the corresponding non-negative impact factor η_δ (cf. (9.18)), we get that $\kappa^{\mathrm{c}}_{\alpha,\beta}(\omega) = -\kappa(\beta) + \kappa(\omega) + \max\{0, \kappa(\overline{\alpha}\beta) - \kappa(\overline{\varphi^\circ_{\alpha,\beta}}) + 1\} \leqslant$

$-\kappa(\beta)+\kappa(\omega')+\max\{0,\kappa(\overline{\alpha}\beta)-\kappa(\overline{\varphi^{\circ}_{\alpha,\beta}})+1\} = \kappa^{\mathrm{c}}_{\alpha,\beta}(\omega')$ iff $\kappa(\omega) \leqslant \kappa(\omega')$, since η_δ is a constant. Thus, (BR2) follows for $\preceq_\kappa$ resp. $\preceq_{\kappa^{\circ}_{\alpha,\beta}}$ via (2.9).

(BR3): Let $\omega \models \varphi^{\circ}_{\alpha,\beta}$ and $\omega' \not\models \varphi^{\circ}_{\alpha,\beta}$, then it holds that $\kappa^{\mathrm{c}}_{\alpha,\beta}(\omega) = -\kappa(\beta)+\kappa(\omega)$ and $\kappa^{\mathrm{c}}_{\alpha,\beta}(\omega') = -\kappa(\beta)+\kappa(\omega')+\eta_\delta$ holds, where is the non-negative impact factor from (9.16) with $\eta_\delta > \kappa(\overline{\alpha}\beta)-\kappa(\overline{\varphi^{\circ}_{\alpha,\beta}})$. The following statements hold for each ranking function κ and will be useful in the course of this proof.

$$\kappa(\overline{\varphi^{\circ}_{\alpha,\beta}}) = \kappa(\bigvee_{\tilde{\omega}\not\models\varphi^{\circ}_{\alpha,\beta}} \tilde{\omega}) = \min_{\tilde{\omega}\not\models\varphi^{\circ}_{\alpha,\beta}} \{\kappa(\tilde{\omega})\}. \tag{9.19}$$

$$\text{For } \omega \models \varphi^{\circ}_{\alpha,\beta},\ \text{it holds that } \kappa(\omega) \leqslant \kappa(\overline{\alpha}\beta). \tag{9.20}$$

The first statement follows from $\overline{\varphi^{\circ}_{\alpha,\beta}} \equiv \bigvee_{\tilde{\omega}\not\models\varphi^{\circ}_{\alpha,\beta}} \tilde{\omega}$ and the properties of ranking functions. The second one holds since $\omega \in \Phi^{\circ}_{\alpha,\beta}$.

In order to show (BR3), we distinguish between cases where (i) $\kappa(\overline{\alpha}\beta) < \kappa(\overline{\varphi^{\circ}_{\alpha,\beta}})$ and $\eta_\delta = 0$ due to non-negativity and (ii) $\eta_\delta > \kappa(\overline{\alpha}\beta)-\kappa(\overline{\varphi^{\circ}_{\alpha,\beta}}) \geq 0$.

(i) Presuppose that $(*)$ $\kappa(\overline{\alpha}\beta) < \kappa(\overline{\varphi^{\circ}_{\alpha,\beta}})$ and $\eta_\delta = 0$. Together with (9.20), $(*)$ and (9.19), we get that

$$\kappa(\omega) \leqslant \kappa(\overline{\alpha}\beta) < \kappa(\overline{\varphi^{\circ}_{\alpha,\beta}}) \leqslant \kappa(\omega')$$

holds. Thus,

$$\kappa^{\mathrm{c}}_{\alpha,\beta}(\omega) = -\kappa(\beta)+\kappa(\omega) < -\kappa(\beta)+\kappa(\omega') = \kappa^{\mathrm{c}}_{\alpha,\beta}(\omega')$$

holds.

(ii) Presuppose that $\eta_\delta > \kappa(\overline{\alpha}\beta)-\kappa(\overline{\varphi^{\circ}_{\alpha,\beta}}) \geq 0$. Then, it follows from (9.19) that

$$(**)\ 0 \leqslant \kappa(\omega')-\kappa(\overline{\varphi^{\circ}_{\alpha,\beta}}) < \kappa(\omega')-\kappa(\overline{\varphi^{\circ}_{\alpha,\beta}})+1$$

holds. Together with (9.20) and $(**)$, it holds that

$$\begin{aligned}\kappa^{\mathrm{c}}_{\alpha,\beta}(\omega) &= -\kappa(\beta) + \kappa(\omega) \leqslant -\kappa(\beta) + \kappa(\overline{\alpha}\beta) \\ &< -\kappa(\beta) + \kappa(\overline{\alpha}\beta) + \kappa(\omega') - \kappa(\overline{\varphi^{\circ}_{\alpha,\beta}}) + 1 = \kappa^{\mathrm{c}}_{\alpha,\beta}(\omega').\end{aligned}$$

All in all, (BR3) follows for $\preceq_\kappa$ resp. $\preceq_{\kappa^{\circ}_{\alpha,\beta}}$ via (2.9).

$\square$

The theorem shows that the core conditional encodes the specific decrease of plausibility for all worlds $\omega \not\models \varphi^{\circ}_{\alpha,\beta}$ in a single, easily accessible logical entity, rather than introducing a meta-level to the revision operator itself as it is the case for standard BR operators. Thus, we have shown that the supplementary information provided by the reference sentence in BR can be incorporated naturally into the object level of an existing revision operator. Since $\preceq_{\kappa^{\circ}_{\alpha,\beta}}$ satisfies (BR1) – (BR3), it follows immediately from Theorem 9.3.2 that $\preceq_{\kappa^{\mathrm{c}}_{\alpha,\beta}}$ also satisfies (BR1) – (BR3).

The following theorem shows that $\kappa^{\circ}_{\alpha,\beta}$ corresponds to a special c-revision.

Theorem 9.3.3. *Let κ be a ranking function. It holds that BR for OCFs $\kappa^{\circ}_{\alpha,\beta} = \kappa \circ_\alpha \beta$ as defined in* (9.12) *displays a c-revision with the core conditional $\delta_{\alpha,\beta}$ and impact factor $\eta^{\circ}_{\delta} = \kappa(\overline{\alpha}\beta) + 1$*

Proof. In general, c-revision $\kappa *^{\mathrm{c}} \delta_{\alpha,\beta}$ are defined via (9.16) with an impact factor η_δ with $\eta_\delta > \kappa(\overline{\alpha}\beta) - \kappa(\overline{\varphi^{\circ}_{\alpha,\beta}})$ (cf. (9.18)). We choose $\eta^{\circ}_{\delta} = \kappa(\overline{\alpha}\beta) + 1$. It holds that η°_{δ} satisfies (9.18), because $\eta^{\circ}_{\delta} = \kappa(\overline{\alpha}\beta) + 1 > \kappa(\overline{\alpha}\beta) - \kappa(\overline{\varphi^{\circ}_{\alpha,\beta}})$, since $\kappa(\overline{\varphi^{\circ}_{\alpha,\beta}}) \geq 0$. Together with (9.12) and (9.13) and (9.17), it holds that

$$\begin{aligned}\kappa^{\circ}_{\alpha,\beta}(\omega) &= -\kappa(\beta) + \kappa(\omega) + \begin{cases} \kappa(\overline{\alpha}\beta) + 1, & \text{for } \omega \not\models \varphi^{\circ}_{\alpha,\beta} \\ 0, & \text{else} \end{cases} \\ &= -\kappa(\beta) + \kappa(\omega) + \begin{cases} \eta^{\circ}_{\delta}, & \text{for } \omega \not\models \varphi^{\circ}_{\alpha,\beta} \\ 0, & \text{else} \end{cases}\end{aligned}$$

is a c-revision of $\delta_{\alpha,\beta}$. $\square$

General c-revisions are capable of revising with sets of conditionals. This is primarily due to their ability to choose flexible impact factors defined by inequalities. However, the solution to this set of inequalities is not unique, and the interactions in between those inequalities make it impossible to define even a unique (Pareto-)minimal solution[2]. In the case of BR, we need to c-revise with only a single conditional $\delta_{\alpha,\beta}$ and therefore do not have interactions with other conditionals to be adopted. Thus, it is straightforward to single out a unique, minimal non-negative impact factor from (9.18) defined as follows

$$\eta_\delta^{\min} = \max\{0, \kappa(\overline{\alpha}\beta) - \kappa(\overline{\varphi^\circ_{\alpha,\beta}}) + 1\}. \tag{9.21}$$

The maximum ensures that $\eta_\delta^{\min}$ is always non-negative. Note that, in the special case of a convex ranking function κ, we can omit the maximum since $\kappa(\overline{\alpha}\beta) - \kappa(\overline{\varphi^\circ_{\alpha,\beta}}) + 1$ is always non-negative.

Proposition 9.3.4. *For a convex OCF κ, it holds that $\kappa(\overline{\alpha}\beta) - \kappa(\overline{\varphi^\circ_{\alpha,\beta}}) + 1 \geq 0$.*

Proof. Let $\omega_{\overline{\varphi}}$ be a minimal model of $\overline{\varphi^\circ_{\alpha,\beta}}$ and κ be a convex κ.
It holds that either (i) ($\omega_{\overline{\varphi}} \models \overline{\beta}$ and $\kappa(\omega_{\overline{\varphi}}) \leqslant \kappa(\overline{\alpha}\beta)$) or (ii) $\kappa(\omega_{\overline{\varphi}}) < \kappa(\overline{\alpha}\beta)$. We show that due to the convexity of κ, we get for case (ii) that $\kappa(\omega_{\overline{\varphi}}) = \kappa(\overline{\alpha}\beta) + 1$ holds. If case (i) does not apply, we can conclude for all worlds ω less plausible than $\overline{\alpha}\beta$, i.e., $\kappa(\overline{\alpha}\beta) < \kappa(\omega)$, that $\omega \models \overline{\varphi^\circ_{\alpha,\beta}}$. Thus, due to the convexity of κ and the minimality of ranks, we get for case (ii) that $\kappa(\omega_{\overline{\varphi}}) = \kappa(\overline{\alpha}\beta) + 1$. And therefore, it holds in general that $\kappa(\overline{\varphi^\circ_{\alpha,\beta}}) \leqslant \kappa(\overline{\alpha}\beta) + 1$. Thus,

$$\kappa(\overline{\alpha}\beta) - \kappa(\overline{\varphi^\circ_{\alpha,\beta}}) + 1 \geq \kappa(\overline{\alpha}\beta) - (\kappa(\overline{\alpha}\beta) + 1) + 1 = 0.$$

□

We substantiate the c-revision $\kappa *^{\mathrm{c}} \delta_{\alpha,\beta}$ from (9.16) via the normalization constant $\kappa_0 = -\kappa(\beta)$ in (9.17) and the minimal impact factor $\eta_\delta^{\min}$ in (9.21) and get the following minimal c-revision $\kappa_{\alpha,\beta}^{\mathrm{c,min}}$ (in terms of the impact factor) with $\delta_{\alpha,\beta}$.

[2] This is why we employ strategies in Section 6.4 in order to achieve coherence across different revision scenarios.

Definition 9.3.2. *Let κ be an OCF and $\delta_{\alpha,\beta}$ the core conditional of BR by β w.r.t. α. We define the minimal c-revision $\kappa^{\mathrm{c,min}}_{\alpha,\beta} = \kappa *^{\mathrm{c}} \delta_{\alpha,\beta}$ via* (9.17) *and* (9.21) *as follows*

$$\kappa^{\mathrm{c,min}}_{\alpha,\beta}(\omega) = -\,\kappa(\beta) + \kappa(\omega) + \begin{cases} \max\{0, \kappa(\overline{\alpha}\beta) - \kappa(\overline{\varphi^{\circ}_{\alpha,\beta}}) + 1\}, & \text{for } \omega \not\models \varphi^{\circ}_{\alpha,\beta} \\ 0, & \text{else} \end{cases} \tag{9.22}$$

Having a unique minimal c-revision is useful when we want to compare different Bounded Revision scenarios, especially in the following examples, we use minimal c-revisions $\kappa^{\mathrm{c,min}}_{\alpha,\beta}$.

The following Proposition compares BR for OCFs $\kappa^{\circ}_{\alpha,\beta}$ from (9.12) with $\kappa^{\mathrm{c,min}}_{\alpha,\beta}$.

Proposition 9.3.5. *Let κ be a ranking function. Then it holds for $\kappa^{\circ}_{\alpha,\beta} = \kappa \circ_{\alpha} \beta$ as defined in* (9.12) *and the minimal c-revision $\kappa^{\mathrm{c,min}}_{\alpha,\beta}$ from* (9.22) *that*

$$\kappa^{\mathrm{c,min}}_{\alpha,\beta}(\omega) \leqslant \kappa^{\circ}_{\alpha,\beta}(\omega) \text{ for all } \omega \in \Omega.$$

And $\kappa^{\circ}_{\alpha,\beta}(\omega) = \kappa^{\mathrm{c}}_{\alpha,\beta}(\omega)$ holds if and only if $\kappa(\overline{\varphi^{\circ}_{\alpha,\beta}}) = \kappa(\overline{\beta}) = 0$

Proof. For the minimal impact factor $\eta^{\min}_{\delta} = \max\{0, \kappa(\overline{\alpha}\beta) - \kappa(\overline{\varphi^{\circ}_{\alpha,\beta}}) + 1\}$ in (9.21) defining the minimal c-revision $\kappa^{\mathrm{c,min}}_{\alpha,\beta}$ and the impact factor η°_{δ} defining the c-revision that leads to $\kappa^{\circ}_{\alpha,\beta}$ according to Theorem 9.3.3, it holds that

$$\eta^{\min}_{\delta} = \max\{0, \kappa(\overline{\alpha}\beta) - \kappa(\overline{\varphi^{\circ}_{\alpha,\beta}}) + 1\} \leqslant \kappa(\overline{\alpha}\beta) + 1 = \eta^{\circ}_{\delta}$$

since $\kappa(\overline{\varphi^{\circ}_{\alpha,\beta}}) \geq 0$ and thus, we can conclude that $\kappa^{\mathrm{c,min}}_{\alpha,\beta}(\omega) \leqslant \kappa^{\circ}_{\alpha,\beta}(\omega)$ for all $\omega \in \Omega$.

And therefore, $\kappa^{\circ}_{\alpha,\beta}(\omega) = \kappa^{\mathrm{c,min}}_{\alpha,\beta}(\omega)$ holds iff $\eta^{\circ}_{\delta} = \eta^{\min}_{\delta}$, i.e., for $\kappa(\overline{\varphi^{\circ}_{\alpha,\beta}}) = 0$. □

Note that, Proposition 9.3.5 follows immediately from Theorem 9.3.3 and the (impact-factor wise) minimality of $\kappa^{\mathrm{c,min}}_{\alpha,\beta}$. As we can see, $\kappa^{\circ}_{\alpha,\beta}$ from (9.12) may implement a non-minimal c-revision which may introduce empty layers in the posterior OCF. We illustrate our results via the following example employing minimal c-revisions from (9.22). But first, we transfer the prior TPO $\preceq_{\kappa}$ from Example 9.1.1 in Figure 9.3a to a ranking function by assigning each layer of $\preceq_{\kappa}$ a plausibility rank

$\omega \in \Omega$	κ	$\kappa^{\circ}_{a,b}$	$\kappa^{\mathrm{c,min}}_{a,b}$	$\kappa^{\circ}_{b,c}$	$\kappa^{\mathrm{c,min}}_{b,c}$	$\kappa^{\circ}_{a,\bar{d}}$	$\kappa^{\mathrm{c,min}}_{a,\bar{d}}$
$a\bar{b}c\bar{d}$	0	1	1	0	0	0	0
$a\bar{b}\bar{c}d$	1	2	2	2	1	3	2
$\bar{a}\bar{b}\bar{c}\bar{d}$	1	2	2	2	1	1	1
$\bar{a}bcd$	2	0	0	3	2	4	3
$\bar{a}\bar{b}c\bar{d}$	2	3	3	3	2	4	3
$abcd$	3	4	4	4	3	5	4
$\bar{a}bc\bar{d}$	3	4	4	4	3	5	4
$ab\bar{c}\bar{d}$	4	5	5	5	4	6	5
$\bar{\Omega}$	5	6	6	6	5	7	6

Table 9.1: Prior κ and the BR revised $\kappa^{\circ}_{\alpha,\beta}$ resp. c-revised $\kappa^{\mathrm{c,min}}_{\alpha,\beta}$.

starting with zero for the lowermost layer. This straightforward transformation towards a convex ranking function enables us to compare our results.

Example 9.3.1. *In Table 9.1 the prior convex ranking function κ which corresponds to $\preceq_\kappa$ from Example 9.1.1 is depicted, alongside, the two ranking functions $\kappa^{\circ}_{\alpha,\beta} = \kappa \circ_\alpha \beta$ resp. $\kappa^{\mathrm{c,min}}_{\alpha,\beta} = \kappa \circ \delta_{\alpha,\beta}$ for each of the BR operations.*

For the first BR by $\beta = b$ w.r.t. $\alpha = a$, it holds that $\kappa^{\circ}_{a,b}(\omega) = \kappa^{\mathrm{c,min}}_{a,b}(\omega)$, i.e., BR for OCFs as defined in Definition 9.3.1 and the c-revision with the corresponding conditional yield the same result. This is due to the fact that $\kappa(\overline{\varphi_{a,b}}) = \kappa(a\bar{b}c\bar{d}) = 0$ and therefore $\eta_\delta = \kappa(\bar{a}b) + 1 = 3$ for $\kappa^{\mathrm{c,min}}_{a,b}$, i.e., the increase of ranks implemented by $\kappa^{\circ}_{a,b}$ and $\kappa^{\mathrm{c,min}}_{a,b}$ is the same.

For the next BR with reference $\alpha = b$ and input sentence $\beta = c$, it holds that $c \in Bel(\kappa)$ and $\Phi_{b,c} = \{a\bar{b}c\bar{d}\}$. So, the input is already accepted in the prior ranking function, and (BR3) is satisfied. Here, BR for OCFs $\kappa^{\circ}_{b,c}$ strengthens the input c via introducing an empty layer, s.t., $\kappa^{\circ}_{b,c}(c) = 0 < 2 = \kappa^{\circ}_{b,c}(\bar{c})$, instead of $\kappa(c) = 0 < 1 = \kappa(\bar{c})$ as before. In contrast, c-revisions as in $\kappa^{\mathrm{c,min}}_{b,c}$ employ minimal change since (BR3) is already satisfied, and therefore the prior OCF is kept.

For the next BR with $\alpha = a$ and $\beta = \bar{d}$, κ accepts the input $\bar{d}$, but (BR3) is not satisfied, since $\Phi_{a,\bar{d}} = \{a\bar{b}c\bar{d}, \bar{a}\bar{b}\bar{c}\bar{d}\}$ and $\kappa(\bar{a}\bar{b}\bar{c}\bar{d}) = \kappa(a\bar{b}\bar{c}d)$ holds. Again, the change implemented by $\kappa^{\mathrm{c,min}}_{a,\bar{d}}$ is minimal, in the sense that the ranks of all worlds outside the core set are increased by exactly the stretch needed to satisfy (BR3), namely

$\kappa(\bar{a}\bar{d}) - \kappa(\overline{\varphi}_{a,\bar{d}}) + 1 = 1$. *BR for OCFs* $\kappa^{\circ}_{a,\bar{d}}$ *leads to a strengthening of the input* $\bar{d}$*, while satisfying (BR1) – (BR3).*

Note that it holds for all BR revised and c-revised ranking functions that their corresponding plausibilistic TPOs coincide with the BR revised TPOs from Figures 9.3b, 9.6a and 9.6b.

The example shows that minimal c-revisions do not change the prior ordering when it is not necessary. This corresponds to a minimal implementation of (BR1) and (BR2), in the sense that the specific ranks of worlds in and outside the core set of BR are kept as long as $\delta_{\alpha,\beta}$ holds. On the other hand, the non-minimal change in BR for OCFs $\kappa^{\circ}_{\alpha,\beta}$ is useful in cases where one wants to strengthen the input belief β by decreasing the plausibility level of worlds outside the core, even though β is already accepted. Note that BR generally increases the number of layers in the posterior ordering, making the belief state more fine-grained.

Intermediate Summary for Part II

We investigated two frameworks for parameterized belief revision, Revision by Comparison and Bounded Revision. These operators take as input not only input information $\beta \in \mathcal{L}$ but also a reference sentence $\alpha \in \mathcal{L}$. Starting from the original works [34] for RbC and [100] for BR, we clarified the role of the parameter α in the framework of plausibilistic TPOs and provided (methodological) implementations for both mechanisms in the framework of conditional revision via c-revisions. Thus, we fully captured the parameterized belief change character of both operations by relying on the internal strengths of conditionals in Belief Revision.

Revision by Comparison. The main goal of Revision by Comparison is to make the input sentence β at least as plausible as the reference sentence α. In Chapter 8, we provided an elegant semantic characterization of RbC for plausibilistic TPOs via three postulates in Theorem 8.1.4, which fully captures the mechanism of RbC. A unique feature of RbC is its hybrid belief change character. It depends on the prior relative positioning of β versus α, whether RbC revises the initial belief state with β or whether RbC leads to the contraction of the former belief α.

The characterization of RbC via two designated formulas, the penalty formula $\psi^{\circledcirc}_{\alpha,\beta}$ and the indirect reward formula $\theta^{\circledcirc}_{\alpha,\beta}$, and the postulates (RbC1) – (RbC3) provided grounds for the investigation of RbC in the framework of OCFs, which lead to the implementation of RbC as a c-revision with a set of weak conditionals $\Delta^{\circledcirc}_{\alpha,\beta}$. Thus, revealing the direct correspondence between RbC and a special iterated contraction operator for conditionals.

Bounded Revision. The main goal of Bounded Revision is to accept the input sentence β as long as the reference sentence α holds and just a little further. In Chapter 9, we presented a Representation Theorem 9.1.3 for Bounded Revision which provides us with three postulates (BR1) – (BR3) that fully characterize the change mechanism of BR via a core formula $\varphi^{\circ}_{\alpha,\beta}$. It holds that BR uniformly shifts all worlds that do not satisfy $\varphi^{\circ}_{\alpha,\beta}$ to plausibility levels strictly higher than one of the least plausible models of $\varphi^{\circ}_{\alpha,\beta}$. Thus, BR tends to increase the number of plausibility levels. Also, it displays an iterated belief revision operator, s.t. β is always accepted independent from the choice of α. Following the line of Rott from [100], we characterized BR as a lexicographic revision with a designated core formula in Theorem 9.2.2 and showed that for special choices of α BR performs a lexicographic resp. natural revision by β.

In Section 9.3, we transferred our results to the framework of ranking functions and showed that BR can be realized as a c-revision with a single core conditional, which fully captures the underlying change mechanism. Implementing the formerly parameterized revision mechanism via a conditional revision operator with a single (non-parameterized) input makes the ensuing application of BR easily accessible.

Chapter 10

Revision by Comparison Versus Bounded Revision

In this chapter, we comprehensively compare two parameterized belief change mechanisms: Revision by Comparison (RbC) and Bounded Revision (BR), both of which we have individually and thoroughly investigated in the preceding chapters. We aim to contrast these mechanisms on the qualitative and semi-quantitative levels and identify their strengths and weaknesses.

In Section 10.1, we initiate our comparison by presenting an illustrative example that serves as a guide for the subsequent analysis. We examine the commonalities and differences in the respective change mechanisms and highlight the characteristic properties that justify these differences. Furthermore, in Section 10.2, we examine the realizations of RbC and BR in the framework of ranking functions. We transfer our illustrative example to the OCF framework and also illustrate the corresponding conditional revisions for RbC and BR. In particular, we focus on the implementation via c-revisions, which provide us with valuable insights into the change mechanisms. We discuss these insights at the end of the section, highlighting the implications of these findings for our analysis.

Moreover, in this context, we also clarify the impact of empty layers in the prior ranking functions, shedding light on how these affect the behavior of the two mechanisms. This comparative analysis provides meaningful insights into the differences between RbC and BR and their respective strengths and limitations.

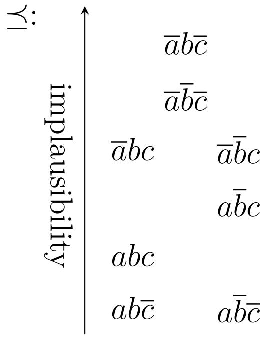

a) Plausibilistic TPO $\preceq$.

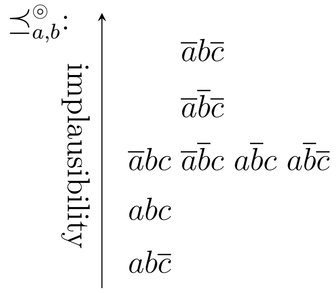

b) RbC-revised TPO $\preceq \circledcirc_a b = \preceq^{\circledcirc}_{a,b}$.

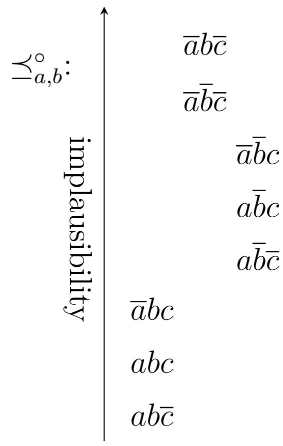

c) BR revised TPO $\preceq \circ_a b = \preceq^{\circ}_{a,b}$.

Figure 10.1: Revision by Comparison versus Bounded Revision (cf. Example 10.1.1).

Bibliographic Remark. This chapter presents a comparison between Revision by Comparison and Bounded Revision, based on the thorough investigations of both revision formalisms presented in the joint work with Gabriele Kern-Isberner [109, 110]. However, the results and conclusions of this comparison are relatively new, although some ideas and approaches are inspired by Rott's work [100].

10.1 Qualitative Approaches

This section compares the mechanisms of RbC and BR in the qualitative framework.

First, we present an example of RbC resp. BR of a plausibilistic TPO side by side. This is helpful to illustrate the main commonalities and differences, which we discuss in the following paragraphs more thoroughly.

Example 10.1.1. *In Figure 10.1a a plausibilistic TPO* $\preceq$ *with signature* $\Sigma = \{a, b, c\}$ *is given. We perform a Revision by Comparison* $\preceq \circledcirc_a b$ *in Figure 10.1b and a Bounded Revision* $\preceq \circ_a b$ *in Figure 10.1c with input* $\beta = b$ *and reference sentence* $\alpha = a$.

For Revision by Comparison we get the following crucial formulas and sets, the penalty formula and set resp. the reward formula resp. set from Definition (8.1.2) *on page 157, defining the change mechanism.*

- *Penalty formula* $\psi^{\circledcirc}_{a,\beta} = \overline{b} \wedge (\bigvee_{\omega \prec \overline{a}} \omega)$ *with the corresponding penalty set* $\Psi^{\circledcirc}_{a,b} = \{a\overline{b}\overline{c}, a\overline{b}c\}$
- *Indirect reward formula* $\theta^{\circledcirc}_{a,b} = b \wedge (\bigvee_{\omega \prec \overline{a}} \omega)$ *with the corresponding reward set* $\Theta^{\circledcirc}_{a,b} = \{abc, ab\overline{c}\}$

And for the core formula resp. the core set from Definition 9.1.2 on page 192 defining the change mechanism Bounded Revision the following holds:

- *Core formula* $\varphi^{\circ}_{a,b} = b \wedge (\bigvee_{\omega \preceq \overline{a}b} \omega)$ *with the corresponding core set* $\Phi^{\circ}_{a,b} = \{abc, ab\overline{c}, \overline{a}bc\}$

Revision by Comparison and Bounded Revision both implement parameterized belief change mechanisms where the input information β is accompanied by additional information in the form of a reference sentence α. The reference serves as an anchor point, which specifies to what extent β is to be accepted in the posterior belief state.

The main idea behind Revision by Comparison is that $\overline{\beta}$ shall be at most as plausible as $\overline{\alpha}$ in the posterior ordering. This goal is expressed in the property $(\mathrm{RbC})_{\preceq}$ from page 153.

(RbC)$_{\preceq}$ $\bar{\alpha}$ is at least as plausible as $\overline{\beta}$, i.e., $\bar{\alpha} \preceq^{\circledcirc}_{\alpha,\beta} \overline{\beta}$

For Bounded Revision $\circ_\alpha \beta$ on epistemic entrenchment relations, it holds that β shall be *just a little more entrenched* than α. In terms of plausibility orderings of worlds, we get for the success condition of BR, $(\mathrm{BR})_{\preceq}$ from page 187, in the BR revised state $\preceq \circ_\alpha \beta = \preceq^{\circ}_{\alpha,\beta}$, a strict inequality expressing that $\overline{\alpha}$ shall be strictly more plausible than $\overline{\beta}$.

(BR)$_{\preceq}$ $\overline{\alpha}$ is strictly more plausible than $\overline{\beta}$, i.e., $\overline{\alpha} \prec^{\circ}_{\alpha,\beta} \overline{\beta}$

Thus, while Revision by Comparison follows an *at-least-as strategy* (‘accept β at least as strongly as α’), Bounded Revision opts for an *as-long-as strategy* (‘accept β as long as α holds along with β, and just a little further’). These differences in the reference sentence’s role lead to non-strict resp. strict inequalities concerning the relative positioning of $\overline{\alpha}$ versus $\overline{\beta}$ in the success condition of the corresponding change mechanism.

The posterior TPOs $\preceq^{\circledcirc}_{a,b}$ and $\preceq^{\circ}_{a,b}$ from Example 10.1.1 satisfy (RbC)$_{\preceq}$ resp. (BR)$_{\preceq}$, i.e., it holds that $\overline{a} \preceq^{\circledcirc}_{a,b} \overline{b}$ and $\overline{a} \prec^{\circ}_{a,b} \overline{b}$. The *at-least-as strategy* for RbC is implemented by the penalty formula $\psi^{\circledcirc}_{a,\beta} = \overline{b} \wedge (\bigvee_{\omega \prec \overline{a}} \omega)$. Via the penalty formula $\psi^{\circledcirc}_{a,\beta}$ and due to (RbC3), it holds that all $\overline{b}$-worlds that are strictly more plausible than $\overline{a}$ are punished, s.t. $\overline{a}$ is at least as plausible as $\overline{b}$. This means, in our example, the worlds $a\overline{b}\overline{c}, a\overline{b}c$ are lifted to the plausibility level of the minimal $\overline{a}$-worlds $\overline{a}bc, \overline{a}\overline{b}c$. This property is reflected in the general success condition of RbC, (RbC)$_{\preceq}$, and it leaves worlds satisfying the (indirect) reward formula $\theta^{\circledcirc}_{a,b} = b \wedge (\bigvee_{\omega \prec \overline{a}} \omega)$, i.e., b-worlds strictly less plausible than $\overline{a}$, among the most plausible ones.

Formula $\theta^{\circledcirc}_{a,b}$ resembles the core formula $\varphi^{\circ}_{a,b} = b \wedge (\bigvee_{\omega \preceq \overline{a}b} \omega)$ defining BR $\preceq^{\circ}_{a,b}$. Yet, the core formula $\varphi^{\circ}_{a,b}$ includes minimal worlds satisfying $\overline{a}$ along with the input formula b, i.e., b-worlds more or equally plausible than $\overline{a}b$. This is in line with the *as-long-as strategy* implemented by BR, which is reflected in (BR)$_{\preceq}$. Because in the BR revised TPO, it holds that worlds satisfying the core formula $\varphi^{\circ}_{a,b}$ are strictly more plausible than the remaining worlds, s.t. b is accepted as long as a holds, and just a little further, namely as far as a minimal $\overline{a}b$-world is accepted in the prior ordering.

Moreover, we observe in Example 10.1.1 that while BR tends to refine plausibilistic TPOs, i.e., the number of layers in the plausibility ordering increase, RbC has the opposite effect and tends to coarsen the belief state, i.e., the number of layers decrease.

In contrast to BR, RbC embodies not exclusively an operation of belief revision but sometimes also a contraction of beliefs. We discussed this hybrid belief change character in-depth in Section 8.2. This fundamental difference leads to differences in the level of postulates satisfied by RbC resp. BR. So, it holds that due to the same beliefs condition $(\mathrm{SBC})_{\preceq}$ from page 187, BR is invariably successful, i.e., the input sentence is always accepted independently of the choice of the reference sentence. In Figure 8.6b, the α-contraction case of RbC is depicted, i.e., the case where RbC of β w.r.t. α leads to a posterior state where $\alpha \notin Bel(\preceq \circledcirc_{\alpha} \beta)$, while the acceptance of the input β is not guaranteed. This example shows that RbC, in contrast to BR, does not incorporate the input sentence β into the agent's belief set successfully in general, but only in cases where $\overline{\alpha}$ is strictly less plausible than β. RbC is a non-prioritized belief change mechanism since it does not prioritize new input over previous beliefs. However, the same-beliefs condition $(\mathrm{SBC})_{\preceq}$ is satisfied by RbC, at least in special cases, as was shown in [99]. The special cases are summarized by the following postulates, which we state using plausibilistic TPOs:

$(\mathbf{SBC})^1_{\preceq}$ If $\beta \in Bel(\preceq \circledcirc_{\alpha} \beta)$ and $\beta \in Bel(\preceq \circledcirc_{\gamma} \beta)$, then $Bel(\preceq \circledcirc_{\alpha} \beta) = Bel(\preceq \circledcirc_{\gamma} \beta)$

$(\mathbf{SBC})^2_{\preceq}$ If $\overline{\alpha} \approx \overline{\gamma}$, then $Bel(\preceq \circledcirc_{\alpha} \beta) = Bel(\preceq \circledcirc_{\gamma} \beta)$

Postulate $(\mathbf{SBC})^1_E$ deals with the case when β is part of the posterior belief set for both reference sentences α, γ, i.e., it compares the belief set of different successful revisions in terms of incorporating the input sentence β. $(\mathbf{SBC})^2_E$ claims that the resulting belief set of RbC is the same for reference sentences from the same plausibility level, i.e., if $\overline{\alpha} \approx \overline{\gamma}$ [99].

Bounded Revision was motivated in [100] by the same concerns as Revision by Comparison, combined with the demand to satisfy the DP postulates. It holds that Revision by Comparison fails to satisfy the DP postulates for iterated revision functions [34]. On the other hand, BR was designed to satisfy the DP postulates and, thus, does so naturally. Note that, in the presence of the $(\mathrm{SBC})_{\preceq}$ axiom and if we are ready to tolerate a little sloppiness in the notation, we can simply use the original formulation of the DP postulates from Section 2.3 even though BR displays a parameterized belief revision[1]. In contrast to BR, RbC fails to satisfy the second DP postulate (DP2) because, in the posterior RbC-revised state, the plausibility distinctions between $\overline{\beta}$ worlds satisfying the penalty formula $\psi^{\circledcirc}_{\alpha,\beta}$ are lost. This can

[1]Rott adapted the DP postulates to the general case of parameterized belief revision. However, he already mentioned that due to $(\mathrm{SBC})_E$, the differences between the original formulation and his parameterized version vanish.

also be observed in our Example 10.1.1, where all worlds in $\Psi^{\circledcirc}_{a,b}$ end up on the same plausibility level, whereas for the BR $\preceq^{\circ}_{a,b}$, the internal plausibility ordering of worlds within and outside of the core set $\Phi^{\circ}_{a,b}$ are kept. Fermé and Rott showed in [34] that iterations of RbC are fairly well-behaved in some special cases, but in general, it is not possible to relate a number of finite applications of $\circledcirc_{\alpha}\,\beta$ for arbitrary input and reference sentences to a single one in a rational way. Note that this makes sense if we take into account that RbC tends to decrease the number of plausibility levels. Thus, it is possible that with a repeated application with the right choice of input and reference sentences for the RbC-operation, the prior state finally collapses to a plausibilistic TPO, which consists only of a single layer.

However, the idea of iterated change in the sense of Darwiche and Pearl should not be given up at this point, and we provided new insights in Section 8.3 where we characterized RbC as a c-revisions with a set of weak conditionals in (8.17) on page 179. Since weak conditionals implement negative information, i.e., correspond to the notion that the respective negated standard conditional does not hold (cf. paragraph on weak conditionals in Section 2.1 on page 19), the revision with Δ^{w} from Definition 8.3.2 on page 175 can be seen as c-contraction of the corresponding negated standard conditionals. Taking this perspective into account leads to the conclusion that RbC actually displays an iterated contraction operator as discussed in Section 2.5.4 and thus is not only applicable to belief states but also respects the principle of conditional preservation at least w.r.t. $\Delta^{\circledcirc}_{\alpha,\beta}$ and the corresponding negated standard conditionals.

10.2 Realization via Ranking Functions

This section investigates and compares the OCF realizations of RbC and BR for input β and reference sentence α. We start with the comparison of the straightforward implementation of RbC for OCFs (cf. (8.12) on page 173) versus BR for OCFs (cf. (9.12) on page 209) before we turn to the comparison of the corresponding implementations via c-revisions. We conclude by investigating the issue of empty layers in the context of OCFs.

Note that the OCF versions of RbC versus BR reveal a main difference between the change mechanisms due to the employment of the arithmetic inherent to OCFs. While RbC tends to punish worlds within the penalty set by making them less plausible, BR promotes worlds within the core set by making them more plausible. For BR, the implausibility ranks of all worlds outside the core set are increased, while worlds in the core set remain on their previous level (modulo normalization). This is not clearly seen when considering the TPO versions of these operations solely.

Furthermore, it is obvious that for the OCF versions given in Definition 8.3.1 for $\kappa^{\circledcirc}_{\alpha,\beta}$ on page 173 and in Definition 9.3.1 for $\kappa^{\circ}_{\alpha,\beta}$ on page 209 the calculation comes at linear costs since we add a constant factor to worlds within a certain set, namely the penalty set $\Psi^{\circledcirc}_{\alpha,\beta}$ for RbC, resp. outside a certain set, namely the core set $\Phi^{\circ}_{\alpha,\beta}$ for BR. Yet, the calculation of these sets is of the same order as the SAT test since we have to check whether $\psi^{\circledcirc}_{\alpha,\beta}$ resp. $\varphi^{\circ}_{\alpha,\beta}$ holds for each world $\omega \in \Omega$.

We transfer Example 10.1.1 to the framework of ranking functions by assigning ranks to each layer of $\preceq$ from Figure 10.1a. Here, we first give a straightforward translation of $\preceq$ to a convex ranking function κ. And then, we consider a ranking function $\tilde{\kappa}$, which leads to the same TPO as $\preceq$ via the translation (2.9) but has some empty layers. Taking these two OCFs as a basis, we perform an RbC resp. BR of b w.r.t. a, first via RbC for OCFs from (8.12) resp. BR for OCFs from (9.12) and then via the corresponding c-revisions.

Example 10.2.1. *In Table 10.1 two OCFs κ and $\tilde{\kappa}$ are given which correspond to $\preceq$ from Figure 10.1a via* (2.9). *Note that, while κ is a convex OCF $\tilde{\kappa}$ has some empty layers.*

We perform an RbC of b w.r.t. a for both κ and $\tilde{\kappa}$. It holds that $\Psi^{\circledcirc}_{a,b} = \{a\overline{b}c, a\overline{b}\overline{c}\}$. For the normalization constants we get $\kappa_0 = -\min\{\kappa(\overline{a}), \kappa(b)\} = 0$ for κ and $\kappa_0 = -\min\{\kappa(\overline{a}), \kappa(b)\} = 0$ for $\tilde{\kappa}$. So, we get for RbC for OCFs from Definition 8.3.1

$$\kappa^{\circledcirc}_{a,b}(\omega) = 0 + \begin{cases} 3, & \omega \in \Psi^{\circledcirc}_{a,b} \\ \kappa(\omega) & \textit{othw.} \end{cases} \quad \textit{and} \quad \tilde{\kappa}^{\circledcirc}_{a,b}(\omega) = 0 + \begin{cases} 5, & \omega \in \Psi^{\circledcirc}_{a,b} \\ \kappa(\omega) & \textit{othw.} \end{cases},$$

because $\kappa(\overline{a}) = 3$ resp. $\tilde{\kappa}(\overline{a}) = 5$ holds. Both OCFs $\kappa^{\circledcirc}_{a,b}$ and $\tilde{\kappa}^{\circledcirc}_{a,b}$ are depicted in Table 10.1 in the column for Revision by Comparison. For RbC as a c-revision, we get the following RbC base $\Delta^{\circledcirc}_{a,b} = \{(|\overline{a} \vee a\overline{b}c|a|), (|\overline{a} \vee a\overline{b}\overline{c}|a|)\}$ with the following *minimal impact factors*

$$\eta^{\min}_{a\bar{b}c} = \kappa(\overline{a}) - \kappa(a\bar{b}c) = 3 - 2 = 1 \textit{ and } \tilde{\eta}^{\min}_{a\bar{b}c} = \tilde{\kappa}(\overline{a}) - \tilde{\kappa}(a\bar{b}c) = 5 - 2 = 3$$
$$\eta^{\min}_{a\bar{b}\bar{c}} = \kappa(\overline{a}) - \kappa(a\bar{b}\bar{c}) = 3 - 0 = 3 \textit{ and } \tilde{\eta}^{\min}_{a\bar{b}\bar{c}} = \tilde{\kappa}(\overline{a}) - \tilde{\kappa}(a\bar{b}\bar{c}) = 5 - 0 = 5.$$

Via $\eta_{a\bar{b}c}$ resp. $\eta_{a\bar{b}\bar{c}}$ and (8.22), we get for $\kappa^{\mathrm{c,min}}_{a,b}$ resp. $\tilde{\kappa}^{\mathrm{c,min}}_{a,b}$

$$\kappa^{\mathrm{c,min}}_{a,b}(\omega) = 0 + \begin{cases} 3, & \omega \in \Psi^{\circledcirc}_{a,b} \\ \kappa(\omega) & \textit{othw.} \end{cases} \quad \textit{and} \quad \tilde{\kappa}^{\mathrm{c,min}}_{a,b}(\omega) = 0 + \begin{cases} 5, & \omega \in \Psi^{\circledcirc}_{a,b} \\ \kappa(\omega) & \textit{othw.} \end{cases},$$

Both c-revised OCFs $\kappa^{\mathrm{c,min}}_{a,b}$ and $\tilde{\kappa}^{\mathrm{c,min}}_{a,b}$ are depicted in Table 10.1 in the column for Revision by Comparison. Note that, as in RbC for OCFs, the normalization constant vanishes.

In Table 10.1, we can see that both $\kappa^{\circledcirc}_{a,b}$ and $\kappa^{\mathrm{c,min}}$ are not convex, even though the prior OCF κ was. This is because if we conditionalize with $\overline{\psi^{\circledcirc}_{\alpha,\beta}}$ we get $\kappa|\overline{\psi^{\circledcirc}_{\alpha,\beta}}(ab\bar{c}) = 0$, $\kappa|\overline{\psi^{\circledcirc}_{\alpha,\beta}}(abc) = 1$, $\kappa|\overline{\psi^{\circledcirc}_{\alpha,\beta}}(\bar{a}bc) = \kappa|\overline{\psi^{\circledcirc}_{\alpha,\beta}}(\bar{a}\bar{b}c) = 3$, $\kappa|\overline{\psi^{\circledcirc}_{\alpha,\beta}}(\bar{a}\bar{b}\bar{c}) = 4$, and $\kappa|\overline{\psi^{\circledcirc}_{\alpha,\beta}}(\bar{a}b\bar{c}) = 5$ which is not convex. Both $\kappa^{\circledcirc}_{a,b}$ and $\kappa^{\mathrm{c,min}}$ keep the distance relations for all worlds satisfying $\overline{\psi^{\circledcirc}_{\alpha,\beta}}$, i.e., the empty layers between worlds in $Mod(\theta^{\circledcirc}_{a,b} \vee (\bigvee_{\min(\bar{a},\kappa)})$ are preserved.

For BR of b w.r.t. to a for κ resp. $\tilde{\kappa}$, it holds that $\Phi^{\circ}_{a,b} = \{abc, ab\bar{c}, \bar{a}bc\}$ and the normalization constants equal $\kappa_0 = -\kappa(b) = 0$ for κ and $\kappa_0 = -\kappa(b) = 0$ for $\tilde{\kappa}$. Hence, we get for BR for OCFs from Definition 9.3.1

$$\kappa^{\circ}_{a,b}(\omega) = 0 + \begin{cases} \kappa(\omega) + 3, & \omega \in \Phi^{\circ}_{a,b} \\ \kappa(\omega), & \textit{othw.} \end{cases} \quad \textit{and} \quad \tilde{\kappa}^{\circ}_{a,b}(\omega) = 0 + \begin{cases} \kappa(\omega) + 5, & \omega \in \Phi^{\circ}_{a,b} \\ \kappa(\omega), & \textit{othw.} \end{cases},$$

because $\kappa(\bar{a}b) = 3$ resp. $\tilde{\kappa}(\bar{a}b) = 5$. Both OCFs $\kappa^{\circ}_{a,b}$ and $\tilde{\kappa}^{\circ}_{a,b}$ are depicted in Table 10.1 in the column for Bounded Revision.

For BR as a c-revision, we get the core conditional $\delta^{\circ}_{a,b} = (\varphi^{\circ}_{a,b}|\overline{\varphi^{\circ}_{a,b}} \vee (\bar{a}b))$ with the corresponding minimal impact factors

$$\eta^{\min}_{\delta} = \max\{0, \kappa(\bar{a}b) - \kappa(\overline{\varphi^{\circ}_{a,b}}) + 1\} = \max\{0, 3 - 0 + 1\} = 4$$
$$\textit{and } \tilde{\eta_{\delta}}^{\min} = \max\{0, \tilde{\kappa}(\bar{a}b) - \tilde{\kappa}(\overline{\varphi^{\circ}_{a,b}}) + 1\} = \max\{0, 5 - 0 + 1\} = 6$$

			Revision by Comparison				Bounded Revision			
$\omega \in \Omega$	κ	$\tilde{\kappa}$	$\kappa^{\circledcirc}_{a,b}$	$\kappa^{\mathrm{c,min}}_{a,b}$	$\tilde{\kappa}^{\circledcirc}_{a,b}$	$\tilde{\kappa}^{\mathrm{c,min}}_{a,b}$	$\kappa^{\circ}_{a,b}$	$\kappa^{\mathrm{c,min}}_{a,b}$	$\tilde{\kappa}^{\circ}_{a,b}$	$\tilde{\kappa}^{\mathrm{c,min}}_{a,b}$
abc	1	1	1	1	1	1	1	1	1	1
$ab\overline{c}$	0	0	0	0	0	0	0	0	0	0
$a\overline{b}c$	2	2	3	3	5	5	6	6	8	8
$a\overline{b}\overline{c}$	0	0	3	3	5	5	4	4	6	6
$\overline{a}bc$	3	5	3	3	5	5	3	3	5	5
$\overline{a}b\overline{c}$	5	7	5	5	7	7	9	9	13	13
$\overline{a}\overline{b}c$	3	5	3	3	5	5	7	7	11	11
$\overline{a}\overline{b}\overline{c}$	4	6	4	4	6	6	8	8	12	12

Table 10.1: RbC vs. BR for OCFs resp. as c-revisions (cf. Example 10.2.1).

Via $\eta_\delta^{\min}$ resp. $\tilde{\eta_\delta}^{\min}$ and (9.22), we get for $\kappa^{\mathrm{c,min}}_{a,b}$ resp. $\tilde{\kappa}^{\mathrm{c,min}}_{a,b}$

$$\kappa^{\mathrm{c,min}}_{a,b}(\omega) = 0 + \kappa(\omega) + \begin{cases} 4, & \omega \in \Phi^{\circ}_{a,b} \\ 0, & \mathit{othw.} \end{cases} \quad \mathit{and} \quad \tilde{\kappa}^{\mathrm{c,min}}(\omega) = 0 + \begin{cases} 6, & \omega \in \Phi^{\circ}_{a,b} \\ 0, & \mathit{othw.} \end{cases}$$

Both c-revised OCFs $\kappa^{\mathrm{c,min}}_{a,b}$ and $\tilde{\kappa}^{\mathrm{c,min}}_{a,b}$ are depicted in Table 10.1 in the column for Bounded Revision.

As for RbC, the convexity of κ is not preserved by BR. Even though κ is convex the conditionalized OCF $\kappa|\varphi^{\circ}_{\alpha,\beta} : \Phi^{\circ}_{\alpha,\beta} \to \mathbb{N}$ is not. It holds that $\kappa|\varphi^{\circ}_{\alpha,\beta}(ab\overline{c}) = 0$, $\kappa|\varphi^{\circ}_{\alpha,\beta}(abc) = 1$ and $\kappa|\varphi^{\circ}_{\alpha,\beta}(\overline{a}bc) = 3$. Both BR for OCFs $\kappa^{\circ}_{a,b}$ and the (minimally) c-revised $\kappa^{\mathrm{c,min}}_{a,b}$ preserve these empty layers because the relative distances between worlds within the core set are kept, and we get an empty layer for rank 2 in $\kappa^{\circ}_{a,b}$ resp. $\kappa^{\mathrm{c,min}}_{a,b}$. Note that the same holds for the counterpart $\kappa|\overline{\varphi^{\circ}_{\alpha,\beta}}$, also here empty layers are preserved, which is why we have an empty layer for rank 5 for $\kappa^{\mathrm{c,min}}_{a,b}$ resp. rank 7 for $\kappa^{\circ}_{a,b}$.

The difference between RbC and BR becomes most apparent when looking at the c-revision for RbC versus the c-revision defining BR. RbC lifts worlds within the penalty set $\Psi^{\circledcirc}_{\alpha,\beta}$ on the $\overline{\alpha}$-level of plausibility. Thus, the shift of worlds in the penalty set $\omega \in \Psi^{\circledcirc}_{\alpha,\beta}$, and therefore the corresponding impact factor η_ω, depends on the rank-wise plausibility of ω and the distance between $\kappa(\omega)$ and $\kappa(\overline{\alpha})$. In BR, all worlds outside the core set are incremented by the same value, i.e., the value of the impact factor is the same for each shifted world. So, we can define a single core conditional where the corresponding impact factor adds the same value to all worlds falsifying it. Example 10.2.1 illustrates these different c-revisions for RbC resp. BR. Also, we see from Example 10.2.1 that for ranking functions with empty layers, like $\tilde{\kappa}$, RbC and BR for OCFs and also the corresponding c-revisions preserve the prior empty layers.

In general, it is possible that the realizations of RbC $\kappa^{\circledcirc}_{a,b}$ resp. BR $\kappa^{\circ}_{a,b}$ for OCFs introduce additional empty layers since the plausibility distances within and outside the penalty set resp. the core set are kept. This effect occurs for both convex OCFs, like κ in Example 10.2.1, and OCFs with empty layers, like $\tilde{\kappa}$ in Example 10.2.1, and is due to the distribution of worlds within the indirect reward set for RbC resp. within and outside the core set for BR. For RbC, it holds that if there are empty layers in the conditionalized OCF $\kappa|(\overline{\psi^{\circledcirc}_{\alpha,\beta}}) : \Omega \setminus \psi^{\circledcirc}_{\alpha,\beta} \rightarrow \mathbb{N}$ with $\kappa|\overline{\psi^{\circledcirc}_{\alpha,\beta}}(\omega) = \kappa(\omega) - \kappa(\overline{\psi^{\circledcirc}_{\alpha,\beta}})$ these empty layers are preserved during RbC, which follows immediately from (RbC1) because RbC does not change the relative distances between worlds outside $\Psi^{\circledcirc}_{\alpha,\beta}$. For BR, we can observe the same effect for both worlds within and outside the core set, i.e., for $\kappa|\varphi^{\circ}_{\alpha,\beta} : \Phi^{\circ}_{\alpha,\beta} \rightarrow \mathbb{N}$ with $\kappa|\varphi^{\circ}_{\alpha,\beta}(\omega) = \kappa(\omega) - \kappa(\varphi^{\circ}_{\alpha,\beta})$ resp. the counterpart $\kappa|\overline{\varphi^{\circ}_{\alpha,\beta}} : \overline{\Phi^{\circ}_{\alpha,\beta}} \rightarrow \mathbb{N}$ with $\kappa|\overline{\varphi^{\circ}_{\alpha,\beta}}(\omega) = \kappa(\omega) - \kappa(\overline{\varphi^{\circ}_{\alpha,\beta}})$ empty layers are preserved during BR, which follows from (BR1) resp. (BR2).

10.3 Conclusion

We conclude this section with a brief summary of our comparison results. RbC and BR are both parameterized belief change operators motivated by similar concerns, and both provide a solution for the parameterized Belief Revision problem (ParamRev) posed in the introduction of this part. However, their corresponding change mechanisms differ fundamentally, which is reflected in their qualitative as well as their semi-quantitative implementations. In the qualitative framework, our thorough investigations have shown that a significant part of these differences is because RbC displays a non-prioritized revision operator, which does not satisfy the DP postulates, while BR is an iterated belief revision operator in the sense of

Darwiche and Pearl. Taking on a conditional perspective and implementing RbC resp. BR as a conditional revision with a set of weak conditionals resp. a standard conditional revealed deeper insights and further clarified our previous findings. The non-prioritized change mechanism of RbC corresponds to an iterated contraction operator, which is implemented via the revision with a set of weak conditionals. For BR, on the other hand, it is enough to c-revise with a single conditional which relies on the relative positioning of the reference versus the input sentence in the prior ordering, i.e., BR is directly related to changing conditional beliefs in such a way that not only the DP postulates but also the more general principle of conditional preservation (cf. Section 2.5.2) is satisfied.

Chapter 11

Conclusion and Future Work

In this final chapter, we recapitulate the essential role of conditionals in Belief Revision and evaluate the results of our research with regard to the conditional perspective we discussed so far. This chapter serves as a reflection of our findings and a pointer for future research questions and further investigations.

One of the most influential extensions to the approach to belief change by Alchourrón, Gärdenfors, and Makinson (AGM) is the well-known theory by Darwiche and Pearl (DP) on iterated belief change [29]. One distinctive feature in the DP framework is the employment of belief states that do not only consist of plain beliefs. Belief states, in the sense of Halpern [48], encode not only what the agent deems to be certain but also provide the necessary, richer structure to implement an agent's conditional beliefs, i.e., the beliefs an agent is ready to accept in the light of new information. The essential meta-structure which enables us to provide coherence over iterated applications of belief revision operators are total preorders related to an agent's preferences [60, 29]. Thus, conditional beliefs and iterated Belief Revision in the sense of Darwiche and Pearl are inherently linked to one another, and one cannot talk about one without meaning the other. Therefore, the adequate treatment of conditional beliefs resides at the core of Belief Revision.

Darwiche and Pearl recognized that the key to rationally revising belief states lies in the competent preservation of conditional beliefs. They provided vague guidelines on how this goal can be achieved. In [63, 64], Kern-Isberner provided a complete axiomatization of a *principle of conditional preservation* (PCP), which subsumes all of the postulates proposed in the DP framework for minimizing changes in conditional beliefs rationally. This fundamental principle directs us to a powerful generalization of the classical AGM framework by providing a method for handling revision by sets of conditionals simultaneously, called *c-revisions* [64]. C-revisions are a gateway for a holistic view of conditionals in Belief Revision since they provide the necessary means to monitor and preserve conditional beliefs by obeying the (PCP) on the one hand and, on the other hand, open up the classical Belief Revision framework to operators that can revise with conditional information directly. Conditional revision operators allow us to monitor and manipulate conditional beliefs as a crucial yet subtle basis of the epistemic state. Thus provide an integral part of the conditional perspective of Belief Revision discussed in this thesis.

Our results decompose naturally into two parts concerned with different aspects of conditionals in Belief Revision. In the first part, the revision with sets of conditionals constitutes the general framework for our investigations. Here, we focus on a specific dynamic that introduces a locality notion via exploiting conditional revision's special features. In the second part, propositional revision with respect to an additional meta-information provides ground for our research. Here, we showed that the richer structure of conditionals enables us to reduce the parameterized operators presented in this part to conditional c-revisions. In the following, we recapitulate and answer the general research questions posed in the introduction of this thesis.

We started our investigations by considering the following research question concerning the special features provided by conditionals as input for a belief revision:

> *How does the specific context of information affect the revision task, and how can we benefit from the inclusion of exclusive contexts in the revision process?*

We precisely answered this question by providing and subsequently thoroughly investigating a Kinematics principle for belief revision w.r.t. to conditional information coming from exclusive contexts. In the case of c-revisions for OCFs and the unique setting of the Kinematics principle, our investigations lead to an intuitive solution for the merging problem that concerns rational ways to set up a globally revised state employing the locally revised substates. This problem is known primarily in the context of Bayesian networks [88]. For the qualitative framework, our main achievements lie in defining a well-behaved conditionalization operator w.r.t. to OCF conditionalization, transforming c-revisions from ranking functions, and therefore providing a qualitative revision operator for sets of conditionals. These concepts are meaningful extensions to the current state of the art in Belief Revision.

Second, we considered belief revision operators that take into account propositional input information that is accompanied by a reference sentence which gave rise to the following research question:

> *How can we incorporate the parameterized information into the framework of (conditional) belief revision so that the relation between input and reference information during the change process is evident?*

The solution presented in this thesis was to incorporate the parameterized reference sentence into conditional information in a way that the relation between the input and reference sentence crucially affects the revision. We showed that the parameterized information is treated adequately by exploiting the expressiveness of conditional beliefs. As a foundation for our investigation, we used two parameterized belief revision operators, whose underlying change mechanisms we clarified, leading to simple yet elegant representation theorems. Ultimately, we extended the research question to a comparison illustrating the versatility of conditional Belief Revision, which can express different notions of parameterized belief revision operators.

Although we answered many questions from different angles about the role of conditionals in Belief Revision and provided a detailed look at the specific aspects we studied, some open questions and possible connections to other research activities remain. We highlight some of them briefly.

Conditionals as Input for Different Belief Revision Operators. Throughout this thesis, we have employed conditionals as input for belief revision and illustrated their expressiveness and adaptability in situations that exceed the classical AGM frameworks. Investigations in that direction can be extended to other generalizations of the classical AGM framework, such as non-prioritized Belief Revision [51, 33], i.e., revision operators that do not guarantee the acceptance of the input information. In particular, our work on parameterized Belief Revision provides ground for further investigations in this direction. Revision by Comparison (RbC) presented in Chapter 8 displays a non-prioritized belief revision operator, and its distinctive features put it in the vicinity of credibility-limited belief revision operators presented in [54, 11]. We believe that it is a fruitful path of research to discuss how conditionals as input information could be used to decide which information should be accepted in the posterior state and which should instead be discarded.

Notions of Locality in Belief Revision. The Kinematics principle allows us to focus on local cases when revising a belief state with sets of conditionals. Thus the Kinematics principle introduces a semantic-based notion of locality which employs conditionalization of epistemic states. In Belief Revision, a similar notion of locality based on syntax exists, called *Syntax Splitting*. Starting from Parikh's Axiom (P) [86] and the works of Peppas et al. [92] on relevance in belief revision, Kern-Isberner and Brewka transferred in [68] the notion of Syntax Splitting to revising TPOs. Briefly, Syntax Splitting aims to capture the intuition that whenever beliefs are revised with a new piece of information A, only those beliefs should be affected that are (syntactically) relevant to A, i.e., definable over the same sub-signature. If we compare this to our Kinematics principle, which implements a "revision by (exhaustive and exclusive) cases" it becomes clear that the Kinematics principle and Syntax Splitting are orthogonal. They both express relevance and independency assertions for belief revision but with different perspectives: Syntax splitting splits the set of worlds *vertically* according to sub-signatures, while the Kinematics principle splits the worlds *horizontally* according to cases. Marginalization is the primary tool for syntax splitting, causing vertical splitting, while the Kinematics principle uses conditionalization which causes horizontal splitting according to cases. As we can see, both Syntax Splitting and the Kinematics principles allow for reducing the semantic space of models to relevant parts of syntax and semantics. It is an open yet promising task to develop an approach for combining these notions of locality.

Synergies Between Strengths of Conditionals. While this thesis's first part was dedicated to investigating a powerful axiom for conditional revision, the second part deals with the advantages of employing conditionals in non-classical revision scenarios. Now that we have characterized different notions of parameterized belief revision as conditional revision, the question arises of whether we can employ the

Kinematics principle in this setting.

It holds that the conditional revision for RbC, presented in Chapter 8.3 employs a set of weak conditionals, i.e., we need to define an extension of the Kinematics principle for weak conditionals. Luckily, this is straightforward, and due to the flexibility of c-revision, this Kinematics principle for weak conditionals is likely to be also satisfied by weak c-revisions. Next, we need to define a case splitting for the weak conditionals in the corresponding RbC base $\Delta^{\circledcirc}_{\alpha,\beta}$. Here, a skillful redefinition of the premises of conditionals in $\Delta^{\circledcirc}_{\alpha,\beta}$ w.r.t. to conditional dependencies or a generalization of the Kinematics principle with weakened prerequisites is necessary. Both of these approaches are promising lines of research. For Bounded Revision, c-revision with a single conditional is sufficient to characterize the operation. Thus, we can employ our results on c-revisions with a single conditional from Section 6.5.1 to simplify the revision task. Generally, i.e., apart from the parameterized revision operator presented in this thesis, linking belief revision w.r.t. to additional information to conditional revision is promising because conditionals (naturally) provide a specific context in which the new input information comes into play.

We close with a final remark on a conditional perspective for Belief Revision. Conditionals and the adequate treatment of conditional beliefs are an integral part of most belief change processes due to their fundamental role as subtle yet powerful guidelines for the cognitive state of an agent in the light of new information. The propositions and theorems in this thesis as well as the discussed further research questions, endorse that the close examination and rigorous employment of conditionals is and will be a relevant instrument for the theory of Belief Revision.

Bibliography

[1] E. Adams. The logic of conditionals. *Inquiry: An Interdisciplinary Journal of Philosophy*, 8(1-4):166–197, 1965.

[2] C. E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *J. Symb. Log.*, 50(2):510–530, 1985.

[3] Y. Ammar and H. O. Ismail. Trust is all you need: From belief revision to information revision. In W. Faber, G. Friedrich, M. Gebser, and M. Morak, editors, *Logics in Artificial Intelligence - 17th European Conference, JELIA 2021, Virtual Event, May 17-20, 2021, Proceedings*, volume 12678 of *Lecture Notes in Computer Science*, pages 50–65. Springer, 2021.

[4] N. B. Amor, Z. E. Khalfi, H. Fargier, and R. Sabbadin. Lexicographic refinements in stationary possibilistic markov decision processes. *Int. J. Approx. Reason.*, 103:343–363, 2018.

[5] C. Beierle, C. Eichhorn, and G. Kern-Isberner. A transformation system for unique minimal normal forms of conditional knowledge bases. In A. Antonucci, L. Cholvy, and O. Papini, editors, *Symbolic and Quantitative Approaches to Reasoning with Uncertainty - 14th European Conference, ECSQARU 2017, Lugano, Switzerland, July 10-14, 2017, Proceedings*, volume 10369 of *Lecture Notes in Computer Science*, pages 236–245, Berlin, Heidelberg, 2017. Springer.

[6] C. Beierle, C. Eichhorn, G. Kern-Isberner, and S. Kutsch. Properties of skeptical c-inference for conditional knowledge bases and its realization as a constraint satisfaction problem. *Ann. Math. Artif. Intell.*, 83(3-4):247–275, 2018.

[7] C. Beierle, C. Eichhorn, G. Kern-Isberner, and S. Kutsch. Properties and interrelationships of skeptical, weakly skeptical, and credulous inference induced by classes of minimal models. *Artificial Intelligence*, 297:103489, Aug. 2021.

[8] C. Beierle and G. Kern-Isberner. *Methoden wissensbasierter Systeme - Grundlagen, Algorithmen, Anwendungen, 5. Auflage.* Computational intelligence. SpringerVieweg, 2014.

[9] C. Beierle and G. Kern-Isberner. Selection strategies for inductive reasoning from conditional belief bases and for belief change respecting the principle of conditional preservation. In E. Bell and F. Keshtkar, editors, *Proceedings of the 34th International Florida Artificial Intelligence Research Society Conference (FLAIRS-34)*, pages 563–568, 2021.

[10] S. Benferhat, D. Dubois, H. Prade, and M.-A. Williams. A framework for iterated belief revision using possibilistic counterparts to jeffrey's rule. *Fundam. Inform.*, 99:147–168, 2010.

[11] R. Booth, E. Fermé, S. Konieczny, and R. Pino Pérez. Credibility-limited revision operators in propositional logic. *13th International Conference on the Principles of Knowledge Representation and Reasoning, KR 2012*, 01 2012.

[12] R. Booth and A. Hunter. Trust as a precursor to belief revision. *J. Artif. Intell. Res.*, 61:699–722, 2018.

[13] R. Booth and T. A. Meyer. Admissible and restrained revision. *J. Artif. Intell. Res.*, 26:127–151, 2006.

[14] R. Booth and J. Paris. A note on the rational closure of knowledge bases with both positive and negative knowledge. *Journal of Logic, Language and Information*, 7(2):165–190, 1998.

[15] C. Boutilier. Revision sequences and nested conditionals. In R. Bajcsy, editor, *Proceedings of the 13th International Joint Conference on Artificial Intelligence. Chambéry, France, August 28 - September 3, 1993*, pages 519–525. Morgan Kaufmann, 1993.

[16] C. Boutilier. Unifying default reasoning and belief revision in a modal framework. *Artif. Intell.*, 68(1):33–85, 1994.

[17] C. Boutilier. Iterated revision and minimal change of conditional beliefs. *J. Philos. Log.*, 25(3):263–305, 1996.

[18] C. Boutilier and M. Goldszmidt. Revision by conditional beliefs. In R. Fikes and W. G. Lehnert, editors, *Proceedings of the 11th National Conference on Artificial Intelligence. Washington, DC, USA, July 11-15, 1993*, pages 649–654. AAAI Press / The MIT Press, 1993.

[19] R. J. Brachman and H. J. Levesque. *Knowledge Representation and Reasoning.* Elsevier, 2004.

[20] M. E. Bratman, D. J. Israel, and M. E. Pollack. Plans and resource-bounded practical reasoning. *Comput. Intell.*, 4:349–355, 1988.

[21] G. Brewka. Nichtmonotones schließen. In G. Görz, U. Schmid, and T. Braun, editors, *Handbuch der Künstlichen Intelligenz, 6. Auflage*, pages 189–226. De Gruyter, 2020.

[22] R. M. Byrne and P. Johnson-Laird. 'if' and the problems of conditional reasoning. *Trends in Cognitive Sciences*, 13(7):282–287, 2009.

[23] P. G. Calabrese. A theory of conditional information with applications. *IEEE Trans. Syst. Man Cybern. Syst.*, 24(12):1676–1684, 1994.

[24] T. Caridroit, S. Konieczny, and P. Marquis. Contraction in propositional logic. *Int. J. Approx. Reason.*, 80:428–442, 2017.

[25] H. Chan and A. Darwiche. On the revision of probabilistic beliefs using uncertain evidence. *Artif. Intell.*, 163(1):67–90, 2005.

[26] J. Chandler and R. Booth. Elementary iterated revision and the levi identity. *CoRR*, abs/1907.01224, 2019.

[27] J. Chandler and R. Booth. Revision by Conditionals: From Hook to Arrow. In *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning*, pages 233–242, 9 2020.

[28] S. Chopra, R. Parikh, and R. Wassermann. Approximate belief revision. *Logic Journal of the IGPL*, 9(6):755–768, 2001.

[29] A. Darwiche and J. Pearl. On the logic of iterated belief revision. *Artif. Intell.*, 89(1-2):1–29, 1997.

[30] B. De Finetti. *Theory of Probability: A critical introductory treatment*. Wiley Series in Probability and Statistics. Wiley, 1975.

[31] P. Egré and H. Rott. The Logic of Conditionals. In E. N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Winter 2021 edition, 2021.

[32] C. Eichhorn, G. Kern-Isberner, and M. Ragni. Rational inference patterns based on conditional logic. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18)*, pages 1827–1834. AAAI Press, 2018.

[33] E. L. Fermé and S. O. Hansson. *Belief Change - Introduction and Overview*. Springer Briefs in Intelligent Systems. Springer, 2018.

[34] E. L. Fermé and H. Rott. Revision by comparison. *Artif. Intell.*, 157(1-2):5–47, 2004.

[35] M. Freund. Preferential Orders and Plausibility Measures. *Journal of Logic and Computation*, 8(2):147–158, 04 1998.

[36] M. Freund and D. Lehmann. Belief revision and rational inference. *CoRR*, cs.AI/0204032, 2002.

[37] N. Friedman and J. Y. Halpern. Plausibility measures and default reasoning. *CoRR*, cs.AI/9808007, 1998.

[38] D. M. Gabbay. Theoretical foundations for non-monotonic reasoning part 2: Structured non-monotonic theories. In *Third Scandinavian Conference on Artificial Intelligence, SCAI 1991*, volume 12 of *Frontiers in Artificial Intelligence and Applications*, pages 19–39. IOS Press, 1991.

[39] P. Gärdenfors and D. Makinson. Revisions of knowledge systems using epistemic entrenchment. In M. Y. Vardi, editor, *Proceedings of the 2nd Conference on Theoretical Aspects of Reasoning about Knowledge, Pacific Grove, CA, USA, March 1988*, pages 83–95. Morgan Kaufmann, 1988.

[40] M. Goldszmidt and J. Pearl. Qualitative probabilities for default reasoning, belief revision, and causal modeling. *Artificial Intelligence*, 84(1):57 – 112, 1996.

[41] I. R. Goodman, R. P. S. Mahler, and H. T. Nguyen. *Three Particular Conditional Event Algebras*. Springer Netherlands, Dordrecht, 1997.

[42] A. Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17(2):157–170, 1988.

[43] P. Gärdenfors. Epistemic importance and minimal changes of belief. *Australasian Journal of Philosophy*, 62(2):136–157, 1984.

[44] P. Gärdenfors. Belief revisions and the ramsey test for conditionals. *The Philosophical Review*, 95(1):81–93, 1986.

[45] P. Gärdenfors. *Knowledge in flux : modeling the dynamics of epistemic states.* MIT Press Cambridge, 1988.

[46] P. Gärdenfors and D. Makinson. Nonmonotonic inference based on expectations. *Artificial Intelligence*, 65(2):197–245, 1994.

[47] J. Haldimann, C. Beierle, and G. Kern-Isberner. Syntax splitting for iterated contractions, ignorations, and revisions on ranking functions using selection strategies. In W. Faber, G. Friedrich, M. Gebser, and M. Morak, editors, *Logics in Artificial Intelligence - 17th European Conference, JELIA 2021,*

Virtual Event, May 17-20, 2021, Proceedings, volume 12678 of *LNCS*, pages 85–100. Springer, 2021.

[48] J. Halpern. *Reasoning about Uncertainty.* MIT Press, Cambridge, MA, 2003.

[49] S. O. Hansson. A dyadic representation of belief. *Belief revision*, 29:89–121, 1992.

[50] S. O. Hansson. In defense of the ramsey test. *The Journal of Philosophy*, 89(10):522–540, 1992.

[51] S. O. Hansson. What's new isn't always best. *Theoria*, 63(1-2):1–13, 1997.

[52] S. O. Hansson. *A textbook of belief dynamics - theory change and database updating*, volume 11 of *Applied logic series.* Kluwer, 1999.

[53] S. O. Hansson. Multiple and iterated contraction reduced to single-step single-sentence contraction. *Synth.*, 173(2):153–177, 2010.

[54] S. O. Hansson, E. L. Fermé, J. Cantwell, and M. A. Falappa. Credibility limited revision. *The Journal of Symbolic Logic*, 66(4):1581–1596, 2001.

[55] J. Heyninck, G. Kern-Isberner, and T. A. Meyer. Lexicographic entailment, syntax splitting and the drowning problem. In L. D. Raedt, editor, *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, pages 2662–2668. ijcai.org, 2022.

[56] A. Hunter. Building trust for belief revision. In D. N. Pham, T. Theeramunkong, G. Governatori, and F. Liu, editors, *PRICAI 2021: Trends in Artificial Intelligence - 18th Pacific Rim International Conference on Artificial Intelligence, PRICAI 2021, Hanoi, Vietnam, November 8-12, 2021, Proceedings, Part I*, volume 13031 of *Lecture Notes in Computer Science*, pages 543–555. Springer, 2021.

[57] R. C. Jeffrey. *The Logic of Decision.* University of Chicago Press, 1965.

[58] D. Kahneman. *Thinking, Fast and Slow.* Farrar, Straus and Giroux, New York, 2011.

[59] H. Katsuno and A. O. Mendelzon. A unified view of propositional knowledge base updates. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence - Volume 2*, IJCAI'89, page 1413–1419, San Francisco, CA, USA, 1989. Morgan Kaufmann Publishers Inc.

[60] H. Katsuno and A. O. Mendelzon. On the difference between updating a knowledge base and revising it. In J. F. Allen, R. Fikes, and E. Sandewall, editors, *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning (KR'91). Cambridge, MA, USA, April 22-25, 1991*, pages 387–394. Morgan Kaufmann, 1991.

[61] J. Kemeny and J. L. Snell. *Mathematical Models in the Social Sciences.* Introductions to higher mathematics. Ginn, 1962.

[62] G. Kern-Isberner. Postulates for conditional belief revision. In *Proceedings Sixteenth International Joint Conference on Artificial Intelligence, IJCAI-99*, pages 186–191. Morgan Kaufmann, 1999.

[63] G. Kern-Isberner. *Conditionals in Nonmonotonic Reasoning and Belief Revision - Considering Conditionals as Agents*, volume 2087 of *Lecture Notes in Computer Science*. Springer, Berlin, 2001.

[64] G. Kern-Isberner. A thorough axiomatization of a principle of conditional preservation in belief revision. *Annals of Mathematics and Artificial Intelligence*, 40(1-2):127–164, 2004.

[65] G. Kern-Isberner. Axiomatizing a qualitative principle of conditional preservation for iterated belief change. In M. Thielscher, F. Toni, and F. Wolter, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference, KR 2018, Tempe, Arizona, 30 October - 2 November 2018*, pages 248–256. AAAI Press, 2018.

[66] G. Kern-Isberner, C. Beierle, and G. Brewka. Syntax splitting = relevance + independence: New postulates for nonmonotonic reasoning from conditional

belief bases. In D. Calvanese, E. Erdem, and M. Thielscher, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the 17th International Conference, KR 2020*, pages 560–571. IJCAI Organization, 2020.

[67] G. Kern-Isberner, T. Bock, K. Sauerwald, and C. Beierle. Iterated contraction of propositions and conditionals under the principle of conditional preservation. In C. Benzm\"uller, C. Lisetti, and M. Theobald, editors, *GCAI 2017. 3rd Global Conference on Artificial Intelligence*, volume 50 of *EPiC Series in Computing*, pages 78–92. EasyChair, 2017.

[68] G. Kern-Isberner and G. Brewka. Strong syntax splitting for iterated belief revision. In C. Sierra, editor, *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, pages 1131–1137. ijcai.org, 2017.

[69] G. Kern-Isberner, J. Heyninck, and C. Beierle. Conditional independence for iterated belief revision. In L. D. Raedt, editor, *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, pages 2690–2696. ijcai.org, 2022.

[70] G. Kern-Isberner, M. Sezgin, and C. Beierle. A kinematics principle for iterated revision. *Artif. Intell.*, 314:103827, 2023.

[71] S. Konieczny and R. P. Pérez. On iterated contraction: Syntactic characterization, representation theorem and limitations of the levi identity. In S. Moral, O. Pivert, D. Sánchez, and N. Marín, editors, *Scalable Uncertainty Management - 11th International Conference, SUM 2017, Granada, Spain, October 4-6, 2017, Proceedings*, volume 10564 of *Lecture Notes in Computer Science*, pages 348–362. Springer, 2017.

[72] S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44(1):167–207, 1990.

[73] D. Lehmann and M. Magidor. What does a conditional knowledge base entail? *Artificial Intelligence*, 55(1):1–60, 1992.

[74] I. Levi. Subjunctives, dispositions and chances. *Synthese*, 34(4):423–455, 1977.

[75] I. Levi. Contracting from epistemic hell is routine. *Synthese*, 135(1):141–164, 2003.

[76] C. I. Lewis. Implication and the algebra of logic. *Mind*, 21(84):522–531, 1912.

[77] D. K. Lewis. *Counterfactuals*. Cambridge, MA, USA: Blackwell, 1973.

[78] D. Makinson. General theory of cumulative inference. In *Non-Monotonic Reasoning*, pages 1–18, Berlin, Heidelberg, 1989. Springer Berlin Heidelberg.

[79] A. C. Nayak. Iterated belief change based on epistemic entrenchment. *Erkenntnis (1975-)*, 41(3):353–390, 1994.

[80] A. C. Nayak, R. Goebel, M. A. Orgun, and T. Pham. Taking levi identity seriously: A plea for iterated belief contraction. In J. Lang, F. Lin, and J. Wang, editors, *Knowledge Science, Engineering and Management, First International Conference, KSEM 2006, Guilin, China, August 5-8, 2006, Proceedings*, volume 4092 of *Lecture Notes in Computer Science*, pages 305–317. Springer, 2006.

[81] A. C. Nayak, M. Pagnucco, N. Y. Foo, and P. Peppas. Learning from conditionals: Judy benjamin's other problems. In W. Wahlster, editor, *12th European Conference on Artificial Intelligence, Budapest, Hungary, August 11-16, 1996, Proceedings*, pages 75–79. John Wiley and Sons, Chichester, 1996.

[82] A. C. Nayak, M. Pagnucco, and P. Peppas. Dynamic belief revision operators. *Artificial Intelligence*, 146(2):193–228, 2003.

[83] D. Nute. *Topics in Conditional Logic*. Philosophical Studies Series. Springer Dordrecht, 1980.

[84] D. Nute. Conditional logic. *Journal of Symbolic Logic*, 54(4):1477–1479, 1989.

[85] M. Oaksford and N. Chater. *Bayesian Rationality: The Probabilistic Approach to Human Reasoning*. Bayesian Rationality: The Probabilistic Approach to Human Reasoning. OUP Oxford, 2007.

[86] R. Parikh. Beliefs, belief revision, and splitting languages. *Logic, Language, and Computation*, 2:266–278, 1999.

[87] J. B. Paris. *The Uncertain Reasoner's Companion: A Mathematical Perspective*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 1995.

[88] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann series in representation and reasoning. Elsevier Science, 1988.

[89] J. Pearl. System Z: A natural ordering of defaults with tractable applications to nonmonotonic reasoning. In R. Parikh, editor, *Proceedings of the 3rd Conference on Theoretical Aspects of Reasoning about Knowledge, Pacific Grove, CA, USA, March 1990*, pages 121–135, 1990.

[90] P. Peppas. Chapter 8 belief revision. In F. van Harmelen, V. Lifschitz, and B. Porter, editors, *Handbook of Knowledge Representation*, volume 3 of *Foundations of Artificial Intelligence*, pages 317–359. Elsevier, 2008.

[91] P. Peppas and M.-A. Williams. Constructive Modelings for Theory Change. *Notre Dame Journal of Formal Logic*, 36(1):120 – 133, 1995.

[92] P. Peppas, M.-A. Williams, S. Chopra, and N. Y. Foo. Relevance in belief revision. *Artificial Intelligence*, 229((1-2)):126–138, 2015.

[93] C. Pornpitakpan. The persuasiveness of source credibility: A critical review of five decades' evidence. *Journal of Applied Social Psychology*, 34:243 – 281, 02 2004.

[94] R. Ramachandran, A. C. Nayak, and M. A. Orgun. Three approaches to iterated belief contraction. *J. Philos. Log.*, 41(1):115–142, 2012.

[95] F. P. Ramsey. General propositions and causality. In F. P. Ramsey, editor, *The Foundations of Mathematics and other Logical Essays*, pages 237–255. Kegan Paul, Trench, Trübner, 1929.

[96] H. Rott. *On the logic of theory change: More maps between different kinds of contraction functions*, page 122–141. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 1992.

[97] H. Rott. Preferential belief change using generalized epistemic entrenchment. *Journal of Logic, Language, and Information*, 1(1):45–78, 1992.

[98] H. Rott. Coherence and Conservatism in the Dynamics of Belief II: Iterated Belief Change without Dispositional Coherence. *Journal of Logic and Computation*, 13(1):111–145, 02 2003.

[99] H. Rott. Two-Dimensional Belief Change. In G. Bonanno, J. Delgrande, J. Lang, and H. Rott, editors, *Formal Models of Belief Change in Rational Agents*, volume 7351 of *Dagstuhl Seminar Proceedings (DagSemProc)*, pages 1–27, Dagstuhl, Germany, 2007. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[100] H. Rott. Bounded revision: Two-dimensional belief change between conservative and moderate revision. *J. Philos. Log.*, 41(1):173–200, 2012.

[101] S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach (4th Edition).* Pearson, 2020.

[102] W. Rödder. Conditional logic and the principle of entropy. *Artificial Intelligence*, 117(1):83–106, 2000.

[103] K. Sauerwald, G. Kern-Isberner, and C. Beierle. A conditional perspective for iterated belief contraction. In *Proceedings of the 24th European Conference on Artificial Intelligence, ECAI 2020*, volume 325 of *Frontiers in Artificial Intelligence and Applications*, pages 889–896. IOS Press, 2020.

[104] F. Schwartz, J. Epinat-Duclos, J. Léone, and J. Prado. The neural development of conditional reasoning in children: Different mechanisms for assessing the logical validity and likelihood of conclusions. *NeuroImage*, 163:264–275, 2017.

[105] N. Schwind, S. Konieczny, and P. Marquis. On belief promotion. In M. Thielscher, F. Toni, and F. Wolter, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference, KR 2018, Tempe, Arizona, 30 October - 2 November 2018*, pages 297–307. AAAI Press, 2018.

[106] N. Schwind, S. Konieczny, and R. P. Pérez. On the representation of darwiche and pearl's epistemic states for iterated belief revision. In G. Kern-Isberner, G. Lakemeyer, and T. Meyer, editors, *Proceedings of the 19th International Conference on Principles of Knowledge Representation and Reasoning, KR 2022, Haifa, Israel. July 31 - August 5, 2022*, 2022.

[107] M. Sezgin and G. Kern-Isberner. Generalized ranking kinematics for iterated belief revision. In R. Barták and E. Bell, editors, *Proceedings of the Thirty-Third International Florida Artificial Intelligence Research Society Conference, Originally to be held in North Miami Beach, Florida, USA, May 17-20, 2020*, pages 587–592. AAAI Press, 2020.

[108] M. Sezgin and G. Kern-Isberner. System Z for conditional belief bases with positive and negative information. In E. Bell and F. Keshtkar, editors, *Proceedings of the Thirty-Fourth International Florida Artificial Intelligence Research Society Conference, North Miami Beach, Florida, USA, May 17-19, 2021*, 2021.

[109] M. Sezgin and G. Kern-Isberner. Revision by comparison for ranking functions. In L. D. Raedt, editor, *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, pages 2734–2740. ijcai.org, 2022.

[110] M. Sezgin and G. Kern-Isberner. Implementing bounded revision via lexicographic revision and c-revision. In *The Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, EAAI 2023, Washington, D.C., USA, February 7-12, 2023*, pages 2822–2829. AAAI Press, 2023.

[111] M. Sezgin, G. Kern-Isberner, and C. Beierle. Ranking kinematics for revising by contextual information. *Ann. Math. Artif. Intell.*, 89(10-11):1101–1131, 2021.

[112] M. Sezgin, G. Kern-Isberner, and H. Rott. Inductive reasoning with difference-making conditionals. In *18th International Workshop on Non-Monotonic Reasoning (NMR 2020 Workshop Notes)*, pages 83–92, 2020.

[113] J. E. Shore and R. W. Johnson. Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy. *IEEE Transactions on Information Theory*, 26(1):26–37, 1980.

[114] B. Skyrms. *The dynamics of rational deliberation / Brian Skyrms.* Harvard University Press Cambridge, Mass, 1990.

[115] P. Smets. Jeffrey's rule of conditioning generalized to belief functions. *CoRR*, abs/1303.1514, 2013.

[116] J. Sparks and D. Rapp. Unreliable and anomalous: How the credibility of data affects belief revision. In *Proceedings of the Annual Meeting of the Cognitive Science Society, 33*, pages 741–746, 07 2011.

[117] W. Spohn. A general non-probabilistic theory of inductive reasoning. In R. D. Shachter, T. S. Levitt, L. N. Kanal, and J. F. Lemmer, editors, *UAI '88: Proceedings of the Fourth Annual Conference on Uncertainty in Artificial Intelligence, Minneapolis, MN, USA, July 10-12, 1988*, pages 149–158. North-Holland, 1988.

[118] W. Spohn. Measuring ranks via the complete laws of iterated contraction. In G. Bonanno, J. P. Delgrande, J. Lang, and H. Rott, editors, *Formal Models of Belief Change in Rational Agents, 26.08. - 30.08.2007*, volume 07351 of *Dagstuhl Seminar Proceedings*. Internationales Begegnungs- und Forschungszentrum fuer Informatik (IBFI), Schloss Dagstuhl, Germany, 2007.

[119] W. Spohn. *The Laws of Belief - Ranking Theory and Its Philosophical Applications.* Oxford University Press, 2014.

[120] R. Stalnaker. A theory of conditionals. In N. Rescher, editor, *Studies in Logical Theory (American Philosophical Quarterly Monographs 2)*, pages 98–112. Oxford: Blackwell, 1968.

[121] R. Stalnaker. Iterated belief revision. *Erkenntnis (1975-)*, 70(2):189–209, 2009.

[122] P. Teller. *Conditionalization, Observation, and Change of Preference*. Springer Netherlands, Dordrecht, 1976.

[123] C. G. Wagner. Generalized probability kinematics. *Erkenntnis*, 36(2):245–257, Mar 1992.

[124] Y. Wang and M. P. Singh. Formal trust model for multiagent systems. In M. M. Veloso, editor, *IJCAI 2007, Proceedings of the 20th International Joint Conference on Artificial Intelligence, Hyderabad, India, January 6-12, 2007*, pages 1551–1556, 2007.

[125] A. G. Wolf, S. Rieger, and M. Knauff. The effects of source trustworthiness and inference type on human belief revision. *Thinking & Reasoning*, 18(4):417–440, 2012.